AN ANALYSIS OF THE PROBLEMS OF R&D FINANCE FOR VACCINES

AND AN APPRAISAL OF ADVANCE PURCHASE COMMITMENTS

Andrew Farlow
University of Oxford
Department of Economics, and Oriel College

This is based on an April 2004 paper with some fairly minor changes.
Copies of that older version are available from the author on request.

This paper is written on the presumption that there is an ongoing debate about how to derive efficient, cost effective, mechanisms to develop the highest possible quality of vaccines in the shortest possible time. All comments are very gratefully received.

andrew.farlow@economics.ox.ac.uk

© Andrew Farlow 2004, 2005

1. INTRODUCTION

1.1. An Overview of the Problem

Infectious and parasitic diseases kill 40,000 people every day, mostly in developing countries – many of them young people and children. Every year hundreds of millions experience the debilitating consequences of these diseases, and the trauma of the suffering and loss of family and friends. Malaria, tuberculosis, and African strains of AIDS alone kill almost 5 million each year. As the AIDS pandemic worsens – especially in Russia, China, and India, but in many other places too – these figures are set to rise even higher. There is a desperate need for more research and development of vaccines and of drugs to treat a wide variety of neglected diseases.

Table 1 shows figures for those diseases that almost exclusively hit the poorest of countries[1].

Table 1: Diseases for Which 99% or More of the Global Burden Fell on Low- and Middle-Income Countries in 1999[2]

	Disability Adjusted Life Years (millions 2000)	Deaths per Year (2000)
Chagas disease	0.68	21,299
Dengue	0.433	12,037
Ancylostomiasis and necatoriasis (hookworm)	1.829	5,650
Japanese encephalitis	0.426	3,502
Lymphatic filariasis	5.549	404
Malaria	40.213	1,079,877
Onchocerciasis (river blindness)	0.951	-
Schistosomiasis	1.713	11,473
Tetanus	9.766	308,662
Trachoma	1.181	14
Trichuriasis	1.64	2,123
Trypanosomiasis	1.585	49,668
Leischmaniasis	1.81	40,913
Measles	27.549	776,626
Poliomyelitis	0.184	675
Syphilis	5.574	196,533
Diphtheria	0.114	3,394
Leprosy	0.141	2,268
Pertussis	12.768	296,099
Diarrhoeal diseases	62.227	2,124,032
TOTAL	176.333	4,935,249

One of the options currently being considered as a mechanism for creating finance for the research and development of vaccines – and increasingly being argued by some as

[1] HIV/AIDS and TB and many other diseases do not appear since they do not meet the 99% cut-off.

[2] Source: Kremmer 2003, based on Global Burden from WHO (1996), quoted in Lanjouw and Cockburn (2001, Table 1). Figures updated from Lanjouw and Cockburn (2001), using WHO (2001).

a means for tackling drugs and techniques for fighting disease in general[3] – is that of an Advanced Purchase Commitment, a commitment to purchase specified 'technologies' in specified 'quantities' in the 'future' at a 'guaranteed' unit 'price' (readers might like to read footnote 3, since it explains how papers will be references throughout this paper). The words 'technologies', 'quantities', 'prices', and 'future' are in quotation marks since an APC turns out *not* to refer to a *unique* technology, quantity, price, or point in time for payment, nor a unique number of sellers. Rather, it refers to a set of 'rules' that determine *sets* of quantities and prices, periods of time over which payments *may* be made, who the buyers and sellers will be, and layers of institutions to administer and police the 'rules'. The word 'rules' is in quotation marks since the rules turn out to contain potentially important degrees of discretion. And this is why the word 'guaranteed' is in quotation marks.

This paper analyses APCs principally for early-stage vaccines, but aims to do so in the context of neglected diseases and R&D issues in general. 'Early-stage' vaccines – such as those for HIV, malaria, and TB – are those for which either no viable vaccines are on the horizon or the current candidates fall well short of 100% effectiveness, and many of the scientific difficulties have yet to be resolved. This is a very different situation from already existing but underused vaccines and 'late-stage' vaccines, where a viable vaccine is close to development.

1.2. Defining an Advance Purchase Commitment, or Precommitment, for Stimulating Vaccine R&D

The following section has been taken from Farlow 2005 Section 1.4, since some readers explained that a description of the object of interest would help early on in the discussion. Readers who have read that, can skip this section.

The phrase 'advance purchase commitments' has come to have varying degrees of strictness in both interpretation and application. At one extreme it has been interpreted as just a generalized notion of 'willingness to pay' for vaccines. However, at the other extreme, there is a benchmark for when such devices are used to stimulate private R&D, and it is worth setting that out exactly, so that we can compare and contrast *that* with real-world enactments. It is sobering to think that we have never had an advance purchase commitment meeting conditions even remotely approaching the benchmark criteria for even the most simple of drug or vaccine cases. And recent policy pronouncements for early-stage vaccines (malaria and HIV in particular) do not begin to approach the benchmark either. How far they fall short, and the implications of this for vaccine development, is an interesting policy issue in its own right[4].

[3] K7:39. This refers to Michael Kremer, Appendix 7, page 39 on the No. 10 Policy Unit website, and is the way all citations will be made below. All papers to be found at http://www.pm.gov.uk/output/Page3704.asp. The quote at K7:39 is: "Potentially, advance purchase commitments could be used to encourage research not only on vaccines, but also on other techniques for fighting disease, including drugs, diagnostic devices, and insecticides against mosquitoes which transmit malaria." Similarly, K7:40: "In principle, purchase commitments are appropriate for both drugs and vaccines, but if a choice had to be made for budgetary reasons, vaccines are probably a slightly higher priority, since distortions in vaccine markets are more severe."

[4] The worst case is when they promise the level of payments supposedly based on an application of the benchmark idealised model (i.e. a large 'pot') but then don't actually enact any of the rest of the framework (though this paper argues that they could not, in all likelihood, enact much of the theoretical framework even if they wanted to).

Advance purchase commitments for vaccines are *legally binding* contracts (on only the funders in the case of early-stage vaccines[5]) that *commit for ever* a sum of money for the purchase of a vaccine or vaccines for a particular disease. According to the literature, this would be anything in the region of $3bn-$10bn (when this was first typed, the bottom of the range was $6bn; but it has kept falling and is now $3bn) per major early-stage vaccine, though the eventual sum is not clear and could be a great deal higher. Pitching to the lower end of this range (indeed pushing the lower range ever lower) has become popular just recently, but we will later see that this is very damaging behaviour if the true requirement is much higher. This is not the whole cost of developing a vaccine. The overall cost includes all public funding needed outside of the mechanism in order to make it work, as well as subsidies, tax-breaks, and other benefits private firms are granted for their research (to the extent that a large *multiple* of these is not removed later from payments, as explained below).

The size of the fund (and its distribution over developers) must be set *precisely* high enough to re-create the *precise* size of *additional* 'blockbuster' market needed to encourage the entry of the *precise* amount of venture capital and stock market finance needed for the *remaining* research and development needed to produce a 'high quality' vaccine or series of vaccines (that will be needed over time, especially in the case of malaria and HIV). This finance would then be fully repaid through the purchase of a *successful* vaccine or several vaccines in a particular period in time (if there are several meeting eligibility conditions in any one period of time), or series of vaccines over time (to combat resistance), and *only* the successful vaccine/s or series of vaccines. Payment would come from the taxpayers of richer countries, by foundations such as the Bill and Melinda Gates Foundation, and through co-payments made by developing countries tied, in advance, to the mechanism.

Observe the multiple directions for decisions about eligible vaccines – across vaccines at a given time and across vaccines over time – with all expected decision rules set in the terms of the 'contract' at the start. In order to overcome any risks (as perceived by developers) that buyers will bid prices down after development, the funds are legally committed *in advance* to pay for those (and *only* those) vaccines generated in response to the mechanism on the basis of the pre-agreed rules. This is important, since one of the key justifications of the mechanism is to solve the 'time inconsistency' problem, that describes what happens when firms have sunk R&D, and then buyers have the power to bid prices down to levels that do not fully cover those collective R&D costs, and, knowing this in advance, no individual firm will perform R&D in the first place. We will see that 'time inconsistency' continues to be an extremely difficult issue to get around under an advance purchase precommitment. Indeed, it turns out to be intractable whatever the mechanism used to stimulate early-stage vaccine R&D, but especially those mechanisms concentrating on payment in the end period. The more complex the science, the greater the ex post discretion, the greater the time inconsistency. Time inconsistency can only be removed by stripping out all hints of scientific complexity (as is done in Kremer Appendix 3), but that is just a convenient 'fix'.

[5] Though, funders also have an opt-out if the contracts fail to stimulate '*enough*' research.

What the 'winners' get

The 'winning' vaccine developer or developers would be paid the value of *all* the privately-funded (and *only* the privately-funded) R&D costs (including *all* capital costs too) of *all* firms (both the successful and the unsuccessful, not just of itself) and *only* the private firms, who used such private funding on R&D towards the vaccine since the time the purchase commitment had been announced (and *only* since the announcement) and *only* for markets covered by the mechanism. Throughout this article, 'capital costs' refers exclusively to the costs of the finance used, and includes the required return to cover *all* risk being borne, including the potentially high risk created by the mechanism itself (i.e. 'capital cost' does not refer to physical real capital investment but to the costs of finance). The winner gets all the vaccine IP.

A blockbuster-style model

As with the blockbuster drug-development model, an individual firm treats its vaccine R&D as a lottery with a very large 'prize' that just makes it a fair risk-adjusted gamble. Individual firms calculate the expected value to it of the 'prize' on the basis of the *privately-funded* R&D activity of all other private firms. If others, not firm *i*, do more R&D then this will reduce the chance that firm *i* will win the contract and hence the expected value to firm *i* of its investment. 'Others' should refer only to other firms working under this mechanism, and not to any other researchers working under any other mechanism. We will see that this proves fiendishly difficult to achieve in areas of complex science involving an interplay of many different funding mechanisms and a complex mix of public and private researchers (Appendix 3 removes all of this by presuming only one mechanism and only one type of researcher is actually present). Worryingly for firm *i*, 'others' could refer to those being paid for under other R&D mechanisms if these other mechanisms are not factored out of payments (since Kremer Appendix 3 factors other mechanisms out, by default, this issue never arises).

A first look at some *very* vague 'size' figures

To frame the thinking, it might help to have a quick overview of possible scenarios, though we also recognise that insufficient evidence has so far been presented to properly analyze early-stage vaccines, and so the figures are necessarily very rough.

When it 'wins' the contract, it turns out that a firm's ex post out-of-pocket costs (including capital costs) are a tiny fraction of the contract size. For example, if 10 firms put in equal effort on an early-stage HIV vaccine (this paper repeatedly indicates that the contracts in practice would work to benefit just one or two large players, but we can maintain the fiction of competition for now) , and we presume that this is the optimal number of firms (we can't), and that (because of all the risks and because of the high cost form of finance being used) they face an expected 70% of capital costs[6] by the time a product is developed (and we ignore all crowding out for now), and we presume for the moment that only one firm wins (though, in most cases there would, supposedly, be a complicated split over time and across firms), then a $6.25bn 'purchase commitment' will go to a firm having spent, in present discounted (2005) terms, less than $200m, on private out-of-pocket research costs. Incidentally, the response of one pharmaceutical executive when this was spelled out precisely was that winning such a contract for HIV would be a "PR disaster". Throw in some of the discretionary elements (discussed below) that the firm would have to

[6] We can only guess at these figures since none have been calculated. See the discussion below.

ex post very publicly fight over in order to get a fair return in *the ex ante sense*, and it would be a "complete PR disaster", and much worse for such firms than some alternative approaches.

In this case, if there was no 'crowding out' (explained in more detail below), the $6.25bn fund would 'pay for' $1.875bn of out-of-pocket R&D costs across all firms and $4.375bn of capital costs (i.e. the cost of equity finance used). If there is crowding out and other inefficiencies, the ratio of 'payout' to the out-of-pocket private costs could be even more extreme. In this simple case, if there was 50% crowding out, the $6.25bn fund would pay for about $900m of new out-of-pocket research costs, or about 9 months' worth of what those working on the Global HIV vaccine enterprise says is actually needed. The most likely short-run response of firms to such an incentive would be no response at all.

But for HIV it would need a mega-blockbuster precommitment

Indeed, if it really is the case that HIV vaccines might take 15 years to develop and need $1.2bn per year of out of pocket research and trial costs, as those working on the Global HIV Vaccine Enterprise argue, then replacing this $1.2bn per-year flow for 15 years with an advance purchase precommitment at the end of the process, would require an advance purchase precommitment of about $85bn to $130bn (based on real required nominal rates of return of 20% to 25% per year, and if we presume no crowding out at all, and that the uncertainty about ever getting a vaccine is embedded in capital costs).

Maybe this is why private firms spend so little on HIV vaccine research? It is hard to believe that rich markets would not pay $25-$50 or so per course of treatment, generating a multi-billion dollars market there for an HIV vaccine. Maybe that is simply not large enough to cover all the risks faced by developers and the mega-blockbuster price tag they need to justify the risks? Maybe it also has something to do with the problem being more than just creating a single vaccine? Maybe if those lobbying for advance purchase precommitments were to work out the potential size of *any* high-value market for HIV vaccines, and take one look at the pitifully low levels of private vaccine R&D funding for *that* market, they might come to a quite different conclusion to the simple 'lack of a market' argument?

Even simple maths makes a mockery of the notion that an advance purchase precommitment is the thing that "has been so desperately lacking"[7] and that if only we had one in place, all would be well. An HIV vaccine advance purchase precommitment – if that is the route chosen – would have to be a mega-blockbuster, and a great deal higher than anything currently being proposed for it. The best a $3bn (the latest figure apparently) advance purchase precommitment would do in such a situation (with all the crowding out discussed below) would be to allow one big, influential firm, at the end of the whole process to maneuver itself to claim all the IP. Even big firms might, ex ante, prefer some other approach to avoid being put in such a position.

For vaccine purchases of currently existing vaccines, these proportions would, naturally enough, be completely the converse (low capital costs because of low risk,

[7] http://www.cid.harvard.edu/books/kremer04_strongmedicine.html.

no crowding out because of the ability to use competitive tenders, much more easily set terms, etc.).

Specifying vaccine characteristics

Each purchase commitment would try to specify in advance – on the basis of expected science, expected difficulty of development, and costs of production and distribution – the characteristics of a vaccine that would be acceptable for those countries covered by the scheme. In truth, this could not be remotely set in advance for conditions such as malaria, TB and HIV (observe how it is not just the characteristics of the vaccine itself that enter the decision process), and there would have to be a great deal of discretion in the terms set. A contract might, for example, specify 250 million treatments for a malaria vaccine at $25 per course of treatment (making $6.25bn overall[8]), with distribution thereafter to those covered by the mechanism at cost-plus pricing.

There would be one, or supposedly several, big winners of the contract with decisions about winners and losers and allocations made by a committee, based on a mix of rules and discretion. In the literature, this has come to be called an 'Independent Adjudication Committee', or 'IAC'. We use the same nomenclature here, but make no *a priori* presumption about its independence since this is highly unlikely to be the case, or, more importantly, highly unlikely to be *expected* to be the case at horizons of interest.

In the above 'best-case' scenario (of no crowding out, though high capital costs), a vaccine costing $25 for the first 200 to 250 million treatments might compose $1-$2 for production and distribution, $6-$7 for out-of-pocket R&D costs of all firms (not just the winning firm), and $16-$18 for the cost of the finance (again of all firms). With 50% crowding out, only about $3 of the $25 would go towards fresh out-of-pocket R&D costs. Incidentally, it is not at all clear that the first few tens of millions of an HIV vaccine could be manufactured *that* cheaply (especially if there is no competition between manufacturers to drive production prices *that* low). We will discuss this more later (in Section 2.14) when ex ante worries about this can undermine incentives to do R&D in the first place.

Competition, supposedly

Freedom of entry and exit in the R&D process and competition to try to win the $6.25bn contract will, we are told, lead to the 'optimal' number of firms working on vaccine trials and hence the optimal speed of development. However, 'competition' is essentially driven by the expected behaviour of the committee, as well as expectations (and worries) about the behaviour of other firms with respect to the committee. The number of firms in equilibrium is dictated by the initial size of the 'pot' of funds, so that having an optimal number of firms requires that the size of the 'pot' be chosen optimally at the start, which requires knowledge of both the science and likely costs of developing and producing a vaccine. If the 'pot' is too small there will be too few firms and progress will be too slow and chances of discovery low. If it is set 'too large' there will be 'too many' (showing up in overlap, waste, lack of cooperation,

[8] 'Making Markets' p 61 ($20-$25x250 million treatments). Recently this has been trimmed to $15 a treatment and 200 million treatments (i.e. less than half the $6.25bn).

rent seeking, efforts to capture the mechanisms, etc. with some of this showing up in harm to other parts of an overall mechanism).

The notion is that if, for any given 'pot' size, there are *too many* firms 'competing', then the chances of any individual firm winning the pot, or a part of the pot, are too low, the risk-adjusted rewards will be too low, and firms will leave (or they simply will not enter in the first place). But if there are *too few* firms, then the chances of being a winning firm are higher[9] and more firms will enter. In *both* cases, the laws of motion supposedly push in the direction of the optimal number of firms working on research leads[10]. That these laws of motion work, requires huge amounts of assumed competition. If terms could be permanently set in advance, firms would supposedly form their optimal strategies on the basis of their expectations of the strategies of other firms, and never on the behaviour of the distributor of the 'pot'. When terms cannot be known in advance, ex ante competition between vaccine developers is policed via the expected ex post behaviour of the committee (very unlike a standard competitive tendering).

Prices of vaccines to those not covered by the mechanism

Populations *not* covered by the mechanism (say Russians purchasing HIV vaccines for their non-covered program) would continue to pay monopoly pricing, since their market is treated as separate. This is an important feature in the case of an HIV vaccine, but, given the recent evidence of the more widespread nature of malaria, it may also be an increasingly important feature in the case of malaria vaccines too. However, given the presence of the advance contracts in poorer markets, this could mean that the prices faced by those not covered by the mechanism in 'richer' markets would be higher than they would have been without it in place[11].

From now on, this is the benchmark model against which all remarks in this paper will be directed. It will be argued below that advance contracting and commitments of various sorts are useful devices, and that late-stage vaccine work can be helped by contracts that commit public funders to pay for 'performance'. But these have to be very clearly separated in the reader's mind from the notion being suggested (though none of the actual mechanism is laid down) in 'Making Markets' and 'Strong Medicine' for early-stage vaccines which is based on the notion of recreating, from the very start of the process, a precisely sized *additional* blockbuster market, and a precise set of rules (though, still, large elements of discretion), based on the notion that this will drive a large amount of the development costs of vaccines. Clearly, purchasing commitments for currently available and cheap vaccines are a degenerate case of the above mechanism, since most of the features described above have collapsed to zero. Such contracts are not capable of telling us a great deal about the above mechanism.

[9] The individual chance may be low, but given how few other firms there are, if one firm wins, the greater the chance it will be oneself.

[10] In practice, leading advocates have not hidden the fact that a few big companies are seen as driving everything.

[11] See Farlow 2004 Section 7.16. The notion is that control over IP generated by the mechanism for the poorer market, and the market segmentation, strengthens ability to price higher in the 'richer' market.

1.3. The Scope of This Paper and the Story so Far

It should be pointed out at the very start that the issue in this paper is not APCs *per se*, but rather APCs when presented as mechanisms for tackling lengthy, complicated, information-rich, technological processes leading to vaccine and drug development and production – even as the implications of this complex process are largely ignored in their calculations. It has been argued, for example, that "there is a clear rationale for focusing initially on HIV/AIDS, TB and malaria"[12]. Carefully-designed APCs may play a rôle in some circumstances; they are, for example, being analysed with respect to meningococcal A conjugate, pneumococcal conjugate, quadra- or penta-valent DTP-based combinations, and rotavirus. These are not the sort of projects that are the principal target of this paper. A follow-on paper looks at these, and other, cases in much more detail. Rather, in this paper, APCs are analysed with an eye principally to the much longer and more complex processes typical of HIV/AIDS, TB, and malaria, since it is in relation to these diseases that they are being most heavily promoted at the moment. As this paper will show these are the very diseases that APCs struggle most to deal with.

Should an APC be found to work in simpler circumstances, it would not imply that an APC would work in these much more complex circumstances. Similarly, even if it were the case that APCs were only being suggested as devices to deal with the last stage of development of a complex vaccine such as for HIV (which they are not), it rather begs the questions of how problems at earlier stages are to be tackled, and of whether it makes sense to concentrate quite so exclusively on the APC part given that problems with the earlier parts of the process feed into the efficiency and cost of the APC. And it is not even clear that these currently early-stage vaccines are not harmed at their late stages by the presence of APCs at their early stages.

Even with respect to those situations where APCs have a more obvious rôle, many of the points made in this paper nevertheless apply, though with reduced severity. Of particular note are the issues of the ultimate ownership of IPR, quality, market power, the efficient setting of the APC terms in the first place, 'crowding out', and the possible bias towards certain forms of pharmaceutical firms over others. Each of these is a crucially important issue largely ignored in the extant APC literature.

The story so far

A considerable amount of material in support of vaccine APCs for HIV/AIDS, TB and malaria has been generated *for* (but largely not *by*) the World Bank AIDS Vaccine Task Force, the International Aids Vaccine Initiative (IAVA), the UK Government and the UK's Department for International Development, the Dutch and Italian governments, and others. If the UK government's websites are to be believed, APCs appear (at least until recently) to be the central plank of the UK's plan for a 'Global Fund for Health'. In particular, a dozen or more APC papers, of very limited authorship, containing hundreds of pages of text, equations, and figures, can be found on the website of the UK's No.10 Policy Unit, with the seeming support of the British Prime Minister[13] – though without a word of balancing argument, never mind any critical assessment. This paper makes extensive reference to these papers, as

[12] Kremer at http://www.number-10.gov.uk/su/health/annex02/content07.htm.
[13] http://www.number-10.gov.uk/su/health/default.htm.

explained in footnote 3. It is not clear that much progress has been made on the 'Global Fund for Health' recently.

Due to the extremely stripped-down and idealised nature of the APC models used in calculations, the many unsupported assumptions regarding how they would work, and the extremely non-idealised modelling of contending approaches, it is not surprising that the figures so far produced favour the APC approach. Much of the current paper is an attempt to check the robustness of those figures to anything less than the perfect world being presumed when they were calculated and to the removal of the many deliberate biases. The figures produced in support of APCs for HIV, TB, and Malaria vaccines should be treated with a great deal of care – as will become clearer as the following sections unfold. The evidence presented for these diseases is almost entirely the work of one school of thought (in fact of mostly one person). When reference is made here to 'the APC calculations' or to 'the APC figures', it is in reference to the figures *so far presented* for HIV/AIDS, TB and malaria, on the understanding that a completely different set of figures would be calculated under a different set of assumptions.

Before embarking on a radical, and possibly extremely costly, 'experiment' (a word, incidentally, used by some supporters of APCs themselves) one might expect the APC framework to go through a much more thorough critique and careful analysis than it has so far been subjected to. Hopefully this paper will go some way towards encouraging this, and rebalancing the debate. It will become clear that a case for an APC to cover all stages of research and development or even just the final stages of a HIV, TB, or malaria vaccine has not been made in any of these papers.

1.4. The Need For a Broader Debate

A further motivation for this paper is the need to challenge policymakers to broaden the discussion about the mechanisms for creating finance and incentives for investment into R&D for vaccines, and for drugs for neglected diseases in general. Much pressure is currently being exerted to limit even the simple act of discussion. In principle the World Health Organisation's report into IPR and health – due at the end of 2006 – should be wide-ranging, but it has come under intense pressure, especially from some in the US, to rule out even the mere discussion of open collaborative research methods or an R&D Treaty[14]. And when the World Intellectual Property Organization, WIPO, announced – after a request from over 60 leading scientists, economists, legal and health experts – a conference during 2004 to open a broader debate into issues including more 'open' collaborative research methods, the US (mainly under pressure from a handful of large corporations who benefit strongly from the current 'closed' system) crushed it at the start. An uncritically accepted interpretation of APCs – a heavy IP and secrecy-enforcing approach when applied to

[14] The US terms of reference for the Commission went to the trouble of repeating this several times. Incidentally, the author's comments about 'open' science or 'treaties' in no way means the author condones any of these ideas. The issue is the need for an open and critical debate, and this is not possible if whole areas are closed off from us from the start. Besides the author is curious. He does not presume that he knows much about these and certainly not enough to discount them without seeing the evidence. For some reason, some choose to view remarks about keeping the debate open as meaning that somehow one 'advocates' these approaches. It is ironic that those most likely to argue this are least likely to have evidence for their own approach.

early-stage vaccines – detracts attention even more from this broader debate. By demonstrating that APCs are not the panacea that they are claimed to be even as they lead to ever-tighter IPR, it will be become clearer that we cannot hope to analyse alternative mechanisms for generating healthcare innovations without putting IPR issues, and fundamental questions about the way we do research, at the centre of analysis.

And there are good reasons for thoroughly investigating IPR issues *before* instigating APCs, rather than after. A key underlying assumption – the glue holding early-stage APCs together – is that of extremely tight IPR that has *no* negative consequences whatsoever for complicated scientific research. Indeed, there are strong ideological roots to this assumption. As Kremer[15], the chief advocate of APCs for HIV/AIDS, TB, and malaria, comments, echoing leading voices in the large-firm pharmaceutical industry: "While intellectual property rights undoubtedly prevent some from obtaining needed pharmaceuticals, eliminating these rights would not help the majority of those without access to drugs,"[16] and APCs enable access "without imposing price controls or eroding intellectual property rights."[17] This may be slightly ironic given that Kremer himself bases, and biases, his calculations of APC effectiveness on the assumption of highly restrictive access caused by patents (see section 8.1.1. below). Given that it is such a key assumption, one might imagine that those supporting APCs would actively encourage discussion of IPR, happy that this strong assertion holds. And it would seem sensible for those thinking of paying for APCs to check that such a key assumption holds before embarking on an irreversible programme based on them.

If more open collaborative research and Treaty alternatives are so transparently bad, it is not clear anyway why so much effort should go into preventing analysis and free discussion of them. And if APCs are so transparently good, it is even less clear why so much of their actual workings should still be hidden from public gaze after so long. The case for thoroughly reviewing the consequences of IPR is even stronger if APCs are *only* being proposed to cover very late stage developments. The irony is that those promoting extremely tight-patent APCs for vaccines can do so without any critical debate, while those requesting a critical debate about alternatives are told that they are being unreasonable.

Though more time and money has been spent on promoting APCs than on analysing many alternatives, hopefully this paper will at least convince the reader that whatever may or may not come of a broader debate, there is nothing so overwhelming about APCs as to rule out this wider debate.

Finally, though this paper is indeed critical of the APC approach, hopefully it is written in a spirit that recognises that *all* proposed mechanisms – including

[15] Kremer, M. (2002), p 68. I have to excuse the constant reference to Michael Kremer. He just happens to be the source for nearly all of the modelling and calculations for APCs.
[16] This in spite of the fact that the world economy spends well over $400bn per year on drugs (nearly $half-a-trillion if government research support is factored in), and yet only a little over 10% of the $400bn goes back into R&D, with typically just over 3% into research deemed by the FDA and other regulatory bodies as 'highly innovative', and less than 1% into R&D to tackle the diseases that cause 90% of the global disease burden. This half a trillion dollars is a tax on the global economy – the mechanism chosen, based on high IPR, to support pharmaceutical R&D.
[17] K8:4.

'collaborative' research, R&D Treaties, and the like – *should* go through a critical grilling. The issues are far too serious for sentimental attachment or for the easy acceptance of *any* mechanism without putting it through a challenging selection process. The mechanisms chosen to tackle such a large and pressing problem have got to work in reality and not just in rarefied settings on paper.

1.5. Core Issues to Guide the Choice of a Mechanism – A Brief Summary of How APCs Fare

This paper is not just about APCs however. Its other justification is that, hopefully, many general principles will come out that will help inform our thinking on how to create mechanisms to finance and provide incentives for vaccine research and development, and how to choose between those mechanisms being offered.

In particular, since APCs share many similarities with prize funds, this paper can also be seen as a commentary on some of the problems that prize funds would also face, and the ways in which this would affect *their* cost-effectiveness[18]. The historical record of prize funds is extremely limited, and there is no systematic analysis of the relative success of the approach. Whether prize funds are appropriate for technology of a more cumulative and complementary nature, in a high IPR environment, with later rounds of technology 'reading off' dozens, even hundreds, of earlier rounds of technology and patents, is not immediately clear. Hopefully this paper will cast some thoughts on this issue too.

The paper is also interested in a number of other core issues. Each is described briefly now, followed by a brief evaluation of how the APC fares against it.

1) The practical difficulties of generating incentives towards creating 'higher-quality' vaccines rather than 'lower-quality' vaccines. It turns out that the APC mechanism struggles greatly with this in the case of early-stage vaccines, with a tendency towards lower quality, especially when dealing with potentially very long processes.

2) How the complications and distortions of highly cumulative and complementary technology feed through to incentive mechanisms and overall costs. This is ignored in the APC calculations, since typical APC models assume that all research projects are strict substitutes for each other, that is they are totally independent 'gambles'[19]. No HIV vaccine project benefits from the presence of any other HIV vaccine project. In a follow-on paper, this will be used as a device to compare and to contrast the APC approach with a more open collaborative approach, where technology is more complementary and more information sharing is taking place.

3) The rôle of the system of IPR as it interacts with the chosen finance mechanism, and the implications of this for the costs of vaccine and drug research in general (including that of non-APC vaccines and drugs). Real-world early-stage APCs would

[18] See also Farlow 2005 "Prize Funds for Drugs and Vaccines: Principles and Problems," shortly forthcoming.

[19] The notion of a gamble is quite a good one. In a lottery, if you buy more tickets (K3:8) it increases your chance of the prize, but only on account of those tickets. No old ticket benefits from the purchase of any new ticket.

deliberately work to make the IPR system tighter. This is needed to keep down the costs of APCs[20]. Meanwhile APCs themselves suffer from the higher research costs imposed by tighter IPR, but IPR is not even modelled in APC calculations and this extra cost is therefore ruled out. Indeed, the favourable APC figures that have been produced are based on idealised, open source, low IPR, information structures, even as tight IPR would be put back in to make them work in practice. Counter-intuitively, the APC figures are calculated such that the *more* harm done in the past by systems based on high IPR, then the *better* the APC and high IPR systems perform in the future and the *worse* direct government R&D and low IPR systems perform in comparisons[21]. In addition, being a high IPR mechanism, there are cases where APCs may be used to segment the market in ways that increase prices both for those in the program and for those outside the program, and times when they may be used to deny or delay access[22].

4) The rôle of fixed and sunk costs. The pharmaceutical industry is renowned for the large sunk costs that build up over time. Certain stages of vaccine research require large fixed investments. How does this feed in to the performance of each finance mechanism? In particular, we will find important effects on the structure of the industry and the strategic behaviour of firms, capital costs, and the risks to firms generated by the APC mechanism itself. The removal of sunk costs (and indeed all fixed costs) in APC models generates serious, and potentially very misleading, results.

5) The nature of capital costs. The ignoring of this, turns out to be an extremely significant shortfall of APC calculations. APCs: i) concentrate in the use of modes of finance that already involve high costs of capital; ii) generate many extra options-based capital costs; iii) generate extra costs of capital through the risks and distortions the APC mechanism generates for those taking part in it. These extra costs are higher the more complicated the technology. High capital costs reduce the discounted value of the end value of any project, requiring a higher end value in the first place – in this case a higher APC price. Together these layers of capital costs mean that the majority of the cost of an APC for HIV of the sort being currently argued would almost certainly be capital costs, with the proportion of capital costs rising the more complicated the technology. The cost of generating a vaccine for HIV/AIDS through an APC might contain only a small proportion of out-of-pocket trials costs and a huge proportion of capital costs. Mechanisms that generate low capital costs are to be favoured over those that create great capital costs.

6) The nature of the interactions of publicly-funded research (other than APC research, that, of course, is also publicly-funded) with APC-funded research. The APC calculations presume *perfect* interaction at *all* times;

7) The different layers of possible 'crowding out'. The APC, in spite of the rhetoric, suffers from many crowding out difficulties. Not one of these is factored into the APC cost-comparison figures. This under-reports the *global* public-funding costs of APCs. Together with the high capital costs, this makes APCs much less cost-effective than is being claimed;

[20] It will be seen below that this is only meant in one fairly limited sense; it is offset in many other senses.
[21] See section 8.1.1 below.
[22] See section 7.16 below.

8) Issues relating to the distribution of firms and researchers being targeted by different finance mechanisms. The incentives created by APCs are tipped in favour of vaccine research in large pharmaceutical firms, to the detriment of those in the not-for-profit sector, small and new biotechs, developing country, and university-based research departments (though this would be masked by a very high APC). Mechanisms that, for the same outlay of public and foundation resources, would be more favourable to the latter are to be favoured if that is where the cheapest and most innovative research takes place, and if it helps technology transfer, and efforts to get production costs down;

9) How asymmetric information issues complicate incentives and distort mechanisms to finance vaccine research. In particular, we are interested in the exact informational assumptions underlying alternative mechanisms and the nature of informational processing going on. In the APC calculations, the amount, quality, and symmetry of information being presumed is surprising high, and is often presumed to be perfect;

10) The rôle and costs of institutions and political processes in the working of different mechanisms. The APC turns out to be layered with institutions and committees with potentially large amounts of discretion (the more complicated the technology, the more the discretion). This leaves it open to all kinds of institutional failure (or costly mechanisms to prevent it), political pressures, dynamic inconsistencies, and coordination failures. This is largely ignored in the APC models, and none of the costs of this are fed into the APC cost comparison figures. Finance mechanisms that would avoid these problems would work out cheaper, with more rapid vaccine development.

11) The nature of strategic behaviour at many different levels and the implications for the efficiency and cost of mechanisms – especially the way that costs are *always* higher than in idealised strategy-free models and can be driven way above the costs in such models. The APC figures are calculated strategy-free, with perfect competition at all times. It is, put quite bluntly, outrageous to simply assume 'perfect competition' and not analyse whether or not this would actually be the case. It is argued here that one of the main faults of the early-stage APC approach is that it does not generate enough competition because it too easily collapses down to a very few companies. This is bad for speed and quality of vaccine development.

12) The nature of financial market problems and how the various vaccine finance alternatives cope with these, and how this feeds their effectiveness. This, in part, touches upon the types of finance that firms use but also upon the nature of the type of firms being benefited by the types of finance used. Again, the APC favours large pharmaceutical firms over the not-for-profit sector, and small and new biotechs. There are no financial market problems factored into the APC calculations. Financial markets work perfectly always.

13) The difficulties of setting optimal values for any finance mechanism. No convincing evidence is provided that the APC price would, or could, *ever* be set even remotely optimally. The more the APC price deviates from the optimal setting, the less effective the mechanism. An auction is mentioned as a solution, but would be unlikely to work given the nature of the asymmetric information problem at hand. In

reality, APCs would be set with reference to the evidence provided from the large pharmaceutical firms it was targeting. All APC figures have so far been calculated on the presumption that a mechanism has already set the terms optimally (though the figures are only "rough rule of thumb" industry figures). Notions of cost-effectiveness are thus biased.

14) The nature of delay built into any mechanism, its costs and dangers. Many of the proposed mechanisms for getting an APC to function efficiently involve important elements of delay. This not only creates costs in terms of welfare loss – that should rightly be priced in to overall cost-effectiveness – but also creates rounds of further problems, including the dangers of reduced private R&D in advance of APCs, even for diseases for which APCs never take place at all.

15) The potential dangers, and the expected costs, of particularly bad outcomes. Each finance and incentive mechanism tends to generate a distribution of possible outcomes. While we might be interested in the average of these, we should also be concerned about the outcomes in the tails of the distribution. Expected social welfare is higher if a mechanism is more likely to generate an outcome with a distribution with less in its tails, and that is more likely to avoid *really* bad outcomes altogether. It is argued here that the APC may generate distributions with more in the tails, and may contain some particularly spectacular worse-case scenarios[23] including cases of collapse, spiralling costs, delay, and even cases where vaccines are developed and withheld from the program. The expected social costs of these are not priced in to any APC figures.

The reader can probably guess that the author's view is that the APC calculations so far presented ignore many extremely important issues, and need to be recalculated in light of them.

[23] This indicates that VaR analysis of mechanisms might be a valid exercise, though probably extremely difficult to do.

2. REASONS GIVEN FOR VACINNE R&D FAILURE – A FEW PRELIMINARY THOUGHTS

2.1. Ten Reasons for Failure – But only one Drives the APC Result

Various reasons are typically given for the failure to invest in R&D for vaccines. It is useful to see which of these is at the core of the APC case:

1) The purchasing power of the poor is too low. Many of those countries that might benefit from vaccines spend less than $10-$20 per capita per year on health. A vaccine program would be too expensive compared to the payments that could ever possibly be charged given such a low health budget. This leads to the 'lack of a market' for vaccines argument. For example, even though 70% of HIV-AIDS infections currently take place in Africa, where clade C is the most common type, commercial developers concentrate on clade B, common in the US and Europe where health budgets per capita are many times higher;

2) There is a free-rider problem amongst developing countries (or those acting on their behalf). They would each rather prefer that other countries pay for the R&D underlying the development of a vaccine than pay for it themselves. Following vaccine development, they will each bargain for their own price to be lower. In this uncoordinated 'prisoners' dilemma' they would collectively bargain the price too low to recover R&D costs. They might like to 'pretend' *ex ante* that they would not bid prices down *ex post*, but this is not credible. This is one aspect of the 'dynamic inconsistency' argument;

3) There is an incomplete set of markets. For example, children are unable to trade the future value of life saved by taking a vaccine now, even though the cost of the vaccine would be trivial compared to the value of that future (even just the wage component, never mind the non-monetary component). In a more developed economy, with social insurance markets, this externality would be factored in. This argument is therefore really about the failure of social insurance markets in developing countries – which is another way of saying that the victims are poor;

4) Consumers are more willing to pay for treatments than for vaccinations. The bad outcome averted is never seen, whereas the bad outcome that is treated and made better is. Many potential customers are illiterate, and many place little faith in public pronouncements anyway, including about the benefits of vaccinations. They wait to see benefits by seeing what happens to others. Even when vaccines do change a situation, the time until the benefit is seen is considerable, many who do not take the vaccine do not get ill, and the vaccination program itself may reduce the visibility of the 'alternative' terrible outcome, lulling others into a false sense of security.

This problem gets worse the higher the level of vaccination. For example, if 90% of the population has been vaccinated against HIV/AIDS, and of those who have not been vaccinated, 20% have HIV/AIDS, the visibility of HIV/AIDS victims is 2% in the population, and probably even lower if there is stigma attached to it (this simplifies greatly the nature of the steady state of the system at the start, since this is

for illustrative purposes only). Those who are not vaccinated are much more vulnerable than they believe, but they are being lulled into a false sense of security by the 90% who are 'safe'. This is especially so if the vaccine is 'only' a 'therapeutic' vaccine, that is a vaccine that does not prevent acquisition of the virus but does delay the need for treatment (a therapeutic vaccine will still reduce rates of transmission by reducing the viral load in the victims).

5) There are health externalities. The private benefit of a vaccine is always less than the social benefit. Taking an HIV vaccine reduces not only one's own risk of mortality, but also reduces the risk of mortality for others. If there are any risks to taking the vaccine itself, this problem is made worse. By not internalising this benefit, most of those taking the vaccine, it is claimed, would not be willing to pay the full social value of the vaccine. Again, it seems to suggest a failure of social insurance to place any value on this externality – another side to the fact that the victims are poor;

6) Designing around patents is too easy. This puts developers off from doing vaccine research.

7) The science and technology of vaccine research is difficult and unpredictable, and serious failures happen because of this. It may be that this is increasingly being aggravated by IPR-related problems;

8) In part because of the difficulty and unpredictability of the science, but also because of the public good nature of important parts of the science, financial markets struggle to finance R&D into vaccines;

9) Vaccines replace profitable treatments. Firms have less incentive to invest in vaccine R&D when they know that it simply 'replaces themselves'. A one-off, cheap, HIV vaccine is much less profitable than a stream of profits from long-term more expensive treatments. Incidentally, this is not supposed to be casting aspersions on the executives of large pharmaceutical firms. Financial markets feed these constraints to them. Those firms working on programs that would risk replacing other profitable programs, find their capital costs are higher (stock markets are supposed to price in all future expected discounted flows of profit, so that the mere *possibility* of reduced overall profit flows is enough to send capital costs higher, for this and all other research). A different set of financial market conditions would feed a different set of constraints, as will be explored below.

In a competitive pharmaceutical industry (where, also, the IPR system would allow entrant firms to acquire technology that might undermine current firms), one might expect that those companies developing vaccines would still have a strong incentive to do so, since vaccines would replace the treatments of *other* companies. But a system heavily dependent on the *same* few companies for both treatments and vaccines – and able through tight IPR to restrict access to information that might undermine their competitive positions – generates a much larger 'replacement effect' and much less of an incentive to develop vaccines. This problem is reinforced if biotechs and not-for-profit firms cannot raise finance to take a vaccine 'all the way', since the only viable market for their output is firms that face a 'replacement effect', thus feeding the 'replacement effect' onto the biotechs. One of the solutions is a

mechanism that allows for more players in the market, not bigger incentives for the same few players.

10) Other ways to perform vaccine research are inefficient. Indeed much vaccine research is very wasteful, reducing the productivity of current research efforts, and the chances of discovery;

Failure on the first five points does not, *per se*, favour one method of financing vaccine R&D over another. The poverty of the victims is often mentioned in the APC literature, and sometime made the chief justification. However, since Foundations and rich country governments pay *whatever* the mechanism, it should be the efficiency of the mechanism and not the poverty of the victims that matters. And it is not, anyway, central to the APC cost-comparison results.

APCs do not directly target the sixth problem. But, since they are even more dependent on tight patents than the current patent-based system to work, they lead to even tighter patents[24].

Although industry frequently mentions the scientific difficulties (point 7) this is expunged from the APC calculations. And all IPR-based issues in such difficult scientific settings are ignored too. Indeed, recent APC cost estimates treat APC-generated vaccines on a par with other products in a typical pharmaceutical firm portfolio[25]. Since the market size issue is the flip-side to the R&D issue, the underlying presumption is that the R&D problem is no more difficult than typical. Hence, the scientific issue is not part of the current APC debate.

We extensively analyse financial market issues (point 8) in Section 12, though problems with these are also removed from the Kremer model, replaced instead by a perfectly-performing idealised set of financial markets.

In certain markets, the 'replacement effect' (point 9) may be important, but these are ruled out in the APC calculations too.

The principle justification for APCs is point 10, the alleged failure of all other approaches, and, in addition, the perfect application of APCs. We will find in Section 8 that the evidence given for the first of these claims is very thin and heavily selected. And the rest of this paper will show that the second claim can't be upheld either.

2.2. What Does 'Lack of a Market' Mean Anyway? The Rôle of Institutions and IPR

What does 'lack of a market' mean exactly anyway? *All* mechanisms currently being considered for encouraging finance for vaccine research visualise the majority of

[24] This has led some, such as the American Enterprise Institute, to question the need for APCs, when the patent system could simply be made tighter anyway, and to argue that part of the benefit being picked up from an APC could have been picked up equally as well without it.

[25] "Advanced Markets for a Malaria Vaccine: Estimating Costs and Effectiveness," Berndt, E.R., Glennerster, R., Kremer, M.R., Lee, J., Levine, R., Weizsäcker, G., and Williams, H. 2005.

expenditure coming from tax-payers in rich economies. This includes the finance for APCs as much as for any other mechanism. The ultimate effective 'market' for vaccines is therefore large groups of developed economies via some institution like the World Bank or the 'Global Fund for Health' being proposed by the UK government, or a large philanthropic foundation.

Rather than thinking in terms of the purchasing power of the poor, the problem is that investors have come to simply not trust that *large donors* will pay a reasonable enough price to cover R&D costs. The 'lack of a market' turns out to refer, in part, to this previous lack of resolve to purchase the results of vaccine R&D investments targeted on the diseases mostly of the poor (partly itself also based on poor health infrastructure which reduces the value of such purchases, and on financial market difficulties, etc.).

This shows up in the fact that we already have vaccines that are heavily underutilised in poor countries. The hepatitis B vaccine became available in the US in 1981, but did not get purchased by UNICEF until 1994. At least 40% of children in Sub-Saharan Africa still do not receive the vaccine. And it took 11 years (till 1998) from the introduction of the first Hib vaccine in the US, for it to be bought in bulk for developing countries. Even now, usage is heavily skewed towards rich countries, with tiny percentages of coverage in poor countries. The yellow fever vaccine costs as little as four US cents per shot to manufacture and yet millions of children do not receive it. Even some very low priced vaccines have not stimulated their universal availability.

Mercer Management Consultants conclude that "larger companies, supplying 20-40% of today's vaccines, attached little or no commercial value to developing country markets. These firms noted the slow up-take and lack of funding for current 'priority' vaccines despite their cost-effectiveness and health value" and "there were very few, if any, projects motivated primarily by the needs of developing countries." Mercer estimated that fewer than 200 scientists in the private sector are dedicated to HIV vaccine related work, and judged that this number probably exaggerated the resources devoted by the private sector, since some of these scientists were grant-supported by the public sector.

This is all in spite of the fact that vaccines are an extremely cost-effective measure. The World Bank[26] refers to health interventions costing $25 to $150 per DALY saved as 'highly cost-effective' (in the US the figure is $50,000-$100,000 per year[27]). Compared to treatment programmes, most vaccines are relatively uncomplicated interventions that do not require monitoring or follow-up visits to health care workers[28]. Kremer in his calculations (on assumption that there is competition) sets the marginal manufacturing cost of a vaccine at about 40-50 cents[29].

[26] 1993 World Development Report (p 8, 64, 68).
[27] Neumann et. al. 2000. Incidentally, this is the sort of figure typically used to calculate the value added by 'me-too' drugs.
[28] As an example, a typical course of tuberculosis *treatment* requires the availability of x-rays and /or microscopes and at least two months of monitoring. All this would be avoided (though, in many cases such treatment simply doesn't happen now anyway). However, this may not be the case in, at least, the early stages of rolling out a therapeutic HIV vaccine program.
[29] K5:13

How IPR impacts on measures of 'lack of a market'

Measures of the 'lack of a market' are also complicated by the state of IPR. In APC calculations, the size of the missing 'needed market' to stimulate R&D is calculated *relative* to 'current market' size, which is deemed a function of IPR[30]. In particular, low–income countries are presumed in APC calculations to be unable to access vaccines during the first 10 years of production (because of high prices and tight IPR), and "after this period it will be at a price that is not a significant source of profits for the vaccine developer"[31] (which also depresses 'current market size'). Indeed, if one thinks of the discounting process that pharmaceutical firms perform (nominal 15+ per cent per year), sales after ten years have such a low present discounted value that it would have made hardly any difference to the firms bottom line whether these sales had been made after ten years or at a very low price at the start of the ten years.

So, the worse the past effects of IPR and the greater the access problems, the greater the measured 'lack of a market', the higher the social surplus generated by an APC, and the higher the APC price – which is then fed to the high IPR system that fed the 'lack of a market' in the first place. Should past IPR problems be *rewarded* with *higher* public expenditure via an APC or should the IPR be analysed to see how much it was part of the problem in the first place? Such figures for 'lack of a market' are also not so useful when calculating the impact of alternative mechanisms that are based on lower IPR.

The APC solution to this past failure – including IPR-related failures – in the face of overwhelming evidence of cost-effectiveness of vaccines, is a promise to resolve to spend at some unspecified date in the future and, meanwhile, for private finance to be sunk in research efforts with all expected capital costs to be repaid by taxpayers and foundations. However, even if 'tight' contracts are used to enforce this resolve, this past history *will* affect the private capital costs component, especially if there remains any discretion in the system (it is argued below that any reasonable modelling of APCs will lead to flexible, discretionary, elements, and therefore will lead to higher costs of capital than if the system did not have this history to contend with). And it is all enacted without any analysis of those aspects of IPR that may have aggravated the problem in the first place.

[30] K1.

[31] K2:2. This is based on the experience of Hepatitis B and Hib.

3. GENERAL PRINCIPLES FOR COMPARING MECHANISMS FOR FINANCING VACCINE R&D

There are a number of general principles against which *any* mechanism for financing vaccine R&D should be judged.

3.1. It Should be a Comprehensive Solution

In *economic* parlance, the mechanism should be weighed for its impact using 'general equilibrium', and not just 'partial equilibrium' analysis. In *non-economic* parlance, any solution should be 'comprehensive'. Or, at least, any solution whereby *only* one or two vaccines are targeted, should not conflict with a comprehensive solution in which *all* vaccines (and drugs) are targeted. Kremer makes no bones of the fact that the APC calculations are based on only partial equilibrium analysis, and that the general equilibrium case has not yet been made for the APC. Judging such 'cherry picking' programs on the basis of just their 'comparative static' features generates potentially very misleading results.

A general equilibrium approach would be interested to know what impact any tightening of IPR needed for an APC to work for one or two vaccines might have elsewhere on all the other vaccines and drugs not covered by APCs. If compulsory licensing on the many drugs is quashed for the sake of the few, what is the impact on health in general? If a partial solution creates incentives that distort research away from other diseases, what is the extra cost imposed (unless efficient mechanisms can be set up to avoid these distortions)? And, while it has been argued that the APC involves no resources up front, what if the use of an APC program to tackle *multiple* vaccines and drugs at once *does* require some financial reserves up front, and how are the costs of this factored into the general equilibrium cost comparisons? The more vaccines and drugs put into the program, the higher the level of reserves needed to help make the program credible and to keep private capital costs down, but the higher the general equilibrium costs to tax payers. Since there is a finite limit to the reserves available, and since these reserves impose costs, many potential vaccines would have to be left out, and costs will be higher than simple partial analysis would suggest. If an alternative, non reserve-based mechanism could have kept *all* vaccines in and kept capital costs down, this makes the APC less favoured on general equilibrium grounds.

Similarly, an APC designed only for the late stages of vaccine development should not conflict with mechanisms designed to tackle earlier stage problems, especially if that is where many of the difficulties lie. A full general equilibrium analysis of the APC would want to know whether any problems it creates at earlier stages might feed back *on itself* at later stages. The repeated danger in much of the APC analysis is that APCs are modelled with the limited assumptions of a late stage device, only then to be compared with alternatives modelled as tackling a much longer process. For fair comparison, APCs need also to be modelled as dealing with this much longer process.

In a general equilibrium solution, the total number of calculated DALYs saved as a result of success on one particular vaccine must be offset by any DALYs lost due to a

failure to tackle other vaccines *if this failure is itself a side-effect of the imposition of the program rather than another*[32].

3.2. Comparison Should be Based on Relative Distortions

"...if the government funds only worthwhile research projects and researchers focus all their energies on developing a vaccine, the expected discounted cost of developing a vaccine is likely to be similar in net present value terms whether research is financed at the front end, through government grants or induced by payments for a successful vaccine at the back end"[Kremer]

Kremer's admission that perfectly enacted front-loaded and perfectly-enacted end–loaded mechanisms have equal efficacy is revelatory, especially when one realises that the comment is based on the extremely simplified notion of technology underlying the APC end-loaded mechanism. A large chunk of what follows seeks to explore how distortions might arise under APCs and how this might alter their effectiveness. The distortions we are most worried about fall into two categories: i) those that are difficult or expensive to design around, although theoretically it might be possible, and ii) those distortions that cannot logically be avoided.

The proponents of APCs have many times referred to the potential to create distortions and perverse incentives through an APC. However, after 7 years of acknowledging this fact, little has been spelled out in public discourse. That the problems are being dealt with 'in private' is supposed to reassure. Meanwhile, all calculations presented in favour of the APC have been based on simple, idealised, distortion-free APCs even as the alternatives are modelled in their distortion-full states[33]. Since distortions raise costs, we should not be surprised if the APC comes out cheapest by this methodology.

Practical policy requires that mechanisms be weighed up on the basis of *realistic predictions* of their distortions in real-world settings. Most of the rest of the paper is essentially devoted to exploring the nature of the distortions that APCs might both create and struggle to deal with. It argues that real-world applications of APCs contain layer upon layer of distortions that are simply missing in the Kremer calculations.

It is argued below and elsewhere[34] that, in principle, a well-designed much more open collaborative research framework, though not without serious problems itself, is less prone to many of these distortions, and is a potentially more cost-effective way to deal with others. This further justifies a broader debate to include models of more open forms of, collaborative, research.

---

[32] All based on ceteris paribus assumptions for any given level of funding.

[33] Stiglizt, J. (2003) p 251, refers to this as the typical tactic of those who preach what he calls the "market mantra". With reference to the disastrous policy errors that led to Enron and many other scandals: "They would have us compare an imperfect regulated economy with an idealized free market, rather than an imperfect regulated economy with an even more imperfect unregulated one."

[34] Farlow 2005 "Collaborative Research Methods for Drugs and Vaccines: Principles and Problems," shortly forthcoming.

3.3. It Should be the Least Expensive, and Quickest, Option

Throughout we should presume that there *is* a global budget constraint. Clearly, if $10bn was thrown at a vaccine development program for one disease, there would be a much greater chance of discovery than if only $1billion was thrown at it. But with many vaccines and drugs being in need of discovery, the opportunity cost of developing a vaccine for one disease is the new vaccines and drugs and treatment programs that cannot be afforded given this global budget constraint. In the case of an APC, this includes also vaccines and treatments for other diseases in the future when the current APC comes up for payment if the APC turns out to be much more expensive than presumed. Similarly, looked at from another angle, the less expensive an option is the more quickly a vaccine is developed for any given outlay of resources.

This is sometimes treated in a rather off-hand fashion by APC supporters. The range in which may fall the cost of an APC for any vaccine is bound below by the need to motivate R&D and above by the 'cost-effectiveness' of the vaccine once it is developed. Kremer claims that the upper bound "is not likely a problem" since "it would be hard to imagine a situation in which purchasing vaccines for malaria, tuberculosis, and AIDS would not be cost-effective"[35]. We do not take issue here with the notion that the social value of vaccines, and therefore the upper bound, is extremely high. But, it *is* argued that keeping costs for a vaccine as close to the lower bound as possible is the main criterion for choosing an approach to R&D finance, given that research into so many other vaccines and neglected diseases is already chronically low. Realistic APCs may have a struggle trying to keep the price close to the lower bound, or even away from the upper bound. It is not particularly helpful to argue that since the top of the range of social surplus is so very high, that not breaching it – but nevertheless being highly wasteful[36] – is any measure of success. And it is not clear that comparisons with alternatives mean anything if the favoured mechanism is then allowed to become as expensive as it likes.

[35] K7:47.

[36] Another way to think of this is that for any given level of social surplus sacrificed to motivate R&D for any particular vaccine or drug, development of this particular vaccine or drug could have been quicker, and more other vaccines and drugs could have been created.

4. THE FORMIDABLE SCIENCE OF VACCINE RESEARCH AND THE SIMPLIFICATIONS OF APC MODELS

4.1. The Formidable Science of Vaccine Research

It is acknowledged widely in the scientific literature that vaccine research faces formidable scientific and technological obstacles. A study for the World Bank testifies to this[37]:

"*The low levels of investment in an HIV vaccine can partially be explained by the inability of companies to see a realistic commercial return. However, all companies cited the scientific barrier as an equally important barrier...the low probability of success of any given candidate and high profile of some failures has had a significant impact on corporate thinking*."

"*The absence of a scientific consensus is evident from the number of approaches still being considered, more than a decade after the effort started. The scientific barriers arise from numerous factors*."

"*This scientific uncertainty is further compounded by the limited understanding of the virus, a lack of correlates of immunity and lack of animal models – all of which contribute to a much higher degree of uncertainty about the potential efficacy of a vaccine candidate than is typical when considering investing in expensive Phase III trials."*

"*A vicious circle is created in which Phase III efficacy trials are needed to rapidly advance knowledge to identify and develop an efficacious product. However, very few companies have sufficient confidence in the probable efficacy of their candidates to risk the required amounts of money. Not surprisingly, the uncertainty surrounding the science and future markets for an HIV vaccine has translated into difficulty in raising private funds to address the issues*."

This vicious circle can be illustrated diagrammatically:

[37] Mercer Management Consultancy, (March 2000) Quotes from pages 6-7.

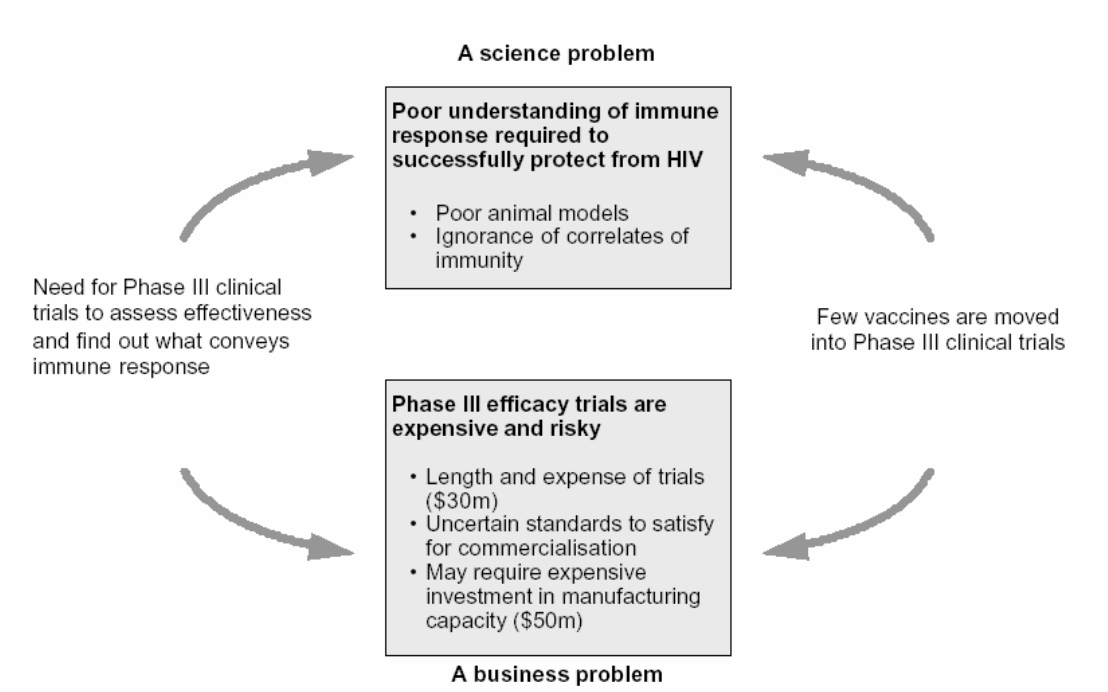

4.2. The Simplified Technology in APC Calculations

4.2.1. The basic APC model

Technology in Kremer is modelled as essentially (strictly) 'stationary', that is the technology is immovably fixed, and repeated the same each period. Firms have access to different research opportunities (presuming it is scientifically possible) each indexed by its instantaneous probability of success (from now on, for simplicity, this will be referred to as 'p'). Probability is per unit of time (from now on we will use the word 'probability' when we really mean 'instantaneous probability per unit of time'). Access to this technology is a function of the human capital of the firm's scientists but also of the firm's ownership of intellectual property. IPR is modelled as always functioning perfectly, with no asymmetric information problems, hold-ups, financial constraints on IPR ownership, or strategic behaviour of any sort. In fact IPR works so perfectly that it is simply ignored. All firms have equal access to all IPR at a market price set perfectly competitively.

Research projects are drawn from a probability distribution over research projects such that the probability of each research project is independent of every other research project. There are no complementarities across projects. Projects are 'perfect substitutes' and incapable of generating information of any use to any other project. As more firms enter to carry out R&D then the greater the collective probability of discovery (there are more research leads being followed), but the lower the individual probability that any one firm will make the vaccine discovery (the more other firms there are in the market, the lower the probability expected by any one firm that when the vaccine is discovered it will be its project that discovered it).

It is then claimed by Kremer that if the distribution of the probability of discovery varies over the product development cycle (which we know it does), then that variation takes place in a known manner, and that the average of this variation can be

taken as 'representative' of the whole pattern of variation over the cycle. With this notion in place, the model can be reduced to that of a *constant* distribution of the probability of discovery per period of time. Game theoretically speaking, technology is reduced to the same stage game repeated over and over again – with the same firm-level probability of discovery distribution both across firms but also for an individual firm across all time periods, and the same collective probability of discovery distribution. And all this is symmetrically known and common knowledge in advance[38].

The model can then be solved for its (unique) equilibrium. Essentially, the larger the APC price then the more research leads will be followed, and this will lead to faster vaccine discovery. But with many more firms pursuing leads, the smaller the chance for any particular lead (hence any particular firm) that it will be one to get the APC. Firms therefore pursue R&D so long as the R&D costs are less than the "probability of success of the project times the value of a successful project discounted at an appropriate discount rate," with the discount rate set exogenously to the APC (repeatedly below we argue that it is actually endogenous).

A value for the APC can then be set in order to motivate the *optimal* amount of R&D. One might like to think of this as the 'Goldilocks' problem; finding a value for the APC payment that is neither too high nor too low. If it is set too high it motivates wasteful forms of R&D, is wasteful of public resources (they have to be raised from taxes so this causes deadweight loses elsewhere) and reduces resources available for other vaccines and treatments; in these circumstances it might have been better to have used an alternative mechanism if it could have avoided this waste. But, if the APC is set too low it motivates too little R&D; the probability of discovery is below the optimal rate, this unnecessarily delays vaccine discovery, and the wastefulness is in terms of the suffering and lives lost on account of delay. Once set too low, it can be very destructive to try to revise upwards the level of payment later. Think of the investment logic (specifically the option cost) if firms know that if they hold back, the level of payment may be set a great deal higher later; early R&D becomes a great deal more expensive. Setting the APC size too low is as bad if not worse than setting it too high.

Some big simplifications

The device of using a distribution of p that is itself constant over time, is really only suited to a situation of a constant state of science at both basic and applied levels[39], perfectly-functioning totally efficient basic science[40], perfectly operating IPR[41], no

[38] For example there is never any secrecy about potential projects.

[39] Intuitively, if the number of viable projects naturally changes and evolves over time, then the shape and position of the curves would naturally 'move about' (See Section 5), and if there are technology 'shocks', the p distributions would 'move back and forth' discontinuously (See Section 6). Both of these, including the 'possibility' of these, are ruled out in APC calculations.

[40] The state of basic science dictates the position of the p distribution, such that the 'better' the performance of basic science, the 'further out' the p distribution. Clearly if anything harms the performance of basic science, the p distribution is moved 'further in'. The APC calculations have stripped out any worries that the negative impact of the APC on the early stages might have 'pulled in' the position of the probability distribution at later stages.

[41] So that the p distribution of one player cannot be 'held back' by any other player's ownership of IPR, nor indeed can any risks related to the actual or potential ownership of IPR 'hold back' the position of any firm's p distribution (in the expected sense).

strategic behaviour[42], and simple comparative static analysis[43]. Kremer[44] concedes all of this: "Because we infer the current value of p and the shape of its distribution (f(p)), this model is best suited for comparing different policies under consistent assumptions about the state of technology," (by which he really means "assumptions about the constant state of technology"[45]). Furthermore, he states that "Our estimate of the time until a vaccine is developed under any single policy is only as accurate as our estimate of the distribution of p"[46]. Anyone who read the last half a dozen or so sub-notes will realise just how limiting these foundations are. As soon as the probability distribution becomes as complicated as we will shortly suggest, the estimates become potentially wildly off course.

So Kremer concedes technological complications, but totally expunges them from calculations. It may be that this is because APCs are visualised as *sufficiently* late-stage that the issue does not arise. But even this is hard to reconcile with the scientific literature[47] that seems to indicate that even late-stage research on some (though not all) vaccines suffers from technological complications, never mind early-stage vaccines such as HIV and malaria. Kremer claims that the state of basic science is fixed (and is also not harmed in any way by the APC) and that therefore so is the position of the distribution of p, and that "more basic science" has very little impact on the distribution of p in the "relevant range of interest", and simply "increases the number of low probability speculative projects...that would not be pursued by applied researchers"[48]. This is empirically verifiable but is *not* verified in this literature. Nor does it seem to square with those working on vaccine research in areas such as HIV/AIDS; all the HIV vaccine scientists *this* author has talked to disagree with the assertion. Neither is it clear that even if the state of basic science is fixed, the position of the p distribution is bound to be invariant, as we will see.

---

[42] Strategic behaviour refers to any acts that deliberately work to boost the position of one's own p distribution at the cost of someone else's p distribution.

[43] For example, all of the above assumptions remove most of the notion of 'expectations' – other than the knowledge that the p distribution always was and always will be the same. And there is no asymmetric information of any sort (or rather there is, but only when public-funders try to choose projects).

[44] K3:25.

[45] The phrase "consistent assumptions about the state of technology" is linguistically consistent with the notion that technology can be very inconsistent, and that it is just the assumptions that are consistent. K: Executive Summary.1. spells it out as: "a consistent set of assumptions about the scientific difficulty", which similarly suggests that the scientific difficulty can be variable, though the assumptions about it are consistent; and it strains to avoid saying "a set of assumptions about the consistent scientific difficulty". For some reason, Kremer *never* does use the phrases "constant state of science" or "consistent scientific difficulty" though that is *exactly* how the model is set up. The assumption about the "constant state of science" is as much about the constant state of basic science as about the constant state of the part of science covered by the APC. The constant state of basic science means that *nothing* can be done via basic science to shift the position of the p distribution. This slightly contradicts the claim made elsewhere that: "At this stage, it is difficult to model the responsiveness to basic science of opportunities for applied R&D for products for malaria, tuberculosis, and HIV." (K2:8). The sensible approach would seem not to rule out *any* responsiveness.

[46] And it might be added, only as correct as the *ex ante* estimates of the outcomes of strategic interactions and a host of other features.

[47] And the private comments and correspondences with the author from scientists working on vaccines.

[48] K2:8. I.e. more basic science does not alter the position of the distribution; it just 'stretches it' – at one end, that is the low-quality end. Kremer happily asserts such things, but does not back them up with any evidence.

Essentially, it turns out, the APC economic models developed so far boil down to the finance and selection of trials of vaccines where the science has become relatively straightforward. Hence, criticisms of alternatives must boil down to criticisms over their abilities to select amongst vaccine trials at a level of science that is relatively unsophisticated. It does not help when trying to determine how to try to tackle a much longer and complex research process. And it is harder to see how comparisons can be made with alternative frameworks that try to tackle this much larger and more complex problem.

Furthermore, in all calculations, Kremer assumes (that is, he calibrates from industry figures) the probability distributions (he allows for a range of shapes), but the underlying data has nothing to do with *vaccine* research costs *per se* or, indeed, with any independently verified set of data on the cost of doing vaccine research. We will return to this issue later.

4.2.2. All dependencies removed, with striking consequences – including the use of open source reasoning

By this technological assumption, all dependencies (technological and strategic), across periods, across players (public and private), and across projects, and all informational problems, are stripped out by *a priori* assumption.

Is this a sensible way to model the technology of something as complicated as HIV or malaria vaccine development? Does it matter, for example, that p might vary, for scientific reasons, over the cycle of a product? Or that p might be a function of changes (even shock changes) in the distribution of technology generally? Or that p might be a function of the changing state of basic science? Or that p might itself be endogenous to the acts of the players and policymakers – meaning that the probability structures themselves can be shifted by the collective behaviour of the players or by government policy? Or that information spillovers might matter?

There are very striking consequences. In particular, the APC solution is thus reduced to one unique value, invariant over time. Set *too high* and it motivates wasteful R&D. Set *too low* and R&D is too slow. Both mistakes impose costs. But, in Kremer, since the APC *never* has to struggle with the vagaries of technological and strategic interdependence, or any informational or IPR-generated problems whatsoever, the value never has to vary and is *always set correctly*. There are *never* any costs of getting the value wrong. This generates cost comparison data that is obviously bound to be biased in favour of the APC.

Most striking of all – and highly paradoxical – the chief result in favour of the APC turns out to be based on the absence of patents on anything other than the final product, i.e. it is based on open source information structures[49]. Only as the technology deviates from the highly simplified version used in such APC calculations do patents and secrecy become an issue. Is it reasonable to take advantage of open source reasoning to get the positive results one desires to support the APC mechanism, only then to put non-open source structures back in to get it to actually work? Even worse, to ignore any deleterious consequences of that reinsertion? And is

[49] Actually, in Kremer there is no information structure across time or across players in the model, so, it is open source in a very default fashion.

it possible to compare with more open alternatives if the potential benefits of the more open approach have, by default, therefore been screened out *a priori*? And is it fair to criticise more open approaches when using a method that itself relied upon such methods to derive its positive results?

4.2.3. Multiple directions to be explored

Several directions can be explored – all of which reduce the measured benefit of the APC over alternative approaches:

1) Keep the stage games independent but have the probability distribution evolve over the development cycle in a known way, but ignore 'shocks' to the distribution;

2) Allow for projects to have externality effects on each other, so that the probability distribution underlying one project is a function of that of other projects. We would be especially interested in cases where projects are complementary, that is the movement 'outwards' of one project's probability distribution 'pulled' the distribution of others' with it;

3) Keep the stage games independent but have the probability distribution evolve over the product development cycle *and* allow potential technological 'shocks' in the distribution;

4) Allow for the probability at one stage to be a function of other stages (i.e. make probability – that is technology – endogenous) which we know is nearer the real-world situation. So, for example, if basic science deteriorates, allow the whole distribution in later periods to 'move inwards';

5) Keep the stage game and probability distributions independent but allow strategic interaction (at multiple levels explained below). The APC solution will clearly generate multiple equilibria, some less socially optimal than others.

In particular, one of the oddest things about the Kremer technology, and potentially one of its most misleading features, is that it contains no sunk or fixed costs at all, even though it is being used to analyse an industry that is very heavily dominated by these features. This assumption rules out important strategic affects, and the crucially important build up of capital costs over time (especially the strategic impact of this build up).

6) Explore how asymmetric information of players regarding other players (including regulators and governments) feeds into each player's strategy and into the APC value (note also that this introduces the issues of uncertainty and risk that must result in a higher cost APC, and, again, a non-unique value);

7) Combine all these features. We find that we get multiple equilibria with 'options' values as part of the cost of a typical APC, and potentially important discretionary elements, uncertainties, multiple institutional mechanisms, extra capital costs, and more.

In conclusion, the APC price is neither unique nor simple, except in the modelling scenarios favoured by Kremer. As soon as we drift even slightly away from the

simple APC approach, the APC price rises and side-effects start to appear that require a more general equilibrium approach.

5. TECHNOLOGY 1 – VARYING PROBABILITY AND EXTERNALITIES

5.1. 'Technology Paths' and the Difficulty of Achieving Optimal Research Intensity

Allowing probability to vary over the development cycle

We start with a very simple change to the underlying technology. We allow the instantaneous probability of discovery to vary over the product development cycle *for a given underlying technology* (i.e. probability is not 'averaged out' as in Kremer) and we allow projects to have externality effects on each other. Even with this simple change a host of new problems enter, every one of which leads to the APC being naturally *less efficient* than in the Kremer calculations, and the Kremer APC is therefore slower at arriving at a vaccine. The size of the Kremer APC is always just the lower bound in the size of the needed APC.

Discovery in science is, after all, cumulative. Newton once described scientific understanding as like "standing on the shoulders of giants" (though he was being rude[50]). There are long and indeterminate periods of very low, even zero, probability of discovery, followed by periods where probability is good, or even high – with a few projects perhaps 'pulling ahead of the field'. Even just considering drug trial attrition rates, fewer early trials go on to later stages of trials, whereas later trials have much lower attrition rates. Pre-trial studies have a similar characteristic. Even then the probability distribution may be highly uncertain[51]. In Kremer, all of this is replaced by the same probability distribution repeated over and over again.

If, for example, for a given underlying technology, early stages of vaccine development have very low probabilities of instantaneous success (even none at all), the APC mechanism has to ensure that enough firms do research at this stage; failure to do so holds back later research and slows the speed of vaccine discovery. Later, when there might be a much higher instantaneous probability of success for a given cost of research, the mechanism has to ensure that not 'too many' firms engage in research at this stage, and do not (wastefully) duplicate each other in an effort to get the vaccine. It might be argued that this is not so bad, since it will speed up vaccine development. But this is not so. The whole of the above is presaged on there being a certain budget to spend on vaccine R&D, so if intensity is 'too strong' with 'too many firms' at the late stage, the expected discounted profits for 'winners' is driven too low. *Ex ante*, worries about this will hold back those at early stages who are not able to earn back enough from the results they produce since they will not be able to sell those results at a high enough price to later firms who do not now expect to make a high enough profit from them. For a *given* APC budget, the presence (and the worry about the presence) of 'too many' firms at the end of the process will mean that *overall* research intensity is lower, and vaccines are delayed and more expensive.

[50] This was written in a letter to Hook, for whom he had intense dislike. Hook was also large and somewhat bent-double, so the complement was something of an insult too.

[51] If research has not been done on the area in previous periods there may be little prior knowledge to draw off to form probability distributions. The more 'new' the area the worse the problem.

Technology externality effects

In practice, externality effects between projects affect the probability distribution facing any firm and the collective probability distribution facing all. In Kremer, for example, the p* project, the marginal project, the project that is *just* profitable enough to do[52], is always in a sense the 'private-non-externality-optimal-probability-project' as it were (excuse the mouthful). However, once we allow for projects to have externality effects on each other then the Kremer-p* turns out to be set too high for some projects and set too low for others. Intuitively – allowing positive externalities – there will be projects that are themselves individually low probability projects that would never meet the Kremer-p* condition, but *would* get done if they were internalised with other projects that learn/feed off them. The correct non-Kremer-p* in these cases is lower than the Kremer-p*, meaning that research intensity should be higher for these activities compared to using the Kremer-p*, and vaccine development will be faster and cheaper[53]. Indeed, some *very low* individual p projects should, in equilibrium, get done.

Similarly, at other times, the p* of projects should be set higher than the Kremer-p* – that is intensity should be lower than a strict adherence to the Kremer-p* at that stage would suggest, in order to ensure dynamic consistency and incentives to do research at the stages where p*< Kremer-p* (there is a prisoners' dilemma and a danger of too much intensity at p* > Kremer-p* stages, which would feed back on and destroy the incentives to do research at stages where p* < Kremer-p*)

The more the APC fails in achieving correlation of the actual path of R&D with the path of theoretically 'optimal' R&D, the more it will reduce the *average* performance of the actual APC compared to the idealised APC – and the less favourable the APC mechanism compares, *ceteris paribus*, to alternatives. Observe, however, that failure to perfectly track has to feed through to a compensatingly *higher* APC price if vaccine development is not to be slowed[54]. The APC price under the simplified Kremer assumptions is always the very lower envelope of these possibilities.

A dynamic optimisation problem, but with asymmetric information

The problem at any given point in time is actually part of a generalised multi-period problem – a dynamic optimisation problem. A planner would look for the whole way through 'technology space', and, intuitively, would allow some low individual p projects since they would contribute to an overall high p project. The Kremer framework presumes (though the presumption is actually hidden in the fact that the problem is removed) that in essence developers can *coordinate* to pick off all the projects on the dynamic path through the 'technology space' even if some parts, on their own, would not look optimal.

But note that this ignores potentially important asymmetric information problems (assumed away by Kremer). In particular, since any one individual firm has low resources *vis à vis* the whole system, they only have access to a partial set of

[52] Higher values of p* mean that fewer projects are pursued; at the margin, the probability of success required from the marginal project has risen.

[53] Contrary to the Kremer calculations, where it would work out more expensive.

[54] The larger APC price ensures that, on average, in spite of all the waste, the intensity of actual R&D matches the theoretically optimal level of intensity even allowing for the lack of perfect tracking.

information (for the sake of ensuring efficiency of the auction mechanism described later, and to rule out a host of strategic behaviours, multiple firms must naturally be part of the APC set-up; each *has to be small* relative to the potential size of the vaccine market). Firms can never *individually* pick off the complete optimal path. Instead the system has to rely on efficiently-working markets – in particular financial and IPR markets – to enable many individuals to *collectively coordinate* to pick off the optimal path.

The averaging-of-probability idea of Kremer is therefore based on the notion that firms are *not* constrained in their access to information at any point in time or over time, and that they *can* internalise all the probability benefits and losses they impose on each other via a *complete* set of markets.

The only ways that this can be done are either through secrecy or through a heavily-enforced extremely comprehensive system of patents (meaning that the thing being patented could be described perfectly so that the contract could be 'complete'). Even in the latter case, it is doubtful the system would be tight enough to prevent leakages of information between players, or remove huge incentives towards secrecy. A whole range of situations could not be locked down in contracts. Let us look at the consequence of these both for the APC and for research in general.

5.2. Secrecy

One way to hold together a 'non-idealised' APC – as narrowly defined in the previous sub-section – is with secrecy. Firms must be able to keep secret all information from earlier stages of R&D to avoid competitors using it on later-stage activities without having themselves sunk early-stage costs[55]. In a competitive industry (the device used by Kremer, though highly unlikely to hold in practical applications of large early-stage vaccine APCs) early information discoverers will be at a competitive disadvantage compared to those who do *not* bother to invest in information discovery but get the information anyway. All competitive firms would seek to free-ride on those doing the early work – so *nobody* would do the early work (in fact those working to find 'early information' will not even be able to get the finance to do it). A classic prisoners' dilemma would infect the early stages of research and the overall costs of the APC would rise, compared to the idealised model. This would be further compounded if information spillovers between projects are an important part of the R&D process.

But secrecy undermines scientific progress[56], and has inefficiencies in a world where a piece of information leads to multiple possible paths that are worth exploring, as is the case in vaccine R&D. In a world of uncertainty and limited resources, under secrecy many routes go unexplored (especially if the industry is concentrated in just a few firms). There is a negative externality effect facing each firm; the firm knows that as it pursues one route, it increases the chances of that route being the one to succeed, but this reduces the usefulness of success on other routes[57]. The firm is in effect

[55] In the language above, firms working on p*<Kremer p* projects have to keep the results secret to stand the chance of being repaid their costs.
[56] Rosenberg, S.A. (1996), Blumenthal, D. et al. (1996), Benowitz, S. (1996).
[57] This is much the same way as in models of monopoly, one disincentive to R&D is that 'new' inventions replace markets for 'old' inventions of the same firm so that the private benefit to the latest

competing 'against itself'. Collectively, firms under-use the information compared to if the information were public knowledge.

Secrecy may also lead to more firms or fewer firms at various stages of the vaccine discovery process than would have otherwise been optimal in a non-secret framework. For example, if a firm wishing to do later-stage research cannot 'buy' from those at an earlier stage, the firm might have to do more early-stage research than it otherwise would have done. Allowing strategic behaviour back in to the framework, larger players may even have good strategic reasons not to sell (hence reveal) information developed at earlier stages, even if this behaviour is globally inefficient. This would tend to indicate that early-stage research problems might lead to a less-than-competitive industrial structure, contrary to that modelled in Kremer.

Secrecy is missing in the simple APC model

Secrecy and externality problems are ruled out in simple APC models because the reduced-form technological assumption assumes that the current stage game is just a repetition of the previous stage game, and there is, in a sense, only *one* route for each firm for each piece of information, and it is the same piece of information as last period's piece of information. Technology is reduced to what might be described as 'single-route technology'. Indeed Kremer quotes sources in support of the notion that what is needed to make vaccine research work are more "individually distinct" projects, suggesting projects with low or zero information spillover. At the same time there are so many individual firms present that they can effectively ignore negative externality effects anyway[58] (though this contradicts somewhat the insistence that it is large pharmaceutical firms that need to be encouraged to follow vaccine leads). Obviously, 'single-route technology' combined with the assumed perfect competition, would make the mechanism work better, but no evidence is provided to support the notion that this is a good description of the actual technology of vaccine development and not just wishful thinking. It is an empirical issue not an issue open to selective assumption to conveniently support a framework.

The damage caused by secrecy

Secrecy clearly works against efficiency in complicated processes like vaccine discovery – even more so if the industry has only a few players. Vaccines do not come out of multiple isolated programs working in secret from each other, but rather out of multiple approaches with sharing of information and 'synergy'. More open collaborative research methods have transparently obvious advantages in these situations. The 'more eyeballs' that can look at and explore a piece of information, the more possible ways will be found to use an early piece of information, and the more competition there is then the more leads will be explored, with under-use mitigated[59].

---

invention is smaller than would be the case in perfect competition. Perfectly competitive firms would ignore this 'replacement effect'.

[58] In the same way as in perfect competition, the marginal revenue curve comes to coincide with the average revenue curve.

[59] One thing to explore with more open collaborative research frameworks is whether they enable more competitive industrial structures with more research-active firms than a program based on APCs that tends (we find later) to narrow the number of players down compared to other mechanisms, for any given level of public funding.

Neither does secrecy sit well with the very justifications given for the efficiency of the APC – that of the power of financial markets to discipline researchers. Given the importance of financial markets in the Kremer framework (see Section 12), information cannot be kept secret from the outside world without leading to inefficient levels of R&D. Secrecy must then give way to a system of comprehensive and tight patents if financial discipline is to operate. Yet, we find that this cannot be done without negative side-effects, and that secrecy remains anyway, aggravating the job of finance.

We also find that secrecy creates a paradox for the mechanism. The nature of payment creates the incentive *not* to reveal information about the progress of the underlying endeavour, yet there is the notion that those running the mechanism are able to pool that knowledge to work out how the mechanism is performing. Creating a model where secrecy is ruled out, helps to expunge a range of important research guidance issues.

5.3. Tight Intellectual Property Rights

The alternative to secrecy[60] as a way to hold a 'non-idealised' APC together (under a perfectly competitive industry structure and assuming away strategic use of IPR for now) is an extremely tight and fully enforceable regime of patents on *all* the progeny flowing from *all* information. In a world of non-constant probability but with secrecy ruled out, all those doing early stage research *must* be able to get a 'fair' price by selling the results on to the next level of the research process. And that price is 'fair' if those working at later stages have ability to sell on to yet later stages at a 'fair' price, etc. In other words, fair price and optimal research incentives (and an 'efficient' price in the APC) are driven by a chain of tight property rights. This condition also means that if the originator of a piece of information plans to use the information for the next stage themselves, they can stop others from using it.

If there are multiple routes for exploring a piece of information but inability to internalise the benefits of any one route for the originator of the information via a complete set of property rights, then the developer of the piece of information might prefer that other firms *not* explore other routes while it explores some of the routes, since the success of others on their route has a negative externality effect on itself (this bites particularly badly if the originator has heavy sunk costs, as in vaccine research). If the original information cannot be patented (and, indeed licensed) it will be kept secret – though this still carries risks should it be revealed. As an example, if the research project of a firm has led to several compounds ready for trial, it might be socially efficient to run trials on all of them – in case of, for example, slightly better efficacy or fewer side-effects on one than another. But then there must be a complete set of patent-based contracts and efficient payments for all subsequent results. Otherwise, the originator will not view favourably other firms or researchers trialling one of the compounds at the same time as it trials another, and yet the originator themselves may be unwilling to trial more than one.

Tight IPR in a less than competitive world

[60] Though, we will see, it still does not rule secrecy out.

All of the problems in the above passages intensify as the industry structure becomes less competitive, since the firm faces even stronger 'replacement effects' from researching multiple routes flowing from a single piece of information. One possibility is that firms employ option-price investment logic, delaying exploring one route (trials on one compound), to see what happens first on another. This is privately optimal even if not socially optimal. If the industry is concentrated – and especially if the sunk costs of exploring any route are high – this may be even worse, since firms will hold off when they know that other firms are exploring routes, to see what happens for other firms first. Overall R&D intensity is lower, discovery takes longer, and the overall cost of vaccine research is higher. If an APC is established in such a setting, the price of the APC is bound to be constrained to be higher.

Again, Kremer rules this out by presuming, in effect, that no firm has more than one route to explore from any given piece of information – his 'single route' technology assumption – and perfect competition (again we observe the contradiction of a model that requires perfect competition to remove these problems but that then specifically targets incentives towards a few large firms that suffer especially from these problems).

5.3.1. Adverse selection of projects

Even if firms can sell the results of early research under a system of patents, those working on early-stage research projects may nevertheless have a better understanding of the true worth of a project, and may be unwilling (or unable) to sell the really valuable results. On average, therefore, the later buyers may only get the average of possible projects, that is they face an 'adverse selection' when buying the results of early research[61]. In consequence, more firms than is optimal may need to work on early stage research in order to overcome this problem for their later stage activities. Again, the APC price has to be set higher to compensate for this inefficiency.

5.3.2. Secrecy even under tight IPR

There still remain strong incentives to be secret even if patents are tight, in order to maintain proprietary claims and competitive advantage.

And it is not just positive information that there will be good economic incentives to hide; the information that a certain approach simply *does not work* has economic value too, if its revelation would encourage competitors to change the target of their investments in ways that will increase *their* chances of success and hence reduce one's own chance of making the first discovery and gaining the APC. The heavier the sunk investment costs so far, the greater the problem. Since this 'it didn't work' information is non-patentable, the only way its value can be kept from competitors is by secrecy.

The APC therefore still involves some important levels of secrecy even if IPR is ultra-tight. The more that vaccines require the cooperative sharing of information, the more this secrecy works against the efficiency of an APC mechanism targeted at that sort of activity.

[61] A precursor of this can be found in the Kremer APC modelling framework but with reference to publicly-financed research; he presumes the problem away for privately-financed research.

5.4. The Purposes and Consequences of Tight IPR – The Troublesome Trade-off for End-Loaded Mechanisms

If payment for R&D is via 'price times quantity' in the end market (either under the current patent-based system or under an APC extension of that system) this creates a trade-off for firms in the way IPR affects them.

On the one hand, tighter IPR helps. Under the 'non-idealised' technology just described, weak IPR generates inefficiencies and losses for firms both at early and late stages of research. This (just as with any other *ex ante* believed distortion) will have to be priced into the APC[62]. Stronger IPR helps to mitigate these inefficiencies and generate a lower APC price.

But tighter IPR has negative consequences – with implications both for the APC itself and for those activities not covered by the APC. If tighter patenting makes basic research and more complicated research (drawing on hundreds of patents) more difficult, it will feed higher research costs overall and create many new risks that will have to feed into the capital cost component of the APC[63]. Since the positions of firms' p distributions are a function of basic research, at each level of private activity the level of p for each firm will be lower. And problems with more complicated research show up in a general drifting of all firms' p distributions in the direction of lower values. Under both cases the flow cost of generating any given 'market enhancement', ΔV in Kremer, is higher – and the needed APC price is higher. If IPR itself becomes a strategic instrument (it will) this intensifies the problem even further (see Section 10).

This trade-off is a 'prisoners' dilemma', a co-ordination failure. In one respect, collectively all 'players' (pharmaceutical firms, financial investors, developed countries who have pledged, and developing countries who have signed contracts and deposited co-payments, etc.) would be better off from lower overall research costs if IPR was weaker. But if IPR were weaker, they would individually be better off if IPR were stronger. Since they cannot co-ordinate over the collective IPR problem, only see their private benefits, and do not internalise the externalities they impose on each other, they each push in the individually rational direction of tighter IPR[64]. Observe that, traditionally, so long as drugs prices are sufficiently flexible upwards, this shows up in higher drugs prices to compensate. The analogy here would be a higher APC price.

---

[62] I.e. The APC price has to be higher to ensure that research intensity is maintained even in spite of the inefficiencies and extra costs.

[63] See Farlow, A.W.K. 'Costs of Monopoly Pricing Under Patent Protection', especially slides 47-54 on 'genetic concerns', Columbia University, New York, December 2003, 'Access to Medicines and the Financing of Innovations in Health Care' conference: http://www.earthinstitute.columbia.edu/cgsd/accesstomedicines_papers.html. See also Section 10 below.

[64] The logic of this, incidentally, underlies much of the current push towards globally tighter IP – though aggravated by the fact that there are more obvious individual 'winners' and 'losers', such that even if the whole game is negative-sum, the winners make positive-sum payouts (This is not so clear at the country level however if the trade consequences are fully worked through; temporary country winners become long-term country losers, but, even within 'country losers', some firms may be long-term winners for much longer, and even indefinitely, and may heavily push for tighter IP).

5.4.1. The trade-off is damaging for complicated vaccine research

When adjusting this trade-off, we cannot say *a priori* if slightly tighter IPR will make the APC price slightly higher or lower. If the benefits of tightening offset the deleterious consequences, then the APC price will fall. If the deleterious consequences offset the benefits, then the APC price will rise. We *can* say several things however about the *likely* affects on the level of the APC price[65].

First, that costs are *always* higher than if these problems did not exist.

Second, that a great deal of evidence suggests that the deleterious consequences in the case of vaccines is very high and is therefore more likely to swamp the beneficial effects.

Third, that the problem *always* intensifies as the technology becomes more complicated, so that the APC price *always* rises with the complications of technology. The flip-side to this is that more open research methods might appeal more as technology becomes *more* complicated. The reader may have spotted that as the technology gets simpler, both the APC and the more open mechanisms work better, and that as technology gets more complicated both mechanisms find it more difficult to perform efficiently. The question is whether they become less efficient at the same or different rates. The open collaborative approach does not need to do perfectly as technology gets more complicated; it 'only' needs to do 'well enough' to 'beat' the now higher APC price. Intuitively, one would imagine that the more open approach could be made to work relatively better in the complicated science setting. One further possibility is that the high level of sunk costs – and consequent rôle of capital costs – might push to very high levels of tough IP in the case of an APC and that this affects the relative ability of an APC to perform compared to more open methods.

Fourth, that as more costs are sunk (or, indeed, in order to create the environment where greater levels of costs *can* be sunk) the incentive to be tough on IPR intensifies in order to keep overall capital costs down[66] (as will become clearer in later sections) even if the deleterious effects of this are intensifying. This is another reason why APCs may not be so wise if technology goes through long gestation periods in an industry characterised by high levels of sunk costs and high-cost forms of finance; at some point the sunk costs are so large, that highly negative side-effects of tight IPR might be tolerated that would not have been tolerated in a world with much lower sunk costs.

Fifthly, that even if the trade-off is positive (unlikely) for an individual APC-covered vaccine or drug, the negative effects for non-APC-covered vaccines will not be offset by any positive benefit, so that the aggregate effect for all vaccines and drugs is deleterious. The paradox is that whilst enforcing the APC to keep individual firms' costs down, the potential inefficiencies of very tight IP would have to be imposed, raising research costs for all, and yet, the APC cannot be targeted at the difficult technology that is becoming more expensive to do in a world of tight IP. Why aggravate the latter in order to support the former (especially if there is an

[65] All this reasoning can be translated into insights about drugs prices in general.
[66] This is a prisoners' dilemma too, since in one direction capital costs are being held down, while in another direction, all the extra risks are pushing capital costs up.

alternative)? When calculating the costs of the former, why do so with no concern for the potentially deleterious costs of the latter?

The Kremer technological assumptions strip patents and all IP issues out of the system on all but the end-products. Statements about the constant state of science turn out, in part, therefore to be statements about the non-affect of IP on the state of science or on the ability to do research (and the suggestion that even if this problem pushes up prices in the drugs industry in general, it will not do so in the APC-financed sector). This should be checked and not assumed.

This discussion suggests that APCs tend to work better the more simple the science, but that more open collaborative approaches can be made to work better the more difficult the science. It does not greatly help therefore to model, as all of the APC papers do, with only very simple science. This also suggests that if one of the reasons why so little private research takes place for HIV and malaria is because of the prevalence of the 'wrong model', the non-open, non-collaborative model, simply making the size of funds available to the wrong model – and indeed reinforcing it – does not help.

5.4.2. Pressures for tighter IPR

Pressure for tighter IPR comes from various quarters under an APC. All 'players' must believe that those running the APC will not renege on their promises. They must all believe that the whole patent environment is supportive of the agreement. As the APC progresses and more and more investments are sunk, this incentive to be tough on property rights will get ever stronger. This is analogous to the way a country that has built up huge debts has to provide credible demonstration of its intent not to default on those debts, even if this demonstration can lead to welfare losses.

This will also necessitate tough measures such as the clamping down on compulsory licensing in general, including on other drugs, so as not to 'frighten' financial markets and pharmaceutical firms into thinking that it might happen on the drugs covered by the APC. Since *all* risk has to be priced into the eventual APC price – via capital costs – the authorities can correctly argue that if the APC price is not to be overly-high and vaccine discovery not overly-slow, then they *have* to be tough in other areas of drug pricing and procurement. Authorities will argue, rightly, that those countries relying on APCs for their vaccines will suffer on their vaccine component of healthcare if compulsory licensing is allowed elsewhere. Furthermore, the bargaining powers of the various sides during TRIPS negotiations will be tipped further against those arguing for compulsory licensing. The costs of suppressed compulsory licensing (based on a counterfactual of what the compulsory licensing would have been without the APCs in place) should be part of the costs of the APC approach.

Should an auction mechanism be used to set the APC price (presuming it works[67]), then all distortions would be priced into the APC price that firms would be willing to accept – to avoid slowing the intensity of vaccine research.

One of the reasons that Kremer models in a way that removes all of these patent-based problems *a priori*, is that he presumes (see the quote in Section 1) that ever-

[67] The section on auctions below doubts this.

tighter patents are not an issue, so simplification by removing them from the framework in the first place is not an issue. However, in a proper assessment of alternative mechanisms, the possible distortionary affects of patents should not be ruled out *a priori*. This is the source of one of the claims made for the open collaborative research model – that if it could avoid some or all of these distortions then it would be a cheaper, more efficient, mechanism to raise finance for investment in vaccine development than that of an APC.

5.5. APCs for Simple Technology Only

This discussion of property rights suggests that the APC is inappropriate to stimulate all but very specific kinds of research. This is accepted by the proponents of the APC. Since those at later stages of research have to buy at a fair price the technology developed from earlier stages, strong property rights imply the need for some clarity as to what a particular piece of information is worth. The APC mechanism does not work well in environments where technology is excessively complicated and/or the market value of a piece of information is hard to assess or cannot be priced due to its public good nature [not just cases where the value is *never* assessable but cases where the value is not assessable *at the point in time when an investment decision must be made* (even if subsequently it proves easier to assess its worth)]. This is why the APC is no good for basic research, or research with large positive non-marketable externalities. If this early research is not very 'own-able', the APC mechanism will not encourage it. Though the proponents of APCs argue that it is all about creating private incentives for research, it will only, at best, do so at the later, fairly un-technical, stages of research.

6. TECHNOLOGY 2 – TECHNOLOGY SHOCKS

The reader may prefer to skim-read this section and jump to section 7. This section is rather long and somewhat tongue-twisted precisely because of the contortions needed to get APCs to work in anything but limited technological settings. It also points up the drastic removal of important features by Kremer in his calculations. This makes for somewhat laborious, and not particularly exciting, reading.

As with the last section, the aim is to set up a framework for the relative evaluation of different mechanisms on the basis of how they cope or fail to cope with the same set of technological difficulties. We find many more reasons in this section for concluding that the size of the Kremer APC is the lower bound of what an APC would actually cost, and many more reasons for concluding that a given APC would be slower than claimed at arriving at a vaccine.

It might be that all of the following section can be ignored if APCs really are targeted only at the easy later stages of simple technology. But, first, the calibration exercise performed by Kremer (and all the press statements) is not presuming quite such a limited rôle for APCs. Second, it is not clear – certainly in the case of complicated vaccines like those for HIV/AIDS – that even limiting APCs to late-stage research will resolve many of the most intransient difficulties. Third, even if the rôle for APCs was that limited, it would again raise the issue of how to deal with all those parts of the vaccine development process that are not simple science and are not being covered by APCs but that might be made more difficult as a result of some of the side-effects of APCs. The technology shocks covered in this section simply make such problems that much more difficult to resolve.

6.1. Uncertainty and Technology Shocks

The second big way in which the technology of vaccine development may be very different from that used in the simple APC models is that not only might the instantaneous probability of discovery vary over the R&D process, but the whole structure of instantaneous probability might shift due to a technological 'shock' such as a biotechnology revolution (such as the discovery of the malaria genome). There has been a tendency in the past for biotechnology to go through such periods of technological 'shock', with the technology of discovery turning out to be much more or much less productive than had initially been presumed.

How is the 'optimal' APC price to be set in such an environment? Though there might be very great uncertainties about the state of current and future technologies, a simple APC would have to be set at a point in time before much of this uncertainty had been resolved and information revealed. For the classic case often used to illustrate a prize fund – that of a method for determining longitude – at the time the prize was set, most scientists believed that longitude would best be determined by astronomical observation. It turned out to be through the development of a sufficiently accurate clock. As Kremer states: "Pre-specifying an astronomical solution would have been a mistake."[68]

[68] K7:40.

This problem is recognised by the proponents of the APC, even though they base all calculations on a constant state of technology. Kremer argues, for example, that it might be sensible to pre-announce that if by some date a vaccine had not been developed that the price would then start to rise, but that it is "probably better to let future decision makers choose whether or not to increase the price, since in some scenarios it would be optimal not to increase price. For example, there would be no need to increase price if the general technological advances in biology reduced the expected cost of developing a vaccine sufficiently that many firms would pursue vaccines."[69] Of course, this could well necessitate that the APC price be *reduced*[70], but Kremer does not discuss that; it is, after all, much easier to raise than it is to lower APC prices. It is also left deliberately vague (even though this is an added risk for firms) because it is difficult to visualise cases where price rises can be expected in advance in a way that does not itself have the deleterious consequence of putting off early risky investment.

It is extremely difficult to design an APC to efficiently cope with technological 'shocks'. The choices are:

1) A 'fixed price' APC;

2) A flexible APC with either a) a complicated set of deterministic rules allowing the APC price to rise or to fall in response to various contingencies, or b) a discretionary mechanism, allowing for the *ex post* adjustment of price – with suitable institutions and rules to govern this;

3) A set of 'staging point' APCs, with varieties of rules and discretion.

Under this categorisation, the simple APC so far used in the APC literature is a 'fixed price' APC. However, since the technology in that literature is (strictly) stationary the 'fixed price' nature of the APC does not raise any of the problems we will now investigate. In reality technology would lead to difficult tradeoffs, that would almost certainly lead to a more 'flexible price' APC, with a knotty set of strategic and institutional problems, and much higher capital costs.

6.2. Fixed APCs

If the payment offered in the APC is not allowed to vary, and yet technology is known to suffer shocks, then these technological shocks must somehow be predicted – as well as the levels of private and public investment that will be needed in response to them, and presumptions about policymakers' abilities at offsetting them (this will be better explained below in the section on public-funding) – and incorporated into the APC price *ex ante*.

[69] K7:38.

[70] And, of course, by failing to clarify the *rules* for the price increase – but instead leaving it to "future decision makers" to choose "whether or not to increase the price" – private capital costs would have to rise to compensate for the potential dynamic inconsistencies, leading to a higher APC price in the first place.

What would happen if the mechanism to set the fixed APC price were to under- or over-estimate future technological improvements? If it overestimates and sets the APC too low, then at some point research will stop, and no vaccine will be developed. If it underestimates, then the APC price will be set too high; this particular vaccine is developed at some point, but at greater loss in terms of other vaccines and treatments compared to an alternative method of finance that might have been able to avoid overpaying (since the budget remaining for these other vaccines and treatments is lower than it would have been had the APC price been set correctly)[71]. This is what we meant above by a comprehensive solution rather than a partial solution to the problem. The open collaborative approach immediately suggests itself as a possibly more flexible approach to this particular problem.

The HIV/AIDS debate

To give a flavour of the problem, consider the debate regarding the costs of vaccine development for HIV/AIDS. DFID points out[72] that:

"*It is too early to know how much an AIDS vaccine will cost...Most of the AIDS vaccines now in development are designed using advanced technology that may mean prices are higher than for older vaccines. Still, some AIDS vaccines in development could be quite economical to produce. For example, a vaccine being developed by the Institute of Human Virology in Baltimore, Maryland, USA and funded by the International AIDS Vaccine Initiative (IAVI) would harness weakened Salmonella bacteria to deliver genetic material. Bacteria based vaccines are expected to be comparatively inexpensive to manufacture*."

[71] We observe that, ordinarily, when private firms contract to supply services or goods to a government at a fixed price, then the government will subsequently turn out to have paid 'too much' or 'too little' (though usually there are terms in contracts to allow for unforeseen circumstance, which are the analogy here of the flexible APC), but the firm would still be *contractually obliged* in both circumstances *to provide* the promised services or goods. Under the fixed price contract, all the risks fall onto the company (and onto financial markets where the risks are, in theory, diversified away). Even then, if the risks are great, it may turn out 'expensive' for the firms (in financial contracting costs) to operate under a fixed contract, but this will be passed on to the government in the contract price. The government operating on a fixed price contract is, in a sense 'insured', and pays an 'insurance' premium as part of the price. Again, if the risks are great, the premium might have to be large. The setting of the price and the premium require some knowledge of the distributions of possible outcomes (in analogy here to the need of those setting up the APC to have some notion of what the technological possibilities are). The more risky the technology, then the higher the premium. If the government is less risk averse than the private sector or (much the same thing) has much better access to credit markets, then it may make *more* sense for the government to bear the risks than for the private firms to bear the risks (in much the same analogy to the way, under an APC, firms might rather prefer the less risk-averse government to bear the risks). The justification for forcing risk-bearing onto the private firms however is the usual requirement to create incentives (especially if there is asymmetric information), in this case probably the incentive to produce the goods or services cheaply (with plenty of contract terms to make sure that the quality is not sacrificed). Observe how usually there would have been some sort of bidding process before the contract are signed and before the firm invests to satisfy the contract. Under an APC, things are slightly different, but in a way that has very significant implications. All these risks and 'premiums' have to be considered, but, since there are no 'contracts' with private firms until a vaccine is developed, firms are not obliged to keep going with an inefficient contract and always have the option to pull out, and they also always have to worry whether those operating the other side of the 'implicit' contract will renege. Firms also have to engage in sunk expenditure to try to win the contract; there is *no* competitive stage *before* investments are sunk. The scheme has to be adapted to avoid these eventualities – at a cost. And the costs are higher the more risky the technology.

[72] DFID web site.

And with reference to a further issue that will impinge on costs:

"*The question of whether AIDS vaccines will have to be designed to fit particular HIV clades predominant in different countries remains unanswered*." Indeed, we still do not know if a vaccine designed for one particular HIV clade would prove cross-reactive against other clades.

Yet, if fifteen or twenty years ago an APC had been set up for an HIV/AIDS vaccine, with a fixed payment, and if firms had under- or over-estimated the technology of discovery, this would have fed into the APC price set then. If the genome revolution that has taken place since then and the invention of rapid sequencing machines and many other scientific discoveries and inventions had speeded up the discovery process and reduced the costs of discovery (including costs related to the risks of failure) compared to what had been expected, this would necessitate a lowering of the APC price (if we could presume that the concomitant explosion in gene patents had not harmed the costs of research, which we know it has).

There is much talk of a genomic revolution of which we are only now getting the first fruits. If the APC terms today are set on the basis of the current technology, then they would turn out to be way more expensive than a payment scheme that is allowed to adapt to the technology. The optimal 'fixed price' APC would seem to require setting the APC value lower today than would be dictated by current technology, in the expectation that technological improvements will make it optimal. And all players would need to believe this, and continue investing under today's technological structure, rest-assured that in later periods their seemingly unprofitable current investments will turn out to be profitable. But for such a fixed APC program to work, we would need to know *now* where this technology is going in the future, and price it in to the APC, and then rely on tight patents to encourage early-period R&D.

Conversely, what would have happened if firms (indeed, investors) had overestimated the ease with which a vaccine could be developed, and had accepted a low APC price[73], but had subsequently realised how tough it was going to be? If the fixed payment was immoveable, research efforts would have ground to a halt. Early researchers would have been stung – and might by now have mounted legal challenges to a mechanism that had failed to price correctly the true costs of developing a vaccine.

Going to Mars

We can think of this further with a parable of what would happen if a fixed-price APC was used to stimulate investment towards getting a man or woman to Mars and safely back. If, subsequently, the technology turned out to be much more difficult than had originally been thought when the APC was priced, the endeavour would collapse[74]. This may not be totally irrational; cutting the project before it involves escalating waste of resources may make economic sense (even then it would not be a *costless* waste of resources, since shareholders – many of them ordinary citizens via their pensions and investments – would have to shoulder the loss; whether they pay it

[73] We presume – though it is wrong to do so – that an auction mechanism could have optimally derived this. More on this issue in Section 11 below.

[74] The corny pun I was going to use was "would not get off the ground". Also refer to an earlier sub-note as to why the APC would not be binding on private firms to provide the service.

through taxes or through loss on their equity holdings is irrelevant). But for a vaccine, this may be deemed unethical and would be ruled out. And it is certainly not a result that we would feel happy about.

Observe also that foreknowledge of the potential deterioration in technology would have to be priced into the Mars APC, and if we wished to make sure that a man or woman would get to Mars and back 'come what may', then a fixed APC might not be the most sensible mechanism to use, since the price might have to be set impossibly high at the start. Maybe this is why the US administration in recently announcing its plan to get earthlings to Mars did not suggest using an APC to pay for it?

6.2.1. A technology option must be priced into a fixed APC

A 'fixed price' APC is not very useful, since, although it is a guaranteed liability on developed government finances, its very lack of contingency means that it cannot respond well to technological 'shocks'. The price is *bound* to be higher that it would be without the presence of shocks (Here we just look at technology, but in later sections we will see that all kinds of shocks will feed the mechanism we describe).

On the one hand, price might be set high to avoid the risk of bad shocks causing collapse of the APC. On the other hand, if the risk of collapse is accepted, and an attempt is made to set the APC price 'low', then the extra financial risk facing those investing towards an APC will force up their private capital costs, which will force the setting of a higher APC price to compensate anyway[75]. Either way, the APC price has to be set higher, with the degree of 'extra price' a positive function of the degree of shocks in the system.

In fact, if the APC price is not allowed to be *ex post* flexible, and if the mechanism is not to risk producing *no* vaccine at all, the price set in-advance has to cater for the worst-case scenario. The more uncertain the technology, the more the mechanism has to err on the side of being overly-generous. Intuitively, in the 'simple' APC setting we do not wish at any point for the APC price to turn out to be so low as to prevent the continuation of the game to the next stage. The Kremer fixed-price APC excludes this extra option-based cost by ruling it out technologically. Again, the Kremer ultra-simple technology assumption generates the lower bound of the set of all prices that might come out.

This logic will push in the direction of a more flexible APC if technology is liable to experience 'shocks'. Meanwhile we will explore the remaining way open to policymakers to respond to shocks under a 'fixed price' APC – the use of other publicly-funded research.

[75] Alternatively, the intensity of R&D effort is even lower at the 'low' APC price once this risk is factored in. If the same intensity of effort is to be targeted, so as not to delay the development of the vaccine, then the APC has to be set higher to offset the risk. Of course, as the APC price rises it also reduces the chances of hitting the bad outcome, so the extra part to compensate for the risks tapers out. Repeatedly, throughout this paper, we find that the simple technology methodology of Kremer has enabled lots of complicated risk-related aspects to be ignored. 'Technology risk' and 'mechanism risk' cannot simply be ignored.

6.2.2. Where do technology shocks fall and how are they offset in a fixed APC?

How the mechanism deals with technology shocks under a fixed APC also depends on *where* the technology 'shock' falls (or, indeed, is likely to fall). We have already seen that the APC is not designed to be directed at basic research; this is left to publicly-funded research. We will see, in Section 8 below, that each APC, if it is to match the efficiency suggested by Kremer, must be based on an exact, flexible, *ex ante,* contracted proportion of non-APC publicly-funded research. Whether the technology 'shock' falls in the private or public part of the technology may therefore determine the mechanism's ability to cope with the shock and achieve efficient vaccine discovery. Again, since non-APC parts of the whole process are stripped out in the Kremer model, this issue does not arise there. For now, let's go through the possibilities (and this is where the comment above about 'contortions needed to make the mechanism work' bites especially, and that some readers may simply wish to skip assured that this section explains even more of the difficulties when the assumption of extremely simple science is removed).

6.2.2.1. Technology shocks in the privately-funded APC-covered sector – the use of a public sector 'correction facility'

If the technology shock falls in the later part of the research process being paid for through an APC, and if the APC 'price' is immoveable, then the only way the shock can be mitigated is by a counter-investment in the non-APC-funded front-end part of the process in an attempt to push returns in the private back-end part of the process back to the level required for the APC to work.

If the shock is negative, and if the publicly-funded[76] part of the research process was set at the efficient level before the shock, then the public counter-investment would have to encroach (inefficiently) onto otherwise private sector activity.

Quite how the public-sector counter-investment might look if the shock to the private sector is positive is not clear; if the publicly-funded part of the process was originally set efficiently, this would seem to suggest that it would require an *inefficient* cutting back on the publicly-funded part of the process.

In both scenarios, the use of a public sector 'correction facility' is always inefficient. However, if publicly-funded research is as (extremely) inefficient as Kremer argues, then the costs of vaccine development have to be correspondingly *even* higher using a fixed price APC with a 'correction facility', such that it may be a very expensive way to conduct research into vaccines (observe how, as the volatility of the technology increases, this observation strengthens).

To avoid using such an expensive 'correction facility', the APC designer might design an APC contract that involves lower expected use of the 'correction facility'. This would be done by setting an even higher APC price than would be strictly 'optimal' (i.e. higher than what it would have been without the wasteful public research). The higher the APC price, the less often the mechanism will have recourse to the publicly-financed 'correction facility'. The optimal fixed APC price would then depend on the

[76] Throughout, this means the publicly-funded non-APC part. After all, the APC is also publicly-funded.

relative cost-benefit analysis of accepting a higher APC price in the first place versus the expected costs of recourse to the 'correction facility'. Clearly, if publicly-financed research is as inefficient as Kremer argues, the APC price should be set very high to avoid recourse to it. But this is very wasteful too[77]. Once again the Kremer APC price turns out to be only a lower bound on the eventual costs of the program, paradoxically, the more this is so the more Kremer is right about public sector inefficiencies. And again it raises the question of why attempts would not be made to improve the efficiency of the non-APC-covered parts of the research process before first instigating the APC.

Unfortunately, the setting of the fixed APC becomes intractable if the underlying technology is truly uncertain[78] since the optimal APC (and its concomitant government commitment) cannot be determined.

6.2.2.2. Technology shocks to the non-APC-covered sector

Technology shocks to the ***non-APC-covered part*** of the technological process must also be offset to avoid passing negative costs and positive benefits on to the APC-covered part of the process; the first runs the risk of not getting the vaccine but also raises risks of developers, while the latter leads to vaccine costs being higher than is efficient.

A negative shock would have to be mitigated by an expansion in publicly-funded investment (again, with this reaction contracted in advance). Failure to do so, and the value of the private part of the process would be reduced[79].

If the shock is positive, however, there would seem to be an argument that the publicly-funded part of the process should be cut so as to not make the reward to the APC-covered private part of the process over-generous. But, if the publicly-funded part of the process constitutes that part of the process containing the positive information externalities, then this is a much less efficient approach than simply cutting the APC payment. Intuitively, the positive shock to the publicly-funded early part of the process risks passing positive benefits onto the privately-funded APC-covered part of the process for which the private parts had made no contribution. Tax payers face a welfare loss since they would rather like to internalise this positive shock into lower tax payments via a lower APC price but are unable to do so. Once again the Kremer fixed payment is only a lower bound.

---

[77] Thinking of the technology for a moment, this is also clearly wasteful. Given technological externalities and the 'public good' aspects of 'basic research', there would be an optimal level of public and an optimal level of private R&D. This higher APC price, to avoid recourse to the 'correction facility', must be presuming that those in the private sector are being encouraged to work on things that would ordinarily be done in the public sector. Since it is expensive for the private sector to do this (they cannot internalise the benefits of the research) using this mechanism to avoid use of the 'correction facility' must itself be expensive.

[78] In the Knightian sense.

[79] Observe how the shock might well necessitate a change in the distribution of activity, with (the usual case) more activity in the private sector. Since the APC cannot rise to encourage this, even more activity has to be done in the public sector to compensate.

6.2.2.3. Rules to prevent corruption, and the problems of co-ordination, dynamic inconsistency, and capture

All of these four eventualities (positive and negative technology shocks to publicly-funded and privately-funded APC-covered parts of the technology), combinations of such shocks, and the public sector's response to them, would – under a fixed-price APC – have to be built into the *ex ante* terms of the APC and the 'contracts' of the publicly-funded parts of the process. The commitment to do this, and the exact rules dictating when publicly-funded actions are triggered, and who should carry them out, would require clear (non-asymmetric) information, yet another layer of institutional detail, and global co-ordination. To the degree that this is not understood, can't be exactly done (for example, if all contingencies cannot be covered in contracts), and can't credibly be contracted in advance, it raises the uncertainties of the private players, raising their capital costs and vaccine development costs, and hence the needed APC 'price' *ex ante*.

And it creates a new dynamic inconsistency problem emanating from the way publicly-funded parts of the R&D process are being used to help support privately-funded parts of the R&D process. Over time, privately-funded research costs are sunk, and the private sector must worry that the publicly-funded adjustments required in response to shocks (whether private or public) would only take in to account the incentives to create sufficient private sector investment in the 'continuation game', i.e. incentives for investments from that point forward only. Worry about this dynamic inconsistency will, *ex ante*, agitate against sufficient private sector investment in early periods of research. Again, the APC price has to rise to compensate for this. The Kremer APC price turns out, yet again, to be only a lower bound.[80]

We observe that an APC is part of a combined package with a publicly-funded program, but it looks very different from the Kremer model that de-emphasises totally publicly-funded research.

6.2.2.4. Multiple APCs, reputation, and credibility

This propels us in the direction of 'rules and institutions', something that will creep up again and again in what follows, even though it is not at all apparent in the literature accompanying the simple APCs.

It also suggests that multiple APCs might help create badly-needed credibility. Since 'reneging' on one APC harms reputation on other APCs, it is avoided. However, this is only useful if technology shocks to one APC are asynchronous to the shocks to other APCs. If technology shocks were synchronous across all APCs at a late point in a multiple vaccine development program, then there might be incentive to 'renege' in part on all or some APCs, in the sense of failing to adjust the publicly-financed part of the technology in a way that takes into account all previously sunk private investments (or, at least, market participants might worry that this would be the case and price this into capital costs ex ante). Observe how 'reneging' is not being done by the APC setters, but by public bodies outside of the APC setters. In a world of shocks and a fixed APC there is a need for credible, water-tight, non-APC contracts on non-APC players too. Failure to achieve these contracts will result in higher risks and

[80] This is getting to be a boring refrain. Just add this comment to the end of every section below and it won't be far off the mark.

capital costs to those working under APCs (none of these notions is mentioned in the APC literature).

This suggests *not* to have too many APCs initially if they are likely to be hit by synchronous technology shocks, so that those 'involved' in early APCs have an incentive to react to technology shocks in ways that create reputation for later, not yet initiated, APCs. The problems with this however are that: this militates against the APC mechanism as a *general* solution to many vaccines *at the same time*; it involves delay in action on some vaccines or drugs, which has to be priced into the social costs of using the mechanism; it does not work if the mechanism for each APC is run independent of every other – maybe via different committees – so that this reputational argument cannot be run over from one APC to another; it still does not satisfactorily explain if and how co-ordination of non-APC players is achieved.

Furthermore, worries about other APCs might lead to those setting the terms of the first APCs (if they had not been set by auction) to set the price too high in order to avoid an early failure; early APCs have to secure good returns to ensure takers for later APCs[81].

6.2.3. Financial options logic to early developers in a fixed APC

Note also that there is option logic to private investment here too, that dictates that 'early investors' hold back investments to see whether the technology is capable of producing the result, what the shocks are likely to be, and how the publicly-funded parts of the process will react. If the APC price is being set in an auction mechanism, this would distort the APC price upwards (in a sense, the auction would get it wrong since it would read information generated by option-based investment decisions as if it were based on non-option based decisions). Repeatedly we find that the price in the simple APC presented to governments and international organisations is only the lower bound to the price to be expected.

6.3. Flexible APCs

The alternative to having a fixed APC with an 'option' element built in and an (inefficient) publicly-financed 'correction facility' is to have a flexible APC, that allows for adjustment of the APC payment *ex post* in the light of technology shocks. The choices are:
1) A flexible APC price and quantity, but based on pre-agreed rules and no element of *ex post* discretion;
2) A flexible APC with *ex post* discretion over price and quantity.

In Section 7, we will see that APCs also need flexibility built in to the quantities and prices of purchases of vaccines in order to generate incentives to create high quality rather than low quality vaccines. Here we concentrate on technology reasons for

[81] This is much the same as what happens when early privatisations and IPOs are set at terms that are excessively generous to initial investors, so as to ensure that there will be takers for future privatisations and IPOs. Early 'failed' privatisations/IPOs, reduce the potential price of later privatizations/IPOs. Low-priced IPOs were also just another element of the corruption playing out on Wall Street in the 1990s (See Stiglitz, J. ibid., plus many others).

flexibility, but all the same reasoning feeds through to the problems covered in Section 7.

6.3.1. Flexible APCs with fixed non-discretionary rules

If the first route is followed – with a flexible APC with non-discretionary rules – the APC mechanism will need a set of precise rules for how the APC payment will adapt to changing technology. And this set of (potentially elaborate) rules will need to be fixed *ex ante* in all contracts, so as not to cause distortions, and be set in such a way that contracts are utterly credible. As before – and as always – any uncertainty about how firms will be treated *ex post*, will factor into higher private capital costs, and a higher APC price *ex ante*, higher vaccine development costs and slower vaccine development.

However, the setting of such rules in a non-discretionary flexible system will require:
1) A very good sense *now* of the underlying science *possibilities* in future periods – even if the science is very uncertain – so that all possibilities can be covered in the terms of the contract, and such that private investors can feel certain that all technological possibilities have been covered in tight legally binding language;
2) Good dispute resolution procedures to avoid the danger of expensive litigation;

However:
i) Nevertheless there will inevitably be inability to specify all contingencies in contracts, either because it is too costly to do so, or because not all eventualities will be foreseen;
ii) There will be potentially large transactions costs;
iii) Even if these terms are all included, it is still not clear, given the expense of research and the costs of contracting, that the terms will be credible;
iv) There is still the risk of *ex post* strategic bargaining games over terms, especially given the sunk nature of many costs;
v) There are risks of setting terms in an inefficient and distortionary fashion, giving investors and researchers perverse incentives;
vi) The APC price will need to be set, *ex ante*, higher to compensate for all this.

6.3.2. Flexible APCs with discretion

If not all eventualities are contractible, or if the state of science is poorly understood, there may be some sense in allowing discretionary *ex post* adjustments to the APC price via various committees of 'experts', rather than rules.

Nevertheless rules for doing this could, for example, be based on principles that do not, *ex ante*, require technological knowledge. For example, a specified action regarding the APC price could be based on the assessment of an independent group of scientists that a particular form of discovery had reduced the costs of a particular process by a certain amount.

It is argued by APC advocates that specifications for a vaccine might be relatively easy to settle in advance, hence that discretion can be reduced somewhat. However, it is argued here that once technology moves away from the simple technology of Kremer (with its unique solution), and once the time profile of R&D and problems with information and intellectual property along the path to discovery is properly considered, and once technology shocks are taken seriously (such that information is

needed on technology in order to optimally adjust the APC price) the emphasis in flexible APCs shifts away from the ability to describe the end product to the ability to describe intermediate products and technological processes. It is misleading to suggest, as Kremer does, that simply specifying the end product is sufficient to drive an efficient result. Even if the end product might be describable in advance (while it is claimed that this is relatively easy for vaccines, we will see below that it is somewhat more problematic), many intermediate products and processes are *not describable in advance* (the consequences of this disappear in the Kremer calculations since it is ruled out in the technology assumptions at the start).

6.3.3. Flexible APCs, dynamic inconsistency, and the ratchet-effect

There is anyway – just as with fixed APCs and publicly-funded 'correction facilities' – a new form of dynamic inconsistency creeping in. At all points in time, research costs in previous periods are sunk. If previous research has raised the probabilities of getting a result from current research, there is always an *ex post* incentive (the incentive increases with the stage of development) to take advantage of this. But the danger of the misuse of discretion will raise the risks of research, capital costs, and the APC price.

Even if the APC setters do not seek, *ex post*, to take advantage of sunk R&D, nevertheless if technology turns out better than expected, and the APC setters exercise discretion by reducing the APC payment, those who have sunk research costs may mount legal challenges. Legal challenges, after all, are a form of rent-seeking, and are likely to rise with the level of previous sunk investments. And since nobody mounts legal challenges if the APC payment is raised – and since there are higher political costs to cutting the APC *ex post* than raising it – there is a ratchet effect, that is a bias towards a higher expected APC price on average. As in any situation involving a regulator, the bias is towards favouring the regulated, who are able to spend up to the value of the sought-after economic rent in fighting a financially constrained regulator[82]

The APC model, it turns out, seeks to solve one time inconsistency problem by replacing it with many others. The original problem targeted by the APC was the risk that buyers – governments and large organisations like the World Bank – would bid down the value of research by bidding down the price of vaccines once they were developed. In the new problem, *with the same set of players*, as the mechanism progresses, the risk is that governments will fail to adjust public-research in a dynamically consistent fashion and that the power of large players might be brought to bear on the APC regulator to bid down the value of research by lowering the APC price after the costs of research have been sunk. This again would necessitate, *ex ante*, a more generous APC to cover the higher capital costs thus generated. This all suggests once again the importance of very tight contracting – but this time it is *outside of the APC itself* – and global agreements, maybe even treaties.

[82] And the regulator is unable to internalise the benefits of good regulatory actions, giving them an extra disadvantage (Any benefit the firm gets from rent-seeking shows up in its share price which reflects well on *it*. No similar mechanism exists to 'reflect well' on regulators).

One way to help weaken this type of time inconsistency problem is to tighten IPR to strengthen the hand of those investing in R&D and facing the APC 'regulator' (or just capturing the 'regulator')

6.3.4. Flexible APCs and the difficulty of improving technology

To complicate matters, we need to consider the dynamics of how APC payments affect the development of *new* technology. This does not arise in the Kremer model since the state of technology is given for all time. A discretionary mechanism may harm the process of creating new technology. Realising that the payments may be 'less generous' if the technology gets 'easier', the APC setter has to be careful not to create distortions that lead firms towards *not* investing in improving the technology in ways that would make it 'easier' for future generations of researchers to develop vaccines (not just this vaccine but other vaccines and drugs too)[83]. It is likely difficult to separate changes in technology caused by those being regulated by an APC from those changes caused by factors outside of their control. The more 'noisy' the technology, the harder it will be. What the exact consequences of this are is unclear, though in all reasonable cases the APC price is higher. If firms worry that they will not get 'credit' for technology-improving activity in the face of an 'unfair' opportunistic APC setter, or simply an error-prone APC setter, the APC price will have to be set higher *ex ante* to encourage them to engage in such activity. Similarly if the APC regulator is soft and malleable, the APC will be higher anyway. It is hard to see the APC price being anything but higher.

In addition, if the technology improvement is caused by *basic* researchers and not by private researchers, then the payment *should* be reduced (unless the basic researchers can charge the profit maximising price to the later APC-paid firms for the improved technology, but this is difficult if the knowledge is a public good). But, if the technological improvements are developed by the *private* researchers, then the APC should not be adjusted downwards (or at least that part accounted for by private researchers). This requires (during the assessment process for the flexible APC) knowledge of what proportion of technological advance was caused by private and what proportion by basic researchers. When Kremer mentions that "there would be no need to increase price if the general technological advances in biology reduced the expected cost of developing a vaccine sufficiently that many firms would pursue vaccines,"[84] he seems to presume easy ability to separate out general advances in technology from advances brought on by the private companies themselves.

Clearly, discretionary systems have potentially severe problems dealing with technological improvement. The less easy it is for the discretionary system to avoid these faults, the more favourable the fixed system (and other systems) starts to look. But, the fixed system carries serious cost distortions in its turn.

These problems are ignored in the Kremer APC figures since the state of technology is always presumed fixed. Issues about improvement of technology don't arise.

---

[83] As in many other places in this paper, this is not to cast aspersions. In a system based on equity markets, the expected returns to technology-improving investment are lower if the markets suspect that APC payments will be altered to take unfair advantage of the technology improvements. Raising the capital in the first place to do the technology improvements becomes more difficult and expensive.

[84] K7:38

6.4. Staging-Point APCs

One further possibility is to allow a series of APCs to get to various 'staging points' in technology, with payments based on the first to reach each point, with an opportunity to review/reset the terms of the next APC. But again this takes the information structure nearer to open collaborative research, while creating many of the same dynamic inconsistency problems *en route*. All the problems mentioned above regarding flexibility and discretion, uncertainty, discount factors and cost of capital still apply, *and* all the layers of problems highlighted in future sections also still affect how we view each staging-point APC.

6.5. Flexible versus Fixed APCs – Higher Costs Either Way

We saw how a fixed APC comes with a set of extra costs and distortions. Now we have seen how a flexible APC carries a different set of extra costs and distortions. All the gyrations in rules, and the presence of institutions, are based on the need – in order to achieve maximally efficiency – to constantly track the optimal intensity of private R&D in a world of technology subject to varying probability distributions and 'shocks'[85]. The reality of this is ruled out in the Kremer technology, and hence all cost comparisons.

The fundamental problem in choosing between fixed and flexible APCs is the trade-off between credibility and flexibility. The more binding and less flexible the APC commitment is, the more 'credible', and the stronger the incentives for vaccine developers. But this requires large extra 'options' components to be built into the APC price to cover the non-stationary technology. The more risky the technology, the larger these extra components need to be. In some cases, where there absolutely is *no* opportunity to change the terms of the APC *ex post*, the extra component even has to cover the worse-case scenario.

Flexibility enables these extra components to be left out, but the imperfect commitment thus generated will increase the risks to developers, which will require large extra components to be built into the APC price to cover extra capital costs. The more risky the technology (and the less credible are mechanisms) the larger these extra components will need to be. Kremer discusses an example of this, though not in the context of technology[86]. A potential vaccine developer has to guess the degree of commitment of potential donors. If a piece of contractual language is missing such that there is a 90% chance of purchasing at the agreed price of $1bn and a 10% chance of reneging and renegotiating the price down to half that, this yields on average 95% of the promised payments, or an expected payment of $950m (Section 7 explains why situations like this are a very real possibility). Kremer states that an "imperfect commitment reduces both the expected revenue of vaccine developers and the expected costs for the sponsor in the same proportion". However, this ignores capital costs. If $1bn was the risk-adjusted figure required to generate optimal research intensity, and if we wish not for vaccine development to be deterred, and if vaccine developers are risk-neutral, then the promised payment by the sponsor has to

[85] Or, looked at another way, simply to achieve the level of efficiency implied in the Kremer calculations.

[86] K7:8. But, as commented before, the same logic goes through on any part of the APC model involving flexibility/discretion.

rise to $1.053bn[87] to ensure the same intensity of research effort. If the vaccine developers are risk-averse, the figure must be correspondingly even higher to reflect this.

Even the small chance of a big renegotiation can have significant impact on the needed APC price. In this simple case, if private capital costs were already 44% of total APC costs[88], presumed for the moment as the very lower bound of capital costs, then the 10% chance of a 50% lower price, results in over 12% more private capital costs (and will be much higher if risk aversion and the probability changes discussed in the last-but-one sub-note are also factored in) [89]. In such cases fixed rules might seem sensible even with all their inefficient consequences.

Whichever way the trade-off goes, the costs are always higher. When the APC-setter gains on the one it loses on the other. The problem intensifies – and the extra costs imposed rise – the more complicated and risky the technology (and indeed the more risk-averse the players and the harder it is to diversify the risks[90]). The exact gains and losses of choosing fixed over flexible systems are not *a priori* clear. A full assessment would require a consideration of all impacts across all APCs both current and prospective (for example, use of early discretion in negative ways will make later APCs more expensive). The problem is also compounded by the fact that different parts of an APC are under the auspices of different committees. So, for example, a choice might be made at the level of one committee to go for a fixed system to maximise credibility – and stomach the extra costs caused by the lack of flexibility – only to be offset by the overuse of flexibility by another committee (or a set of committees unable to coordinate) at another level (say those funding research at an earlier stage of development), leading to higher levels of inefficiency overall.

In addition, there may be a worry that even a fixed APC may not be fully credible, raising the capital costs of fixed APCs anyway. This follows much the same logic as used to describe fixed exchange rate regimes, where currency markets have to price in the small chances of large devaluations or revaluations, and where this can lead to self-fulfilling pressures on the system to revalue or devalue. Knowing *ex ante* that this possibility exists, complicates enormously the analysis of private firms engaged in vaccine research who have to factor it in to capital costs under a fixed APC. Like exchange rate regimes, there may be some technologies that are *so* risky, that – apart from setting the APC price so very high that the chances of collapse can be avoided (with huge waste in the meantime) – it might be better to avoid a fixed regime

[87] x such that 0.9x + .0.1*0.5x = $1bn, presuming the probabilities are not altered in the process, which they no doubt would be; the probability of reneging is likely to rise with the APC price, necessitating adjustment of the APC price even further upwards to compensate. Price would settle at the stationary point in this reasoning process. All of this would have to be adjusted upwards in proportion to the degree of risk aversion.

[88] This is the TUFTS calculation of the proportion of capital cost in developing a 'typical' drug (though see the literature for discussions of the veracity of this figure).

[89] A real-world calculation would have to work out the whole probability distribution over reneging. And indeed, the distribution evolves over time, leading to the possibility of instability. Avoiding this instability would be another reason for setting the APC extremely high to start with.

[90] It is very likely that the risks will be difficult to diversify away.

altogether and adopt a more flexible regime and put up with the high capital costs consequent on the risks of the latter[91]. Again, either way, capital costs are higher.

The simple constant-instantaneous-probability technology framework used in the Kremer APC models strips out these problems and costs. The more complicated the technology, and the greater the distortion from the idealised model, the higher these extra costs, and the more biased are the Kremer figures.

Phrased another way, in the Kremer model the trade-off has been set to zero; technology cannot impose higher costs on the APC via options elements, neither can it impose extra costs by discretionary elements. In that model, the only trade-off of any relevance is related to the choice of quality, and that is supposedly solved by good description of the end product. The above sections clearly indicate however, and we will see it much more clearly below, that even the apparently simple quality trade-off is related to technology, since the reason that *that* problem can be reduced to the mere description over only an end product is itself driven by the same technological assumptions that stripped out the above layers of problems. And the simplification of the quality rules over end-products is related to the simplifications of assumptions over technology.

One solution would be to set such a high APC that the need to adjust afterwards, and any dynamic inconsistency problems would be avoided entirely, but – as well as being very wasteful of limited resources – this questions the whole point of doing cost comparisons of mechanisms, if the plan was to throw everything at the APC mechanism in the end anyway. And what if another approach, say a more open collaborative approach, could have allowed this heavy cost to be avoided?

[91] It is not even clear if this solution works if the APC has to be set very high indeed, say for a HIV vaccine.

7. WHAT HAPPENS WHEN VACCINES ARE DEVELOPED?

7.1. The Difficulty of Generating Quality Vaccines

This section is all about 'quality' issues in the APC setting. But the flip-side to quality is cost. The driving logic of this section is that if, at any given APC price, the expected quality of vaccine being produced is lower, then in order to generate any given quality of vaccine the APC price will need to be higher. The overriding conclusion is that since there *are* strong forces operating under a typical APC mechanism pushing in the direction of lower-quality, the public costs of developing good-quality vaccines will be higher. Indeed, without good knowledge of the underlying technology, it may not be possible at any price.

The principals discussed in this chapter are pretty general to any mechanism that controls quality entirely through the end market (which is what the rawest form of the APC does). Therefore, many of these issues relate to the patent system generally. Issues of control of quality en route are also therefore of great interest, although they are expunged from the APC models presented so far for HIV, malaria, and TB.

At the heart of the problem is the way that expected behaviour at late stages of the program feeds back – in backwards induction fashion – onto behaviour at earlier stages. A belief that poor vaccines may be tolerated increases the risks and reduces the profits of those researching good quality vaccines. 'Poor quality' could simply refer to the notion of not investing enough in investigating a wide enough range of vaccine leads, that in consequence, in a probabilistic sense, decreases the average quality of vaccines even if there is no particular control over the quality of any particular vaccine lead and/or the quality of each lead is purely stochastic. But it could also refer to choices over what lead to follow where some leads are more expensive or risky to develop even though they might produce a better result (say two boosters instead of three). If those researching 'good quality' vaccines are not to face overly-high capital costs and be deterred from doing research, then the rules of the system have to be designed to favour *them*. There is also a *self-fulfilling* aspect, since if few other firms have concentrated on quality – believing that the program's rules are not credible, and that poor-quality vaccines will be accepted – this makes it more likely the program will accept poorer-quality results anyway, thus making *not* concentrating on quality the rational, profit maximising, strategy in the first place. The system has to constantly battle against this self-fulfilling tendency[92]. Aggravating all this is the cost

[92] Incidentally, the current patent-based pharmaceutical industry also has to battle against this, but without many of the system features present in APCs. Maybe this is one of the reasons why so many 'me-too' drugs and so few truly innovative drugs are produced by it? Comments in previous and later sections on the issues of patents also suggest possible reasons for why the current dearth of truly innovative new drugs is concomitant with an explosion in costs. Much of the analysis of this section can be reinterpreted for pharmaceuticals in general. Tentatively it suggests it might be worth exploring how the various systems rank. The problems here would seem to suggest that the order is: open collaborative; APC/prizes with some control over quality; patent with market.

of delay, which is potentially high in the case of vaccines, adding to the pressure towards lower-quality[93].

Rules for steady states

Key in *trying* to solve these problems is the pre-agreed minimal quality of vaccine acceptable for purchase, an utterly credible set of rules regarding this, and the ability of the APC to, somehow, optimally redistribute itself post-vaccine-development so as to encourage research on higher-quality vaccines (naturally entailing restriction on purchases of vaccines meeting the minimal quality). We immediately observe that this pre-agreed, and supposedly optimally-set, minimal condition is only unique for any system once the state of underlying technology and epidemiological conditions are in steady state. If the state of the underlying technology and/or epidemiological conditions changes, the chosen pre-agreed minimal condition is no longer optimal, and the rules are increasingly pulling against the constraint of a 'wrongly chosen' minimal condition[94]. Maybe this is why Kremer runs his model on the assumption of constant underlying technological conditions. This is an assumption, not a fact. The more wrong it is (and it is very wrong) the more biased the figures.

There are potentially large problems with creating *quality* vaccines, and preventing inefficient vaccine purchases, if the eligibility and pricing rules of the APC are not set vary carefully. This potentially faces vaccine developers with a great deal of risk[95].

The damage when the eligibility rules are wrong

Given the underlying technological and epidemiological conditions, if the minimal conditions are set too toughly, they may prove too difficult to achieve, discouraging firms from pursuing research leads in the first place; either no vaccines are developed or vaccine development is delayed and expensive. If, for example, the rule requires 90% efficacy against *all* strains of a virus, then firms may avoid leads that might give 99% against some strains but only 85% against others. And overly-tough minimal conditions would *delay* the roll out of vaccines even as products that could have satisfied more minimal conditions would otherwise have been ready. There could be situations where a vaccine has been developed that does not satisfy the minimal condition but the program is not prepared to let it be used since it would ruin the market for the 'legitimate' drugs being aimed at by the program. It is not clear though that the use of a vaccine failing to meet the minimal conditions would, or ever could, be ruled out. How much longer would policymakers (and developing countries) be prepared to wait for an AIDS vaccine meeting 100% of the minimal conditions if there were currently a vaccine meeting 'only' 90% of the minimal conditions?

---

[93] Think of it this way: If a poor quality vaccine is produced, and there are no costs to delay, the system could hold out for the better quality vaccine. Given the knowledge of this *ex ante*, no firm would ever work on the poor quality vaccine. Similarly, if a poor-quality vaccine is produced, and there is no imminently available good-quality vaccine, then the costs of delay by *not* using the poor-quality vaccine make it an *ex post* more favourable proposition to use it, increasing the incentives to work on it in the first place. If there are self-fulfilling aspects to the problem, this might make this the only logical outcome.

[94] In a sense, all of the systems that we henceforth describe become less and less optimal the further away the system drifts from the underlying technological/epidemiological conditions that generated the originally set minimal condition.

[95] And, as everywhere else in this analysis, this pushes up the proportion of the APC accounted for by private capital costs.

Naturally this feeds back into the risks of those developers working on the vaccines meeting 100% of the conditions[96].

Conversely, if minimal conditions are set too weakly, the issuer of the APC might, for example, be obliged to buy a vaccine that provides only temporary protection or interferes with the development of natural immunity. Even if one is developed, and no other vaccine is available, it would be difficult to resist using it. And there would need to be more restriction on purchases of vaccines that do meet the minimal quality so as to leave room, and thereby create incentives, for developers to be rewarded for better quality vaccines.

Perfectly-set quality rules need perfect knowledge of technology

Clearly, the proper rules for minimal acceptable quality and the rules for APC redistribution have to be based on the current and all future expected states of the science and costs of vaccine development at all stages of development, public and private, and on all expected future epidemiological changes. An intuitive way to think of this, is that the APC mechanism has to try to do something that curiously *looks* a bit like price discrimination, to extract a certain amount of social surplus. On the one hand, the overall generosity of the APC price has to be just high enough to maximise incentives towards research, but it has to try to set terms higher for more innovative outcomes (with these terms a function of the expected costs of generating such outcomes), that is to 'discriminate' in their favour and create incentives for high-quality firms to try to hit those outcomes. All periods are connected dynamically, and the terms for optimally extracting surplus in one part of the state-space cannot be set without having a handle on what the other parts of the state-space look like. To do this perfectly (as Kremer's calculations presume) requires huge knowledge about the underlying technology – including about future states of technology. Yet again, Kremer's claim that an APC economises on the need for policymakers to gather and efficiently process a great deal of information is found to be wanting, and is masked by the use of a technology device that strips the problem out anyway *ex ante*.

This is a major headache for mechanisms like APCs, and again suggests that a more collaborative research framework is likely to be more flexible and efficient. Once set in an APC, the options value of being able to re-set the minimal acceptable quality in the light of technological or epidemiological change is lost (unless yet more *ex ante* rules can be derived to allow *ex post* change of the minimal quality, and such that dynamic inconsistency problems can be avoided). The system is left struggling to adapt around a pre-determined minimal condition, when it might be more efficient to change the pre-agreed condition (maybe this is why the Kremer framework ignores changes in underlying technological and epidemiological conditions and concentrates on what happens once these condition are set?). A more flexible approach to vaccine development would have a clear advantage over permanently fixing the minimal conditions. Whatever the potential downsides to more open approaches, we repeatedly

[96] There is some myth about the way APCs work. The acceptance of a poor quality vaccine *is* a risk to those working on better quality vaccines, because there is nothing in the APC system that guarantees a particular market size for the later better quality vaccine. Once the inferior vaccine has been accepted, the funds remaining in the APC, and the market size available to later, better, vaccines is reduced. And as the 'left over' market shrinks so do the incentives to work on better quality vaccine since the largest potential market available to them has shrunk.

see the way it 'wins back' these many options values, which ordinarily would have to be priced into the APC.

7.2. The Problems and Risks for Early and Late Vaccine Developers

This potentially complicated post-APC problem is in a sense an extension of a problem that already exists under a patent-based pharmaceutical system – tying rewards for research to 'price times the quantity of vaccine used' even if superseded by a better vaccine. And it relates to the way in which sunk costs of R&D build up in the pharmaceutical industry without any repayment until a drug is patented, or, in this case, until an APC is awarded. Those doing innovative research have an unenviable trade-off between the need to aim for the first APC contracts to repay heavy sunk costs (including a large element of capital costs) and the desire to create quality vaccines but risk losing all sunk costs if beaten by a competitor.

Often (but, importantly, not always) the most innovative and difficult research takes place leading to the *first* drugs in a class. Later, *better*, drugs come along building on this earlier research (for example, they have fewer side-effects, or require shorter courses of treatment or fewer pills per day). Rationally, the inferior first drug should be dropped, with the first drug developer still compensated for their highly innovative research. However, the patent system only allows the first developer to extract compensation through the sale of the drug, and so the first drug has to stay on the market at a high enough price and for long enough to earn back the resources expended on its research (in fact to earn back several times its research costs, to compensate, in 'blockbuster' fashion for research on non-successful drugs too). Later drugs, if they are cleared for sale, then compete with the first for market share. Both the first and later drugs are forced to engage in heavy marketing (and the knowledge of this marketing is *ex ante* priced into choice of research projects and the intensity of research effort) to try and gain, or simply to keep, market share – when the rational arrangement would have been to clear the market for the later drug. We cannot escape the fact that the first developer needs to be rewarded for the science breakthroughs they achieved even if we might not want to use much of their drug. Tying the reward to the quantity of drug used in order to repay them for their innovative research is inefficient and forces use of inferior products. Dropping drugs as they are superseded would be perfectly natural under a more open research system for vaccine development.

At the same time it is not necessarily the *first* vaccine that comes along that does contain *all* the innovative science and that should therefore only be rewarded! For example, if two vaccines are near to completion, but one has fewer side-effects but will take longer to develop, the terms of the APC must not punish it and deter its development in the first place. *Ex ante* the mechanism has to be set up in such a way as to be able to, credibly, redistribute the APC payment post-development in a way that rewards fully the vaccine that takes longer to develop but is 'better' (that is the developers of later 'better' vaccines should not just be rewarded according to what they 'add' to social surplus since that will never enable them to recover costs). Otherwise, under 'winner takes all (or nearly all)', the poor products hit the market first every time and – knowing this *ex ante* – no firm will ever bother to invest in R&D leading towards 'better' products that take longer to develop (i.e. it will not be part of the dominant strategy equilibrium). The Kremer technology removes this

second type of developer, i.e. the probability distributions of Kremer are not capturing all reasonable possibilities. Modelling as if they do, risks generating poorer-quality vaccines[97].

7.3. Orphan Drug Reasoning – First and Late Developer Risks

Various orphan drug extensions – granting market exclusivity of some sort – are suggested by Kremer to solve the APC redistribution problem and to try to tip incentives towards quality vaccines. In no case does he convincingly argue that lower-quality, and more expensive, vaccines will not crowd out better quality and cheaper vaccines. In all cases, the inability to set and then target terms perfectly, leads to cost estimates higher than those derived by Kremer based on idealised conditions.

Multiple Leads and Multiple Developers

Unfortunately, market exclusivity also creates a problem that may reduce the number of leads being followed. It might be thought that if several companies develop a similar vaccine around the same time, the very first to win approval should be given all of the APC. But market exclusivity to the first one is a risk to all those working on the others – and the knowledge of this destroys *ex ante* incentives for multiple developers to follow multiple vaccine leads. If following multiple leads is one of the main driving forces for quick vaccine development, it is not clear that allowing only one winner is ever the risk-adjusted optimal result. An AIDS vaccine, for example, is likely to come out of a process involving multiple, in certain ways sometimes similar, leads. It would not make sense to punish all but the very first of a set of leads to make the breakthrough, thus reducing the number of leads being followed in the first place, thus slowing down the whole AIDS vaccine initiative.

In addition, in section 11 below we see that we need multiple competing developers in order for an auction mechanism to efficiently set the APC 'price' (if this is the mechanism chosen), so that anything that deters multiple leads will feed a higher APC 'price'. We cannot rely on a mechanism that requires large numbers of 'individually distinct' projects in order to work but then punishes them. No rules can really get around this problem.

Limiting Purchases to a Period

It could be that a limit could be set to the purchase of vaccines developed under the program within a certain period (a year or two say) following licensing of the first acceptable vaccine, unless a subsequent vaccine is clinically superior. Though this would reduce risk for firms in a tight race to develop a vaccine, it would reduce the chances of 'me too' vaccines reducing the sales of the first developer. However, on the down-side, this might put pressure on followers to do unnecessarily quick and expensive research (after all, their costs up to this point are sunk, and all that matters is the return from the continuation game) with pressure on them to manipulate figures. All of this has to be factored into the APC price *ex ante*.

[97] The reader will have spotted in some of the above the ways in which prize systems may have superior ability to reward innovation, even if subsequent products make most of the sales. This is being explored in another paper: Farlow 2005, 'Prize Funds for Drugs and Vaccines: Principles and Problems'.

Linking Payment to Marginal Improvement

Kremer[98] suggests another possibility for a modified orphan drug framework, one where "the price would be related to the marginal improvement the subsequent vaccine represents over the original vaccine, and the original vaccine developers would continue to receive compensation in line with the social value of their work." What might this actually entail in practice:

1) Does the first developer continue to distribute the inferior vaccine (they are paid through sales after all). Or does this mean that they stop selling and get compensated through the sales of the follow-on drug in proportion to the first vaccine's social worth, and simply stop selling the first vaccine?;

2) Do developing countries use any of the first inferior drug beyond the development of the second drug, since the price mechanism is 'twigged' to dictate this outcome? We cannot escape the fact that *no* efficient system should *ever* keep using inferior drugs – achieved by manipulating the price mechanism – simply because R&D has to be paid for through the price of the drugs that actually get sold and consumed[99];

3) The follower vaccine in a sense costs more though those who develop it are paid only in proportion to the value they add, but, in a sense part of what they get paid for their vaccines has to get passed on to the developer of the first vaccine;

4) It is not clear that it solves the problem that earlier poorer-quality vaccines crowd out better-quality later vaccines. If such quality vaccines tend to take longer to develop than inferior vaccines, why should they be paid for through 'marginal' payments on top of payments going to inferior drugs that supposedly somehow were 'first' to develop the technology? Those developing these better vaccines would rightly expect to extract all of their costs by a rule relating to the vaccine's *total* social worth and not relative to its *marginal* social worth. Failure to sort this problem will force incentives towards the early low-quality vaccines. As with all inefficiencies created during its operation, this feeds into a higher *ex ante* APC price to stand a chance of getting the good quality vaccines produced.

5) How is the 'marginal improvement' measured anyway? What if it is wrong? Who measures it? What if it suffers from many of the information distortions and asymmetric information problems mentioned elsewhere in this paper (see more below)?

6) What political pressures could possibly force countries to use inferior drugs?

7) What are the dangers of using inferior drugs – if this is indeed the suggestion – in terms of resistance for example?

8) The phrase 'marginal improvement' not only refers to marginal improvement in 'quality' but also to the 'marginal populations' targeted by later vaccines. To the extent that these later vaccines do not flow from the first vaccine, the marginal populations might actually require a 'new' APC (to ensure that all sunk costs of

[98] K7:29

[99] This argument is even stronger for vaccines, given the issue of the build-up of resistance.

development for the sub-group are covered), rather than relying on bonuses based on older APCs. Otherwise there is a risk of not getting a vaccine to cover the population sub-group[100].

In addition, to the extent that these later vaccines do flow from the old vaccine, the use of a bonus presumes that *all* firms have access to the technology of the firm that developed the first vaccine and can compete to develop the vaccine for the new population. If this is not the case – and it *is never* the case – the later groups are tied in to relying on the first vaccine developer to be bothered to develop the vaccine for them. And society is tied in to a monopolistic situation requiring monopoly-size bonuses and yet still does not get the follow-on products.

Buy-Outs to Encourage Quality

Kremer suggests yet another possible extension of the orphan drug framework to try and get around this problem – to "give the developer of the original vaccine incentives to buy out the technology of the second producer."[101] Somehow, the bonus payments (explained in a section below) attaching to the second developer give the incentive for this. This suggests that the original inefficient drug can be dropped. But this has problems:

1) It still does not give incentives to longer-to-develop but better quality vaccines. The vaccine developer who sticks with the 'better' vaccine that naturally takes longer to develop, will still not get fully rewarded for innovative research if they only get paid for the value they add on top of the first inferior vaccine.

2) At the point of 'buy-out', with so much investment sunk and facing a bilateral bargaining situation, only a 'regulator' could possibly determine the terms of the buy-out. Worries about this in advance put off firms from sinking investment in the first place. Once again, Kremer's ruling out of any sunk investments (in an industry dominated by sunk investments) leads to misleading results.

3) We see later that these terms will be highly unlikely to be efficient. In particular, the efficient pricing of the buy-out relies on bonus payments being 'correct' at all times, whereas they are often heavily distorted and poor instruments for determining value.

4) How does the system cope with bonuses that are spread over multiple periods and multiple buy-outs with consequent 'layers' of bonuses, as would be the case for a product with a long development cycle?

5) What happens if the firm doing the 'buy-out' is simply buying-out the stream of bonuses to a product that is not yet in mass production – such that it will have to invest in large fixed capital to produce it? If it already has large fixed capital devoted to its own current vaccine production, this would now become obsolete[102]. This

[100] The new APC for the subgroup might, for example, have to be set at a price that is higher than the 'implicit' APC for the subgroup that had been contained in a previous APC that had embedded the subgroup amongst various subgroups that – it had been hoped – might have been covered by the previous APC.

[101] K7:29.

[102] To the extent that it could not be adapted to the new vaccine.

obsolescence reduces the value to it of the firm it is buying out and leads to a lower offering price. Similarly, to avoid forced obsolescence the firm will have an incentive to keep current production of the inferior vaccine going. The cost conditions of Kremer rule this danger out, since they contain no sunk or fixed cost elements.

6) Many of these problems are aggravated by the fact that production of a vaccine is not just about capital but about know-how. This tends to suggest that the developer of a vaccine should be the one to produce it. It makes buy-out more difficult.

7) It presumes that financial markets work equally well in providing finance for all firms to perform buy-outs (one can visualise scenarios where larger firms have greater access to capital markets in order to 'buy out' smaller later firms, and scenarios where less credit-worthy smaller firms are unable to buy out the technology of larger firms, etc.). Failure on this raises the capital cost component of the APC price.

Late Developers to Pay Early Developers

Another version is to pay for the newer vaccine based on its efficacy, but require the developer to pay the original developer an amount equal to the price of the original vaccine, less an allowance related to the production cost of the new vaccine. But this is hard to administer and expensive, it creates all kinds of perverse incentives (like affecting incentives relating to any factor that will push up the fee paid to the first developer), and it still does not solve the problem of making sure that developers who take the greater risk of sticking with the 'better technology that takes longer' get properly compensated.

The problem of resistance and the need for newer vaccines

All of these problems are aggravated if the build up of resistance to vaccines and medicines in general necessitates a rolling stock of fresh vaccines built on a rolling stock of new APCs. The setting of the terms for the new APCs must not be distorted by the 'old' APCs. For example, those working on vaccines to deal with new strains would want all of their sunk research costs covered. They would not want to be paid *as if* they had lost out to previous vaccines on previous strains. Yet, separating out which firms were responding to the need for new vaccines (or indeed drugs) and those who had simply failed on previous vaccines, is very difficult to determine. The danger is that those working on new vaccines to tackle resistance are put off from doing the research if they think that they will not be able to be separated out and get *all* their private costs covered. Vaccine research to catch up and get on top of an evolving, resistant virus, is slowed down.

7.4. The Problems of Rewarding Innovation

Interestingly, the ultimate goal even under an APC, is clearly to reward innovation. All of these gyrations are simply the result of presaging everything on the notion that 'price times the quantity of vaccine sold' is *the way* to compensate for the value of innovation. So, every time something 'new' comes along that 'adds value', the method of payment via the quantity of sales of vaccines and their price has to be 'rejigged' to incorporate it so that this 'price times quantity sold' condition still holds, and such that there is incentive to 'add value' in the first place. If this cannot be done *perfectly*, good vaccines do not result. It does not matter how much Kremer tries to create layers of mechanisms onto the orphan drug framework and add ways to try to

get over this problem, he cannot guarantee to generate incentives that lead to the best quality vaccine being developed. And it rather begs the question of why a mechanism might not be sought that simply rewarded innovation and allowed a totally flexible approach to the level of usage of particular vaccines.

It's a bit Like Orphan Drug, but not as we know it

Part of the constraint here is that the situation is markedly different from that of typical orphan drugs. In orphan drug cases, once the first drug has been approved, no individual or institution sets the size of sales or the ultimate profitability of the drug. That is left to the free market. The company markets the drug and tries to build up a following. When any follow-on drug is approved, it competes freely in the market with the first, and purchasers get to choose whether or not to use or to switch to it. There is no legal wrangling over size of sales, and so forth. There remain private incentives to the developers of the second drug since if it is much better, it will take more of the market. It may have to engage in marketing and other costly expenses that will result in overall drugs prices being much higher – but it does not face, what might seem to its developers, the vagaries of committees setting prices and quantities. The first developer has to calculate whether it is worth investing in the somewhat less useful drug knowing that its exclusivity may be lost with the development of a new drug. Developmental decisions must price in these considerations. Since the APC *is the market*, the size and profitability of sales over all the developers is set by a regulator. The regulator has to emulate what the environment for investment decisions that the vaccine developer would face were it a traditional orphan drug. And the vaccine developer has to face the uncertainty of the regulator.

The APC perpetuates the problem already in the system – that reward is through the sales of pills and injections and not for the innovative research that went into them. Attempts to avoid this basic problem lead to the creation of layers of committees, complicated mechanisms, and lots of risk layered on to developers. And all risk has to be priced in to capital costs and the required APC.

7.5. Minimal Conditions and the Acceptance of Imperfect Vaccines

One of the ways around the problems with minimal conditions, it is suggested by Kremer, would be a commitment to purchase 'imperfect' vaccines and to base calculations of how much a vaccine developer would get on what a commitment to purchase an ideal vaccine would pay[103]. The terms would have to be very precise to ensure developers still concentrated efforts to develop 'perfect' vaccines; if, given the lower costs of developing an inferior vaccine, the conditions are set such that the inferior vaccine is more profitable at the margin, then developers will concentrate on that (this requires knowledge of costs of developing both perfect and imperfect vaccines).

[103] K4:131 More precisely one calculates "number of vaccinations that would be cost-effective given any set of vaccine characteristics selected. It then calculates the total social surplus generated by such a vaccine. Under the approach to vaccine pricing implemented [here] the total commitment is then set to equal the commitment size for an ideal vaccine times the ratio of social surplus generated by a vaccine with the selected characteristics to that generated by the ideal vaccine." The phrase 'information?' (or even 'perfect information'?) should be going through the reader's mind at this point. And the fact that these rules would have to be set in stone for credibility.

And it would require a huge amount of knowledge about the current and future state of the science of vaccines. Let's say that policymakers at some point conclude that on the basis of the current state of science there will *not*, it turns out, be a vaccine satisfying the minimal conditions soon. If they choose to use up more of the APC than stipulated in the terms of the original APC on what is currently available, how much of the APC should be used up on the imperfect vaccine? If they still think that there will be a vaccine meeting minimal conditions *some* time, they will not want to push up the imperfect 'allowance' too far since that will leave less of the APC for the vaccine meeting minimal conditions and hence reduce the incentives to develop that vaccine. And they face litigation in the meantime from those who were working towards the more perfect vaccine.

The Self-Fulfilling Outcome

Could policymakers credibly commit to the minimal conditions and avoid the dynamic inconsistency of accepting too high a level of sales of vaccines meeting less than these conditions? Probably not. Of course they would like to commit (in order to create maximum incentives) not to raise the imperfect allowance in cases of there being no perfect vaccine, but they have (under probably strong pressure) the *ex post* incentive to change this rule. Of course, inability to credibly commit, makes the inferior vaccine more likely as a self-fulfilling outcome: If developers conclude that policymakers might at some point conclude that there will not be a vaccine meeting the minimal conditions (given asymmetric information, they may not have the information needed to conclude otherwise) and that they will then purchase imperfect vaccines, the rational response of developers may be not to work on the more perfect vaccine in the first place. But this self-fulfillingly feeds the original conclusion that no vaccine is being developed meeting the minimal conditions.

The Negative Views of the Industry

Mercer Management Consultants conclude that "Some manufacturers believe there is a low probability that an HIV vaccine will ever be universally recommended in industrial countries. These firms noted that the minimum performance characteristics of an 'acceptable' vaccine is significantly higher today than it was ten years ago. Combined with the emotional and moral viewpoints charging HIV/AIDS control, this vaccine could be even more controversial to introduce." How does all this feed through into the APC price? Any uncertainty over all of this will raise the risks of researching the 'perfect vaccine' and hence raise the required APC price to compensate.

7.6. The Use of Payment Bonuses to Developers

If information on the number of lives saved by a particular vaccine and its quality is revealed only slowly, it might be better to condition purchases on long-run outcomes, and not pay all in one up-front price anyway. Does a particular vaccine provide temporary or permanent protection? To what extent does a particular vaccine prevent secondary infections? To what extent does resistance to the vaccine in a particular population build up? And so on. Another suggestion of Kremer is to use earlier vaccines but to hold back on APC payments for them and use later additional 'bonus' payments based on the realization of DALYs as a consequence of the use of these vaccines. There is some sense in this. In trials it is usually much easier to make sure that delivery protocols are followed than might be the case in a 'real world' setting.

To encourage firms to work on vaccines that work in the real world, it might be a good idea to link rewards to actual lives saved/DALYs, and not just to the results of clinical trials. Bonuses also make it easier to force purchase committees to pay a remunerative price. Before a vaccine is used they might be able to claim that there are potential problems with the vaccine. But once the vaccine has been used, they might find it more difficult to argue that it is ineffective.

The drawbacks of bonuses

The use of bonuses has several important drawbacks (common to any system that might involve bonuses):

1) It would mean that capital costs would have to accumulate on unpaid parts of the APC set aside for bonuses. If many of the inefficiencies and risks mentioned in this paper already hold, this may be at a very high rate. This might be hard to stomach if it were transparent.

2) How does the system ensure incentives to report correctly (and gather such reports) on the effectiveness of vaccines if it is known that this will trigger higher costs? Or indeed, the more disturbing 'principal-agent' problem on the ground: What is the incentive to administer the use of vaccines correctly if it will trigger a bonus[104]? Independent outside observers (more committees) would be needed[105]. If pharmaceutical firms can influence DALY calculations through politically-connected actions rather than through research, they will put effort into this and not into research towards good vaccines.

3) There are already – given a patent-based system – insufficient incentives to research the adverse effects of drugs (or to identify populations that should not use them) leading up to the patenting of a new drug. The use of bonuses extends this effect even more after patenting. Similarly, if the build up of resistance to a vaccine leads to reduced or even 'negative bonuses' (though the latter is not discussed in Kremer) what are the ramifications of this for the incentives to research the level of resistance? If the growth of resistance to vaccines necessitates a rolling stock of new vaccines, the knowledge of the degree of build up of resistance should feed into the setting of the APC terms for the next vaccines in the series. How is it ensured that this information is timely and not distorted? Observe how interventionary the APC turns out to be yet again.

4) How are the incentives of vaccine developers affected by the risk that their bonus payments may in part be the result of those administering the vaccines and, possibly, on the politics of a particular participating country? Basing payments on lives saved and on delivery costs, opens developers to many risks when outcomes are based in part on the actions of others. For example, payments may depend on the ability of those on the ground to control cold chains. Should vaccine developers get involved in this part of the process to reduce their risks? But maybe this is not such a bad thing if

[104] This is another way of saying that the marginal cost of better administration of vaccines is higher if bonuses are present. For efficiency and social welfare it might be better to keep this marginal cost as low as possible.

[105] As a follow-on to the previous sub-note: If such independent outside observers are going to be needed anyway, the need for bonuses might be mitigated anyway.

it encourages vaccine developers to design heat-stable vaccines for imperfect on-the-ground systems in the first place rather than designing for an idealised system.

5) Bonuses create incentives to strategically extract payments from vaccine developers. What if health ministries (for example if they are under political pressure to achieve numbers of treatments quickly) behave strategically seeking to extract payments (which could also simply refer to more favourable terms on other pharmaceutical contracts) from a vaccine developer in exchange for agreeing to distribute efficiently?

6) What if the committee simply makes mistakes? Even if mistakes might go either way, if vaccine developers are risk averse then this does face them with risk, and *does* have to feed through to their capital costs[106]. And early mistakes by committees can create later distortions.

As always, all extra risks have to be priced into capital costs and hence the APC. It would be better to base payment on indicators of 'likely' lives saved/DALYs rather than to use bonuses. But this leads back to an inefficient non-bonus based mechanism, and a great deal more need of *ex ante* information.

Kremer finds in favour of APCs compared to patent buy-outs (another pull program) partly because of the risk that the government will purchase the patent for a product that will "turn out not to be useable due to side-effects that were not immediately apparent...[and that] It may be difficult to recover the funds at this point"[107]. But this comparison seems to be presuming either an APC with a (perfectly-working) bonus system built in, and that somehow such bonus systems are not available to patent buy-outs, or that APCs simply never generate these quality issues in the first place (the calculations seem to suggest the latter).

The notion of bonuses raises many unanswered issues. All of this has been ruled out in the Kremer calculations.

7.7. Delaying and Committing to Delay the Release of Vaccines, and More Dynamic Inconsistencies

One of the politically most difficult consequences of trying to avoid dynamic inconsistency and enforce quality vaccine research via an APC mechanism, is the need to force countries in the program not to use vaccines that do not meet minimal conditions, and indeed to limit the use of vaccines that *do* meet minimal conditions.

[106] That Kremer repeatedly ignores issues related to capital costs can be found in the way he treats this. He states (K7:25) that if these mistakes "do not systematically tend to underestimate or overestimate the actual effects of the vaccine, then the potential profit from developing a vaccine could as easily be increased or decreased by the uncertainties in calculations of DALYs or lives saves. The attractiveness of investment in vaccines would be reduced, but only to the extent that vaccine developers are not willing to take gambles that could turn out to help as easily as to hurt them." But the financial risk of these *new* gambles still has to be bourn by someone. Mistakes increase capital costs.

[107] The other reason is that the product might be difficult to produce, leaving the developer with an effective monopoly, even without the patent. K2:5

But (not for the first time in this analysis) the program then has to be set up *ex ante* on the basis of some notion of where the science is going as well as what other products are near completion. For example, there may be a cost to holding back if really there will be *no* superior products in the following five years – but this has to be known and written into the rules of the program even before any of the research for the first drugs is sunk.

And there are difficult issues regarding holding back or pushing forward with vaccination programs. It may be too risky to hold back on rolling out the first drug in order for the second drug to have a decent amount of APC payment left over to pay for it. The cost-benefit analysis might indicate that HIV/AIDS has got so out of hand in certain countries that it is better to push out a 90% qualifying AIDS vaccine to avoid later need for even more vaccines and treatment programs, than to wait till there is a vaccine meeting 100% of the minimal conditions.

There are also issues relating to monitoring after the release of vaccines to ensure that resistance to the vaccine had not developed and spread. It might be difficult to demand repayment of an APC payment from a drug that subsequently is found to be much lower-quality than originally thought. At the same time, the solution of holding back its release so as not to 'spend' too much APC allowance on the drug may not be feasible.

Note that 'delay' and the 'commitment to delay' is yet another source of dynamic inconsistency, the presence of which will raise capital costs. If the APC setters commit to delay, then the more that vaccine researchers believe the promise and act accordingly, then the higher the *ex post* incentive of the APC setters to renege on the original commitment to delay. This might bite more if either there are very few APCs in operation at the time, or if the payoffs for those administering one APC are very independent of the payoffs of another APC (so that they fail to internalise the knock-on effects onto other APC programs of reneging on this one), and if there are strong political pressures not to delay even if it has been promised. The whole thing unravels in a self-fulfilling fashion so that vaccine researchers do not trust that the APC setters will in fact ever delay, and they concentrate research efforts instead on lower-quality vaccines.

The mechanism may even have to commit to delay use of a vaccine covered by one APC program even though it is becoming clear that there is little chance of further 'better' products on the way – just so that it can be understood in the case of *other* vaccines and drugs that the APCs for *them* will delay payments, so as to create incentives for research on later 'better' vaccines and drugs based on *those* APCs. If the developer of the first vaccine on the first APC is not disciplined in this way, then in a world of asymmetric information, firms working towards other APC contracts may not believe that it will happen in their cases, and this will distort incentives away from better products in these other vaccine areas. So the degree of hold back is also important when considering the general equilibrium of a *system* of APC programs.

Incidentally all these delays have to be worked into the cost-effectiveness comparisons since they *do* impose social costs and they *are* a part of the APC mechanism[108].

7.8. 'Adjudicating Committees'

This leads to the creation of 'adjudicating committees', that determine *ex ante* eligibility and pricing conditions and *ex post* adjustments to these *ex ante* rules, and how to police and punish those who deviate from them.

Eligibility issues for *ex ante* consideration might include:

1) Vaccine efficacy:

i) against some strains rather than others;
ii) in some regions rather than others;
iii) for some age groups rather than others;
iv) for some groups more than others. For example, a low efficacy AIDS vaccine that can be used to target core groups in some regions that will 'break the chain' that would otherwise lead to greater spread in the general population, may still be more useful than waiting much longer for a vaccine more efficacious against the general population. But how does one incentivise the former without weakening the later? Or incentivise the latter without weakening the former?;
v) against severe symptoms, but not minor symptoms;
vi) if various sub-types of a disease are present in a region, then efficacy against more sub-types in one drug is better than efficacy against just the one sub-type in one drug;
vii) different percentages of efficacy (related to the difficulty of the science of achieving different rates of efficacy), etc.;
viii) with all these changing over time as the epidemiology changes as a virus is eradicated. See the discussion of malaria in Farlow 2005.

2) The number of doses needed. Generally, as the number of doses increases so does the cost of delivery, and as a result, the percentage of targeted groups that actually get the vaccine declines, possibly substantially – hence the impact on DALY's saved and total discounted cost per DALY is much reduced.

3) Side-effects, that may also differ across sub-populations, age-groups, etc. What about side-effects of those who do not comply fully with the prefect delivery protocol? And how does delivery protocol interact with the side-effects issue? If a vaccine has contra-indications this has to be built in as an extra cost for screening/testing, follow-up, etc.

4) Time over which protection lasts. How many booster shots does it need? Each booster shot adds a layer of cost, but efficacy rates also fall if booster rates are not 100%. For example if, in certain difficult regions or especially poor settings, boosters are only taken by 80% of those who took the previous shot, the need for three boosters reduces the efficacy of those initially successfully targeted to 50%. A one-

[108] This is not done in the Kremer figures.

shot vaccine has huge advantages in such circumstances. But in other regions the cost of reducing the number of booster shots by one, may not be worth it.

5) How rigorous should field trials be? What size of samples of the population, and how long would they be tracked until determining the length of protection given by a vaccine? How many studies in how many different countries/environments are needed to assess efficacy against different strains?

6) How long does resistance take to build up? Does this aggravate the difficulties of generating follow-on vaccines, and how is that priced into the early vaccines? There are also cases where a vaccine might weaken the limited immunity built up in childhood.

7) Eligibility issues across APCs. When setting the terms for one APC, care is needed to take into consideration the effect on other APCs (in fact this is a general observation relevant to many other aspects of the APC mechanism). For example, setting 'lenient' terms for an APC covering a vaccine that at a low rate of efficacy might still be cost-effective, might nevertheless undermine confidence in what would be accepted on another APC, so it might make sense to set the efficacy requirement of the first APC higher than strictly optimal. This is also another reason why even when setting up just one or two APCs, care has to be taken not to think too much in terms of partial equilibrium analysis, and, why, if there is a plan to have multiple APCs it makes more sense to create them all at once. Observe how this requires common knowledge *ex ante* of coordination across APCs *ex post*, so as not to run the risk of dynamic inconsistency – made worse now by the fact that the inconsistency flows 'across' from APC to APC. This is the source of a further observation, that the APC mechanism would not be light on international treaties as well as committees.

8) Terms of eligibility need to be set always mindful of the need for follow-on vaccines, to counter the build-up of resistance. Again this is an informationally-heavy process.

Membership of Adjudicating Committees

How should such committees be constituted? Kremer argues that, to enhance the credibility of a vaccine purchase commitment, "appropriate decision makers" on such committee should include "members who have worked in the pharmaceutical industry."[109] This would help "convince potential vaccine developers that the committee would not impose unreasonable conditions after they developed a vaccine." Recently, few emerging-market developers have had any involvement in discussions about setting up APCS. It might be thought that the above logic applies to all vaccine developers and manufacturers, but, if so, what are we to make of the lack of involvement of many of those who will be impacted by an APC mechanism? While mentioning the importance of insulating such committees from "political pressures" and political capture, the notion that industry pressures, and even industry capture, might be an issue are not even considered. Yet, this cannot be regarded as an insignificant possibility. At the time of adjudication decisions, large pharmaceutical firms (the emphasis is placed by Kremer on these rather than on small biotechs or not-for-profit initiatives) will already have heavy sunk costs and – largely based on this

[109] K7:2, and K7:12.

but also based on the need to be tough in other product markets – have large bargaining and rent-seeking incentives. At the same time, due to past antagonism between vaccine purchasers like the PanAfrican Health Organisation and UNICEF, these other institutions "therefore, might have difficulty administering a program designed to increase private-sector incentives for vaccine development."[110]

Rules versus discretion again

Committees in this model are rather like *central banks* – they need time to develop a reputation. They also face the same dilemma of discretion-versus-rules talked about above. On the one hand, discretion allows flexibility in the light of changing conditions, but risks reputational costs and runs the risk of capture. Since all of this feeds into higher capital costs, it leads to a higher APC price. To avoid this problem the system might instead try to specify exact rules, *ex ante*, for how 'discretion' will operate, but this has many of the consequences of fixed rules. Fixed rules remove the reputational and capture elements, but may lock in suboptimal choices, and raise the APC price this way. Both discretion and rules therefore lead to a higher APC price than if these problems did not exist[111].

7.9. 'Purchase Decision Committees' and Developing Country Co-Payments

Each developing country is visualised as having a 'purchase decision committee' vested with the job of making 'efficient' purchasing decisions, and having the power to release resources from the country's sub-account within the program.

These countries (though *not* those *not* covered by the APC program) are required to contribute co-payments towards the purchase of qualifying vaccines as part of a "market test"[112]. Such countries (who supposedly have full flexibility to choose when and how to spend their co-payments) are "forced to consider the suitability" of any vaccine that comes along to "ensure that they felt that the vaccines were useful given the conditions in their countries" and (subject to a barrage of information from developers, all of which would have to get priced into their costs) to consider when they feel that they do "not expect a superior vaccine to come on the market shortly" (apparently scientific communities don't know these things, but most developing countries do). If they thought a better vaccine was on the way "they would be better off saving the funds in their sub-account" and not signing deals[113].

Co-payments would be set in each country's case "just below a country's estimated willingness to pay for vaccines"[114] (something else that would require a great deal of information in advance), since this "maximises incentives to develop vaccines". Get this right for all countries and the overall incentives to develop vaccines (for given overall payments) will "correspond to the aggregate willingness to pay for vaccines," and be maximised. Get it wrong, and incentives are less than maximal and costs of development are higher (Kremer models on the assumption that it is always set right).

---

[110] K7:12.
[111] Observe again, how this trade-off has been greatly reduced, indeed removed, in Kremer due to the simple technological assumptions chosen.
[112] K7:18, K7:47.
[113] K7:19.
[114] K4:10.

Co-payments also *need* to be set exactly correctly, in order to remove developers' temptations to try to extract extra payments from the purchasing countries following vaccine development. The developing country might be willing to agree to these supplementary payments if the co-payment was set below the estimated willingness to pay, but that would just lead to the whole program costing much more than claimed, and the original APC price would be misleading as a measure of the overall price of the program.

A 'purchase decision committee' draws on zero-interest deposits, previously made by the country into the sub-account created when the APC was first initiated. The size of these zero-interest deposits would rise with "each country's respective GNP per capita,"[115] generating a form of tiered pricing. At the same time, as Kremer points out, "vaccine developers need not take the politically damaging step of revealing their willingness to produce additional doses at low cost, thus risking generating enhanced pressures for price regulation"[116] (To optimally make their co-payment purchase decisions, the 'purchase decision committees' would, it seems, have the necessary information to be able to choose between actual and potential vaccines, even though the producers would not be obliged to reveal important sets of information to help them make that choice).

With the bulk of payments into the APC program being from tax payers of richer countries, and with co-payment funds from developing countries linked to income, the overall distribution of costs of an APC program is potentially similar to that of an R&D Treaty – but without the benefits of openness and information sharing of alternatives like more open collaborative research funded from such a Treaty. An R&D Treaty could certainly incorporate an element that allowed countries to choose whether or not to 'purchase' a particular product by drawing down on co-payment funds, but might avoid the secretiveness of technology and tightness of patents that the APC approach would rely upon.

7.10. Co-Payment Problems

Co-payments, however, introduce multiple problems:

7.10.1. Another route to the self-fulfilling development of low-quality vaccines

First, how is it guaranteed that each and every 'purchase decision committee' always acts optimally, and that coordination across all purchase committees is achieved at all times, pushing incentives towards good-quality vaccines? What if several countries face pressure to go for a short-term gain by accepting a vaccine that does not meet the minimal conditions, or even meets the minimal conditions but is short of the perfect vaccine hoped for (and we remember that they are making decisions in an environment of extremely poor-quality information involving expectations over future vaccine as well as current vaccines)? The chance of a later good-quality vaccine falls (in fact those working on the short-term projects have an incentive to feed this belief,

[115] K4:11, K4:53

[116] K7:16. We will see this logic in action again later when we analyse the way that the APC scheme itself might enable market segmentation and the charging of higher prices.

and it is unrealistic to expect them to provide the transparent information Kremer presumes of them) since fewer firms will work on these vaccines since they expect less payment following their development (the global co-payment fund now falls because of the inefficient decision of the purchase decision committees). Even those countries not facing such pressures come to realise that, *given these pressures*, the development of an earlier less-than-minimal vaccine at the cost of delaying a 'better' vaccine is looking the most likely outcome, and it becomes rational for them to join in (especially if they worry that allocations are limited either by the rules or by production capacity). The less-than-minimal vaccine becomes a self-fulfilling equilibrium. Indeed, this is self-confirming of the committees' choices; consumers need never know the committees were wrong, but those who did not join in the choice of the poor-quality vaccine would look 'wrong' and as if they had mismanaged the choice. Short-term gains for some are offset by a longer-term loss for all. In a multi-APC setting, this would even knock on, in a sort of contagion, to other vaccines, with bad vaccines driving out good vaccines.

7.10.2. Would all developing countries agree to deposit?

What is the mechanism to ensure that all developing countries coordinate by depositing at the start of the APCs? Would those who refuse to deposit be barred from using any vaccines resulting from APCs? If it was known that they would not be barred, the incentive would be to hold off depositing and free-ride on the deposits of other developing countries into the APC (in the knowledge that the APC price, and the call on developed country taxpayers, would be allowed to rise to make up for the loss of deposits). A multi-country Treaty agreement would be needed in advance of the setting up of an APC to avoid any free-riding.

7.10.3. Problems when accounts pay zero interest

Interest must not be paid on co-payment accounts (to "prevent bad purchase decisions") otherwise developing countries would be under less time pressure to use the account and reach agreement with vaccine companies (vaccine companies are automatically under time pressure to sign deals since their patents are time-limited). But this means that developing countries are effectively being forced to pay for the high cost finance that goes into pharmaceutical research (which includes any extra capital costs caused by the use of APCs) even as they are denied any interest on their deposits themselves. They certainly would not be allowed to invest at the rates that those investing directly in the pharmaceutical research would be expecting to get via an APC. For developing countries, might it not be cheaper in expected discounted terms to commit to a front-ended R&D flow (via an R&D Treaty perhaps) in support of vaccine R&D than to an expensive co-payment fund?

What happens if countries disagree with a vaccine adjudication committee's decision not to allow a vaccine to be used within their borders, given the zero return on the co-payment account while they are forced to wait until the adjudication committee changes its mind or agrees to some other vaccine being made available?

The argument made by Kremer is that if interest was paid this would give developing countries bargaining strength *vis à vis* vaccine developers and prevent developers recovering their R&D costs. However, it is not clear that the bargaining game that developing countries would face would not then strengthen pharmaceutical companies

(particularly when we move away from Kremer's assumption that the industry is perfectly competitive).

And how politically acceptable is it to require impoverished developing countries to deposit into zero-interest earning accounts while arguing that other countries not covered by the program but benefiting from it need not, and that those in rich economies paying into the program only need to 'pledge' their contributions?[117]

7.10.4. Corruption and political pressures to use (or not to use) the co-payment accounts

Co-payments have to be held in an account. This raises the issues of who controls the account and of what their incentives might be. Plenty of possibilities arise for corruption and 'pressure' onto purchase decision committees both from pharmaceutical companies themselves and from the countries where pharmaceutical firms are based. Since the bulk of payment for the vaccine would come from the vaccine purchase program, and only a small amount from co-payments, large deals can be leveraged by relatively small enticements (for example a 5% co-payment generates a 20:1 ratio). Either a vaccine developer could offer a kickback on the purchase price – like offering flexible antibiotic prices if a country should happen to purchase a particular vaccine that happens to have been developed by the maker of the antibiotics –, or tied deals, or even bribes. These would generally favour larger rather than smaller pharmaceutical companies and biotechs (especially if the biotechs only concentrate their efforts on developing this vaccine). Several of these practices might even be technically legal and hard to disentangle from obviously corrupt practices. Implicit tied deals can be very hard to detect. Or countries might add relatively minor enticements to aid packages that serve to tilt the recipient in favour of one deal over another. Such payments from third parties are very difficult to regulate. The fact that the alternative to an early deal may be a much later deal (especially if there are limitations to the total allocations of a particular vaccine), reduces even further the expected discounted cost of enticements.

Again, all of this has to be worked into the expectations of the investors into R&D. To the extent that they do not trust that this kind of behaviour can be ruled out, the developers of 'good' vaccines will hold back, *and financers will not finance them*. Similarly, those vaccine players who feel that they will be in a weak position to exploit these strategies may be deterred from investing in the first place, *and financers will not finance them* either. This has certain self-fulfilling aspects to it. This necessitates yet more layers of regulatory organisations. And if the industry is not competitive, the bargaining game involving developing countries and large pharmaceutical companies is strengthened even more in favour of the larger pharmaceutical companies.

Kremer[118] argues that the technical requirement is the "first and most important" line of defence against corrupt practices, but he recognises the need for purchase decision committee "whistle blower procedures...to protect, or even reward, committee members reporting attempts at bribery by vaccine developers." But ruling out corruption and undue pressure is not so easy. Even if there is *ex post* adjudication of

[117] Though we will shortly see that this is not strictly true in a multi-vaccine APC.
[118] K7:21

technical tests, corruption is risked simply by alternative routes, especially by capture of one or several of the APC-generated committees, pressure elsewhere on the trade playing-field, secrecy and lack of transparency (Kremer is extraordinarily naïve about the degree of transparency of information provided by large pharmaceutical firms, transparently that is utterly necessary if purchase decisions are to be efficient). And how many layers of committees and how much policing will be tolerated or enacted in practice over the 20 to 30 year horizon of run up to, through, and out of an HIV vaccine program?

But there are forces working the other way. A decision to make a co-payment usually carries a further financial obligation – to pay for all the other steps necessary to ensure a vaccine is delivered. In other words, the marginal cost of deciding to use the sub-account may be quite large, and very large if the choice is not discrete (for example if agreeing to use a vaccine leads to a large needed program rather than a purely marginal program). If a government is not prepared to pay all these extra costs, it may put non-price obstacles in the way, like requiring more evidence of efficacy, changes in the product, etc. Worries about this feed through to investors and the APC price.

7.10.5. Political short-termism and political reality

Politicians have clear incentives to sign deals that would not be optimal in a multi-period setting. How is it ensured that one generation of policymakers/leaders does not have an incentive to use up the account on visible less efficacious current vaccine programs, to "be seen to be tackling a pressing medical problem" instead of less visible, but better, later programmes (that may fall in the period of office of somebody else anyway)?

Imagine a situation where a country has held out using its sub-account, and yet a superior vaccine has not been developed. What are the political ramifications if neighbouring countries used their sub-accounts on less-superior products and got their populations treated while this country did not? What if the domestic buying committee cannot convey to the general population the scientific information on which it based its decision to hold out (it is an asymmetric information problem after all, and, besides, 'everyone knows' that the purchase committee has a dominant strategy to always reveal information in a way slanted to supporting its decision, even if wrong)? They may not even have the information to provide to the general public, since it is argued that one of the benefits of the co-payment system is to shield pharmaceutical firms from revealing too much information. And what information channels do they have anyway in resource-poor settings?

And even if they could reveal information truthfully, what if they are very risk averse, such that while waiting a bit longer might be the globally optimal thing to do, it would, privately, risk the outcome of not getting a vaccine in a politically timely fashion? Running the risk of drawing from the 'wrong' tail of a superior probability distribution may be just too risky compared to the alternative of getting, for certain, the middle of an inferior distribution. What if we throw in that members of the committee have short horizons (some of them are political appointees and others have career objectives and so on)? What if we throw in the notion that purchase decision committees might tend to think there is safety in numbers (in the same way that professional investors, worried about how they might be judged by the public relative to other professional investors, 'go with the herd')? Not going with the prevailingly

available poorer vaccine chosen by other committees, the 'herd', and waiting for the next vaccine may be a very risky strategy.

There is even the danger in a multi-APC framework where multiple co-payments have been deposited in a fund, that vaccine developers would have the leverage/incentive to extract even more from purchase decision committees for early vaccines but at the loss of funds to make co-payments for other later vaccines. What are the mechanisms to ensure that enough is left for later vaccines and that research intensity on those vaccines is not harmed? Sub-accounts might have to make sure that each disease carried its own allocation, and no pooling was allowed[119].

Fancy adjustment mechanisms – based on the notion that it is inefficient to use up all of the allotted co-payment account on the imperfect first vaccine since this harms incentives to research for better products – may not work as a political reality. To the extent that this is understood in advance, researchers and investors avoid working on vaccine programmes that might have been globally superior. And to the extent that countries break from the APC anyway, the act of taking part in the first place will turn out more expensive than if they had kept out.

7.10.6. Placing time-limits on APC programs: Sunset provisions

Kremer suggests that a sunset provision might be written into an APC program such that if, say, after fifty years, no qualifying vaccine had been developed, or at some earlier point a scientific committee had determined that the burden of, say, malaria had been significantly cut, then countries could have their co-payments returned. But this has problems too. Setting the *optimal length* of the sunset provision is difficult without good knowledge about the underlying current and expected technology – and it *has* to be set optimally since those countries holding co-payments in funds will get zero interest, so setting the sunset provision too long forces costs onto them[120].

Similarly, the sunset provision forces risks onto vaccine developers, which, again, has to show up in their capital costs. The sooner the sunset provision is known to cut in, the greater the risks. Kremer is too dismissive of this, suggesting: "but biotech and pharmaceutical firms routinely have to bear risk that alternative technologies will render the projects they are working on superfluous."[121] This is true. But it does not take anything away form the fact that a *new* risk, additional to their "routinely" bourn risks, has been added to their burden via the sunset provision, and that this *has* to be factored into their capital costs.

Once the program has *any* time limit, there is a risk of winding back to poor quality vaccines as the time limit approaches, and of creating extra risks and capital costs for developers. Indeed, since this is still a patent-based system with payment based on 'price times quantity purchased' it does have a time-limit built in, an 'option'-based time limit that is triggered into existence the first time a vaccine is created anywhere in the world (APC-qualifying or not-APC-qualifying). For example, what happens if a 'better' vaccine is developed but it is 'late' in the life of an APC, because the program is approaching (or is even in) the competitive phase, with earlier-developed vaccine/s

[119] Observe that as with any other risky situation, pooling of funds might have cost savings.
[120] Incidentally, it also feeds into the expected value of the co-payment funds, hence the needed payments from the donors.
[121] K7:26.

(both any that qualified for the program and any that did not) now being produced (or being close to being produced) by generically competitive firms? Are those in the program forced to keep to sales of the newer vaccine satisfying the stricter conditions of the program? Are they stopped from using (or waiting to use) cheaper generic competitive products based on the earlier vaccine/s? It is not clear how this competition might feed back to affect payoffs of those firms working according to the strictures of the program, particularly those working on better-quality vaccines. At the very least it increases their risks and capital costs, and this alone may deflect them from investing too heavily in higher-quality vaccines.

Equally, it might seem 'better' to use up the program than risk failing to distribute the whole amount (and the risks that this might happen, and that the program ends without fully dispersing the 'perfect vaccine'-payments before generic competition hits in, might reduce the certainty of payback and generate a higher APC price *ex ante*).

In addition, if the pricing mechanisms varies inefficiently over first, second, and third vaccines anyway, there may be incentives to delay/not delay if the pricing of later drugs is wrong. Getting the pricing wrong distorts incentives towards delay or wasteful haste. A time-limited program might help correct the former problem, but possibly only at the cost of poorer-quality.

Since 'late' vaccines have trouble recouping their development costs, this pushes us towards using time-unlimited programs, but the risk then is of being tied in to a poor non-working program.

7.10.7. Allowing developing countries to buy non-scheme vaccines but not with their co-payment funds

Interestingly, Kremer does not seem to think that the unregulated sales of poorer-quality vaccines that did not get approved by the adjudicating committee, would be a threat to the vaccines intended to flow from the APC, or therefore that any mechanism or Treaty, with APC or non-APC countries to prevent such sales, is justified. He argues that "if a vaccine turned out to be socially useful, but not good enough to qualify for purchase under the program at the promised price, this would not preclude individual countries from purchasing the vaccine or other donors from purchasing it"[122]. Since, of course, the insiders to the APC program would not be allowed to use their co-payments to pay for the non-qualifying vaccine, the marginal price of purchasing the vaccine is made higher than the marginal price of a vaccine qualifying under the program. This is supposed to stop countries in the program from using it.

If a country were to get itself fully treated with vaccines from outside the APC, say for HIV/AIDS, it would *not* be allowed a rebate on the co-payments it never used, since that would destroy the whole point of the exercise, and the *ex ante* knowledge of the ability of countries to claw back co-payments by breaking the scheme would feed into destroying incentives towards vaccines meeting the programs conditions. Once the co-payments are deposited they are lost to all but qualifying vaccines for ever. Even the sunset provisions would have to be modified, since, to be dynamically consistent, the country should have *no* incentive to renege in the knowledge that they

[122] K7:13.

can get their co-payments back, and a sunset provision would be treated as a guarantee by a country that it could get all its co-payments back (with delay) if they are not spent because of reneging, with the incentives to renege rising as the sunset provision gets closer.

There is even a further danger of encouraging poorer-quality vaccines. Countries who have gone outside the program and treated a proportion of citizens will find themselves still with co-payment funds to spend on those remaining to be treated. If better, but still poorer than optimal, vaccines subsequently come along, they are able to bid more for such vaccines than would have been the case had they not already gone outside of the scheme, with very little reason for firms not to exploit the fact or for the countries to resist. The whole set of post-development rules is wrecked by those going out of the program.

If the country *does* get itself *fully*, or even just partly, treated with poorer-quality vaccines generated outside of the APC, it would end up doing so with inferior drugs than stipulated in the APC, at a higher overall price, and would have to forego or waste co-payments in the process.

Since the APC is a contract with developing countries, would they be legally bound *not* to buy outside of the program? Or is it that they just would not be allowed to use their co-payments to pay for purchases made outside of the program? If they are legally prevented from buying outside the program will this not seem (and be) coercive? If they do buy outside anyway, will it look unfair if they are prevented from having some or all of their co-payment funds returned? If it was the fault of those setting up the APC that a country then chose to go outside of it for its vaccines – because the APC qualifying conditions were set too high, or the APC did not work correctly and was not generating quality vaccines – countries would end up being punished for the mistakes of the APC setters.

7.11. Ways to Extract Even More Developing Country Payments

In the APC models, co-payments are set exactly just below countries' willingnesses to pay. First, to ensure maximal incentives to develop each vaccine. Second, to, supposedly, minimise incentives to extract even more payment from developing countries *ex post*; it is after all still a patent-based system. Third, so as to give a true measure of the cost of an APC program.

Unfortunately, there is always *some* incentive to extract extra payments beyond co-payments, and there may be quite a lot. If not disciplined, this may show up in extra payments related to vaccines themselves (such as lower-quality vaccines sold at higher prices, and the consequent crowding out of higher-quality vaccines) or in components relating to the distribution/follow-up of vaccines, or in side-payments related to other deals, or even through bribes[123].

[123] This is not inconsistent with bribes to get a country to take a particular vaccine, since the one benefiting from the bribe and those being charged for the vaccine do not have to be the same.

7.11.1. First extraction method

If the co-payment simply turns out to be lower than a country's actual willingness to pay, vaccine developers have an incentive to extract extra marginal payments up to the level of social surplus.

This could be because the co-payments were set wrongly. Given the amount of information required to set co-payments correctly, this is quite likely[124]. But it might even have been deliberate. Those setting the co-payment terms might have erred on the side of caution and set them too low, not wanting to put developing countries off from joining. Or – realising that there is anyway some ability (which they may not be able to police) to *ex post* extract extra payments – those setting the payments had less incentive to get them exactly correct or simply wanted to keep the apparent overall costs of the scheme lower. Or, it could be that the epidemiology changed dramatically, and this was not fully predicted, such that the co-payments fell short of the social surplus and eventual willingness to pay. The problem is that whenever, and for whatever reason, the *ex ante* set co-payments fall short of the true social surplus, there is always an incentive to extract extra payments up to the social surplus.

But it could also be that co-payments were set perfectly correctly. If the co-payments were set *ex ante* on the basis of *expected* willingness to pay, and if this expected willingness is the *average* of a distribution, then half the time the *ex post actual* value of willingness will perfectly naturally be less than the *ex ante average*. Therefore, ex ante, half the time there is some extra 'willingness' to exploit, and conversely, half the time the co-payment will turn out to be above the actual willingness and there will be no further 'willingness' to exploit (but it is not clear if some co-payment would be returned in this case). It might be thought that some adjustment mechanism could be worked into the APC program to solve this problem, but there is essentially no way around it except by either: i) allowing *ex post* payment to adjust *up to* the actual *ex post* willingness, and also ensure return of a portion of co-payments if the *ex post* 'willingness' turns out less than the average willingness presumed *ex ante* (though one can imagine the information/bargaining nightmare); or ii) set co-payments at the bottom of the range of expected willingness to pay, and allow top-up later through extra payments. Both these options generate another layer of institutional detail[125]. Failing the use of one of these options, developing countries face a 'one-way' bet. If the willingness turns out to be higher than was initially predicted, there will always be incentive to extract more payment, but if the willingness turns out to be lower, there is no return of co-payment. On average they pay more than the co-payment.

If a stream of drugs is intended to come out of the process, these extraction effects will show up in the early poorer-quality vaccines taking a greater share of the co-payment fund than was strictly optimal, leaving a need later for additional funds to pay for better drugs (the early over-use of the co-payment fund was at zero marginal cost to those dipping into it, even if those who come later have to pay more). There is an obvious danger that these funds may not be so easily forthcoming later. To the extent that it is not, countries get lower-quality vaccines at higher prices and (to the extent that developers believe this to be the case) developers face incentives to work

[124] Observe how messy this becomes if the consequences of each disease overlap; so, also, must the payment schemes in a very precise fashion if uncertainty and distortion are not to be created.
[125] Not to mention problems for all the post-development adjustment rules.

on lower rather than higher-quality vaccines. Observe how this, like many of these affects, is driven by the fact that the price is way above the manufacturing costs of the drugs, with most of the price being the R&D and capital costs.

The APC calculations presume no probability distributions over country 'willingness', and, instead, use point estimates of a system in steady state epidemiologically and population-wise – and they are thus able to ignore this extraction problem. It is not such a sensible assumption if the epidemiology (or population) is less clear and less in steady state – for example in the case of HIV/AIDS.

7.11.2. Second extraction method and the paradox of quality

There is also a natural ability to leverage more than the originally stipulated co-payments, once they are sunk, using the threat of delay – which has a real cost to countries. Potentially this can lead to a great deal of extra surplus extraction.

If a vaccine has been developed, and it is worth exactly the discounted value of social surplus (it could be that 'Case 1' extraction has already taken payment up to the value of social surplus), this might suggest that no *extra* surplus could be extracted. However, this is not the case if the co-payments are sunk, in the sense of only *ever* being available to spend on vaccines for *this* disease, and 'lost for ever' otherwise (we see elsewhere that there *should* be no chance to get co-payments back if the APC mechanism is not to be harmed by the *possibility* of this)[126]. In this situation, any expected delay in the release of vaccines creates negative social surplus to the developing country, and the positive discounted value of this negative social surplus (or, more precisely an amount just under this value) can be extracted above the co-payments, in exchange for immediate access to the vaccines. The extra marginal payment that developers might seek to extract is always lower than the marginal payment they would have had to extract without the co-payment fund in place. It might be argued that those administering a country's co-payment fund would surely take into consideration the *whole* marginal cost of the vaccine, *including* the costs of drawing down the co-payment funds. But this is all presaged on the notion that there are vaccines generated by the APC scheme *other than this one* that they can buy. If this holds, then, indeed, the marginal payment for *this* vaccine includes the drawing down of the co-payment funds. Otherwise it does not.

A simple analogy might help. The discounted value to you of lifetime membership of a health club is $1000. You approach the one health club in the entire world, and the only health club that will *ever* be in existence. How much are you prepared to pay to join it? Let's say that, instead, you put $1000 into an externally-held fund that you can never touch for anything other than health club membership. Now, there are hundreds of equally good health clubs. If you approach any one of them, could any one of them ever take more than the value of your fund from you? Return to the world with the one health club, and where there will only ever be the one health club. When you approach the club, is the situation any different now that they know that you have pre-committed the $1000 and that they are, and will only ever be, the only health club in the world? If they try to make a small extra charge on top of the fund payment in exchange for not making you wait five years (they have others waiting to join after all

[126] We see elsewhere the possibility of 'sunset provisions' but we also find problems with the operation of these.

and places are rationed), would you refuse, leave, wait five years, lose five years' worth of use of the health club, and still have to pay over the total value of the fund to join the club (and maybe face the same request for an extra fee or risk another five year wait)? If they offered you instant membership at $200 (plus the fund) would you make a commitment never to join even if you know that there would never be any other health club you could ever join and though you continue to value membership for life at $1000, and forfeit your fund entirely? Would you regret ever setting up the externally-held fund?

The logic in APC models is that 'potential' vaccines will discipline this behaviour. Indeed, to the extent that other vaccines are 'on the horizon', there is downwards pressure on this extra payment[127]. The more imminent are other vaccines, the less extra surplus can be extracted from *this* vaccine. But it does not remove the effect entirely. And it requires a great deal of knowledge about future vaccine possibilities, especially for those countries buying vaccines[128]. For them, asymmetric information regarding vaccine possibilities (fed also by those developing vaccines) will feed the ability to extract the extra payments. The effect may also be stronger if attempts are being made to limit the supply of vaccines for post-development re-adjustment purposes, since the limitation of supply may impact extra upwards pressure on price[129].

This even bites in situations where co-payments are already greater than the social surplus. The co-payments are sunk (and effectively non-retrievable) and will have to be spent on a vaccine for this disease anyway. The marginal payments extractable from delay are therefore still positive, even if the co-payments already represent overpayment[130].

The paradox of quality

Observe how, paradoxically, this relates to quality. If the perfect vaccine comes out the first time, there *will be no* other vaccines, and this will generate maximal incentive to extract extra surplus. However, given the quality of this vaccine (which we presume, on average, reflects also its costs of development), it is already worth the total of the co-payment funds anyway, so these extra funds will have to be extracted from elsewhere. If it could be sure of no competition, the developer of this vaccine could seek to extract up to the value of a country's social surplus *on top of* that already deposited in the co-payment fund (limited by any limitations on the country's ability to contribute more than already deposited in the co-payment funds). There must therefore be an assumption underlying the APC model that, even in this case, behaviour is disciplined by *competition* from other vaccine developers. The notion seems to be that the ability to extract extra surplus is profitable and would attract entrants or the threat of entry. If this *potential competition* is perfect, then an individual firm could never extract more than the already deposited social surplus. It is no surprise to find that the Kremer framework has no sunk costs and perfect access to IPR and, therefore perfect *potential competition*.

---

[127] Imagine the above analogy, with the added knowledge that a few more health clubs would be created eventually – but maybe after a long delay.

[128] In the analogy, imagine not knowing for certain if any other clubs would ever be built.

[129] In the analogy, the membership numbers at any club are being limited to first-come-first-served.

[130] In the analogy, it really does not matter to the behaviour of the club-owner whether you put $1000 in the fund or accidentally put $1200 in (unless he/she feels some sympathy for your mistake).

But this does not even remotely describe the world of pharmaceutical research, never mind HIV, malaria, or TB vaccine research. As soon as a few large firms and sunk costs enter the equation this reasoning breaks down. The ability to extract these surpluses does indeed attract more investment into vaccine research and, in equilibrium, the expected cost of this activity just offsets the expected value of this extra surplus. But this is bad news. First, if the APC price was set optimally in the first place, then this is pure waste – and it is at the cost of developing countries. Second, this means that the co-payment contributions of developing countries seriously under-reports what they do end up paying (in expected social surplus) for vaccines. To them this extraction is a hidden cost. Third, it means that the APC price under-reports the eventual overall global costs of vaccine development. Fourth, developing countries are worse off after the scheme, and certainly worse off than a scheme such as a global R&D Treaty that could have extracted their social surpluses without tying them in to the use of co-payment funds.

This is part of a general paradox of the way developing country payments relate to the quality of vaccines. Extra extraction is only limited by the possibility of further potential, better, vaccines. Strangely, therefore, poorer-quality is needed at some point to discipline price. As the possibility of better vaccines recedes, the incentive to extract extra payments from what is actually produced rises. This might be because good vaccines have been derived (as just described) but it could also be in cases where poor-quality vaccines have been derived. If poor-quality vaccines are derived (that are worth a much smaller proportion of the co-payment funds than perfect vaccines), but higher-quality vaccines are expected, the pressure to extract might be low. But if poor-quality vaccines are derived *and* are expected, then the pressure to extract might be high. The worse the system does – in the sense of the lower the chance of there being better (or any) vaccines later that a country might want to purchase – the higher the ability to extract extra surplus. This surplus extraction is on top of the low supposed co-payment funds going to the poor vaccines (that cost, on average, less to produce), so, for their producers, this might be quite a profitable outcome. And, unlike the 'best' quality vaccine, the extra payment can be extracted from the co-payment funds themselves (at zero marginal cost to those spending those funds) rather than requiring sources from elsewhere. So, a deteriorating system may actually reward the earlier developers of poorer-quality vaccines (this may itself feed developers' incentives), reduce the remaining co-payment funds available for later vaccines (further feeding those incentives), and make poorer-quality self-fulfilling.

7.11.3. Third extraction method

Even if firms are legally obliged to sell to the program if they meet its conditions (current APC set-ups do not presume this), a vaccine that does not quite make the minimal conditions can *legally* be kept out. If individuals living in countries covered by the program try to buy this vaccine privately (or indeed their governments buy it for them), this will involve payments beyond co-payments.

Therefore, one further profitable way for a firm to extract more payments both from those in the program as well as from those outside of the program, might be to set the quality of a vaccine just below the threshold for the program (this is increasingly profitable – and possible – the more a firm's vaccine is coming in at a quality close to the threshold), supply the outside market and the wealthiest parts of the inside market

for a while at higher prices, use the firm's hold over IPR to suppress the development of vaccines based on this vaccine, and only later, when profitable to do so, or after pressure to do so, push the vaccine over the threshold (this may even simply mean manipulating information in a world where firms have great informational advantage over APC setters). There may be little pressure working against this. Other potential competing vaccine developers may not wish to sink huge costs to develop a competing vaccine if they realise that the incumbent can modify their vaccine (or simply modify information) to compete in the program-based market, and if they believe that the incumbent – for all their refusal to supply the program-based market – will ultimately get that market. Even if the APC rule-makers exercise discretion and reduce the threshold, there is no obvious way that a firm could be legally obliged to supply the market under a new set of rules from those they had been led to expect. And, in fact, lowering the threshold only makes the problem worse since it harms even further the position of other potentially competing vaccine suppliers. Again, we are reminded of how utterly important is the assumption of perfect competition everywhere, and always, in the Kremer model. Once we have just one or two firms, we are in real trouble.

The knowledge that the co-payments are sunk may even feed this behaviour, since such behaviour actually *creates* ability to extract even more payment from developing countries[131]. It might even be the same firm that picks off the wealthiest parts of the program-covered market, that then has the ability to extract more payments from those who remain in the program-covered market, giving the firm even stronger incentives to engage in this behaviour in the first place. It also increases incentives to adapt strategies (including secrecy on IPR) to weaken the threat to this behaviour posed by other vaccine developers. Worryingly, setting a *higher* minimum condition is more dangerous if there is a risk of this situation arising.

This may seem to be very cynical behaviour. However, firms are not charities, and will understand in advance that stock markets will punish them if they do no profit maximise. The lesson is that firms may themselves prefer not to be forced from the start to face situations that will generate certain kinds of behaviour. Firms may prefer to avoid from the start schemes involving co-payments.

7.11.4. Commonalties to these extraction methods

In all of the above cases, the APC price under-reports the true cost of the vaccine.

Clearly, once the co-payments are sunk, there must be a guarantee that this *will* cover the entire country-cost of any vaccine. Observe that if there was competition in the supply of the vaccine, even if there was only the one vaccine, then this extra social surplus could not be extracted in cases 1 and 2. The situation described above arises precisely because there is only one supplier. In addition, with only one supplier, if the time-cost of delay to a country is higher than the opportunity cost of delay to the firm, this puts the country in a weak bargaining position. This would be especially the case in a health crisis.

[131] The more who get treated outside the program, the more APC co-payment per capita is left for those still inside the program. Extraction of extra payments for lower-quality vaccines for those who remain to be treated on the program is now easier.

The third way to extract extra surplus was simply a function of the lack of competition in the market for IPR. Once a firm has a hold on IPR, it can use that to create a lack of competition not only in the supply of the current vaccine but also in the supply of other competing vaccines. In this case it was also helped by the way the APC program has created an 'inside' and an 'outside' market and the legal ability (via IPR) to deny to the inside market products it openly sells to the 'outside' market.

All three problems are related to the way the IPR to vaccines – once developed under an APC – is still held in the hands of the developer. Given the huge dependence of APCs on ultra-tight IPR, how is it guaranteed that this will not backfire with the IPR system being used to delay access?

All of this, incidentally, potentially also feeds the bigger-picture problem of delaying the production and release of a vaccine. If it is known that extra payment can be extracted after development, then if the value of the social surplus is rising greater than anticipated, then so is the value of these extra extracted payments. This acts like a negative discount factor, slowing intensity of research. Again, this is another reason for encouraging multiple suppliers of the vaccine rather than just the one.

In addition, we have seen that there may be some conflict with the post-development APC redistribution system; a system that needs to generate limited quantities and prices of poorer-quality vaccines in order to make the APC actually work, but which might then feed some of the problems of extraction. Intuitively, at the heart of *this* problem is the fact that the post-development redistribution rules are all based on the total pool of actual and *potential* vaccines, when in many real-world scenarios, actors are making judgement only on the basis of those vaccines currently in existence. This naturally risks the break-down of those rules, all the more likely if asymmetric information bites (we see once again how basic assumptions about highly symmetric information are being made to hold the APC approach together).

All of this suggests strict policing of any extra payments, and the benefits of a quicker move to generic manufacture even in a world with co-payment funds (and maybe even more so). And, once again, it alerts us to the dangers of using a system where the property rights to vaccines are not publicly owned once the vaccine is discovered. Maybe there should not even be a co-payment system in the first place? Though co-payments were trumpeted as a way to ensure that developing countries would choose quality vaccines, maybe if there are to be any co-payments they should pay for something like the handover of the IPR itself followed by rapid generic competition to prevent incentives to extract more.

7.12. Forcing Countries Not to Sign Contracts

There are many political and practical difficulties of getting countries covered by the program (and, it will turn out, also those *not* covered by the program) to refrain from signing too many contracts on inferior vaccines that nevertheless meet the minimal conditions, and to refrain from using altogether vaccines that do not meet the minimal conditions.

This will be compounded by significant coordination problems across countries; if it is believed that no other country will do excessive deals on the less-than-minimal

qualifying vaccine, the 'punishment' for the one country doing excessive deals might be more credible, but if countries cannot coordinate or discipline each other then they might *all* try to sign excessive deals knowing that the 'punishment' will be light.

Aware that another drug is on the way, the developers of the first drug would have an incentive to get as many contracts signed as possible (including corruption of the evidence base regarding the efficacy of vaccines, etc. Marketing for this would also get priced into the APC *ex ante*). Asymmetric information helps them. Aware of the 'limited allowance' under the terms of the APC program, unless they can be sure of the second drug's development, individual countries covered by the program might rather get in early on the first drug deals than risk the exhaustion of all of that drug's allowed distribution rights and have to face the risk of having to wait for the second drug to come along. Given a 'life and death' choice over the life of its citizens they may be very risk averse to being left out. This would lead to political pressure to break the APC allowance and also pressure for the use of the countries co-payments in ways that break the APC altogether (including illegally shipping in generic copies from non-APC producers). Signing up to bad deals becomes another self-fulfilling equilibrium where there is worry of getting 'left out' when allocations are being restricted.

Kremer suggests that the first vaccine developer will get a larger share of vaccines anyway since they will sell vaccines to immunize the backlog of un-immunized adults, while follow-up vaccine developers will be limited to new cohorts of children. But, this suggests that those who sign deals with the first developer may find they are prevented from signing deals involving their co-payments with later better products. Being locked into inefficient contracts may be bad, but unavoidable.

And countries would need to be prevented from buying from anyone but the official source. However, since the previous research costs of any firm are sunk, those close to development but who get 'pipped' at the post by other firms who win the official deal, may still carry on to develop and try to sell by whatever means they can. Like those who sell cheap last-minute airline tickets, anything above marginal costs (in this case the extra costs incurred since they lost the APC deal) will be profitable. *Ex ante* this is inefficient and should in principal be prevented. Heavy-handed tactics might be called for, like demanding payment from the co-payment pool to compensate the legitimate vaccine patent holder if a country is found to have worked around the program (though this always risks litigation and political instability). To the extent these heavy-handed techniques are not dynamically consistent, the mechanism is harmed and the APC price has to be higher. In a general equilibrium model, the co-payment pool would pool resources for many different vaccines (hoping to benefit from the 'Law of Large Numbers') and this 'punishment' may harm purchases on other APC vaccine programs.

7.13. When Countries Not 'In the Program' Destroy the Program

Those countries outside of the APC program get benefits from the program via the vaccines that it motivates, yet *they* do not have to hold co-payments in zero-interest bearing accounts. Meanwhile, their behaviour, 'unregulated', outside of the program feeds back onto the workings of the program.

In particular, the value of the APC to researchers, and hence the price of the APC, is based on the market that it *adds* to the 'initial market size'. The 'initial market size' *includes* markets *not* covered by the APC. The guarantee of a minimum price for a vaccine meeting minimal eligibility criteria is necessary for credibility, yet, if an imperfect vaccine is launched that meets less than the minimal conditions, and those countries *not* covered by the APC program buy and use it, this will reduce the value of the *expected* 'initial market size' for any vaccine produced under the program. Knowledge of this possibility reduces *ex ante* incentives to develop the 'good' quality vaccine and requires a higher APC price if such vaccines are to be produced. If no more co-payment is extractable[132], then this will have to come from richer countries contributing to the scheme. The situation may have arisen simply because the non-APC countries failed to co-ordinate, or that they did not have access to the same information as those running (and in) the program (or, indeed, have been fed more marketing information than those in the program). Kremer's assertion that if a poor-quality vaccine is developed, countries both outside the program as well as those inside it (though, in the latter case, without being allowed to use their co-payment funds) should be allowed to buy it, is simply wrong by elementary arithmetic. Buying it contributes to the destruction of the market motivating the good quality vaccines.

Similarly, if the APC is attempting to adjust APC allocations post-development so as not to give all of the APC market to one vaccine, if those outside the program are not doing likewise, and simply increase their usage of the vaccine, it weakens the APC exercise; the suboptimal behaviour of the outside countries destroys 'initial market size' that was to be factored in as part of the market sizes of the later-to-arrive APC vaccines. Intuitively, those outside the scheme are not pricing according to the pool of all potential vaccines that the rules of the scheme are attempting to price to. As always, knowledge of this possibility *ex ante* would reduce incentives to develop the good vaccine and thus require a higher *ex ante* APC price.

7.13.1. The need for an international APC Treaty?

Countering this – in order to keep the APC price down, and, indeed, to keep down the costs of *all* vaccine purchases including those not covered by the APC – might necessitate an international agreement requiring *all* countries, whether in the program or not, to only buy vaccines that satisfy the minimal conditions of the program, to adopt the same post-development adjustment programs, and to police each others' behaviour. Without an international agreement, individual contracts with non-APC members to prevent use of less-than-minimal vaccines might be mutually beneficial, but are highly unlikely since they would be unenforceable, with too many 'cheating' from the socially optimal equilibrium (it's the usual prisoners' dilemma problem). If an APC Treaty would be needed anyway, maybe an R&D Treaty would make more sense? The two should certainly be contrasted and compared.

Not only might countries outside of the program sign deals for their own usage with the unofficial suppliers or sign deals with official suppliers that would take the official suppliers beyond their APC limits, but they might also be used as a route around official suppliers via parallel trade. To keep costs of the APC down (since all

[132] This turns out not to be a redundant phrase. It might seem that since the co-payments are set to extract all developing countries' willingnesses to pay, that no more could be extracted. This is not necessarily so in the situation being here described, as will become clearer below.

uncertainties, including the uncertainty of market size, has to be factored into firms' risks and hence in to the APC price[133]) there will need to be strict banning of parallel imports into those countries covered by the program as well as out of those countries covered by the program if the drug inside the program is cheaper than in non-program countries. Since the APC is supposed to be a 'market enhancing' measure, for it to approach any of the measures of cost-effectiveness claimed for it, the market it is supposedly enhancing has got be protected at all costs. For some vaccines (for example for AIDS) this might be important. Outsiders to the program must pay higher prices than insiders to the program. Will Russia/India/China, etc. be willing to pay such non-program prices if they are kept outside the program? Countering parallel trade might necessitate an international agreement requiring all countries to join, or at least to police it (it would also enhance the APC if this agreement covered drugs not covered by the APC itself). Failure to set this up in advance also has to be priced into the APC price.

7.14. The Problems of 'Existing Market Size' – Another Option Price Component in the APC Price

Just as non-stationary technology causes problems for the optimal APC price, so too does the non-stationary nature of the market size. APC calculations so far have, for example, been calculated on the basis of steady-state levels of prevalence everywhere. The APC price should be set according to the expected prevalence of disease in the country, or even more precisely the prevalence expected at the time of discovery of the vaccine[134] and should also be a function of expected non-APC-generated, 'initial market size'[135] which is a function of prevalence elsewhere.

Predicting prevalence is not easy. For HIV/AIDS, for example, prevalence rates are growing rapidly but in an uncertain fashion so that the non-APC-covered 'initial market size' is also growing in an uncertain fashion. Even for 'established' diseases for which the pattern of prevalence is more stable, programs may be initiated that cut or eradicate the disease (e.g. malaria) in certain areas. Population rates of growth may vary too. How does all this factor into the overall APC price? For example, a scheme initiated to eradicate malaria in large areas might slow vaccine effort if there is no compensating increase in the APC price, and firms might even worry *ex ante* about possible eradication schemes[136], raising uncertainty of research, capital costs, and the needed APC price *ex ante*. Kremer claims that "It is efficient for researchers to consider the possibility that their work will be superseded by other technologies when choosing their research projects." But these sorts of possibilities, especially if they are under the control of organisations that also run APCs, would raise risks to vaccine

[133] Observe, how, as far as firm *i* is concerned, it is not just the risk of parallel trade for itself that matters, but parallel trade for *all* other firms at *all* other points in time, since, given the importance of post-development mechanisms and the behaviour of *other* firms in determining the payout to any particular firm, any uncertainty about how parallel trade will treat any other firm at any other point in time, *will* factor into the *ex ante* decision of firm *i*.

[134] Indeed, the expected time of discovery, expected prevalence at time of discovery, and the needed APC price would need to be solved together, by feeding the laws of motion governing prevalence into the model of the technology that drives discovery (and all players would need to know these laws of motion).

[135] This refers to the market for a vaccine that is *not* brought about by the APC. If the APC is fixed, it generates a fixed 'APC-generated market' size.

[136] Observe, also, the need to adjust co-payment funds to recognise the degree of eradication.

developers and the needed APC price. There is clearly yet another dynamic inconsistency problem, and yet more potential for coordination failure.

As with non-stationary technology, to the extent that an APC price cannot adjust to changes in the non-APC-generated 'initial market size', part of the APC price has to contain an 'option based' component based on the non-APC-generated market size. Again, this potentially increases the value and efficiency of more open collaborative research methods that are more flexible *ex post*, and hence cheaper (and quicker), in response to this non-stationarity.

Also, to the extent firms are risk averse, they need to know the degree of riskiness around the central projection of the non-APC-generated market size, including assessments over how non-APC countries will respond to get the non-APC-generated market size down. For example, the more done to prevent the spread of HIV/AIDS, the lower will be the non-APC-generated market size, and the higher will be the needed APC price for an HIV/AIDS vaccine.

And it is not clear even what the concept of 'existing market' means anyway. All current failures (including the current levels of prices and low consumption of drugs in developing countries, problems created by the current system of patents, poor health infrastructure, and the past failures of large institutions to roll out vaccine programmes elsewhere, etc.) show up in lower 'current market size' for drugs, and hence a higher APC. It is a little perverse that these distorted estimates of current market size are part of the mechanism for determining the needed stimulus, without some analysis of what the 'true' counterfactual market size might be without the presence of these failures. This counterfactual would be different under different mechanisms, again revealing that the Kremer result is based on partial analysis.

7.15. APCs Crowd-Out Some Existing Market Size

It is not clear what the exact mechanism is that will prevent the APC from crowding out some of the 'initial market size'. This may not be great in many countries in Sub-Saharan Africa (though even in this case there would be *some* crowding out). But it might bite if relatively 'richer' countries are put into the program (say a HIV/AIDS vaccine APC), such that a proportion of their populations would have generated a market without the program in place. Since the APC is supposed to be adding 'additional' fresh market onto the 'initial market size', the more of this crowding-out[137], then the less *additional* research created for every dollar spent on the APC, the less well the mechanism performs compared to alternatives that would have directed finance directly at vaccine research, the higher the needed APC, and the lower the resources available for other vaccine and drug treatment programs. In effect there is redistribution from rich country tax-payers, who pay into the program, to better off poor and middle-income country buyers who would have bought anyway, with the deadweight loss of the taxes and the expensive capital costs of the APC incurred in the transfer.

[137] The degree of this form of crowding-out will also depend on government pricing policies in these countries, and also on how these policies have been factored into the co-payments of such countries.

This sort of crowding will vary by disease and population profile, and may actually be worse where there is already a current market that is quite large, but not quite large enough to induce optimal research intensity. It might seem that a small APC could induce optimal research intensity, but, even small amounts of crowding out could raise the cost of the APC considerably.

All Kremer's cost-effectiveness calculations[138] (rather amazingly given all we have been told about the importance of crowding-out issues) fail to incorporate this particular form of crowding out.

Incidentally, *worries* about crowding out, and any uncertainty about the prices that non-APC countries might pay, would raise private pharmaceutical capital costs and feed into a higher APC price.

7.16. The Dangers of Losing Vaccine IPR – Market Segmentation, Denied or Delayed Access, and Higher Prices

The APC introduces at least three features that were not present before. First, countries that are segmented into two groups; those who are in the program and those who are out. At least in principle those in the program can be denied access to the *type* of products that are available to those who are out (though not *vice versa*). This might, in principle, enable the creation of an 'earlier' outside market, a 'later' inside market, and some ability to pick off part of the 'later' market earlier (we saw this at work in extraction method three). Second, vaccine developers who can potentially be divided into those serving the program and those serving the market outside of it (which may be the same firm at different times) with this distinction based on product features. Third, full developing country payments (if the co-payments are set correctly) that have already been irretrievably[139] deposited, which reinforces dependency on the program, and makes it easier to stop purchases based on alternatives to the program. Additionally, being an extension of the patent-based system, the system is characterised by strong IPR. This may even be stronger than before, reinforced by rules specific to the program, barring APC recipient countries from accessing non-APC vaccines and from trading APC vaccines. At least in principle this creates potential dangers, and perverse incentives, that should be fully investigated.

Firms that sell vaccines under an APC, and even those who do not but were motivated by the APC, keep all IPR to the vaccine created and on processes that were patented leading to development and production of that vaccine. Indeed this is what makes the mechanism work as an incentive device. It is not completely obvious that, post-development, a firm would be willing to accept the terms of an APC program if it turned out more profitable to work outside of it (at least for a time). And it is not immediately obvious that extra ability to segment the market and charge overall higher prices has not been created, aided also by strong IPR. Since none of the investment taking place under the APC is within the direct auspices of those 'running'

[138] K2:9

[139] See the section on 'time-limited programs' for caveats.

the APC program[140], they would have no control over this behaviour. At least some of the consequences of this should be explored.

7.16.1. The APC is a financial option

An APC program is, in the investment jargon, a financial 'option'. It is the market firms revert to if *other* markets turn out insufficiently profitable[141]. This 'option' has financial value, will reduce capital costs, and will increase investment into vaccine research, *even if the APC is never supplied with the results*. There is nothing irrational about being motivated, at least in part, by the APC and then not supplying to the program, or supplying to it but only after a delay. *Ex ante*, the APC is part of a set of possibilities, and investment is made according to all these possibilities. The value of investing is therefore boosted by the inclusion of the APC. As history unfolds and a drawing is made from this set of possibilities, decisions about strategy – including who to supply – are made according to what is optimal in the 'continuation game' following each drawing. If the drawing over other possibilities is particularly good, it may be optimal in the continuation game to ignore the APC.

If, for example, HIV/AIDS vaccine markets in Russia, China, and India[142] (or indeed any of these if the others are included, but *it* is not) may one day become highly profitable, but this is not currently known for sure (or the speed at which they become profitable is not known for sure), an APC is the back-up to investments targeted at these hoped-for markets. Option price theory will show a value for this option that will make investing cheaper.

If these other markets indeed turn out highly profitable (with vaccines[143] able to sell in larger quantities at ten or twenty or more times what is promised for the APC program countries, and also with some ability to price discriminate[144]) and if firms (there may be more than one) do not wish to threaten these much more profitable markets, then they may simply refuse, or more likely delay (on top of that discussed above)[145], supply of a vaccine meeting the minimal conditions to those countries inside the APC program. They may, for example, worry that the transparency of the very low prices

[140] Since the APC is supposedly an 'additional' device, investment towards vaccines is fungible; some is accounted for by the APC, and some is not (accounted for by other devices like non-APC markets, tax-breaks, subsidies, etc.), but there is no way to legally tell the difference, so no way to assert rights over *any* of it.

[141] If the 'products' of different markets are not quite the same (different sub-types of a disease, etc.), the exact size of the option value will depend on things like the fungibility of investment across the different sub-types, the expected timing of the revelation of information, the expected timing of the flow of investments, etc.

[142] This is on the presumption that they have not signed up to the co-payment scheme. See comments at the end of the following sub-section.

[143] The argument goes through even for different clades to those that might be sold in the poorer markets.

[144] This may be easier than before if the program countries are being kept from the products over which price discrimination is taking place. Price discrimination may be enhancing the value of the behaviour being described.

[145] This may simply show up in zero effort to try to convince the 'purchase decision committees' of such countries to purchase, or be backed up with arguments about the dangers of misuse/lack of comprehensive distribution programs leading to resistance to the drug being built up, and the need to protect the product, etc.

at which they would effectively be forced to sell to APC countries[146] (with low ability to price discriminate across populations in these countries too[147]) would affect the ability to price highly in the more profitable markets. Or they may be concerned about the risks of parallel trade from countries inside the APC program. If they are not obliged to sell to the program, they will not do so (and this is easier to achieve if they fall below the minimum conditions of the program[148])

The perversions of the growing 'initial market size'

This interferes with part of the optimality condition for determining the APC price. It is supposed to be the case that if the value of the 'initial market size' is larger, the APC price can be set smaller, since the APC is only supposed to be *additional* to this 'initial market size'. But the logic may work in totally the opposite direction. As the value of the non-APC market segment – the 'initial market size' – rises, the relative value of selling to the APC market segment falls, especially if selling to the APC market segment undermines the value of the non-APC market segment in any way. The APC price needs to *rise* rather than fall to encourage firms to sell to the APC market segment. This obviously matters less to established diseases and diseases concentrated amongst the most impoverished, but may be important for diseases such as HIV/AIDS that are still rapidly evolving and that may affect the rich as well as the poor.

The post-development adjustment mechanism may also feed this problem. That mechanism only enforces quality if it can be sure to set higher prices and quantities for better quality vaccines, and lower prices and quantities for poorer-quality vaccines. If a vaccine is produced that does not meet enough of the minimum conditions to get a large purchase or a particularly good price, the firm may refuse to sell to the scheme if the 'bad deal' acts as a bad signal that reduces the price and the market-size outside of the scheme (where there may be more marketing anyway, etc.). Incidentally, this suggests that *ex post* they should be able to extract a higher price than supposedly 'optimal' under the rules of the program.

Two other APC option values to developers

There are two further ways that APCs may act as options, and motivate research even if the results of that research never get sold to countries in the program, or are sold to them with a delay.

The first way is when, ex ante, a firm may not know what '*type*' of vaccine they will get at the end of the process. For example, in HIV/AIDS research some investments may lead to results useful to different clades, some of which will lead to vaccines suitable for wealthier markets, but this may not be clear at the start. The inclusion of other possible suitable markets for results (different clades) will act as an extra option component reducing risks and costs of research, even if the vaccines developed are in the end for the wealthier markets. The presence of the program has value and

[146] This would be aggravated by any need to create heavily discounted prices. Kremer suggests, for example, that given all the uncertainty in calculations of the necessary number of doses, an APC "should cover a smaller number of doses at a somewhat greater price, with an option for the program to buy additional vaccine at a *discounted price* at the program's discretion" K4:5. Italics added.

[147] This also is reinforcing the behaviour being described.

[148] This could be easy to do. If those running the program have set up rules to accept vaccines that are below 100% of the minimal condition, firms may simply not rush to supply.

encourages investment, but once the 'type' is revealed, then in the continuation game it may nevertheless not be rational to sell to the program. This also creates a discontinuity; once the clade for the more profitable market is covered, suddenly research incentives are much lower for the other clade. Similarly, even if the research is concentrating on one particular clade we still do not know if a vaccine designed for one clade would prove cross-reactive against other clades. Once that is revealed research for the poor market becomes of secondary importance. Observe how having a late-stage APC, of a procurement variety, 'up one's sleeve' might be useful in this case, but has been lost in the fixation on the other, early-stage, type of APC.

The second way is when a firm is not sure of the *quality* of a vaccine that will be produced. For example, and related to the segregation of markets and various other issues mentioned in this paper, the program may act as the back-up receptacle for a poorer-quality HIV/AIDs vaccine, if the vaccine did not turn out of high enough quality go get the more profitable bigger markets elsewhere.

All of these options-based problems suggest care when measuring the investments motivated by the program if they might contain options-based elements. They also suggest problems for any auction-based mechanism seeking to generate 'optimal intensity' of research if the actual research intensity is not very revealing of the chances of getting a vaccine for the program-covered countries.

7.16.2. Fourth extraction method – Another way to make developing countries pay more

We already saw three cases of extraction of surplus beyond the co-payments. The third was related to the creation of an 'inside' and an 'outside' market, since, without the program in place, the quality of a vaccine could never be below the threshold of the rule to keep a vaccine out of the APC-covered market.

There is a fourth way to extract extra surplus (there may even be others) that also overlaps with the others. Firms may be in a position to add to those previous extraction possibilities the fact that they now have non-APC market options. The 'drawing of a good outcome' elsewhere simply reinforces their ability to extract more.

In this they are further helped by the co-payment funds. We already saw how the sunk nature of the funds means that the marginal payments that developers would be seeking to extract are lower than the marginal payments they would have to extract without the co-payment fund. We saw that it is wrong to presume that those administering a country's co-payment fund will automatically take into consideration *all* of the cost of drawing down the funds when working out the marginal cost to them of a vaccine. That the administers of a country's co-payment fund price into the marginal cost of the vaccine *any* of the drawing down of co-payments, is presaged on the notion that they know that there will be vaccines generated by the APC scheme that they can buy – if not this one, then another later – and *that the scheme is working perfectly*. As soon as this is not so, they may not price *any* of the drawing down of funds into the marginal cost of a vaccine to them. Given the logic of the previous sub-section, once countries start to realise that i) even *other* vaccines meeting the minimal requirements will not sell to them at the prices stipulated or ii) that the system is not tending to create vaccines targeted at them, or if it is, iii) it is tending to create an incentive to delay the creation of vaccines targeted at them, or iv) that the logic above

regarding the need for a higher APC price applies anyway, *then* they start to judge *any* vaccine on the marginal costs it imposes above the already sunk co-payment funds. Naturally, knowledge of this possibility will ex ante alter incentives of developers[149].

This is also fed by any realisations that those outside the program are, by their uncoordinated acts, breaking the program by over-consuming, as described in a previous section. It was said there that "if no more co-payments could be charged" then extra payments would have to come from richer countries. But if countries reason that the breaking of the program by outsiders decreases the chances of future better vaccines, then they will be prepared to pay marginal payments on top of their sunk co-payments to avoid this delay (indeed, this is another aspect of the scheme being an option to developers; one of the option elements is the ability to encourage the breaking of the scheme by outsiders).

As before, this effect is strengthened if allocations are limited. Limited allocations are supposed to enforce efficiency, but if efficiency is not being produced, these limited allocations will simply create greater ability to extract more surplus (indeed it may be that – according to the information of the APC regulator – efficiency would be produced by restricted allocations, but the APC countries do not have access to, or believe, the information underlying the decision).

The fact that the threat of compulsory licensing would now be much weaker naturally strengthens all abilities to extract more. Ruling out compulsory licensing on this particular vaccine might be part of the *ex ante* conditions that countries would have to sign before taking part in the program at the start – in order to help prevent dynamic inconsistency problems. But it is apparent that this also creates new dangers

7.16.3. APCs segment markets and drive up prices

This is all part of the general problem of the way the APC-program enables (enforced) market segmentation. This creates the potential ability to generate higher prices (than would be the case if the APC was not in place), at least for a time, to countries *outside* the program, and the ability to pick off profits from sections within the APC program. All of this should be (but is not) priced into the costs of an APC.

There may be little incentive for pharmaceutical firms to give countries inside the APC program a low price (or at least a price as low as what they might have got without the APC program in place) on a drug that effectively breaks the APC program since i) it only weakens the price that can be charged currently to those countries outside the program (in fact it might be rational to set a higher price to outside countries precisely so as not to cause those covered by the APC program to agitate for it) and ii) it reduces profits from a later drug that meets the program's requirements. We observe that the behaviour on point i) will lead to a higher APC price for this later drug anyway (especially if the price is set in the auction mechanism described), even if it is of poorer-quality than if there had not been market segmentation.

Some further aggravations

There are further aggravations to the market segmentation problem. In the APC calculations, the individual country co-payment is based on the overall country

[149] There are certain self-fulfilling aspects.

marginal willingness to pay. There are two possibilities. Firstly, if the co-payment was calculated on the notion that no price discrimination would be possible across wealth levels within a country (if the richer sections of the market could not be prevented from trading with the poorer sections), then developers may profit from the segmentation of the program and non-program countries by trying to sell to only the richer sections of program countries (the richer sections will not trade with the poorer sections). It may prove difficult to deny these richer sections access to a product from outside the APC program[150], even as the poorer sections are denied access. Through segmentation, therefore, the program enables some price-discrimination that was not possible before, and this increases the expected profits of those actions that feed this ability to price discriminate.

Secondly, if the co-payment was based on social surplus presuming that the government of the country would price discriminate (probably very imperfectly), then the picking off of the top parts of the wealth distribution by developers, increases the *ex post* cost of the lower sections beyond the social surplus of those sections. This all further reinforces the problem being described here.

What are the dynamic consequences of this? What if potential developers of better-quality vaccines believe that developing poorer-quality vaccines will i) lead to higher payment outside the program consequent on the market segmentation caused by the program, and that ii) weaker vaccines may still undermine the program for better vaccines anyway[151]? This may slant early research efforts towards products meeting less than minimal conditions, and the need for a higher APC price to encourage research into better quality vaccines for those covered by the program. In a sense the APC allows the market to be separated out over time. This needs to be explored in more detail.

7.16.4. The IPR is less likely held in the public domain

Ordinarily, given the huge expenditure of government and foundation resources, their might be some leverage over this behaviour. However, given that the property rights are held by firms and not by the ultimate funders, governments or foundations, on behalf of society as a whole, and given that the APC system has itself made IPR tighter, this is even less likely the case. Similarly, the holding of the IPR would continue to bar others from using information and important 'enabling technology' that, under more accessible IPR systems, would have enabled them to undermine some of the problems just described. When comparing the APC with other mechanisms where the public sector and foundations might get to keep the IPR for the public good, care should be taken to include this benefit in all comparisons. Similarly, more thought could be given to the way that most current push mechanisms, in contrast to pull mechanisms, do not generate sufficient rights to the outcomes, and how this might be rebalanced.

The social surplus of access turns out to form a huge proportion of the value of an APC. The net present value calculated by Kremer of the purchase and delivery costs for covered countries, are $3.2bn for malaria, $5.3bn for tuberculosis, and $4.0bn for

[150] Of poorer-quality than the minimal conditions, if the developer is obliged to sell to the program if the vaccine meets the minimal conditions.

[151] Some of the comments made above about time-limited schemes, etc. suggest reinforcing features to this process.

HIV[152]. But the social surplus lost from failure to address access is $7bn, $5bn, and $14bn respectively[153], with DALYs lost being 284m, 246m, and 589m[154]. Clearly, losing the IPR or concentrating it in a few powerful hands has potentially serious consequences.

7.17. APCs Lead to Tighter Property Rights and Stifling of the IPR Debate

One of the consequences of an APC, found time and time again, is that it goes hand-in-hand with ever-tighter property rights, especially for early-stage vaccines. Tight property rights were needed to overcome the problems of nonstationary technology in Chapter 5. We have now seen how important it is that a typical APC – in order for it to work efficiently – be capable of all kinds of adjustments post-vaccine-development. This *also* needs very tight property rights. As Kremer puts it: "For purchase commitments to spur research, it is essential that intellectual property rights be respected" and that this might entail that "U.S. funds would not be used to purchase vaccines that violate U.S. Patents."[155] Essentially, tight IPR helps to counter the incentives of countries to 'collude' against the APC mechanism[156]. Tight IPR is absolutely essential to reduce risks to developers and to keep capital costs down, even if it drives overall costs higher.

Even if the first vaccine is protected, it might be relatively easy to work around it or imitate it (though this is not necessarily the case). The risk that a marginally superior vaccine will come along based on the first vaccine and take the whole market (this is, of course, efficient) might deter research on the first vaccine (that might, though not always, contain all the difficult science). First-mover advantages such as network effects and brand loyalty are much weaker for vaccines than for just about any other product, since governments are the buyers and the knowledge is easily disembodied from the physical product (though for vaccines there are huge problems with know-how and sunk production costs that limit the usefulness of knowledge of the patents). The weaker are property rights, the higher will the APC terms need to be set *ceteris paribus*.

This leads to several conclusions:

First, efforts to establish a strong APC (even just a couple of APCs) will be concomitant with a strong effort to strengthen the IPR system and, indeed, might even require promises up-front that debates about the IPR system will *not even* be on the agenda. Failure to commit to this would otherwise increase the APC price.

Second, with several APCs in place, efforts to weaken the IPR system later, will run the risk of weakening the ability of the APC to generate its end result. Those seeking

[152] K4:14.
[153] Incidentally, these also give some idea of the amount of social surplus extractable should the APC not discipline price away from the upper bound of the range of possible values.
[154] Observe, that these are boosted also by the nature of past IP-related failures.
[155] K7:26.
[156] Actually, it is just recognition of the fact that the system is an elaboration of the current patents-based system.

alternative IPR systems will find themselves undermining (or being presented as undermining) the few initiatives in place for creating new vaccines.

And third, the APC mechanism may even create the dangers of losing IPR to vaccines altogether, with access and price problems to those covered by APC programs. This should at least be explored in more detail.

7.18. Committees, Information and Risk – and the Need for a 'Global APC Treaty'

The attempt to avoid quality distortions and inefficient vaccine purchases, leads to the creation of layers of committees and rules, and risks for vaccine developers and investors.

The risks are a function of the rules, but the optimal rules are a function of the underlying technology: "The type of technology in question will influence the formation of eligibility and pricing rules."[157] This is a tall order. Intuitively, quality varies over the technology space (distributions over research leads). If the rule treated every part of the technology space equally, all developers would always go to the part of the space that is easiest, i.e. cheapest to reach. The job of the APC setter is to set the rules so that 'effort' towards the more difficult and expensive parts of the space – where the quality lies – is relatively more rewarded. That is, they wish to create incentives that make movement towards these parts of the space a 'dominant strategy'. If the APC setter knew the technology space *exactly* (which includes knowing firms' costs, i.e. the costs of their 'effort') they could set a precise rule with larger rewards the more difficult it was to get to a particular part of the space[158]. If they do not know the space exactly, they can only create a highly imperfect rule, taking great care over where in the technology space rewards are placed in case they cause distortion. They are reduced to generally average rewards everywhere, and, indeed will *never* achieve the highest quality results. Hence, on average, achieving 'quality' is much more expensive.

In addition, because of all the uncertainty to players, the APC setter cannot simply pick out the quality area (even if they knew where it was) with a huge payment compared to the rest of the space (which might seem the most logical thing to do), since this would face players with huge risks should they fall onto other parts of the space where the payments are tiny. So the rule over 'how much the quality rule varies' over the technology space, itself requires knowledge about the characteristics of firms, such as their access to finance, degree of risk aversion, etc. Kremer has ruled all these difficulties out, since he only talks about rules, and presumes they are enacted perfectly – which is a polite way of saying that he presumes the APC setters know everything about technology and about the characteristics of firms (which they do in his model).

The APC was claimed as a solution to the lack of information of policymakers about vaccine technology when publicly funding vaccine research up-front. We yet again

[157] K2:2.
[158] Imagine the dimensions and complications of this technology 'space' if research projects were not independent.

find, ironically, that to stand a chance of being efficient, policymakers, when setting the *ex post* rules for the APC, need to make extraordinarily heavy use of the very information that they are claimed not to have access to. The words 'chicken' and 'egg' do rather come to mind. And it is unclear how optimal rules could ever be set behind a wall of scientific or technological secrecy as encouraged by the APCs (and high IPR) themselves[159].

One solution to the lack of information on which to base rules would be to give policymakers discretion. The proportion of the APC that any one firm and its investors will on average secure will then depend on the exact mix of rules and discretion, and their treatment in the hands of these layers of committees. This creates the usual trade-off between fixed rules and discretion. If all is not transparently obvious, discretion introduces even more uncertainty and risk[160]. And all risk has to be reflected in capital costs, and hence the APC price[161].

Of course it could be that the APC price could be set so extremely high in the first place that many of these committee issues are avoided – helped in part by the sheer paucity of information, such that even if the APC price is set extremely high it is hard to judge that this is taking place. But that hardly suggests that the APC is being thought of as a cost-effective instrument.

If we are prepared to countenance the use of so many, supposedly well-informed, committees and regulators *ex post* to make the APC work, it does rather numb the criticism that more open collaborative approaches and other methods need good scientific knowledge and make use of 'committees'. And, given the high degree of 'interventionism' by committees and others in the APC, it rather nullifies the notion that somehow the APC is light on centralism and solves everything without 'intervention'. Quality under an APC is driven by the mechanism, not by any collective group of scientists as would happen under, for example, more open approaches. But the mechanism still has to be 'set' by someone.

The dangers highlighted in the previous sub-section also indicate the need for all countries that do intend one day to *benefit* from the APC program to be in the program from the start (in the HIV/AIDS example given, these would include Russia, China, and India). This is parallel to the logic earlier that all those who intend to *contribute* to the program should be in the program from the start. Both these arguments suggest the need for a global Treaty, or similar binding agreement, at the time of initiating an APC.

[159] Again, it is worth exploring the degree to which more open collaborative research might remove some of the bricks in this particular wall.
[160] These risks include not just the holding down of price, but also non-price manipulations (that may act to suppress price tacitly) like requiring excessive product testing and improvement.
[161] It is not in the APC papers.

8. NON-APC FUNDED VACCINE RESEARCH, AND THE APC COST-EFFECTIVENESS EVIDENCE

8.1. The Inefficiencies of Non-APC Vaccine Research

In the executive summary of his proposals to the British Government, Kremer claims that APC-motivated private-sector research is *four and a half times more cost-effective* than direct funding of applied research and joint ventures with private companies of research[162] into a HIV vaccine, and three and half times more cost-effective for malaria, and TB[163]. Even roaming exclusivity (normally thought of as a bad method to finance R&D) is *two to three times more cost-effective* in cases such as TB and HIV than direct funding of applied research and joint ventures with private companies. That the APC outshines so much is explained thus – and *only* thus:

> "*The cost-effectiveness of government R&D is limited by the potential of crowding out private R&D, difficulties in picking winners among competing research projects, potential politicization of funding decisions, and difficulties in shutting down unpromising research projects*".[164]

In all APC calculations, publicly-funded initiatives (as opposed to the, equally publicly-funded, APCs) are modelled as being extremely bad at supporting good vaccine research. As the quote reveals, this – and the fact that technological, strategic, finance, auction, capital cost, and many asymmetric information, and institutional difficulties are presumed away – is what drives the result.

There is not room here to analyse all publicly-funded push and pull mechanisms, including tax-breaks, subsidies, patent extensions, government grants, government direct research, the patents system, etc.[165] On many of these (though, conspicuously, not on the last) Kremer is right to draw attention to failures and inefficiencies. Interestingly, many of the criticisms are not echoed amongst large pharmaceutical firms, who are generally encouraging of the notion of a big rôle for government support of the pharmaceutical sector. It is even asserted that one of the major factors drawing large pharmaceutical firms into the US, in spite of the higher costs of doing research there, is the size of the largely publicly-funded science base and the size of the NIH budget[166]. The criticism here is the way in which Kremer deliberately selects from the set of potential failures in ways that bias the case for APCs, exaggerates what he does find in ways that he does not properly support with evidence, and ignores all failures that may impinge on the workings of the APC itself. As Jaffe[167]

---

[162] This is, of course, a small sub-class of all publicly-funded research, so for example it does not refer to many instruments like tax-breaks or non-targeted subsidies, etc.
[163] K Summary.4.
[164] K Summary.2.
[165] Incidentally the patent system is included here as a publicly-funded mechanism to support pharmaceutical R&D since it *is* a form of taxation, and does come largely out of government tax revenues in many countries. Even where it does not come out of government revenues it operates as a tax on companies and employees.
[166] Fulcrum Newsletter 2003, Interview with Frank Fildes, former Senior Vice President and Head of Global Development of AstraZeneca Worldwide.
[167] A.B.Jaffe, Oxford Review of Economic Policy Vol 18, No. 1, 22-34, p23.

comments with respect to publicly-supported research, "Much of the political debate surrounding such programmes remains at the level of ideology...Yet as social scientists we have an obligation to try to bring facts to bear on these debates."

APC supporters should have an independent interest in trying to correct many of these non-APC failures, and in working out which of them are relatively less bad than others. To be optimal, each APC requires a concomitant optimal amount of 'push' research. If this optimal amount of push research falls short, the APC price has to be set higher to compensate. And if publicly-financed or foundation-financed research is as bad as Kremer suggests, then the cost of the APC is pushed even higher than it otherwise need be.

The cost penalty is built up by Kremer in layers:

8.1.1. The first four layers of inefficiency – Ruling out all the good projects

Kremer models public-funding bodies as buying projects from the private sector. Since the funder cannot tell one project from another, and cannot put in place mechanisms to efficiently select between projects, they can only work out the median value of all projects[168], which is the value 'm'. At the same time, nobody owning a project worth greater than m will ever want to sell it to a publicly-funded program, since, allegedly, the seller can find private buyers for their projects prepared to pay greater than m. The public-funders buy a selection of projects on the distribution running up to the m project, but no higher. Kremer thus rules out the top half of the probability distribution (the portion between the median, m, project and the highest, h, project) in *all* publicly-funded vaccine research projects.

In fact, since m is the median project – and not the mathematical average – this turns out to be quite a severe penalty given that the probability distribution may be skewed with only a few really top projects. This is a second layer of inefficiency.

More harmful IPR creates more favourable results for high IPR

But the matter is much worse than this. The median project is not based on all conceivable projects, but is the median project of those projects that are profitable based on the low 'initial market size', i.e. the market without the APC. This leads to a totally counter-intuitive conclusion. Since the 'initial market size' is itself calibrated on the basis of a set of problems caused by high IPR[169], the *more* harm done in the past by systems based on high IPR, the *worse* direct government R&D (and low IPR approaches) performs in comparisons, and the better is the performance of systems based on high IPR, including APCs. This is not just patently wrong, but it makes no sense when trying to work out the relative effectiveness of alternative mechanisms in solving a market failure, to advantage in proportion to that market failure one of the mechanisms that was most implicated in creating that market failure in the first place.

There are several problems with this assumed inefficiency:

[168] They are supposed to know the probability distribution.

[169] The calibration presumes that, on account of high IPR, there is no access to vaccines for ten years and even after that vaccine prices are so low that they are barely profitable; both of these suppress initial market size.

1) If the government suffers a 'lemons' problem, is it reasonable to model the situation as if there is no way to get around it – via, for example, the use of reputation, or contracts that are time-dependent. Most research projects are not one-off investments but a stream of investments. So, does this assumption really boil down to an inability to kill poor projects, rather than primarily an inability to pick good projects? A more open framework for research would enable funders to more openly kill poor projects than the government here is modelled as being capable of doing; in an environment where all information is transparent and secrecy is at a minimum it would be harder for poor projects not to draw attention to themselves.

2) The modelling device presumes that the publicly-financed method is capable of *some*, quite difficult, targeting nevertheless. It visualises funders picking off projects at the top of that part of the probability distribution available to them (the 'truncated-at-m' probability distribution based on the initial market size) even if they are incapable of picking any projects on the part of the distribution above m. It is not clear how they manage to do this. For example, if they are really never prepared to pay greater than m, then no project worth greater than m will offer itself. The average project offered will be worth approximately m/2 (it depends on the exact shape of the distribution). How do they get to work out the 'close-to-m' valued projects from this? And if they are able to judge projects as worth 'close-to-m' from the information given by those presenting themselves, then why can't they use similar methods to judge projects worth 'greater-than-m' by offering a price, or a schedule of prices, 'greater-than-m' and allowing 'greater-than-m' projects to use the same methods as the just-below-m projects to demonstrate their worth?

At the same time as some sort of targeting *is* obviously going on, the public funders are nevertheless visualised as paying p_m (the price of the median m project), even though they know that the average value of any projects offered is m/2, for *all* projects, even including the second tranch of dud projects that will be covered in the next section ("random projects below the p_1 project")[170]. So they seem to be able to work out the value of projects they do get offered, but they still, hopelessly, pay p_m for all. Hence they overpay for all projects. This feeds into the measured higher costs of publicly-funded non-APC vaccine research. This is part of a recurring pattern; every chance Kremer gets to boost the costs of research that is not APC-motivated, he takes it.

3) It is not clear why the public funding bodies are not prepared to offer a price more commensurate with the social worth of developing a vaccine, and then choose between projects offered up.

4) If the APC relies (to hold the path of technology together) on secrecy at various stages of project development, can we be so sure that poor projects do no survive under an APC over these 'secret' stages and that more open approaches are not a better mechanism for picking them off? Secrecy allows poor projects to hide and survive. Openness exposes them.

5) Kremer presumes that vaccine projects between p* and h would be pursued without government assistance, that there would be no failure of private investors to fund

[170] Equation 25 on K3:19.

them, and that those holding such projects would therefore never be prepared to offer them. That this continues to hold when the true alternative is that there is no market for vaccines, and hence no funding for such projects outside of the publicly-funded sector, is not clear. One might expect that anybody with a project worth greater-than-m might have an incentive to reveal their project's worth in such circumstances to whoever is prepared to fund it. In comparisons with the APC, Kremer seems to model on the basis that such projects would hold out for the APC value even if it is not the counterfactual alternative that will actually be available to them.

6) Do private sector investors into APC-seeking projects face any similar problems, or are they somehow able to overcome such problems in ways that the publicly-funded sector cannot? Could these ways, if they exist, be imported into an alternative framework (say a more open framework)? There is more on this issue in Section 12 below on private finance.

Four layers of inefficiency are thus rolled into one: The loss of the top half of the distribution; the use of the median not the mean; the median itself based on the greatly reduced market due to IPR failure; and all projects purchased at the price of the most expensive project. Even without any more layers of inefficiency, the results are already stacking up heavily against alternatives to the APC. When used to generate cost comparisons between alternative ways to fund, say, HIV vaccine research, these assumptions are quite absurd.

This is not to say that the above problems do not exist. It *is* to say that the evidence for them needs to be openly presented and quantified, and that ideology and superstition are not good surrogates for facts.

8.1.2. The fifth layer of inefficiency – 'Crowding out' and the problem of 'additionality'

It might be thought that this ability to select investment projects just below but close to the median, m, project would nevertheless be regarded as somehow 'good', and that this would act in favour of publicly-funded research in comparisons with the APC. Kremer, however, describes even such productively used public resources, as "crowding out private investment", that is of projects that "would have been done anyway." Such projects add nothing to the increased likelihood of the technology being invented, and indeed are modelled as pure waist. Since crowding-out leads to no net increase in resources, and given that such public resources are raised in distortionary ways – such as through taxation that creates 'deadweight losses' – society might actually be worse off overall. The default setting for this in the APC model is 40%.

One moment, publicly-financed projects are penalised because they target badly, failing to hit the top half of the (already repositioned) probability distribution. The next moment good targeting counts against them.

Wallsten

For empirical evidence of this Kremer uses just one study – over and over again – to argue that publicly-funded vaccine research initiatives would simply replace privately-funded vaccine research that "would have found private investors in the

market". He explains: "For an empirical argument that government crowding out effects on private R&D can be close to 100% please refer to Wallsten (2000)."[171]

The Wallsten study looks at the Small Business Innovation Research Program (SBIR) and the way it "seems to have crowded out privately funded research dollar for dollar"

But this evidence is very problematic for our purposes:

1) At the very best, as Jaffe[172] puts it, "results on this question are mixed" with limited evidence on either the effectiveness or lack of effectiveness of programs like the SBIR. It is well known that the issue is plagued by the lack of a 'control' group of firms that did not get funding. Most studies look at retrospective survey data provided by recipients. Practically none tries to address the counterfactual.

Lerner (1999) does try, and comes to completely the opposite conclusion to Wallsten. Lerner looks at 1,453 SBIR awardees and a matched sample of non-awardees over a period of 10 years, and finds that those receiving SBIR grants grow significantly faster after receiving their grant than those that do not[173]. Even then, this is plagued by selection bias – that those firms or academic institutions likely to get funding are likely to have the best ideas anyway, and so would have the incentive to spend their own money, and more ability to attract support from third parties, causing them to grow the most anyway.

Branstetter and Sakakibara (2000), in a study of Japanese research consortia found that the R&D of participating firms increased with public funding of research. In a study of research of commercial firms in Israel, Lach (2000), found that total R&D expenditure rose by $1.41 for every $ of public funding.

David et al (2000), in a survey of econometric evidence, find that many studies at the firm level find 'crowding-out' effects, but studies at higher levels of aggregation mostly find 'crowding-in'. But there are difficult econometric interpretations of what this means exactly.

2) The Wallsten study does not refer to pharmaceutical trials. In fact, Kremer provides no evidence whatsoever, in any article he has ever written on the APC, on the efficacy of trials, publicly or privately-funded, although this is *the* driving force of his argument. And we know that there is strong evidence for efficient publicly-funded trials[174].

3) The study does not relate to publicly-targeted research, but to a general tax credit regardless of the type of R&D being undertaken (other than it be on 'research'). We know – and the author agrees entirely with Kremer on this point – that general tax credits are a very inefficient way to stimulate vaccine R&D. It would not, for example, be sensible to use such tax-credit generated figures to discuss subsidies[175]

---

[171] K2:10, K3:20.
[172] Oxford Review of Economic Policy, 2002, p
[173] For further studies of such programmes also see Spivack, R.N. 2001.
[174] See section 8.3, and references therein.
[175] There is mixed evidence on the social benefits of subsidies on research nearer to commercial application. See David et al (2000), and Klette et al (2000).

and grant-funding (that are typically targeted, but the results of which are often very unclear since much of the benefits are diffuse in basic science and hard to measure) or to make judgement on direct government R&D, or of more open collaborative research methods, etc.

4) Being a straight tax-break, there is no requirement in Wallsten for the co-funding of research proposals. Firms spend their own resources up to the point where the expected marginal return is equal to the costs of funds. After this point the firm cuts its own funding dollar for dollar. But if grant agencies stipulate that for each additional $ of funding the firm also has to contribute from its *own* resources, this *does* reduce the marginal cost of research to firms. A profit maximising firm facing a downward sloping marginal research returns schedule *will* increase total expenditure under such circumstances. The exact degree of crowding out then depends on how rapidly marginal productivity decreases.

5) If funding agencies pick projects that are believed to have large social returns, but such that these projects are far down the private marginal-returns schedule, they would not be undertaken by an unsubsidised firm, but might be if the government funds them. The Wallsten study refers to high-tech projects (in this case, in California) with R&D much less of a public good, the end product not a public good, where there is a market, and where there is plenty of private funding available, and finds, under the methodology used, that government subsidies simply replace private funding (in fact the subheading of the Wallsten article is "Why is government subsidising commercially promising business projects?"). It is clearly misleading to adjust calculations for neglected vaccine research on this basis. Kremer himself points out that the 'social returns to research are typically twice the returns to private developers"[176] but then seems to oddly ignore this.

6) If there are large spillovers – and these are perhaps the true reason for some of these programmes – there may well be a particularly large gap between the private and social rates of return, so, *even if there were crowding-out*, the issue would be the relative size of spillovers versus the crowding-out effect. Kremer's technological assumptions rule out *a priori* any rôle for spillovers, so that there is *only* ever crowding-out. Again, we are interested in measuring the 'general equilibrium' interactions between the funded researchers and the rest of the system.

7) The alternative to the public funding of vaccine initiatives is the APC, and not 'the market'. The APC is *also* publicly funded, with just the timing of payment different. This is unlike the alternatives in the Wallsten study which really are the private finance markets and the profits of the firms concerned. It is slightly tautological to use 'crowding out' to favour the result towards the APC and against the publicly-financed research, when the burden of both ultimately falls on tax-payers. Here once again, failure to do a global cost analysis works in favour of the APC;

8) To the extent that the APC creates distortions and is expensive, it would be optimal for alternatives to crowd *it* out as a policy option. One of the things highlighted in this paper, and ignored by Kremer, is the way that the APC would generate higher capital

[176] K6:16. References are Nadiri, M. I. (1993), and Mansfield, E. et al. (1977).

costs and many distortions itself. The contention here is that these would be high enough to 'crowd out' many alternative forms of funding for vaccine R&D.

9) Even the Wallsten study quotes 82% crowding out. Allowing crowding out of each other, but recognising that they both require tax revenues, the use of the 82% figure would tip the analysis in favour of publicly-funded non-APC based research, ceteris paribus;

10) In the cases covered in the Wallsten study, the researchers may not have been particularly credit-constrained; a public form of finance simply replaced private unconstrained finance. However, if private research is credit-constrained (biotechs for example) then public finance may be cheaper than alternative finance and will 'crowd-in' private finance[177]. One mechanism discussed in the literature is the notion that the screening of R&D proposals for likely success is a costly and uncertain process (and may also suffer from free-riding if there are multiple private funders), so that public funding certifies proposals as 'high quality'. Non-public funders free-ride on this. This 'certification' or 'halo' effect seems important in the US as a way to increase the total research spending of grant recipients.

When Kremer argues, for example, that those holding good projects are not prepared to reveal themselves to the non-APC public-funder, this must mean that they have equally good access to finance from elsewhere. Supposedly, if projects reveal themselves then the government can pick off the best; it's just that the very best don't reveal themselves. Allowing credit-constraints and some ability to pick those projects that do reveal themselves, along with the ability of this to signal to financial markets the quality of projects and hence attract more private finance, and this piece of the logic of Kremer starts to look much weaker;

11) Firms are modelled as facing a 'principal-agent' problem *vis à vis* public funders, but not, it is presumed, *vis à vis* private investors. This will be explored in more detail below in Chapter 12 on finance. Principal-agent problems in the case of private vaccine finance are likely to be unusually high compared to other types of investments;

12) Kremer complains that public finance tends to go to those "already well-known in the field" and that this is inefficient. But, where reputation matters, or where the gathering of minds matters, this, surely, has efficiency-enhancing features too?;

13) Crowding-out is a much more ambiguous affair than Kremer argues. The real issue is the marginal (realistically, though, the average) social product of public expenditures. 'Additionality' of public funds is thus neither necessary nor sufficient to produce a positive social product. It all depends on the opportunity costs of the alternative funds. For example, if public funding crowds out foundation funding, or indeed any other type of funding, and if that funding is pushed to other productive research that we do not observe, then the social product could be large even though all we see is crowding-out in the area we are studying[178]. The degree to which crowding-

[177] See Diamond (1998).
[178] Similarly, the 'certification' or 'halo' effect could also mislead us since it simply redistributes in a zero-sum fashion the available non-government funding.

out or crowding-in effect the social productivity of public research depends entirely on the elasticity of supply of alternative sources of funding.

As Jaffe[179] comments: "Measuring the direct impacts of public funding on the funded entities is only a small piece of the overall evaluation problem...the crucial question is not the extent of additionality *per se*, but rather how, in the presence of the possibility of crowding out or crowding in, can one measure the social product of the public research funding itself...Once one recognises these complexities, the relationship of the additionality question to the underlying public-policy issues becomes ambiguous...Failing to reject a null hypothesis is not the same as showing the null to be true...political processes may ignore these subtleties and misuse research findings no matter how many caveats appear in the papers supporting these findings."

Given the huge importance of the vaccine issue, it really is untenable for the measures of the cost-effectiveness of various financing approaches to be subject to the vagaries and the "misuse of research findings" in this way.

8.1.3. The sixth layer of inefficiency – The random choice of projects

Once the first few 'targeted' projects are picked off, a large proportion of projects are then just randomly chosen. The default setting in the Kremer model is 40%. This adds a heavy penalty in the calculations of cost-effectiveness of publicly-financed vaccine research**.**

8.1.4. The seventh layer of inefficiency – When the science is not ready yet

All non-APC scenarios involve agents chasing tax credits, and government subsidies *even if there is no chance of a breakthrough*, whereas the APCs have the supposed advantage that they only pay if a researcher makes a breakthrough. The default in the Kremer model is 33%.

The possibility that the science is "not ready yet" "increases the appeal of the advance purchase commitments"[180] since the APC, it is claimed, disciplines researchers to be truthful about the true nature of their projects (though there is no explanation as to how this comes about or to the factors leading to the science being *made* ready).

8.1.5. The eighth layer of inefficiency – Projects that should not be pursued

The last tranch of projects chosen are those that are "certain to fail" due to "politicised, corrupt" behaviour, or simple because of "inability to cut off funds from failed projects". The Kremer model starts with the default setting for this of 20%.

USAID 1980s example

The piece of evidence that Kremer cites, repeatedly, as evidence of the "plague"[181] of "politicisation and corruption" leading to publicly-funded vaccine trials that "should not have been pursued" by the NIH, USAIDS, and many others, and that demonstrates that "the risks that grant-funded scientists and research administrators...will overstate the chances of success and divert resources away from vaccine research are far from

[179] Jaffe ibid.
[180] Kremer, M.
[181] Kremer, M., "Pharmaceuticals and the Developing World," *Journal of Economic Perspectives* 16(4), Fall 2002. p82.

hypothetical" is the "sad story" of the US Agency for International Development's (USAID) program to develop a malaria vaccine in the early 1980s

The problem with Kremer's use of this piece of evidence, is not that this is not a good demonstration of corrupt and lax standards of funding, but that it is the *only* piece of evidence that he ever uses, and that he uses it repeatedly. Given the huge importance of this in driving the results, one would expect by now, more than 7 years on, a large body of supporting evidence tightly connected to the *trials* issue at hand. But Kremer uses the same case – nearly a quarter of a century old – over and over again. It is not as if there are not many more contemporaneous trials to judge efficiency and levels of corruption from. In August 2003 the US National Institutes of Health was recruiting patients for 2,832 clinical trials that it was sponsoring directly, including 293 for HIV alone. Other US agencies were recruiting for a further 186 clinical trials. And a further 1,796 were taking place in universities; many funded by US federal grants. A rich source of evidence one would think.

The total costs of this "sad story" came to about $5m. The current NIH budget is about $37billion per year. It is hard to imagine why a case representing maybe a couple of ten thousandths of one percent of the NIH budget since the early 1980s could come to symbolise all that is corrupt in the use of that budget. And it is not clear why the moral rectitude of those running the NIH and other funding bodies for the past twenty or more years should be judged on the basis of this one particularly bad case of corruption.

By the end of 2002, unhappy that this was a rather tiny figure compared to the NIH budget of over $3bn a month, Kremer sought to boost the case by, quite misleadingly, converting this $5m waste into $60m of "trials without success."[182] Given that most trials yield no positive 'result' anyway, however efficiently they are carried out (see the Tufts data on attrition rates that demonstrates just how few private pharmaceutical projects come "with success", especially early stage trials[183]) the bare figure for "trials without success" means little on its own. This desperation does rather suggest difficulty in coming up with sufficient evidence to drive a hugely important assumption in the analysis.

It is not claimed here that public funders do not suffer many problems in allocating funds efficiently or sometimes face perverse incentives or that, under political and interest-group pressure, they do not seriously misallocate funds at times (some would argue, for example, that the NIH and the EU allocates far too high a proportion of its budget to diseases that do not constitute the majority of the global disease burden). All vaccine finance mechanisms face these problems, and need to honestly demonstrate how they can work around them. Might, for example, more open collaborative models of research handle these problems better than current publicly-funded push initiatives? Without an open discussion on the matter, we cannot know. The question is still begged, however, why it is that Kremer does not provide a dossier of evidence to support his modelling choices, instead of using this one example repeatedly?

---

[182] Kremer, ibid.

[183] DiMasi et al., J.S. "The Price of Innovation: New Estimates of Drug Development Costs," Journal of Health Economics, March 2003, p151-185.

It is also ironic – given the concern of APC supporters that the administrators of NIH-funded, foundation-funded, and publicly-funded projects in general might "overstate the chances of success" of projects that they have got attached to and have vested interests in sustaining – that those pushing most for the APC mechanism do not insist on higher standards of evidence for their own case, to avoid the risk that *they* too are overstating the chances of success of their particular project.

8.1.6. The ninth layer of inefficiency – Extra administration costs

There is an extra cost variable imposed in the calculations, to account for the "higher administrative costs associated with government activities." This adds 20% to all government figures in the default setting. Again no evidence is provided for this. This is on top of the distortions listed above.

In spite of multiple layers of institutions and treaties required to make an APC work, there are no administrative costs of this incorporated into the calculations for APCs (maybe this is because the technology of Kremer rules out institutions anyway?).

8.1.7. The tenth layer of inefficiency – Biotechs

Regarding tax liability of biotechs, Kremer points out that 50% of vaccine research is done by biotechs. Biotechs on the whole have no profits, so large pharmaceutical companies have to engage in transactions with them to use their tax advantage (including buying biotechs up) and this has a transaction cost resulting in a further inefficiency. One of the side-effects of this is that it makes tax credits even less effective for vaccine research than standardly is the case[184] since much vaccine research goes on in biotechs. This is added in all calculations as an extra fraction on the tax credit (presumed at 5%).

Incidentally, this is an empirical issue. The value of the unused tax-break should be reflected in the share price of a biotech. To the extent that there are transaction costs for large pharmaceutical firms involved in using up the tax-break, the biotech share price will need to be lower to compensate (i.e. the post-acquisition value of the biotech, including share price and tax-break value, must be just high enough that the biotech is still a 'fair buy' for the large pharmaceutical company). The inefficiency is that the initial tax-break is therefore less valuable to the biotech and, at the margin, the return on its investment must be higher, and the number of leads followed and the intensity of research lower for any given tax-break. This should leave a trail of evidence in share prices as well as in required rates of return for biotechs.

[184] K2:12 gives a good summary of the case against general tax credits for tackling vaccines: i) It is potentially claimed by firms not pursuing research on these diseases; ii) It can still be claimed for research on these diseases that may not be applicable for developing countries; iii) Governments can face these excess claims even in the case in which the desired product is technologically infeasible; iv) It does not address the access issues (firms keep the IPR and are not committed to selling products at a particular price). The caveat is for tax credits for clinical trials and an R&D tax credit triggered by trials. The problems with these though are that firms may inflate costs (including keeping the best ideas out of the tax credit scheme if any stipulations are made on availability, but then taking full advantage with clinical trials when they are less likely to succeed), and that much of the riskiest and most important research may be at the pre-clinical stage anyway especially for clade C HIV and malaria.

8.1.8. "No Probability of Success" – What Does it Mean Anyway?

There is "some dispute about the state of science for some potential projects."[185] In APC calculations, however, firms "know" this and, if privately funded, will not pursue such projects, whereas policymakers do not know it *ex ante*. The default setting for the probability that the vaccine is scientifically impossible is set at 33%.

The whole notion that we can get a handle on those projects that had "no probability of success" *ex ante* is probably wide of the mark. The high attrition rates in studies of drug development costs suggest that, in a sense, most projects had "no probability of success" *ex ante*. Rates of failure are high in private pharmaceutical trials, because the whole process is risky. And, in vaccines, the rates of failure may well be higher than in other areas of pharmaceuticals. The empirical issue is whether different mechanisms can improve this situation.

To make operational Kremer's notion of "No probability of success" projects, clearly we require good empirical evidence on the relative incidence of such projects across sectors and on the past choices of various players – private financiers, the public sector, etc. It is *ex ante* expectations that we are essentially interested in, so that *ex post* failures are no guide to this *per se*; we never really get to know *ex post* whether the actual rate of failure that took place was the 'efficient' rate or not.

The "No probability of success" phrase therefore turns out to have little empirical or behavioural content; it is a device for introducing a parameter for describing poor choice over projects, and is a tautology: It failed because it was the result of a poor *ex ante* choice; it was a poor *ex ante* choice because it failed. When the APC modellers set the "no probability of success" at 20% for public science but 0% for private science, this simply indicates an extreme assumption about the ability of publicly-financed initiatives to target research. Although each layer of inefficiency looks to refer to something different, in a sense, they *all* simply refer to an inherent inability of publicly-funded (non-APC) research to choose good projects and kill bad projects.

8.2. Poor Targeting – A Summary, and The Need to Cut Publicly-Funded Research

All these layers of inefficiency add up. Depending on the shape of the probability distribution (in the sense that the more 'steep' and the more 'drawn out' the good projects are at the top of the distribution, then the more costly are targeting mistakes), adding together the loss of the top half of the distribution containing all the most promising projects (indeed the APC calibrations seem to presume only a few 'good' projects, suggesting this does indeed impose a very high cost penalty[186]), the biased way this is done (including basing it on previous IPR-related problems), the penalty imposed on 'good projects' since they 'crowd out' taxpayer-financed APC-backed projects, the picking of large numbers of dud projects, the randomness of the rest of the projects, the 20% extra administrative penalty applied to all non-APC activity[187], and the 5% penalty to allow for biotechs, it is no wonder that even the best publicly-

[185] K3:3.
[186] K3:7.
[187] Apparently, APCs have no administrative costs. In reality, is likely that APC administrative costs are higher than for many alternatives.

financed and foundation-financed vaccine R&D is visualised as spectacularly hopeless.

In the Kremer calculations, these inefficiencies reduce the number of vaccine leads being followed by push mechanisms compared to pull mechanisms for any given injection of taxpayer funds. Lower vaccine R&D intensity leads to slower vaccine development and a higher eventual cost per vaccine[188]. This is what drives his cost-per-dose figures.

This is not a minor issue. It is one of the two main driving forces behind Kremer's results (the other being his removal of multiple layers of distortions and costs from the APC mechanism itself) and the main justification he himself gives for the APC. At the risk of labouring the point, the quote above read: "Even in the best case, if the government funds only worthwhile research projects and researchers focus all their energies on developing a vaccine, the expected discounted cost of developing a vaccine is likely to be similar in net present value terms whether research is financed at the front end, through government grants or induced by payments for a successful vaccine at the back end."[189]

The favourable results for the APC – that vaccines produced from it are, in some cases, less than a quarter the price to develop of vaccines produced via the public funding of applied research and joint ventures – turns out not so much forced by the wonders of the APC *per se* as by all these perceived failures of applied research and joint ventures. These layers of assumed inefficiency have sealed the fate of anything other than the APC (given that the APC *is* modelled as always working totally efficiently itself). There is much less consensus in the economics literature on these inefficiencies than ever hinted at in these calculations[190].

As a sign of the poor quality of the 'cost-effectiveness' data once used to support the APC agenda for early-stage vaccines for HIV, malaria, and TB, these figures are no longer used in *any* recent publications on the subject. Neither 'Strong Medicine'[191] nor 'Making Markets for Vaccines: a practical plan'[192] make any reference to them.

The APC should be allowed to crowd out other research

One logical consequence of this reasoning is that the roll out of an APC program should be contemporaneous with a cutting back of publicly-funded front-end research in the APC equilibrium. If the back-ended APC really is more than four and a half times more cost-effective in cases like HIV/AIDS, then it would make sense to set the APC higher and restrict the use of front-ended research. Large-firm pharmaceutical

[188] In the auction section below, slower vaccine development – caused by i) the period when the APC price is still rising towards the optimal level; ii) getting the rate of increase in the price wrong; iii) 'experimenting' with some APCs before being able to tackle other (maybe more serious and larger) APCs – is *not* priced into the APC figures even though it *is* part of the APC mechanism, and does impose a cost penalty.

[189] One presumes this is on the basis of the same costs of capital in each case. The phrase 'best case' is, it seems, deemed only applicable to the first.

[190] See Klette et al, ibid.

[191] 'Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases', Kremer, M, and Glennerster, R, Princeton University Press, November, 2004.

[192] 'Making Markets for Vaccines: a practical plan' Centre for Global Development, 2004, http://www.cgdev.org/globalhealth/proj_pull.cfm.

R&D should be allowed to crowd out publicly-funded R&D until, at the margin, the effectiveness of the two is the same.

That Kremer does not make much more of the need to cut front-ended research in equilibrium is either because it is expedient to ignore it at the moment, or that front-ended research is not being presumed to be anywhere as inefficient as supposed in the APC calculations. Indeed, Kremer uses industry figures that presume the already typical level of front-ended support, without making any adjustment for its inefficiency[193]. Kremer even suggests in places that publicly-funded research should also 'chase the APC' – which would be wasteful if it really did involve four or more times the resources[194]. Or maybe even Kremer does not trust his own figures? Maybe he knows that they are not to be taken seriously; they are only a lobbying device even in his interpretation of them?

This is part of a generally contradictory attitude towards inefficiency. Contrary to all the supposed inefficiencies used to drive the APC result, it is concluded that: "Government-directed research programs may be well-suited for basic research..."[195] and that "much of the riskiest, most important investments must be made at the pre-clinical stage"[196] If anything, the monitoring and selection problems should be much worse at those stages, with misrepresentation much more difficult to achieve at later trials stages. Yet, inefficiencies seem only to bite at those stages that are traditionally the preserve of large pharmaceutical companies, and mysteriously disappear at these other, one would expect more difficult, stages.

If the argument is made to semantically revolve only around the inability of publicly-funded vaccine research to pick promising candidates at later trials stages, it leaves all the early stage problems, to the extent they exist, unresolved. Since all the failures in the earlier stages will feed through into the effectiveness of the APC, if the APC is to somehow be declared 'more cost-effective than alternatives' it has to somehow tackle these much earlier problems. The late stage problems may not be the most insurmountable problems. Either way, the APC leaves us wanting.

Before imparting on a vast experiment to create a new mechanism to avoid a "vaccine-development effort that might not be warranted scientifically" perhaps the exact evidence for this assertion could be more adequately spelled out and backed up with evidence? If the APC mechanism carries many distortions and costs of its own, would it not make sense to keep the debate open to continue the investigation of alternative approaches to the selection of projects that might enable easier removal of those that are 'scientifically unwarranted'? A system based on ever-tighter patents,

[193] This affects cost-effectiveness comparisons. The logic of the model into which the figures are then fed suggests that this support should be cut back in equilibrium.
[194] It might be argued that the reason no large-scale cut-back in equilibrium has been suggested is because the kind of research being done by these publicly-supported projects has large public good/externality effects and that large pharmaceutical firms find this difficult to do. The high inefficiency of such projects is tolerated in a trade-off with the costs that large firms would face trying to overcome the problem. However, this logic does not go through. Once the APC is in place, the position of the equilibrium in this trade-off would still move in the direction of *less* front-ended research.
[195] K6:Abstract.
[196] K8:5.

more secrecy, and even less sharing of information does not transparently improve this ability. Perhaps a more open, collaborative structure would[197]?

8.3. Publicly-Financed, Foundation-Financed, and Pharmaceutical-Financed Research

The cumulative effect of these many layers of inefficiency is to bias the results heavily in favour of APCs. Kremer provides practically no evidence to justify this. Given the insistence of Kremer that vaccines will come largely from large pharmaceutical firms and not from small biotechs, not-for-profit, developing country, or government-sponsored research, and given that the APC biases in favour of the former over all of the latter, this is not a trivial issue. In particular, it turns out that the APC model's criticisms are not targeted at all science *per se* but at that part of science often performed by large pharmaceutical firms, which seems a little disingenuous.

The problem is that there is plenty of counter-evidence of good publicly-financed and biotech-financed trials and of poor-quality expensive pharmaceutical-financed trials. Just for a flavour of this counter-evidence, however, the following letter from five government or government-funded researchers reveals the overwhelming role of US government funding in the development of the AIDS drug AZT. AZT is one of several drugs at the centre of many a contentious dispute over access to AIDS medicines, in particular in Sub-Saharan Africa.

New York Times, September 28, 1989.
Heading: Credit Government Scientists with Developing Anti-AIDS Drug.

To the Editor:

The Sept. 16 letter from T.E.Haigler Jr., president of Burroughs Wellcome Company, was astonishing in both substance and tone. Mr. Haigler asserts that azidothymidine, or AZT, was essentially discovered and developed entirely by Burroughs Wellcome with no substantive role of government scientists and Government-supported research. This will be a surprise to the many men and women who have devoted their lives to working for the viral cancer program and developmental therapeutics program of the National Institutes of Health over the last 25 years.

We (associated with the National Cancer Institute and Duke University) make this statement as co-authors of the first publications describing AZT as a drug for treatment of acquired immune deficiency syndrome (Mitsuya, et al. Proceedings of the National Academy of Sciences, 1985, and Yarchoan, et al, The Lancet, 1986). There are few drugs now approved in this country that owe more to Government-sponsored research. In the interests of brevity, perhaps this point can be summarized most efficiently by stating what Mr. Haigler's company did not do.

The company did not perform the first synthesis of AZT. This was done by Dr. Jerome Horowitz of the Michigan Cancer Foundation in 1964, using a Government grant.

[197] At the very least this would enable the pooling of more scientific information, a more up-to-date overall picture of progress, better monitoring of incentives to act in self-interested ways, and those who face the consequence of a funding scheme having more of a rôle in its running via peer review.

The company did not conceive or provide the first demonstration of an effect against animal retoviruses. This was done by Wolfram Ostertag at the max Planck Institute in 1974, using a mouse retrovirus in a test tube. Mr Haigler's implication that his staff "discovered" the antiretroviral potential of AZT in 1984 is noteworthy. What he did not say was that his staff repeated the Ostertag mouse experiments. You cannot 'discover" something published by someone else 10 years earlier.

The company specifically did not develop or provide the first application of the technology for determining whether a drug like AZT can suppress live AIDS virus in human cells, nor did it develop the technology to determine at what concentration such an effect might be achieved in humans. Moreover, it was not first to administer AZT to a human being with AIDS, not did it perform the first clinical pharmacology studies in patients. It also did not perform the immunological and virological studies necessary to infer that the drug might work, and was therefore worth pursuing in further studies.

All of these were accomplished by the staff of the National Cancer Institute working with staff at Duke University. These scientists did not work for the Burroughs Wellcome Company. They were doing investigator-initiated research, which required resources and reprogramming from other important projects in response to a public health emergency. Indeed, one of the key obstacles to the development of AZT was that Burroughs Wellcome did not work with live AIDS virus nor wish to receive samples from AIDS patients.

In a number of specific ways, Government scientists made it possible to take a drug in the public domain with no medical use and make it a practical reality as a new drug therapy for AIDS. It is unlikely that any drug company could have found a better partner than the Government in developing a new product. We believe that the development of this drug in two years, stat to finish, would have been impossible without the substantive commitment of Government scientists and Government technology. It does not serve anyone's interests to nullify the importance of Government-sponsored research in solving problems of American public health.

Hiroaki Mitsuya, M.D.
Kent Weinhold
Robert Yarchoan, M.D.
Dani Bolognesi
Samuel Broder, M.D.

Bethesda, Md., Sept. 20, 1989

The AZT case is just one of a long list of other successful publicly-supported drugs[198]. The case for APCs needs first to assess *this* greater body of evidence carefully to see how to encourage improvements in trial attrition rates whatever sector the trials are conducted within. Incidentally, Burroughs Wellcome also benefited from the Orphan Drug Act, which provided a tax credit of 50 percent towards *their* outlays on clinical

[198] The author would welcome some balancing views from the industry.

trials, reducing even further the proportion of private funding leading to the development of AZT.

Kremer claims (see the quote earlier) that patents are not an issue for access in developing countries. But, in spite of the majority of work on AZT being publicly-funded, the rights to AZT were not held publicly, so that cheap generic copies could not be made and distributed *en masse*. Patents to AZT and other AIDS drugs *were* a problem that contributed to millions of AIDS victims in sub-Saharan Africa being denied access to the drugs that would have saved their lives. The South African Competition Commission concluded so in its ruling in late 2003, which found that the companies involved had acted anti-competitively in setting prices, helped by the nature of the IPR system.

8.4. When the APC Crowds Out the Private Sector

When we discussed the rôle of the non-APC financed parts of the research process above, we made the crucial conclusion that the part of the process that is amenable to the APC is dependent in very specific ways on the part that is not. That – in order to work efficiency – the APC needs the other part to work to a *specific and precise degree*. Otherwise the APC price is higher. Kremer presumes that this coordination is achieved perfectly at all times.

In addition to the inability to co-ordinate perfectly the APC and non-APC parts of the process, there are two much greater dangers if the underlying Kremer analysis is taken at face value.

8.4.1. The dangers and costs when public sectors 'chase' an APC

It is sometimes argued, including by Kremer, that public sectors should try to 'chase' an APC (that is go beyond the *specific and precise degree* of involvement required of them) and that somehow this would be good, since it would drive up the chance of discovery. It is decidedly not so, and would *not* drive up the chance of discovery (at least not if the private sector is reacting efficiently). It would simply encourage the public sector to overpay (possibly considerably), drive out private sector activity, drive up overall costs of vaccine discovery, and wastefully distort research. The Kremer figures become the floor to the expected costs of discovery.

The fault in the reasoning is to forget that an APC is designed specifically to encourage the private-sector research component *only*. The public sectors of all countries involved in vaccine research need to be policed by contracts to pre-agreed (contingent) amounts of research and *not* to try and win the APC[199].

For example, let's say that an APC designed to stimulate the *private* sector is set at $10bn and that the eventual repayments are shared across many countries. Let's presume for a moment that public research and private research are equally efficient (so, for the moment we not only ignore differences in research efficiency as described

[199] Incidentally, this would require some notion of the relative efficiencies of different public sectors and the ability to allocate a precise amount (though with discretionary ability to vary the relative proportions across countries over time) of activity for each to contribute towards supporting the APC.

by Kremer, but also differences in capital costs[200]). Let us say that it turns out that the public sectors of various countries 'chase' the APC, such that on average the global public and foundation sectors, *contrary to what should have happened*, spend as much 'chasing' the APC as the private sector (with public-sector and foundation activity redirected from other areas, especially neglected diseases and vaccines not covered by APCs, since at the margin it is now more profitable than before to chase the APC than to research other areas, especially if these other areas are neglected diseases or other vaccines with no market[201]). Discovery is made when an expected discounted (including capital costs) $10bn has been spent ($5bn in the private sector and $5bn in the public and foundation sector if capital costs are the same[202]). This is all based on the notion that any excess of public sector activity chasing the APC is exactly offset by the private sector cutting back (This is the logic of the Kremer model; privately-financed activity *has* to respond to an excess of publicly-financed activity by cutting back since the value of the continuation game has fallen with all the new 'players' doing research).

If one country wins (we presume an outright winner for now, though in reality this would not be the efficient outcome according to the analysis in Section 7 above), the overall expected discounted costs to the public and foundation sectors of all 'losing' countries, excluding the winner, is the $5bn of expenditures 'chasing' the APC minus the expenditures of the winner, plus the $10bn owed to the winning public or foundation sector. The cost to the 'losing countries' is close to $15bn. But one country is a big winner, 'winning' $10bn minus research costs. Interestingly – ignoring the distributional issues and the deadweight loss of taxation – countries are collectively better off for having chased the vaccine since the overall aggregate public expenditure is $5bn. However, the total cost is still $10bn since those holding investments in pharmaceutical firms take a $5bn hit.

However, what if some firm in the private sector wins? It gets the $10bn APC contract, the public sector pays the $10bn and has its resource costs of 'chasing' the vaccine of $5bn to cover. Countries are collectively worse off since the total cost to the public sector is $15bn (with no big public sector winner).

Averaging across both outcomes however (there is a fifty/fifty chance of either) the average cost to the global public sector is $10bn, and the total expenditure on research outlays is $10bn. It might seem to not matter that the public sector 'chased' the vaccine. However, besides the distributional issues (which are not insignificant) the deadweight losses – caused by taxation – are greater on account of chasing, since $15bn rather than $10bn needs to be raised in tax revenues[203].

If the public sector does research as poor as claimed

[200] The research efficiency part of this is obvious. But if private research involved higher capital costs, then any given amount of private research costs will lead to fewer out-of-pocket research expenditures and lower chances of discover. In equilibrium, for each to have equal chance of discovering the vaccine (as being modelled here) the private-sector has to spend more than half of the total outlay, and the public-sector less than half of the total outlay, with the difference being in proportion to the difference in their relative capital costs.

[201] Remember the general equilibrium requirement laid down much earlier in this paper?

[202] They are not, but we ignore this for now.

[203] This is offset by any 'deadweight gains' in the winning country.

If, however, we take seriously the notion that publicly-financed initiatives are poorly targeted as Kremer argues, such that each public dollar of expenditure reduces the probability of winning of the private sector by less than would be the case had it been a dollar's worth of private funding, the effect is much, much, worse. For example, if publicly-financed research is only a third as effective as privately-financed research (roughly Kremer's claim), one scenario for the $10bn APC is that it attracts, as before, an expected discounted $5bn of private research, but now an expected discounted $15bn of public research (given the inefficiency, this is equal only to an expected discounted $5bn of private activity, so that total 'effective' expected public, foundation, and private spending together is still $10bn). Even if the public sector discovers the vaccine first, the global expected opportunity cost is $20bn ($5bn of private expenditure plus $15bn of public expenditure). One country wins $10bn with all the other countries shouldering $15bn of research costs minus the research costs of the winning country, and the $10bn APC, bringing the costs of the non-winning countries to close to $25bn.

If the private sector discovers the vaccine first, it costs the global public sector an expected $25bn in total – the original $10bn APC price agreed, plus its own spending of $15bn. This is two and a half times what the APC was supposed to have cost global public finances. There are of course even worse scenarios (opportunity costs can approach $30bn and the cost to tax-payers – excluding the winning country – can approach $40bn, though the probabilities of the latter are asymptotically small).

Averaging across the good and bad outcomes, global public expenditure in this latter case is $20bn (0.5*$15bn+0.5*$25bn), twice what the APC was supposed to have cost. The distributional costs are greater than before, and the deadweight cost of taxes even greater (and this time, the deadweight gain to a winning country does not offset the deadweight losses of the others, though it probably did not before either).

The more that those setting up the APC feel that the public sector is inefficient the less they would want to tempt them with an APC and the more they would want to police them not to overdo their expenditure[204].

The argument that overspending in this way is fine since it increases the chances of getting a vaccine is not only wrong, but totally ignores the global budget constraint mentioned earlier. Think what would happen for the funding of other research projects and for global efforts to create other vaccines and spread treatments if countries cut back on programmes that did not have an APC in order to concentrate more on those areas covered by an APC. And what would happen when they then had to cover the excessive public costs of earlier enacted APCs? And it flies in the face of the notion that we are supposed to be comparing mechanisms on the basis of their cost-effectiveness. If we allow overpaying on the APC mechanism it is hypocritical to criticise it on other mechanisms.

[204] Observe that if public sectors did not have delusions of being equally efficient then they would clearly realise that it was never in their interest to chase the APC. But, Kremer claims that they proverbially over-exaggerate their chances and will think that they have an equal chance with the private sector of 'wimming'. Indeed, if Kremer is wrong, and public sector vaccine programs are actually much better than the private sector, it would be rational to chase the APC, and in fact that would make the global costs lower! It is argued here that the side-effects of the APC work against this however.

8.4.2. The danger that already low levels of private research will collapse

There is a second danger. For various reasons, research is already going on in large pharmaceutical firms for vaccines for diseases primarily of the poor, even though the prospective market is very low. One possibility is that this is done for its 'option' value[205]. It may also have some PR/social responsibility value. At the same time, more than 50% of vaccine research is taking place in small biotechs.

For APCs to work efficiently (and presuming no *en masse* creation of APCs in a surprise overnight announcement with no period of discussion leading up to it[206]) firms would have to be assured that they would be *fully* compensated, post-vaccine development, for any R&D costs sunk prior to an APC being set up. The private sector worry (including capital markets) would be that an APC would factor in only what is *marginally* required to get a vaccine developed. Indeed, that is how the auction component of the APC model is deemed to work, and is how the cost-effectiveness of APCs has so far been calculated. This argument bites more, the fewer the firms in the industry and the fewer the firms currently investing in the area of a particular prospective APC[207]. If, unlike the Kremer model with its same-probability-distribution repeated each period, past R&D draws vaccine discovery closer, a lower APC price will be needed to stimulate the remainder of the R&D needed for its development. There is clearly an incentive to hold back on pre-APC R&D if there are worries that the APC price will be negatively related to these expenditures. This effect is worse if the decision as to whether or not to initiate an APC is also dependent on these previous decisions.

Adjusting to account for past research will not solve the problem

It might be thought that an adjustment could be made to the APC terms to reflect this previous sunk R&D expenditure. However, if firms have sunk different amounts before the APC is set up, an average adjustment to the APC to reflect the average will disadvantage those who have sunk costs already and advantage those who have not. This would create a 'prisoners' dilemma', with no firm prepared to sink investments.

Besides, the adjustment (average or otherwise) could not simply show up in a higher APC price, since that would simply operate as a higher APC in the continuation game and encourage excessive behaviour compared to what is optimal in *that* game, leading to higher vaccine development costs, and still fail to compensate for pre-APC sunk costs. Instead, the adjustment for previous levels of research costs would need to be some sort of 'side-payment' on top of the APC[208].

---

[205] Nichols, 1994.

[206] This section recognises similarities between the setting up of a system of APCs to the setting up of a fixed exchange rate regime or even a single currency. Policymakers need to be mindful of incentives created in the period leading up to the fixing of rates of parity between currencies. Surprise overnight announcements that lock-in parity will avoid the dangers of the instability on the path to the creation of parity, but run the risk of fixing the rates at the wrong parity since there is not time to learn and adjust. Long periods leading up to the fixing of parity will enable better learning and may lead to a more correct parity, but will run the risk of perverse incentives and instability. The analogy to the latter is an APC system that is allowed to experiment and adjust. Open source would have much the same set of issues to contend with.

[207] Intuitively, the APC price benefits the many who might now enter, but punishes heavily the few who have already sunk investments, giving them a competitive disadvantage.

[208] Observe the level of information required.

The problem is further aggravated if *delay* is part of the mechanism, as frequently asserted by Kremer. If the idea is to 'experiment' with some APCs before setting the terms of others, then some mechanism has to be in place in advance to ensure that those currently sinking research costs in areas covered by the future APCs will still invest assured of a side-payment to cover these sunk costs. Otherwise, the danger is that by investing now, they lose out by worsening the terms of the APC they face later (another prisoners' dilemma).

Without such a mechanism in place – and handling the problem *efficiently*[209] – the danger would be that even the small amount of private R&D currently taking place would dry up. Quite literally, as far as financial markets would be concerned, this R&D has suddenly become much more expensive at the margin. Indeed, the long gestation period leading up to the setting of terms for an APC will have a depressing effect on research if this mechanism does not exist or does not work well.

What happens to the options value of such pre-APC research is not clear. It might be thought that the options value rises if an APC is 'suspected', but that if all firms reasoned this way and sunk research costs in advance of the APC, and if the APC failed to reflect these costs fully in its payments, then the overall profitability from the market would be lower. Perhaps, the most one can say is that the fewer the number of firms, the greater the chance of the 'collusive' outcome of lower intensity option-based R&D: The 'game' is such that firms would like to collude to hold back so as to get a higher APC price, but this gives incentives to cheat since the options value of research is rising; if the industry is competitive it may be difficult to discipline this cheating, but if the industry has a few large players, the incentives to cheat are lower, and the few large players will manage to 'collude' to hold back on pre-APC R&D. Indeed, they do not really have to 'collude'; this sort of behaviour on non-APC areas of research is forced on them by financial market pressures.

The effect is further reinforced if large pharmaceutical firms are engaged in neglected disease research for partly or even largely PR or social responsibility motives. The presence of APCs in some neglected diseases areas but not in others will distort the relative pattern of activity away from those areas that are not covered; intuitively the firm gets all the PR and social responsibility kudos from working on neglected diseases covered by APCs with the benefit of possible payment via APC. In developed economies the public sector is also often as bad as the private sector at targeting the diseases of the poor, tending to spend by far the greater proportion of funding on, relatively speaking, 'rich economy', diseases. Perverse distortions in the public sector cannot be ruled out either, with those limited public resources spent on the diseases of the poor now seemingly more productively employed working towards APC covered rather than non-APC covered diseases. All these are forms of crowding out, but are hard to quantify.

Even if APCs are not promised for specific diseases, their mere *possibility* will reduce the value of current R&D if a side-payment mechanism does not work. This affects even those diseases for which there *never will* turn out to be APCs. Again, this is

[209] It would be *highly unlikely* to be efficient given the huge informational requirements.

another reason why the APC has to be thought of not in terms of one or two APCs but in terms of its general equilibrium impact[210].

8.5. Does the APC force Global Centralisation of Control over Public Sector Research?

The APC is only supposed to deal with *private* incentives. The APC should have been set bearing in mind the pre-agreed globally-contracted publicly-financed non-APC-based research. There is no reason to encourage public research beyond this (notice how getting up to this pre-agreed amount may need more expenditure than is currently taking place).

The 'chase' equilibrium described earlier is actually the result of a prisoners' dilemma. The public sectors of individual countries have incentives to 'cheat' and spend beyond the pre-agreed amounts, crowding out private expenditure. Discovery does not speed up (this is all in the logic of the way the probabilities and technology are set up in the basic Kremer model). But it does run the risk of the public sectors of the majority of countries overpaying towards the eventual vaccine, and reducing heavily the effectiveness of public spending. Another way to think about this is that, once the APC is in place and all pledged contributions are fixed, the marginal cost for each country of its research is less than its expected average cost (which factors in the pledged contributions).

The APC mechanism would have to *police* public-sector non-APC research – to make sure that *no* public sector player went beyond its *pre-agreed* level of research. This, yet again, indicates that the mechanism is far from being light on institutions and centralised 'control'. Paradoxically, rather than avoiding making decisions about publicly-funded research, the APC setters would have to have a good handle on what was going on in the public sector to make sure that the sector did not overspend (i.e. distort expenditure flows in this direction). The APC does not avoid having to tackle the asymmetric information problems of that sector, as Kremer suggests. The asymmetric information problem just raises its head in a different fashion.

8.6. When the APC Fails 'Additionality' – the *Greater* Use of Public Funds, and the Bias to Large Pharmaceutical Firms and against Not-for-Profit, Biotechs, Developing Country and University Research

In Section 8.5 we looked at the incentives in the *public sector* in response to APCs that can cause crowding out and higher public costs. This section looks at incentives in the *private sector* that can also cause crowding out and public costs to be much higher.

All APC cost comparison figures, so far, have been calculated on the basis that the APC pays *only* for that part of *privately*-financed research that is stimulated, i.e. *additional,* as a result of the APC's presence. Kremer refers to APCs as "market enhancing incentives"[211], and views this as a strong factor in their favour: "Public

[210] Open source would also have to find some way to handle many of these problems too.
[211] K1:9.

funding is generally given to institutions that are already well-known in a particular field, which creates little incentive for new players to invest funds in a different idea or approach. In comparison, a pull program such as a purchase contract or tax credit for purchases is open, and encourages innovation form any participant."[212]

There is nothing in the APC modelling to indicate how this is guaranteed or, indeed, even conceivably possible. For example, the calculations presume that firms whose research was funded by tax-breaks or subsidies would *not* be able to be granted an APC (though they would no doubt object if this were the case), or at least should have the APC price cut in proportion to their tax-breaks and subsidies (observe straight away that this is greater than the tax-break any one of them individually received[213]).

An un-modelled separation mechanism

To avoid totally the inefficiencies of the *other* publicly-funded R&D instruments that Kremer discusses, he must be presuming that either those firms being motivated by APCs are researching towards vaccines relying totally on private capital markets stimulated on the basis of the APC alone and on *no other* publicly-funded research support device, or that there is some, so far un-modelled, mechanism for policing the APC to apply *only* to private research activity that is *additional* to that supported by other public research support devices.

Such a separation mechanism would require a phenomenal amount of information to separate out that part of private activity to be rewarded via the APC from that that should not. This separation would suffer all the usual asymmetric information problems and incentive distortions (e.g. firms distorting research so as to make it look as if their tax-breaks or subsidies had indeed supported activities *not* subsequently covered by an APC). It would also mean different firms being paid differently for essentially the same outcome, resulting in a lack of transparency about what exactly firms were being rewarded for.

To the extent that no such separating mechanism exists, taxpayers would find themselves paying tax both to cover the research tax-breaks and subsidies *and* to cover the costs of the APC (both of these sets of taxes generating deadweight losses), and the APC (for equal expenditure of resources as other mechanisms) generates more crowding-out than the effects of the APC on its own would suggest, thus under-measuring the public costs of developing vaccines via APCs.

Some unpleasant figures

As a quick thought experiment, imagine what the cost-effectiveness would be if the British government or the World Bank initiated an APC for one vaccine, if 50%-60%

---

[212] K2:10

[213] The logic is this: Imagine that ten firms are trying to win the APC, that ten firms is optimal, and they operate at optimal intensity that will lead to them spending the expected discounted value of the APC trying to win it. If they all receive tax-breaks of 20%, then the APC price has to be set 20% lower than without the tax-breaks. If one of them wins the APC, the deduction from the APC is 20% although that firm's proportion of total tax-breaks was only 10% of that, which is 2% of the total APC price. The reduction in APC price is ten times this firm's tax-break. If firms differed, the APC should be adjusted accordingly. If a firm gets 30% tax-break, the APC price needs to be cut by 30% for that firm. How the evidence is gathered to achieve these adjustments is anyone's guess. And large pharmaceutical firms are more likely then other types of researchers to be able to hide information to prevent repayments.

of its costs were capital costs, and – being unable to police firms globally on their sources of funding – firms benefited from tax-breaks, subsidies, joint venture outlays, etc. covering 35%-50% of their costs. A dollar of public funding would translate to at best only 39.4cents and at worst 28.5cents of private out-of pocket research expenditure[214]. This does not sound like a good deal. And this is almost certainly the very best case scenario we might hope for for a vaccine as complex as HIV.

Adjusting to account for other support will not solve the problem

It might be argued that the APC price would be set (lower) to reflect the *average* component of vaccine R&D being covered by tax-breaks and subsidies, etc. But this suggests:

1) The need for a great deal of information about investment decisions of firms and of the levels of tax-breaks, subsidies, etc. in the 'APC-generated equilibrium' (i.e. quite a dynamically complicated calculation to make) so that the APC price can be set correctly[215];

2) That this would require a great deal of transparency of information on all sides, both by firms and by 'APC setters' (given the long gestation periods of R&D investments, much information, even if clear, has unclear meaning anyway. And the APC still encourages a great deal of firm-level secrecy which works against setting the APC optimally);

3) Given that there is no such thing as an 'APC setter' but a whole set of APC institutions and committees, that this would require a huge amount of coordination and policing across disparate institutions and countries;

4) That this would generate 'prisoners' dilemma' situations, and hence break-downs in coordination, leading to a total inability to ensure that the APC is applied globally only to *additional* research activity (countries would have to be banned form introducing new tax-breaks, subsidies, etc.[216]);

5) That, because this is only an average, for those not amenable to tax-breaks (such as not-for-profit companies and non-profitable biotechs) the APC would end up being set too low to achieve *their* maximal incentives (this, yet again, simply indicates that the APC *should* be able to *discriminate* between those companies totally relying on it for incentives and those relying on other public support, so as not to discriminate against the former. This is just another way of saying that part of the APC payments to winning companies should be taken away in proportion to their tax-breaks, and subsidies, etc.);

---

[214] 28.5cents worst-case scenario derived as: $1 leads to 60cents capital costs and 40cents funding for R&D. This pays for 50% of R&D, the rest being from other tax-breaks and subsidies. Total public funding producing 40cents of privately-funded R&D is $1.40. 40cents divided by $1.40 is 28.5cents. The best-case scenario is done similarly using 50% capital costs and 35% of tax-breaks and subsidies.

[215] To the extent that it is set wrong, all the usual comments about the lower cost-effectiveness of APCs compared to alternatives apply. And the more dynamic the issue is, the more likely it *will* be set wrong (for example, how should the APC adjust to changes in various countries' taxation systems?).

[216] Even if not connected to the particular vaccine/drug at hand, since creating generally more favourable conditions for firms may be one of the ways around rules barring specific measures to favour firms.

It biases towards large pharmaceutical companies over others

6) That this would still leave the APC biased towards companies that are able to take advantage of other public support devices such as tax-breaks and subsidies – that is large pharmaceutical companies[217]. Indeed it would most likely distort research in that direction leading to even greater use of public resources, since it is difficult to see how it could be policed. It is yet another 'prisoners' dilemma'; if every other country has policed the use of tax-breaks and subsidies out of its domestic industry working on this particular vaccine, there is a strong incentive for individual countries to cheat by supporting their industry since the marginal gains are now much higher at any given APC price (they could also provide support in hidden ways, including hidden subsidies and favourable tax legislation and even deals). This is perverse given that APCs are supposed to achieve the complete opposite.

Yet again we find that the distribution of firms relying on APCs is biased in the direction of large pharmaceutical firms and those tax-advantaged in other ways. If relatively more innovative research takes place in the not-for-profit or non-profitable biotech companies, this would, *ceteris paribus,* slow the speed of vaccine development compared to an equal expenditure on some other support device that might more directly benefit not-for-profit and biotechs. Of course this could all be masked by such a large APC price that even the not-for-profit or biotechs are happy[218].

7) That the APC is then only *part* of the overall public costs of developing a vaccine and all these other public costs, and concomitant inefficiencies and deadweight losses, should be priced into the overall cost of a mechanism based on an APC.

This suggests that the APC – to be totally efficient – would have to be a closed system, such that activities covered by it would have to be barred from coverage by any other system of public support. A public or foundation research body working on a vaccine with government subsidies or foundation grants should *not* be allowed to win the APC. If they are allowed to win the APC, then the private non-subsidised sector would, *ceteris paribus*, need compensating by being fed a higher APC price to get them to take part in the mechanism in the first place. Also, given the heavy reliance on patents and the build-up of sunk costs, this puts a bar on university researchers – working along-side private researchers from a pharmaceutical firm – from sharing information with other university researchers not on contracts with the firm, since this is an extra risk to that firm and adds to its capital costs.

In a sense there would have to be no 'straying' by any of the players covered by the program away from the program. Might more open collaborative systems be much more capable of preventing such 'straying' and crowding-out, be potentially much more self-sufficient with more sharing of information, and potentially lighter on

[217] 'Not-profitable' biotech firms can take advantage of tax-breaks only to the extent that they can be bought up by larger pharmaceutical companies to 'cash in' on the value of the tax-break (the biotech firm amasses all its unused tax-breaks as an asset until taken over), but Kremer himself argues that this involves a level of inefficiency due to all the transactions costs involved. And unless they are bought up (in sufficient time?) the tax-break goes unused.

[218] Since in the equilibrium described here they *should* be unhappy to use an APC compared to alternatives, if it turns out that they are *not* unhappy, it means the APC is set too high.

public finances? To the extent that choosing an APC – over alternatives like open collaborative approaches – makes the ability to correct many of these crowding-out distortions more difficult, the APC even makes matters worse. This problem has been ruled out by Kremer from the start in all calculations.

8.7. The Dangers of the UK (or any Single Government) Going it Alone

Kremer argues that "An advance purchase commitment from a single government would pave the way towards commitments from a broad range of governments and donors... One possibility would be for the United Kingdom to pledge to purchase the first vaccine developed [from a selection of diseases for which vaccines are being sought] and [then] seek commitments from other countries for subsequent vaccines." He further argues that "It is unlikely that vaccines for all three diseases would be developed simultaneously, but if donors wanted to limit their exposures, they could cap their total promised vaccine spending under the program, for example at $520 million annually."[219] The logic of the APC model and of rudimentary finance not only does *not* support such an assertion, but suggest that this is extremely dangerous and expensive logic.

Terms could never be set optimally

No APC could *ever* be set *optimally* if it was set *unilaterally* to *only* pay for the first one of a selection of vaccines. This would require, from the start, complete certainty in the minds of vaccine developers and financial markets that other governments would pick up later vaccines *and* would cover *all* of the costs of R&D of those later vaccines including those costs incurred between announcement of the original APC package and the date when the APC for the disease that a firm is working on was set up. Otherwise, it will mean one of the following:

i) Having an inefficiently large APC to compensate for the expected costs of failed research projects not only on the successful disease but on all other diseases that are being worked on but which at some point may (if other countries do *not* sign on to the APC[220]) lose the APC option. That the government has the 'option' on only the first one of several diseases from a selection *has* to be priced in financially;

ii) A small and inadequate APC in the hope that other countries will enact APCs to boost the value of the first country's APC, with, meanwhile sub-optimally low research intensity for any one disease, and slow discovery of vaccine for any disease.

The only way the APC can be set 'correctly' and for research intensity to be optimal is if private players are 100% sure that the second, third and further vaccines *will* be picked up by other countries.

Things are made even worse by the logic of dynamic inconsistency – leading to the danger of collapse and of even more expense and delay in developing vaccines. As time goes by, the APC setters of vaccines subsequent to the first one will have incentives not to take into consideration the research costs being sunk on these vaccines (supposedly stimulated by the first APC) that took place before the new APC

[219] K7:35

[220] The probabilities that this may happen get worked into the APC price, though it is hard to see how this works if the APC price is supposedly set in an auction.

was set up, and will offer too low an APC price (if firms really have invested in research into these other vaccines, then the needed APC has indeed fallen[221]). If private players realise this *ex ante*, it will raise the uncertainties and risks of pharmaceutical players *ex ante* about how they will be treated *ex post*, hence raise capital costs of *all* players, numb research intensity on all three diseases[222], and raise the APC price that the first country has to set (it is not clear how this APC price would be set though, since the auction mechanism could not work in this case). At the very least, the idea necessitates that these later APC setters contract now *not* to take advantage of sunk research costs, and commit *now* to setting up APCs. Given that it is not even clear who the countries will be that will offer the APCs, this is ruled out.

Besides, commitments by one country to purchase vaccines will benefit other countries, giving incentives for other countries to free-ride on the initiative. Allowing this problem to intensify – as, one by one, countries join – does not seem a sensible way to proceed. Co-ordination to get over this problem is always much easier when no country has yet signed on to an agreement, i.e. through some sort of Treaty (the private benefits are always higher knowing that one's vote gets to tip the Treaty into existence[223]).

The greater dangers later

The risk that other countries will not come on board later would lead to:

1) The British government having to guarantee at the start to bail out a failed system by paying all costs if the UK fails to secure APCs from other governments. After all, Kremer is suggesting that the British government policy would work by convincing private investors to sink the optimal amount of research expenditure while they were awaiting the setting of the remainder of the APCs. To the extent that private investors believed and trusted the British government, they should be compensated for this failure to deliver.

2) Higher capital costs, and the risk of a self-fulfilling collapse. The increasing risk that failure might happen would feed back into private capital costs and a higher needed APC price[224]. Seeing the way the needed APC price is rising, the incentive for other countries to take part will faill, which (in rounds of the same logic) feeds back to make failure to establish further APCs and destruction of the first APC a self-fulfilling equilibrium[225].

3) Having to put up with slow vaccine development;

4) Possibly reneging on contracts, having to compensate firms, and initiating a non-APC approach. But by then, even that approach may have become more difficult, and

[221] This is ruled out in the simple model of Kremer, since the technology is stationary.
[222] Since at the start the winner is not clear.
[223] Incidentally, this is why sometimes it is better to set these things up to with the possibility of not getting the hoped for initiative 'unless there is unanimity'. This helps to overcome dynamic inconsistency, and helps to bring the thing about in the first place.
[224] If the APC price is being set by an auction, the price would be 'bid up'.
[225] If set in an auction with a totally flexible price, the price would spiral out of control. Though if the auction is a fixed rule regarding the allowable price rise and this is not able to rise fast enough, then research intensity collapses at any particular price.

there will have been unnecessary delay of vaccine development and waste in to the bargain.

It is suboptimal to try to encourage vaccines for three diseases but only be prepared to pay for one, and potentially risks great waste of resources and damage to vaccine R&D incentives. Any APC program has to have *all* vaccines that plan to be covered by it, actually covered from the start, and all countries who plan to take part cooperating from the start. Going it alone and hoping that others will follow, will increase the chance of an expensive failure. Why does Kremer make the claim that countries can go it alone in this fashion, when the economics clearly shows that this is a totally harebrained way to proceed?

9. THE UP-FRONT RESOURCE COSTS OF APCs[226]

It is argued that the APC mechanism does not require resources up front and therefore is not a burden on financial flows until it needs to be repaid: "Such a commitment does not require money now"[227]. And it is claimed that "given appropriate legal language, the key determinant of credibility will therefore be eligibility and pricing rules, rather than whether funds are physically set-aside in separate accounts"[228]. This has some useful angles, but turns out to be much less straightforward.

Clearly, if current push research costs have to come out of current expenditures, it may be easier to finance research from pledges rather than from current budgets. A front-ended R&D mechanism leading to a free, or close to marginal cost, AIDS vaccine would mean maintaining two streams of 'front-loaded' costs in exchange for a stream of zero end-loaded costs. If vaccines replace treatment programs, it might make sense to dovetail the roll-out of vaccine payments with the decline in treatment payments, for a given stream of discounted expected payments. This is a big 'if'. If 'only' a therapeutic HIV vaccine is ever discovered there will continue to be treatment and prevention programs well after the roll-out of the vaccine, with a bulge in costs while the two co-exist. And a less than fully efficacious malaria vaccine will, in the absence of treatment and prevention, leave many, if not most, cases both unprevented and untreated. For example, a 50% efficacious vaccine that achieves less than half coverage in resource-poor settings (for example if it requires too many boosters) leaves up to 75% of potential victims still vulnerable. Worries about resistance might even dent this scenario. It would be unrealistic to base finance on the presumption that one budget simply replaces the other.

But, even if we allow the big 'if', there are a number of counter-arguments:

It's cost in the end – not timing of payments

1) Ultimately, the decision about method of finance should be based on a comparison of present discounted cost – and the cheapest method chosen – and should not be based on the *path* of payments. This is the only really valid criterion. Presuming that the methods work, taxpayers do have to eventually pay an amount equal to *all* research costs (i.e. including all failed research projects, capital costs, and interest on government loans). An expensive, imperfect, APC, with many of the faults listed in sections above and below, is much less favoured than the idealised versions used in calculations. Choosing the APC over another mechanism, on the grounds that it redistributes payments to later and not on the basis of overall cost, is tantamount to the World Bank or the British government borrowing from the future at a very high discount rate.

Other instruments, like the International Financing Facility, IFF, proposal of the British government, that makes use of bond markets, is designed for dealing with two

[226] The author would especially value thoughts on these issues since some of the observations he feels are still slightly speculative.
[227] K2:1
[228] K2:1

streams of 'front-loaded payments' in exchange for a stream of zero end-loaded payments. If the APC turns out more expensive than more open research approaches, then using an IFF-type arrangement with a more open collaborative vaccine initiative would work out ultimately much cheaper than an APC.

A big system needs reserves

2) As Kremer points out, he only considers the partial equilibrium analysis of one vaccine at a time. This ignores important 'general equilibrium' considerations, that might particularly bite if APCs turn out to cost a great deal more than originally claimed.

One of the criteria we set up at the start was that any solution should be judged on its ability to generate a global, and not just a partial, solution to the lack of vaccines and drugs for neglected diseases. This would require a large number of APCs (especially if the side-effect of the presence of a few APCs would be to create difficulties for basic research). APCs are a form of borrowing from the private sector (the private sector finances the flows into pharmaceutical companies in the expectation of getting repaid later). Once a large system of APCs is in place, the *credibility* of the system's ability to repay matters, and private capital costs become a serious issue[229]. Those sinking funds must be assured that the "sponsor of a commitment has sufficient funds to fulfil the commitment"[230]. Crucially, the capital cost component of an APC has to absorb all risks, including also the risks of collapse of the APC.

As more and more APCs are added, the chances of at least one of them being activated – triggering a stream of repayments – rises. In addition, once we move away from Kremer's notion of probability to one with probability of discovery varying over time – and, in particular, with probability rising over the development cycle – the chances of triggering large payments in a multi-APC system rise over time. A financial system has been created with a distribution of risks within it, with instances when the triggering of multiple APCs within a very short time of each other might threaten the whole system of APCs, and yet other instances where the cost of the multiple APCs has become very high with no apparent results.

Default of / threats to the system can be interpreted quite widely to include:
i) Uncertainty about how 'price' and payments might be manipulated if vaccine discovery has happened earlier than expected on some APCs, generating a seemingly excessive stream of payments on them;
ii) Worries about reneging on many of the discretionary elements discussed above, on both current and later APCs if many APCs have been triggered already or if APCs have become costly;
iii) Worries from private capital investors that the program will be wound up early, especially if costs have escalated with little research.

In turn, the capital costs of those investing in APCs is a function of expected default/reneging rates. Firms have to judge how much to invest in research given, in part, the chances of collapse of the mechanism. There is an interesting dynamic problem here, so far unexplored. The higher the APC prices have become, the higher

[229] This reinforces the glaring omission of capital costs in the current APC models. These models will not be able to handle what is about to follow.
[230] K7:5.

the chances of the system 'defaulting'. This leads to even higher risks to private finance and even higher needed APC prices to compensate for this added risk. Pharmaceutical firms hold out from sinking sunk research costs given this worry, but this itself generates an even higher needed APC[231]. All the standard analysis of financial crises and potential self-fulfilling collapse of such a system can be explored.

This leads to a rôle for reserves to keep these private capital costs down and to keep the system stable. There is an optimal level of reserves for every set of APC prices. If insufficient potential reserves are in the system, the above scenario could bite quite nastily. The risk of collapse (the worry that government or backers might pull out) necessitates almost a central bank type rôle for the holder of reserves in preventing collapse.

If the APC mechanism fails in its claims, and APC prices end up being set inefficiently and excessively high anyway (it is argued here that this is a very real danger), a situation may arise where the general system of APCs has generated an excessively high set of future liabilities, with the need for the system to be backed up by high levels of reserves. This consideration alone might rule out the APC as a sensible *general* solution to the global problem of vaccine development.

No such thing as a free lunch

3) Taxpayers "pay nothing if an effective vaccine is not developed". But, they do. Even in the case of research failure with no APC payment, or indeed if the set of APCs collapses, taxpayers *do* pay *all* the costs in their capacity as shareholders and investors in the companies that do vaccine research. There is no such thing as a 'free lunch' – and that includes failed vaccine research, and failed APC programs.

4) A system of APCs that collapsed with loss of private shareholders' wealth, might anyway lead to calls for a bail-out, with recourse to public finances. There are plenty of features of real-world (as opposed to idealised) APCs that suggest areas for litigation. A bail-out might seem easier.

Capital costs of front- and end- loaded expenditure

5) If discount rates for front-loaded expenditure are lower than for end-loaded expenditure, then front-loaded expenditure is favoured. Kremer uses a 4% real discount rate[232]. This is excessively low for private finance. The Tufts study of 'average drug development costs' uses a real cost of capital of 11%, generating a nominal cost of capital of 15%[233]. With all the extra risks of the APC mechanism, this could easily become much higher.

6) There may be slower switch in expenditure flows than Kremer suggests. It is not clear that the development of vaccines does end the first stream of payments at the same rate as the roll out of the vaccine. For example, if a vaccine for AIDS is developed there will still need to be a large and growing programme of treatments in Africa, India, Russia, China, etc. as well as the payment for the stream of vaccines. There is still a 'bulge' in resource requirements, though it is perhaps, spread over a

[231] In the auction mechanism referred to below, price is supposed to rise to encourage higher research intensity, rising until research intensity hits its 'optimal' rate.
[232] K4:17
[233] DiMasi et al, ibid.

longer period. To the extent that policymakers falsely rely on the APC to smooth this bulge, other treatment and research programmes may have to be cut.

7) Repayments of APCs do cause budgetary flow problems at some point. If only a partial APC program is initiated – for example an APC is applied to just one disease like HIV/AIDS – the 'bulge' in financial requirements later will affect programmes for treatments for those diseases not themselves covered by APCs. It will also affect the ability to initiate vaccine programs later should those programs not be based on APC logic.

8) One or two APCs now, even if they prove wasteful – and perhaps even more so if they turn out wasteful – may force the use of APCs later since the resource requirements to support the early APCs reduces the ability to fund front-loaded mechanisms later.

Some argue that it does not matter if one or two expensive APCs are allowed. This is not the case. If APCs turn out to be a more expensive option than claimed (though plenty of evidence would no doubt be produced to justify what, *ex ante*, had seemed an excessive price) they may lead to serious consequences for other vaccines and drug-treatment programmes, and for alternative mechanisms, like open collaborative research, should such mechanisms subsequently be chosen instead.

This is another reason why proponents of APC need to be much more up front with their calculations and re-do them more carefully.

To be optimal they need big changes in other front-ended payments

9) Each APC requires an optimal level of publicly-funded front-ended research to make it work (see the subsection above). Enacting an APC without putting this research into place will lead to long-term inefficiency and higher cost of the APC. The APC is not a panacea for the lack of publicly-financed research on large components of the research process. Efficient APCs no doubt require higher levels of publicly-financed research than is currently taking place, even advance push commitments.

10) The interest costs of co-payments enter as an extra up-front cost. Poor countries relying on co-payments may prefer an F&D Treaty.

11) There is always the chance of being locked into bad contracts. In particular if the legal language of the contract works to stop reneging on what turns out to be a bad contract, excessive resources are needed to support the system, preventing other drugs and treatments form being efficiently used.

10. STRATEGIC INTERACTIONS OF FIRMS

One significant omission from the APC framework used in policy pronouncements is that of the strategic behaviour of firms: "There are no interactions between firms. Preliminary investigations with alternative models suggest that the key qualitative features of comparison with alternative policy instruments would be similar[234]…We chose this structure that eliminates strategic interaction between firms in part for simplicity."[235] Indeed, in the APC models there is no strategic behaviour of *any* firm with respect to *any* variable at *any* time. Ordinarily, this would be regarded as a very strange way to study a topic in industrial organisation.

Kremer's technology modelling device presumes that perfect competition describes typical pharmaceutical research at *all* stage and that all projects are strict substitutes for each other. The probability of a particular firm's research project being awarded the APC is a declining function of the research projects being engaged in by other firms. Overall probability of successfully deriving a vaccine is a positive function of the overall number of projects being pursued. Anything that takes the real world away from the idealised notion of perfect competition will alter these probability functions, and hence the chances and costs of vaccine discovery, and the required APC price. In particular, to produce the lowest possible APC price, the APC mechanism needs to be robust to there every being just a few firms at *any* stage.

In the APC calculations, firms pursue R&D so long as the "R&D costs are less than the probability of success of the project times the value of a successful project discounted at an appropriate discount rate" – which suggests that decisions on the pursuit of R&D can be altered by anything that alters actual or expected 'R&D costs', the 'probability of success', the 'value' of a successful project, and the 'discount rate'. One might add 'expectations' (of all of these) too, since the technological and strategic set-up in the APC models has no serious notion of expectations or asymmetric information. Allowing some of these real-world (and standard industrial organisation) features back in generates quite some room for manoeuvre.

10.1. Strategic Investments

In Kremer, there are no fixed costs, only variable costs. The average and marginal cost of research is constant per unit of time, and is common-knowledge amongst firms

[234] Unfortunately, in rebutting awkward arguments made against the proposal, Kremer has a habit of claiming that the model has been tested against the particular problem at hand, when there is no evidence that this has, or even could have, been done, and claiming that 'the mechanism can be designed to deal with the problem' when there is no evidence it had ever been, or could ever be, thought through. Furthermore, one observes that if the APC price was two or three times higher, the qualitative comparison with the direct government funding of applied research and of joint ventures with private companies of research would not change. If some of the layers of assumed inefficiency in Section 8 were much different, however, the ordering would change. This again reminds us that it is the layers of assumed inefficiency in Section 8 that are driving the result. It would be interesting to see these 'preliminary investigations with alternative models' that 'suggest that the key qualitative features of comparison with alternative policy instruments would be similar.'

[235] K3:8.

(including potential firms), governments, financiers, and all other interested players. Kremer states that: "it would be possible to also assume that projects varied in both cost and probability... Since the algebra would be substantially more cumbersome, however, we made the simplifying assumption that all costs are the same...." Even though in reality cost varies over the development cycle, projects, and firms, he presumes that it can (like the technology) be averaged out as: "the average annual development cost over the full development cycle for the average project". This averaging, advertently, or inadvertently, removes or (to use a word used earlier) 'partials' away strategic behaviour from the underlying probability distributions. When these distributions are then used to derive cost comparison data, this will bias the result in the direction of the lowest possible APC price based on the least possible amount of strategic behaviour – i.e. none. It also helps to generate an instantaneous, unique, APC price, rather than a range (that may vary with time), and a complete lack of a rôle for competition policy or any other standard industrial policy interventions. Highly significantly, from the perspective of the pharmaceutical industry, past sunk costs do not play a rôle in current decisions. Bygones are bygones. Just as importantly, a whole new layer of *risk* has been removed.

In reality, the strategic choices of one firm would alter the costs or opportunities (the probability distributions and the potential value of the market) of other firms. Kremer claims to have considered a toy model where variable costs appeared, but explains that "The strategic aspects introduced in the firm's problem would considerably complicate our analysis. While we did not extend our model to cover the richer scenario, we briefly analyzed a very simplified version, where a number of firms faced the same probability of success in a 'one-shot' model, and found similar basic qualitative results"[236].

This is totally unsatisfactory, especially given that strategic behaviour could *only* ever increase the expected APC price[237], and given that it removes important strategic features that impinge on other parts of the framework, such as the incentive to extract extra payments from developing countries, and the ability to do post-development redistribution of APC payments to target quality. These issues are likely to be very significant in the real world of practical vaccine research and delivery, and can't be simply swept under the carpet because it would 'complicate' the analysis.

The 'commitment' factor

Most vaccine research involves large fixed costs and large 'commitments' of costs (in Kremer each period is a fresh new game; there is *no* notion of commitment) and hence variable average costs. When deciding to commit or not to commit resources to a project, a firm would have to take into consideration the chances of other firms discovering the vaccine first, in order to work out its own likely average costs (based on the spreading of its own fixed cost). This is a function of (amongst other things) the investments/sunk costs of other firms. For example, if firm A assesses that other firms have invested heavily in fixed research factors (for example, production plants to produce significant quantities of trial vaccine candidates) this means that, given the increased probability of the other firms gaining the vaccine sooner, firm A's own

[236] Ditto the last footnote.

[237] Given the above admittance, it is then rather odd for DALY calculations of the value of APCs to be made to a precision of a cent per DALY and for no suggested upwards adjustment to be made to allow for this strategic behaviour. Kremer's figures again are the lower envelope of the possible prices.

investment is likely to spread over fewer periods, raising firm A's *expected* average cost. This is a particularly significant issue in an industry with a large build up of sunk costs and no costs repaid until an 'APC' is 'won'. In the modelling framework used by Kremer it is the 'continuation game' that matters, but it is always the same continuation game. In a situation where there are multiple sunk costs at any point in time, these sunk costs (and the knowledge and IPR stockpile built up as a result of them) can be used to raise the probability of discovery in the continuation game for some whilst reducing it for others, and to reduce (marginal) costs in the continuation game for some whilst increasing it for others. The 'continuation game' is always up for strategic manipulation.

The problem of research leads

If vaccines are best discovered through multiple routes of investigation, the use of fixed costs as strategic entry-exclusion devices has the potential knock-on effect of reducing the number of drug leads being followed, thus slowing the vaccine discovery process. For example, imagine what happens if it looks as if just one or two players have leads that will result in a vaccine? For other firms the risk has risen that their projects will not be the ones to gain the APC[238]. They cut back on their investments (which are now more risky and hence entail higher capital costs, with these costs higher the larger are the typical sunk costs in the industry). This reduces the number of players in the 'game' still further. This increases the original firm's own chances of getting a good APC-generated price even for a lower quality vaccine. Players with research leads that seem to be going somewhere have obvious incentives to make strategic investments/sunk costs and strategic decisions over patents in order to deter others from following leads, and to strategically manipulate information to exaggerate the chances they have. Truthfulness, like everything else, becomes strategic. Importantly, the socially optimal and the privately optimal behaviour of the 'other' firms deviates because of the behaviour of the one or two with the so far successful leads, leading to fewer overall leads being followed compared to the socially optimal level.

We remember also that the investment may not lead to the vaccine, but rather lead to 'something' *en route* to a vaccine that will have to be sold to the next layer in the chain, so the argument applies equally to situations where the vaccine is not yet discovered. In its individual profit maximisation condition this may put an individual firm off making important intermediate investments.

Research intensity is non-linear

There are layers of other possibilities. For example, as a firm tries to increase research intensity for a given fixed factor, at first the costs of research decline but eventually they increase. In Kremer the relationship between intensity and cost is linear, perhaps implying an ability to alter fixed factors continuously and costlessly. The scientific literature (also see the diagram in Section 4) on the development of vaccines, however, suggests that at certain stages a large fixed factor, such as a production plant, has to be put in place before being able to go on to the next stage (we also know from the literature on trial sizes and costs that 'cost of discovery' varies greatly with the trial phase). These large discrete sunk costs introduce highly non-linear segments to research cost curves, sometimes leading to fewer eventual players in the market.

[238] Simplifying to the case where there is just one winner.

They also suggest a large element of option-price thinking in investment decision-making (thinking that is also heavily influenced by what other firms have done in terms of their investments).

Complementary, rather than substitute, research projects

It may also be the case that research projects are not as strictly substitutable for each other as presumed in the APC calculations. It may help to think of this through the language of product differentiation. With strictly substitutable projects, firms choose research leads to differentiate themselves form each other as much as possible. If they choose leads too close to those of other firms then they risk too low a share of the eventual market (remember the way the APC redistributes). But if technology *is* more complementary, and particularly if it is more so in certain parts of the technology space than in others, then society benefits from projects being more similar. This raises two problems. Firstly, it increases the chances of collusion (to break this requires forcing firms to differentiate projects more, but at the cost of the lost benefits of complementarity). Secondly, it suggests that APC policies on sharing markets after vaccine development should adapt. If the rules are based on strictly substitutable technology, they deter complementary research. Thinking of the probability 'space' underlying the model, strictly substitutable technology would lead to post-APC rules that are in a sense equal in their effect everywhere on the space, whereas complementary technology requires the rules to allow 'regions of more favourable treatment' or regions where groups of firms will, in a sense, be treated better on average than groups of firms would on other parts of the technology space. Failure to get this right leads to inefficiency. The assumption of strictly substitutable research is not only a major simplification, but also rules out the losses that would come from a misapplication of post-APC distribution. This leads to a complicated trade-off of the need for similar leads and more generous treatment post-APC development of firms following similar leads, against the dangers of them colluding[239].

Financial 'deep pockets' of large players

There may also be a problem with large players (perhaps large pharmaceutical firms with 'deep pockets' financially) who may have the ability to make strategic choices over fixed investments (and other factors) that are not available to smaller (possibly more innovative) firms. This effect is strengthened by the presence of an APC, with its emphasis on 'deep pockets' finance. The effect on the population distribution of firms between highly innovative and less innovative is not clear. In the Kremer model there is no distinction. This should not be presumed. Less innovate firms can come to crowd out more innovative firms in more realistic strategic settings.

This strategic interaction would be further complicated by the technological consideration described in Sections 5 and 6 above and by the many layers of strategic behaviour that are already taking place *in response to the APC itself*. And the effect is even stronger if there are common costs across vaccines, since one firm might develop an advantage over another in one vaccine market that gives strategic advantage in another market.

[239] Observe also that if the underlying technology space is not known to the APC 'regulators' then they will face a signal extraction problem in trying to detect collusion. Those groups of firms on the technology space that is more complementary look the same as those firms colluding on parts of the space where technology is more substitutable.

We cannot presume that at *all* stages in the research process, technology (and finance – see more on this below) is such that it generates the competitive structure Kremer claims and relies upon to derive his APC cost figures. Indeed he seems to favour an industrial structure with a core rôle for a few large firms. Since none of the above is modelled in the Kremer APC framework, we do not know what the consequences for the efficiency of the APC will be.

Problems with generic follow-on drugs 'after ten years'

Fixed and sunk costs also spoil somewhat the idea that after ten years, competing vaccines are likely to emerge[240]. Some experts have commented privately to the author that given the large discrete sunk cost element of production, and given the importance of 'know-how', there may be too little of a market at too low a price to justify entrants, and the incumbent may have too much of an advantage and yet little incentive to manufacture[241]. It is apparently already the case that some vaccines are under-produced and under-used because it is simply not profitable to produce them after a point (the large firm is no longer interested, but entrants, aware of the sunk cost needed to manufacture and their missing 'know-how', cannot enter anyway[242]). In such a setting it might be more valuable to create a generic market earlier, especially if that enables 'know-how' to be spread more widely more early in the process.

10.2. Strategic Use of Patents

"Many companies have developed large portfolios of patents and other forms of intellectual property but use only a small portion of these intangible assets in their core products or services. The remaining assets effectively sit on the shelf, yet some of them can provide enormous economic [and social] benefits." *Mercer Management Consultancy*[243][bracketed portion added]

10.2.1. Hundreds of patents and highly cumulative research

We realised above that the APC is a very high patent system. To get sensible cost comparison data, we should incorporate the strategic uses of patents at all stages – including at intermediate stages, and any deleterious consequences this has for research – into calculations of the cost of using APCs. As we have seen, the APC methodology has ruled all such problems out *a priori*. Here we give a basic overview.

---

[240] K4:50.

[241] Some counter evidence was presented at the World Bank and Centre for Global Development suggesting that generic manufacturers in China, India and elsewhere would be only too happy to follow on with generic copies of vaccines, and that part of the current problem is the lack of reasonable prices for vaccines once off patent.

[242] Comments made at MSF Malaysia meeting, Feb 2004. One supposes that this relates to the fact that with the incumbent in place satisfying some, but not much, demand, and given the cost disadvantage of entrants, the 'residual demand curve' left to entrants is too low given their now much higher supply curve (higher because of their higher costs).

[243] http://www.mercermc.com/defaultFlash.asp?section=Perspectives. It should be acknowledged that property rights to information can lead to trade in IPR that is socially beneficial, since, at least in theory, it enables the spread of general purpose technology over many industries. The Mercer quote seems to indicate problems with this however. Why would firms sit on their intangible assets if a 'knowledge market' allowed them to efficiently trade in such assets? No doubt the issue is a lot more complicated than it at first appears, with the practical 'knowledge market' differing in important ways from the theoretical 'knowledge market'. And, clearly, at times secrecy and strategic considerations must overpower the desire to sell on the 'knowledge market'.

There are many ways to use patents in strategic ways, to increase risks to other players and to reduce the research intensity of other players. Added to the usual social costs flowing from the creation of temporary monopolies caused by patents, we have two further problems that are increasingly numbing the ability of researchers into vaccines (and drugs generally), that inevitably will feed through to higher APC prices:

i) Patents created in technological fields where a single product 'reads on' hundreds of other patents;

ii) Patents in fields where innovation is not an isolated event, but a highly cumulative and complementary process (Scotchmer, 1999, 2004; Shapiro, 2001). Innovation does not just give benefits to current innovators and consumers but to future innovators and consumers. Granting patents in such cumulative research such that each patent-holder has the right to exclude another, leads to a breakdown in reaching mutually beneficial agreements to share technology if transaction costs are high enough. As Scotchmer points out, the innovator is both buyer and seller of IP, and what they gain on the latter cannot be presumed to match what they lose on the former. In some areas of cumulative innovation, such as biotechnology, innovations may be especially complementary, such that the overall probability of a particular result being achieved within a given time period increases more than in line with the number of research lines adopted by potential innovators. Plenty of studies have found the way that tight property rights stifle research in such areas, with researchers increasingly facing infringement actions for using patented materials, processes or research tools, and the very possibility of this stifling them from action in the first place[244]. Strengthening IPR in such areas (say to enforce the workings of an APC) slows rather than speeds technological progress, and vaccine discovery. It would not make much sense to adopt a device to stimulate investment into vaccine research that requires tight IPR to work, if those carrying out the research were only then forced to face the consequences of that tighter IPR in the research process itself. It would also make the device ultimately a great deal more expensive.

The importance of the intermediate steps

This is a radically different view of the world from that underlying the technology device of Kremer – which strips it out *entirely*. In Kremer, the technology is such that *only* the end discovery/product matters anyway. Since there are no intermediate discoveries, there is no notion of cumulative technology or of the disincentive effects of patents (the other interpretation is that Kremer simply presumes open source for all intermediate discoveries, but that does not seem that likely!). Since research projects essentially refer to the end products paid for in the APC, they are *necessarily* competing and non-complementary. While it may make sense to model the end products of research this way, it makes no sense at all to (implicitly) model all points of the research process leading to the end products in this manner. In Kremer, the overall probability rises when there are more projects, but the probability of any one project falls if another project enters, just as would be expected with competing products. But that is not a good way to model the science of HIV, malaria, or TB vaccines.

[244] Solsted, J. (2004); Heller M.A. and Eisenberg R.S. (1998); Bunk S (1999); Freundlich, N. (1998).

Instead of ignoring these issues – sweeping them away by technological slight of hand – we need to have an open and frank debate on what incentives are actually needed to achieve the optimal amount of 'pioneering research' and to promote the optimal amount of effort to build on that initial research. A model that presumes these questions away from the very start, can obviously provide no answers.

10.2.2. The problems with patents

There is a large body of literature on the problems of patents, summarised (rather roughly) below, with the emphasis on issues that will lead to the APC price being higher than in the Kremer calculations. For now, some of the more positive aspects of patents are not discussed, not because they are not important, but because the purpose here is the rather more limited one of rebalancing the debate in the context of the Kremer figures, and the deliberate decision to remove IPR issues from the method of their calculation. Besides, this is the whole point of the assertion that we need a full and frank public debate about these issues, and about alternative frameworks, such as open collaborative research, that might help to get around some of these problems:

• Evergreening – extending monopoly/marketing dominance (if not monopoly) beyond an original patent. In pharmaceuticals, a lot of research is conducted to extend the life of a product by resetting the 'patent clock'.

• Bargaining games involving litigation costs. Firms with higher litigation costs are less likely to patent in areas where other firms have lower litigation costs (Lerner, 1995).

• Patents to increase bargaining positions in cross-licensing deals (Granstrand 1999, Rivette and Klein, 2000).

• Patent to create 'zones of exclusion' around inventions, so that others cannot exploit their own patents.

• 'Patent thickets' (Shapiro, 2001) 'excessive compounding' (Gangi, 1999) of property rights (also 'clustering', 'bracketing', etc.). The overall strategy is to obtain patents similar to each other in scope so as to make it harder for potential entrants to gain 'know-how' or to acquire the overlapping patents they need before they can introduce new products or processes. 'Know-how' is especially important in late stage vaccine research and in ensuring rapid generic competition. This generates an anticompetitive barrier to entry and raises the costs of entrants and prices to consumers.

• 'Patent portfolios' as a form of currency, especially in biotechs. The aim is to achieve as wide a scope as possible – numbers matter more than quality. The scope tries to anticipate future scientific developments. 'Patent portfolio wars' result from the desire to avoid the breakdown mentioned above in highly cumulative areas of research.

• Patent portfolios as 'bargaining chips' – acquired to negotiate access to important external technologies. As such they are a way to get around hold-up problems in investment caused by important patents being held by outside firms. A study by Hall and Zeidonis (2001) investigating the doubling of patenting in the semiconductors

industry – following changes in patent legislation in 1982 and the creation of the Court of Appeals for the Federal Circuit – found that this was largely due to the fact that inventions in the semiconductors industry use technology covered by hundreds of patents held by numerous firms. Firms increasingly faced litigation and preliminary injunctions and risked 'hold-up' if they did not have cross-licensing agreements in place. Large patent portfolios significantly strengthened firms in negotiating access to technologies developed elsewhere, and added credibility to any threats they made that other firms would be sued for infringement. Many firms were simply engaging in defensive drives to gather as many patents as possible as quickly as possible. Hall and Zeidonis refer to the 'patent paradox' of an industry with an increasing propensity to patent, yet, relying ever more on secrecy, lead time and superior manufacturing and design capabilities rather than the patents as protective devices. Hall comments "This had little to do with encouraging innovation, and in fact looked like a tax on innovative activity."[245]

• Since the APC still maintains patents as the basis for staking a claim, the need to 'be first' can lead to wasteful patent 'portfolio races'[246] and the need to amass as many patents as possible for strategic reasons. This behaviour is especially inefficient if there are large complementarities in information (for example in the search for a vaccine).

• Use of 'Selective patenting' – to control the channels of distribution. Once a country is a potential threat to the manufacture and supply of drugs to neighbouring countries, patents are imposed in that country even if not on the countries nearby. Incidentally, this counters the claim sometimes made that low levels of patents in Sub-Saharan African countries, other than South Africa, proves that patents are not a bar to accessing drugs. South Africa is the main route into Sub-Saharan Africa. As a game – it is perfectly consistent to have low levels of patents in some countries and yet no domestic manufacture or availability. Just the expectation that 'selective patenting' will be the reaction in response to any moves towards domestic manufacture or availability, deters investment *ex ante*.

In conclusion, the strategic use of patents raises the cost of doing R&D. It especially raises costs of doing complicated research on 'difficult' projects like, for example, vaccines, where there is great uncertainty at some stages of development and where sharing of information and synergy of ideas is highly important, and where capital costs are already high. Hence the cost of new drugs is raised, even as the breadth of coverage is reduced.

How this relates to the APC model

All this would play out against the background of an APC. For example, under an APC with non-stationary technology and tight patents (to hold all the stages together) firms – instead of selling intermediate information as they 'should' according to the APC model[247] – may build up patent pools to strengthen their hand at later stages in

[245] Hall, B.H., 2002, p 7.

[246] Hirschleifer, J, and Riley, J.G. (1979); Dasgupta, P, and Stiglitz, J. (1980a and 1980b); Tandon P (1983)

[247] As we have seen, in the Kremer model there is no reasoning really going on about what firms *should* do at intermediate stages, since there are no such stages. This paragraph, instead, refers to the APC model that must be required to make the reduced-form Kremer model work.

order to suppress competition, and also in the hope of keeping their expected capital costs down. The benefit of selling one's discovery at an early stage when the probability of vaccine discovery is low (even zero) may be lower than the benefit of keeping it till later. In this case there is a positive externality – onto the probabilities of this particular firm discovering the vaccine on later parts of the state space, from strategic actions they take with respect to IPR at earlier parts of the state space.

Above, we saw how the non-stationary technology model of an APC would be forced to work either via secrecy or via patents. Clearly, this presumed the non-strategic use of patents. As soon as the strategic use of patents is incorporated, we can show that while the individual probabilities of firms can, on *ceteris paribus* assumptions, be increased by their strategic use of patents, nevertheless the externality effect that this generates on *all other firms* is in the direction of reducing probabilities in the *aggregate*, raising costs in general equilibrium, and slowing *aggregate* vaccine discovery for any given APC price.

The purpose of this section was not to suggest that patents do not have benefits that might also serve to increase R&D, but to suggest that there is an important rôle, in promoting healthcare research, for sectors that can more easily share information. Adopting an APC system that deliberately leads to more secrecy, and that biases its benefits towards large pharmaceutical firms that rely on tighter IPR, should not be done lightly without first exploring any negative consequences. The Kremer model ignores the issue entirely, and feeds this ignorance through into all its cost comparison figures.

10.3. Strategic Manipulation of Information

Entwined with the strategic issues discussed above is the notion that firms will not always be truthful with information – even to financial markets[248] Firms have incentives to try to signal their chances of vaccine discovery to give themselves a competitive advantage, and incentives to be secret about information that increases competitors' advantage. It may be a dominant strategy to signal an exaggerated notion of one's chances of vaccine discovery if it causes others to ease off in their efforts to discover a vaccine or if it makes one's access to finance easier.

Kremer is extraordinarily naïve about the degree of transparency of information provided by large pharmaceutical firms, transparently that itsutterly necessary if purchase decisions are to be efficient.

10.4. All Strategic Behaviour Feeds a Higher APC Price

The fact that the APC price can exist in a large range – bounded below by the need to create incentives for R&D but above by the social surplus generated by a vaccine – means that there is plenty of room for rent-seeking economic activity. As in all situations of rent-seeking, firms will be prepared to 'waste' resources seeking to extract this social surplus (in this case via sunk costs, patent inefficiencies, and the manipulation of the evidence base). 'Waste' is incurred up to the point where the marginal cost of a bit more 'waste' is exactly equal to the marginal benefit of a bit

[248] In Section 12, we will see how this creates problems for that part of the Kremer model too.

more ‘waste’. Nothing in the Kremer model explains how this would be avoided, and why firms would stick close to the lower boundary when the upper boundary may be a great deal higher. Kremer takes the unjustified short cut of presuming perfect competition at all times and at all stages to rid the problem from the model.

Another way to think of this is that the APC price has to rise to cover all ‘waste’ in equilibrium. Otherwise, the *ex ante* risk to players that rent-seeking and strategic behaviour will *not* be compensated for in a higher APC price will slow vaccine discovery. An auction is a potentially ideal mechanism for encouraging this wasteful rent-seeking, especially if there are no limits on how far the price can rise[249]. This will be explored more in the following section on auctions, since the auction is a further stage in which strategic manipulation can be played out.

It would be worth exploring whether an open collaborative research framework is more capable of creating a more competitive industrial structure with more research-active firms and lower incentives towards the strategic use of investments, patents and information, than a framework based on APCs that instead tends to narrow the number of players down for any given level of global public spending on vaccine research, and tends therefore to feed such strategic behaviour. This too could usefully be factored into the APC cost comparison methodology.

[249] This might require yet more adjustment to the auction ‘price-rising’ rule. See Section 11.

11. THE APC AUCTION MECHANISM

11.1. Introduction – The Problems of Chicken and Egg Thinking

What if the technology (i.e. the relevant probability distributions, in the Kremer model) is not known so that the policy maker cannot calculate the exact, optimal, APC price? As Kremer puts it[250]: "There is no single answer to the question of how large a market is needed to spur research."

Yet some mechanism is needed to avoid either paying much more than necessary or offering too little to stimulate research in the first place. The solution originally suggested is the device of an auction – a low-valued contract/deal that is raised until it gets 'sufficient' takers[251]. An auction is a potentially useful device where there is asymmetric information. Vaccine developers (in economic parlance the 'agents') have all the information (cost structures, knowledge of science and of future prospects, etc.). The organiser running the mechanism (in economic parlance, the 'principal') does not have this information. An auction mechanism extracts the information and, in this case, determines the APC terms – or so we are told. Strangely, although Kremer devoted a great deal of attention to auctions initially – suggesting real worries that it would be difficult to set the terms of early-stage APCs remotely efficiently – all thoughts of auctions have been abandoned recently.

Kremer[252] suggests starting with a modest program, "not too expensive", with "an option to increase the value of the program if the original program proved too small to stimulate sufficient research." As long as the vaccine price "is not expected to increase too quickly" firms will not hold off a vaccine from the market in hope of getting a better price. "One way to avoid either paying more than necessary for a vaccine or offering too little to stimulate research would be to offer a relatively modest price initially, and if this price proved insufficient, to raise the promised price gradually *until it proves sufficient to spur vaccine development*."[253] (emphasis added). Given the argument made by Kremer that the main reason for the APC mechanism's superiority over all other mechanisms is that it avoids policymakers having to make decisions about the underlying science (though we have already found this not to be the case), we should hope that the suggested 'auction' mechanism will work efficiently and not entail any informational requirements on the part of policymakers.

Unlike an auction run for normal procurement purposes, private firms under an APC do not sign contracts agreeing to provide a particular level of R&D in the same way a firm might sign contracts to supply a particular number of tanks. Their intensity of R&D is a totally private, un-contracted, choice (indeed, information on it is usually

[250] K7:31
[251] The auction idea is presented in early cases of the APC, though it is less clear how strong the commitment to it is given that all the APC calculations are based on industry claims and that little has been revealed about how the auction would actually work. Recent (March 2005) evidence is that worries about getting the size and terms right are far from on the minds of those most promoting early-stage APCs.
[252] K7:30
[253] K7:4

confidential), and APC contracts are only signed by the winning developers who produce the vaccines.

Which information comes first?

Supposedly the APC price would have to rise if it had so far "proved inadequate to spur sufficient research"[254]. But it is hard to visualise how an auction could work out the moment when optimal, or 'sufficient', *intensity* of research had been achieved given that the only information we have to judge this by is the discovery of the vaccine itself, i.e. the moment "*it proves sufficient to spur vaccine development*". Optimally setting the former (intensity) in order to achieve the latter (a vaccine), is impossible if it requires information provided by the latter that only comes out on average, a very long time *after* the former. It's a chicken and egg situation. R&D intensity cannot be conditioned on the information provided by development of the vaccine itself. Neither is it clear that the optimality of the intensity of R&D going on *could* be judged without knowledge of the latter. This is only made worse by all the hidden information regarding R&D expenditure. Even then, for statistical reasons we would never be able to judge, even after the development of the vaccine, if intensity had indeed been optimal; the result of any APC (the timing of vaccine discovery) is just one possible point on the distribution of times to discovery, so we would need the results many APCs on the same vaccine to derive the distribution over 'times to discovery' to work out whether this particular APC had been set 'efficiently' in any statistical sense. Unfortunately, the auction would *need* to work out an exact, 'optimal', amount of R&D – neither too much nor too little, for the APC to do better than alternatives. It has not yet been explained how this optimal amount will be calculated without a great deal of *ex ante* knowledge of technology.

There is a novel, and somewhat contradictory, attitude to auctions in Kremer. On the one hand there is the assertion that an auction mechanism removes the need for policymakers to have much information. On the other hand there is a repeated reliance on detailed scientific information for a whole set of institutions and mechanisms involved in the running of other parts of the APC mechanism. Perhaps it is being assumed that the revelation of information in the auction is used to derive information for use in other parts of the APC, but no explanation is given of how the linkage is made. On one hand, the figures used to compare alternatives to the APC are based very heavily on the hopelessness of non-APC institutions in deriving and using information. On the other hand, there is a heavy reliance on the information coming from large pharmaceutical firms to work out the APC price.

An auction is no panacea, and bland assertions (like "this mimics an auction, which are typically efficient procurement mechanism in situations in which production costs are unknown"[255]) do not help. Auction mechanisms can go, and have many times gone, badly wrong in a wide variety of circumstances if they are not carefully designed[256] or if (as in this case) the situation is not appropriate to the running of an auction. *Many* countries got the sale of 3G, and other, licences very badly wrong, with the price going too low and in effect giving huge amounts of tax-payer revenues to the

[254] K7:Abstract
[255] K7:8
[256] As a salutary reminder of how badly wrong things can go, the reader is encouraged to read the multiple recent examples of auction 'disasters' to be found in Varian, H. R (2002). The reader should also see Klemperer, P. (2002).

companies who were bidding (compared, say, to the UK where huge revenues were raised for the UK Treasury from the sale). In the case of APC auctions, the analogy would be huge revenue losses for governments and foundations caused by the APC price going too high, and/or serious vaccine quality losses, and/or delays, or combinations of all three.

11.2. Is an Auction a Suitable Mechanism in this Case?

11.2.1. Setting the growth rate of the APC price

The obvious worry is that firms might hold back on research effort as the APC price rises, in the knowledge that the price for the vaccine will be even higher later. Another way to think of this is that increasing vaccine research 'too soon' leads to a price too near to the bottom of the range of possible APC prices, with a large amount of social surplus going to those countries buying vaccines and less expected discounted profit to the developers. Holding back is more profitable. In such cases, reducing the growth rate of the APC price, or even cutting the price, might, paradoxically, increase the rate of vaccine research.

Kremer argues that so long as the growth rate of the price is no greater than the growth rate of the interest rate (under certain technological assumptions) then this situation will not arise. However, this is not the correct condition when the market structure is less than competitive and strategic interactions of various sorts are possible. Then, the correct condition is that the growth rate of *potential profits* is not greater than the growth rate of the interest rate. If competition is low, as this author for one contends would be the more likely case (Kremer seems to concede this at times, given his emphasis on encouraging one or two large pharmaceutical firms), this makes this problem worse, with incentives to slow vaccine research.

11.2.2. Collusion, Suppressed R&D intensity, and contradictions with the technology of vaccine R&D

A very basic conclusion of auction theory is that to prevent collusion against an auction mechanism (which in this case may simply mean tacitly adopting less intensive research strategies), there need to be 'enough firms' to maintain competition. Kremer presumes perfect competition at *all* stages of vaccine development with "many symmetric pharmaceutical firms"[257]. Delay would not be severe "if many firms can potentially compete to develop a vaccine"[258], thus generating "a price very close to the cost of its development."[259]

Under a more realistic model of the technology of vaccine development, however, it may be perfectly reasonable (in fact the most profitable course of action) for vaccine developers to, at particular stages, hold back on their R&D intensity if lower intensity generates a higher APC price. This may hold, for example, if the technology eventually involves (or possibly converges at some point during the auction) a distribution of projects with a few very high probability projects, and a mass of projects at low levels of probability (i.e. some probability distributions over projects

[257] K7:48.
[258] K7:38
[259] K7:48.

are more 'stretched' than Kremer visualises). Know-how[260] also becomes hugely important at late stages of vaccine development, biasing the advantage towards the few incumbents. If by later stages of development there are just a few firms working on the clear favourites for a vaccine and the APC price has not yet been settled, the incentive is to suppress the intensity of research (helped by these firms' holdings of patents on processes and intermediate products). This is increasingly the more profitable strategy, the more the marginal cost of research rises with the intensity of research (this is another reason why it is wrong to presume that the average cost of research is constant). This leads to later vaccines but at higher prices. In addition, *ex ante* knowledge of this possibility will cause strategic behaviour in earlier stages in order to reduce the number of players in the later stages (including holding back on the release of technologies and access to patents) in order to achieve these higher prices (and slower vaccine development).

Multiple research leads all the way?

Not only does the vision of a rising APC price, working its way towards an 'optimal' level, require multiple paths to vaccine research, but it requires that those multiple paths *always* exist on the way to vaccine development. This is an unrealistic description of the technology of vaccine research. There *will* be stages where it is necessary to have many research leads being followed to ensure that what *ex post* turn out to be very good leads are followed (indeed we saw the struggle needed to create rules for the redistribution of the APC after vaccine development, so as not to deter the optimal number of leads from being followed earlier in the process). But it is not obvious that having multiple leads to the bitter end, as *would be needed* to prevent collusion against the auction, would be efficient. At the same time it is important that the expected strategic behaviour on these few last-remaining leads does not, in backwards-induction fashion, feed earlier strategic behaviour that reduces the number of earlier leads being followed.

Pricing of patents in a world with collusion

Similarly, it is not clear how those on 'non-collusive' stretches of the development process should behave. For example, in order that sellers of patented ideas on earlier parts of the technology process do *not* have the incentive to price their patents 'too high', they should be convinced that for all stages of the mechanism there will not be collusion. Otherwise they should either logically hold on to their ideas, or they should incorporate this later collusion into the price-setting of any patent they do sell (to any use, not just to uses within this vaccine process). One can visualise low numbers at one stage of the process extracting collusion profits from the rest of the system. For example, it is a standard notion that a monopolist at one level in a vertical chain can create monopoly prices that feed through to prices at other levels (of course this extracts a higher APC price but also slows vaccine research *ceteris paribus*).

The large numbers of competitive 'bidders' needed to prevent 'collusion' against the mechanism may anyway conflict with the number needed to get optimal costs of vaccine development. Having a large numbers of firms risks duplication of projects and greater collective fixed costs (in Kremer this is ruled out). In fact, there are cases where, even without collusion, "if there are few firms, it is possible to construct

[260] I thank a number of experts for this observation, in particular Christopher Garrison and Mary Moran of MSF, London.

examples in which expected time until a vaccine is produced increases with the growth rate of p."[261] This generates a trade-off between the cost-efficiency of having fewer firms versus the inefficiency caused by the ability these fewer firms have to collude against the auction mechanism.

If another option is chosen – "to pre-announce that if no vaccine had been developed by a certain date, the price would start to grow automatically"[262], this is less likely to succeed with fewer players, since the marginal cost of early research is now effectively higher and, indeed, as the automatic price rise period is approaching the marginal cost of current research is rising (the automatic price rise period acts as an extra discount factor on early research costs). If investment has an option price element, even if firms do not try to actively collude, firms may still hold off investing to see how valuable it is in the light of the actions of others (collectively by holding off, they eventually reach the automatic price rise period[263]).

These affects are even worse if there are common costs across vaccines, since when one firm develops an advantage over another in one vaccine market that gives it strategic advantage in another market, it becomes more difficult to ensure that there will be enough players in each APC in the other markets to ensure that competition generates the optimal APC price in those markets.

11.2.3. What if the APC concentrates on the last stages of development only?

We found earlier that the APC is not really designed for dealing with technologically sophisticated science, and may only therefore concentrate on the last stages of development (leaving the basic science part to other public funding). Paradoxically, the more the APC auction mechanism concentrates on the latter parts of the technological process the more likely it will suffer from collusion and the more likely non-auction mechanisms would be used for setting the terms.

11.3. Strategic Interactions to Drive the APC Price Higher and the Rôle of Competition Authorities

There is also a complex connection between the strategic problems discussed in sections above and the manipulation of the mechanism for setting APC terms. For example, if strategically chosen fixed investments deter entry of other firms, it may also lead to there being too few firms to make the auction mechanism work to achieve a low APC price. Knowing this *ex ante*, there is an extra incentive to engage in such strategic investments in the first place. Similar feedbacks onto the auction mechanism happen in the case of strategic patent and information distortions. In all cases, using strategic choices to reduce the number of other firms creates further rounds of ability to delay and force a higher APC price via the auction part of the process. Observe that strong patent links in a completely non-distorted model (including not distorted by patents) should help prevent firms from holding off in early stages to take advantage

[261] K7:49.
[262] K7:38.
[263] Though there must be some offsetting option price thinking as the period approaches. What happens all depends on how firms interact, and whether they collude to slow research before the price rise period hits in.

of the later more profitable phases. But, as soon as patents themselves become the source of distortions, this function is wrecked.

An important objective of any program to ensure efficient vaccine development would be to ensure competition. This suggests that concomitant with APCs there might need to be a rôle for competition authorities. But it is not clear if 'interference' in this way would be tolerated or possible in a world governed by APCs. First, there are the multiple other layers of APC-generated committees and institutions that the competition authority would have to interrelate with. Second, given the central importance of financial instruments, such as stock options, in the workings of APCs, shareholders may mount legal challenges – based on the 'contractual terms' of the APC – against competition authorities. Third, the competition authorities would be facing large, highly profitable, politically influential, pharmaceutical companies (as envisaged both by Kremer but also in the workings of the APC model[264]). Fourth, there would need to be coordination globally across competition authorities.

11.4. Large Pharmaceutical Firms and the APC Auction

Kremer says it is important that incentives are strong enough to "induce major pharmaceutical firms to pursue several potential leads simultaneously." He argues that small market-enhancing measures might bring in small biotechs, but will not bring in large pharmaceutical companies, and that "a large incentive might bring in a single major pharmaceutical firm, a still larger incentive would bring in more."[265] There is a conflict between Kremer's obsession with the need to get large pharmaceutical firms involved and the need to keep the industry competitive for the sake of the price-setting auction. None of the actual mechanics of the way price would be set with just a few large pharmaceutical firms has yet been demonstrated. The previous section showed the dangers that it would lead to delay, a higher APC price, and greater cost to taxpayers. We know[266] that biotech companies respond more strongly than large pharmaceutical companies to the orphan drug incentive package. An expensive APC that is designed to appeal to large pharmaceutical companies may be inferior to a cheaper alternative (such as a more open collaborative framework perhaps) that is targeted more closely at highly innovative biotech companies. Mercer Management Consultants conclude that: "The perspective of firms on the potential developing country markets were also divided. Primarily smaller concerns such as biotechs with limited experience supplying today's vaccines, assumed a substantial market and adequate funding for the vaccine would exist," while larger pharmaceutical firms were more negative. The tipping of the balance in the direction of large pharmaceutical companies is further reinforced by the financial market implications of APCs, as described in section 12.

Kremer hints at the inefficiencies of relying on large pharmaceutical companies when he points out that many larger firms won't bother to change "corporate strategy" for the sake of small markets. This begs the question of why the APC should have to make up in price for this inefficiency in 'corporate strategy' or culture, especially if smaller biotechs are more willing to do the research.

---

[264] See more below, especially in the section on finance.
[265] K1:9.
[266] For example, see Kettler (2000).

11.5. The $336m-$586m Per Year 'Needed Market Size'

We know that the true figure of 'needed market size' has still to be worked out, since the model that could possibly set it is not in the public domain. In early APC papers, the auction notion was mentioned though no workings were given. In later papers the auction idea has been downplayed in favour of figures derived from large pharmaceutical firms themselves.

The $250-$500m/year for ten years (plus average $86m/year co-payments from developing countries for ten years[267]) is derived after Kremer and others "discussed these issues with market participants and reviewed estimates by others who use this approach," and "reviewed information on the actual sales revenue for drugs and vaccines."[268] The review of the "relevant literature" included the analysis of Mercer Management Consultancy[269] and Blanc (1999), which itself is based on similar surveys of the industry. As pointed out earlier, the methodology of basing figures on the previous levels of sales, runs the danger that the greater are the problems caused by tight IPR the greater is the 'needed market size' and the APC price[270]. We also know that heavily-quoted parts of the literature are based on controversial methodology with little of the underlying data placed in the public domain.[271]

A "rough rule of thumb"

The best that Kremer can muster in favour of the $250m-$500m figure is that it is "a rough rule of thumb in the industry."[272] He points out, similarly vaguely, that "A ten-year purchase commitment would likely be sufficient to motivate research"[273] though he gives no evidence to back this up. The above Mercer study finds that "With only one exception, none of the companies interviewed were able to give a detailed description on the potential market for an HIV preventive vaccine. Views on the market potential were vague and divided". Given that the APC is supposed to obey the condition of 'additionality' – that is add to the initial market size – it is hard to see how it could have been calculated by use of industry data if industry cannot even work out the market-size that an APC needs to be additional to.

These industry figures, as Kremer (mostly) admits, also suffer from distortions:

1) They tend to overstate the needed market since firms think that this is part of a negotiation process;

2) There are political constraints to such sensitive issues that lead to high claims being made as to the 'needed market size'. For example, executives may be reluctant to

---

[267] The $86 is an average across three diseases. Total co-payments over ten years are $128m for Malaria, $251m for tuberculosis, and $442m for AIDS (K4:14).
[268] K1:5.
[269] Mercer Management Consultancy, ibid. See, also, Whitehead, P (1999).
[270] In fact, this is worked into the APC figures by presuming that no vaccines are purchased by low-income countries in the first 10 years after the release of a new vaccine.
[271] DiMasi, et. al. (2003).
[272] Also K7:30. No mention is made of the fact that the types of drugs and market conditions (i.e. me-too drugs, the kinds of technology underlying drugs, marketing costs, etc.) typical in the industry are different from that required for production of vaccines under an APC.
[273] K4:10, K4:50.

explain their lack of interest in vaccines on the basis of the lack of profitability, and find it politically more expedient to explain it on the basis of the low prospects of the science;

3) Answers are based on their own experiences. Most respondents are large pharmaceutical firms; small biotech firms may be more willing to do vaccine research at lower prices[274];

4) (Something not mentioned by Kremer) By using the size of the "market that previously proved sufficient to motivate research" it suffers from the fact that those figures include adjustments to account for the current level of inefficiencies in the large pharmaceutical sector, the failure of the blockbuster model, the levels of me-too research, the levels of marketing (which is more than is spent on research), strategic interactions, any side-effects of tight patenting, any distortions caused by lack of information sharing, etc. It also ignores the tax-breaks and public research that went into previous drugs, and important price-related issues (like the fact that only rich markets are covered). At the same time it ignores the fact that some of the more difficult vaccines may cost a great deal more than typical drugs to develop or manufacture. There is no explanation as to whether the $250-$500m is an optimal figure (it can't be) or just what large industry players have told the author.

Calibration

This does not stop the "rough rule of thumb" 'needed market size' figures being used extensively to 'calibrate' the model, in particular the probability distributions over technology that are then used to work out the time to vaccine development[275]. Kremer calibrates in such a way that "the underlying distribution of projects matches our estimates for two different scenarios describing the market for the product". In other words, if this framework is to be used later in policy advice regarding optimal 'needed' market sizes, the framework feeds a tautology in via the calibration exercise. Alternatives are then compared by feeding the figures into fancier calculations than could possibly have been used to derive the figures in the first place. Being based on such "vague and divided" and "rough rule of thumb" figures, the APC price figures cannot themselves be any more than "vague and divided" and "rough rule of thumb".

Kremer, for example, feeds in "parameters concerning the median time to discovery (assuming the product is technically feasible) under the current market size and under the hypothetical [i.e. "rule of thumb"] market size that would generate intensive research effort[276], as well as the total R&D costs if all feasible projects were

[274] It seems odd for Kremer to state this, but then design a mechanism that works most in the favour of large pharmaceutical firms.

[275] One should usually worry when the phrase 'calibration' is used, since it generally means that there is no data to support the framework and that what is available is used to create the data that would be needed to support the framework's conclusions. In this case, K3:7: "There is clearly uncertainty about γ, as well as h and l" (where γ refers to the elasticity of aggregate research expenditures to the rewards for successful products), and in fact they are not known. Instead they are made to fit the "underlying distribution that matches *our estimates* for the two different scenarios describing the market for the product." (emphasis added). The problems are that the estimates for the market under the two scenarios are themselves based on "rule of thumb" figures, and that they are also a function of the tight IPR in the pre-APC period.

[276] We mentioned earlier the chicken and egg problem of trying to work out optimal intensity from the actual discovery of the vaccine. Here the optimal intensity is based on figures derived from data

undertaken. From that we can recover ProbSuccess." But, by taking industry-generated measures of "hypothetical market that would be needed to generate intensive research effort", all the faults mentioned in point 4) above are conditioned away in the calibration exercise that derives the probability distributions used in the model. So when we later make efficiency comparisons of alternative approaches to vaccine finance, the results for the APC are all based on 'cleaned' probability distributions. No wonder the methodology comes out favourably in favour of the APC when all the faults have been conditioned away and yet all the layers of inefficiency of publicly-funded research have been kept in[277]. A correct methodology would try to work out what the counterfactual probability distributions would be if these faults could be avoided, and then put these faults back in when analysing APCs, but keep them out when analysing various other mechanisms, like for example more open collaborative research approaches.

Given the problems with the notion of using an auction to set the APC price, this does leave us rather at a quandary. If the auction cannot be relied upon to fix the APC price, we must have been hoping to get a fix on the cost of an APC some other way. Now it turns out that the *best* we can hope to achieve if the auction fails us, is a "rough rule of thumb" ball-park figure.

11.6. Adjusting the Figures up When the APC Fails 'Additionality'

In Section 8 we talked about the danger that the APC might fail the 'additionality' test – i.e. the APC is supposed to reward only *additional* privately-financed research not in anyway subsidised by any other form of public funding, such as tax credits and subsidies. If there is no mechanism to achieve separation and to reward firms *additionally* with APC payments, then in order to get a full measure of the public tax-payer costs of an APC, we need to add all the costs of other public funds directed to pharmaceutical firms for research covered by APCs. We cannot, like Kremer, presume that these costs are not there.

Indeed, the typical 'needed market size' quoted by the industry, and used in the APC figures, does *itself* already presume the current levels of subsidies and tax-breaks feeding into a 'typical' drug. The APC calculations must therefore *also* be presuming that firms working towards APCs will be in receipt of tax-breaks and subsidies. The presumed amount needs to be derived and added to find the overall cost to public finances of a typical APC (and adjusted to take account of all the inefficiencies of these other components).

11.7. An Auction will Price all Distortions into the APC, Including Marketing Costs

A running theme here is that for vaccine research not to be slowed, anything that distorts or raises costs of research needs to be factored into a higher APC price so that

---

generated by such discoveries (in this case, mostly drugs). This suggests that the 'auction' mechanism is not about *discovering* optimal intensity where the overall technology and cost of developing a vaccine is not known, but rather starting with a known cost of development and a known intensity and motivating firms to achieve that known level.

[277] All quotes from K2:7-8

those doing vaccine research are compensated enough to carry on participating up to the optimal intensity required. So long as players know that there is a high upper bound to the APC price, and so long as there is no effective constraint on the APC price-setting auction other than this upper bound, then strategic interactions, rent seeking behaviour 'chasing' up the APC price to try to extract the social surplus, and all other 'inefficiencies' will have to feed into a higher APC price. The auction does not reveal the perfectly competitive cost of development, but the cost generated by the less-than-perfectly competitive industry structure. No argument is given as to how non-perfectly competitive firms could be faced with a *constraint* on the APC price (other than the upper bound) without this harming vaccine development efforts. Arguing that the upper bound is so very high as "not to constitute a problem" starts to look worrying rather than reassuring[278].

One way firms may dissipate potential profits is via marketing/lobbying expenditure. If this is understood *ex ante*, and if there is no way to 'police' it out of the system, then it too should rationally be priced into the APC price via the auction mechanism. Observe that policing marketing efforts would be like policing a 'prisoners dilemma'; to the extent that it is not possible to avoid marketing in equilibrium, then not allowing marketing costs in the APC price will simply reduce the incentives to develop vaccines, delay development (again, especially so if the cost of development is related to the intensity of development and not held constant as in Kremer), and reduce the quality of vaccines. For example, if two or more firms have developed vaccines then they will market against each other to encourage the most countries to take their products. If the second or later vaccine is a better product, then why should any country be forced to use the first vaccine? But – given that under some versions of the APC the first vaccine is not removed – if the first vaccine engages in marketing to hold its share, it will be efficient to allow the second product to also engage in marketing, even if this eats into the social surplus supposed to have gone to this vaccine (the social surplus that it had hoped would have been 'left over' for this vaccine after the APC had reallocated after the first vaccine).

If the APC is set without an auction it is not clear by what degree the terms should *not* be adjusted upwards to allow for strategic interaction, including marketing costs. Failure to do so would destroy the chances of getting the desired result (a quick vaccine) if these strategic interactions nevertheless could not be policed out of the system. *A priori* it could be a very serious issue. It is not helpful that Kremer simply presumes the problem away.

11.8. Low Healthcare Infrastructure Generates a Higher APC price

One of the features that reduces the value of the deployment of a vaccine to the developers is the needed expenditure on currently non-existent infrastructure. Where a vaccination program is already under way, the cost of adding an additional vaccine is very low because the delivery cost is a large percent of the overall cost. For example, the World Bank calculates that adding a one-dose yellow fever vaccine and the three

[278] It also makes no sense since one could no doubt say the same thing about just about any intervention (clean water, sanitation, nutrition, etc.). And developed economy customers could never be charged for the use of telephones, computers and the like on the basis of the very high ceiling to the social worth of such products and services.

dose hepatitis B vaccine to the EPI package would only add 15% to the overall cost (including delivery). This is still only about 40 cents per dose.

This leads to a simple, but often painfully missed, observation in the debate about creating incentives for investment in vaccine research; that the current value of a vaccine to its developer is heavily suppressed by this lack of infrastructure. Creating a 'market' artificially, with the price necessarily incorporating a component to compensate for insufficient infrastructure, is indirect and inefficient. Given the positive externality effects anyway of health infrastructure it may, under sufficiently large distortions in APC methods, prove cheaper to expand health infrastructure *first*. As a very general economic principal, if the solutions proposed for financing R&D in vaccines create large distortions (as, it is argued, the APC does), then the most cost-effective response might be to fix the primary distortion first.

Clearly, also, an APC would itself require commitments of funds for infrastructure, if the APC is not to have to be set higher to compensate. In an auction mechanism, the APC price rises if less thought has gone into creating health infrastructure that boosts the value of vaccines. Concentrating on an APC without tackling these issues simply pushes up the APC price and adds to overall costs (or, the flip side, leads to too little R&D and slow vaccine development at any given size of APC).

11.9. The Auction Needs a Great Deal of Information

Even when the APC price is visualised as rising towards 'equilibrium' this still requires *some* notion of the state of technology, in particular that the cost of developing the vaccine is in a *range*. Unlike auctions for say, 3G licenses, the APC 'auction' is going on *at the same time* as the sought-after activity is taking place – so that intensity of activity and delay become issues in ways that would not exist in other auction settings. The designer of the auction would not want to start at a price above the range since that gives unnecessary economic rents to developers, but, similarly, the designer would not want to start at a price below the range since it would unnecessarily delay vaccine development. Even within the range, if the start level for the price is set 'too low' then the slower will be the progress towards the efficient price, while if the start price is set 'too high', then the quicker will be the progress towards the efficient price. But this leads to the conclusion that even with an auction, better *ex ante* information is needed to achieve the most optimal outcome.

11.9.1. Inability to get the auction 'price rise' rule right

The 'auction' would anyway simply consist of a *rule* to let the initial APC price rise. Getting the start price and the rule correct would need huge amounts of information. The optimal speed of price rise is highly sensitive to the underlying state of science and the strategic interaction of the firms.

In the simple technology of Kremer, a simple rising price might eventually be able to hit optimal intensity, but in the underlying technology that Kremer refers to, where the probability of discovery varies over the product development cycle in ways un-modelled by Kremer, this is not so clear. The model was reduced to that of a *constant* distribution of the probability of discovery per period of time by assuming that this variation is known and the average of this variation can be taken as 'representative' of the whole pattern of variation over the cycle. *If* the actual real world distribution were

constant, then a constant rule could be set. But setting the rule in the more realistic world of the varying probability with the imperfections listed in earlier sections would require knowledge of the whole distribution in order to work out the optimal price rule, and this rule would not necessarily be constant.

How would it adjust to the state of technology?

In reality the rule would also have to somehow adjust for the state of technology. If the state of technology got 'worse' there would have to be some way to let the price rise 'a bit extra' to compensate[279]. If the rule for letting the price rise did not allow a rapid enough rise in this period (ditto for a rule requiring a decline in the rate of price rise or even a price fall) then the project would be delayed. The rule about price adjustment when there are 'shocks' to technology is much messier than the simple rule in Kremer, and even more complicated if information regarding it is potentially wrong.

Strategic behaviour over the setting of the start price and rule

And there would be strategic behaviour over the setting of the initial start price and, even, the rule. This strategic behaviour in turn would flow over from APC to APC. Since a higher start price will enable the extraction of more surplus, there are incentives to 'pretend' that a high start price is desirable on early APCs in order to influence start prices on later APCs.

11.9.2. The paradox of the need for information and the need for secrecy

Without knowing the underlying technology, it is hard to know what the optimal level of research intensity should be, the start price, the optimal speed of price rise, and how to make adjustments to the price rise in the light of technological changes. This leads, once again, to the paradox that it is quite likely that the mechanism works best with relatively well-understood technology, but that this defeats the supposed object of the whole exercise of creating a mechanism that can work in an environment of highly asymmetric information.

An alternative interpretation is that the APC advocates (for all their talk about the asymmetric information facing public-sector funders) presume that the public-funders *do* know the underlying technology; they just do not know about the individual projects of firms. In which case the funders can work out the underlying optimal research level and use the auction to stimulate research up to *this* level. But this suggests the need for transparent information if the public-sector is to work out the optimal research intensity. But, we found earlier that in order for the APC to work in a world of non-stationary technology, the APC would need either strong patents or, failing that, secrecy[280]. So we face the paradox that when we most need the auction mechanism to work we cannot rely on the APC setters to have the transparent information they need to set it up in the first place.

There is a yet further paradox. Sometimes, hidden information is good if it helps to make an auction work better. If firms don't know what other firms are doing, this helps to stop them from colluding against the auction. Under totally transparent

[279] And the 'extra bit' of price rise (or fall) should only relate to the bit of technological change that was outside of the control of those taking part in the auction.
[280] And we also found that secrecy might still persist even with patents.

information, they find it easier to discipline each others' behaviour and hence will 'cheat' against each other much less than they would if information was more opaque. Hence, holding collusive behaviour together is much easier with transparent information. But this suggests a problem: On the one hand we presume that the organiser of the mechanism has accurate information on the exact nature of research going on in firms in order to be able to judge when 'optimal intensity' is being achieved[281]. On the other hand, if that requires transparency about everything, then the auction mechanism fails to work well. But this suggests the paradox of allowing secrecy (by, for example, keeping secret the results of research funded by firms) to help make the auction mechanism work better at achieving 'optimal intensity', but then never knowing fully if the auction mechanism was achieving anywhere near optimality anyway! And how do firms work out their optimal levels of research if they have no idea what the aggregate level of research is?

While vaccine trials could not be kept secret, Kremer explains that "research towards patents could be"[282] and that this helps to make it hard to collude against the auction. Implicitly, this seems to suggest that the more the mechanism is targeted towards trials, the easier it will be to collude against it, and so it must therefore be targeted at non-trials research. But this contradicts the assertion then made that the mechanism is mostly about trials-stage research.

11.10. Experimenting and Collapsing APC Auctions

Kremer suggests that it might be possible to select "easier to develop vaccines and drugs as a way to build credibility" and experiment with purchase commitments for these before modifying or extending to other diseases (as of early 2005, having got the ear of policy-makers for malaria and HIV, such nuances have been cast to the wind). But:

1) What is the time-frame of the delay of *other* vaccines/drugs, while this credibility is being built up by experimenting on these few vaccines? This has a cost in terms of lost lives and social welfare (a cost that *should* be factored into the Kremer calculations, since it *is* part of the proposed mechanism).

2) If, as Kremer claims, the average time to vaccine development is ten years, and the post-development period ten years, it is impossible to believe that later vaccines could build on much of an experiment with earlier vaccines without a great deal of delay in application to the later vaccines. Is it realistic to delay an initiative to develop 'harder to develop' vaccines for HIV/AIDS while experimenting with some 'easier to develop' vaccine? In reality, if unnecessary loss of social welfare and lives is to be avoided, the APC would have to deal with 'harder to develop' vaccines with very little experience with earlier vaccines.

3) What is the time-frame of delay for each vaccine? When Kremer says that the price in each auction is first set at a "relatively modest level" and then rises "gradually until

[281] There are other paradoxes. It is presumed in the APC models that private firms might, for example, wish to hide research on a clade of HIV/AIDS that is targeted at profitable developed economies in order to get a tax credit designed for a clade prevalent in a poor country. But it is presumed that similar sorts of behaviour, hiding the true intents of research, does not arise when firms face each other under an auction mechanism.

[282] K7:38

it proved sufficient to spur vaccine development", he is suggesting that the price takes time to find the optimal level, which typically could take many years. But the APC figures are calculated on the basis that this learning process is complete and the price always set optimally. But this learning process for each vaccine is part of the cost of using the mechanism and needs to be priced in. It cannot be presumed instantaneous as in the Kremer cost-effectiveness figures.

4) Experience with 'easier to develop' vaccines may not give much information for 'harder to develop' vaccines anyway. It is likely that the first situations in which APCs would be used would be those that are relatively the simplest and avoid some of the biggest problems mentioned in this paper. This would give a misleading impression of the ability of APCs to deal with the more complicated situations.

5) As with any individual experiment, one would need to take great care with interpreting the results. Thirty experiments might enable some sort of distribution over potential outcomes to be built. A single experiment would represent just one drawing from the distribution of all possible experiments and it runs the risk that – being in one of the tails of the distribution – using it to set the terms of follow-on auctions will lead to the terms of those auctions being set less than optimally[283]. It might make more sense to have a system that can be experimented with but such that the terms of which can be altered as experience comes in. It is less easy to visualise this with an APC than, say, with an open collaborative research approach, since the APC system involves more sunk levels of private finance; the ability to alter terms *ex post* would impose risks on private capital and increase the private capital cost component of the APC.

6) What if it goes wrong and the idea has to be abandoned? We saw above the self-fulfilling aspects of a system of failing APCs. This cannot be ruled out and its aftermath should be explored. What are the costs of delaying an alternative (like, for example, a more open collaborative approach to research)? What are the legal commitments to already sunk resources? Could the APC instigator be forced to pay or keep a failed APC system going?

7) What if we wish, later (if the APC mechanism does not generate the hoped-for results), to introduce an alternative method like, for example, one based on more open collaborative research? A partial collapse might be worse than a total collapse, since we would be left with a tight patent system to support the remaining APCs even though this would harm a more open alternative. A later reversion to more open approaches would require the more open framework to be modified in potentially inefficient and costly ways to take account of those APCs still in place.

8) The period of 'experimenting' is also a period during which firms can adopt strategies that influence future APCs (like delaying research in early APCs in order to create the impression that future APCs should be initiated at much more generous starting prices). What are the extra costs of allowing this manipulation in terms of

[283] If, for example, the distribution gives information on the possible cost of APCs, then just using the result of one experiment will run the risk of over- or under-exaggerating what APCs might typically cost. If not treated with care, in the former case this runs the risk of an extremely wasteful follow-on auction, while in the latter case it ruins the ability of the follow-on auction to achieve an optimal result (or any vaccine at all). The notion of 'experimenting' becomes rather hollow in such situations.

future APC costs? What are the costs in terms of delay to the creation of a general system for stimulating vaccine development?

9) What is the cost of delay if it is ten years before generic competition is allowed to drive down prices?[284] This is an extra cost compared to simply allowing generic competition as soon as the vaccine is discovered.

Clearly it would be less risky to experiment first on a few diseases with a more open framework rather than with an APC framework. The open framework would generate shared information and might be easier to unwind later. And it would not risk left-over APC commitments, and high capital costs, forcing even more APCs onto researchers for no other reason than budgetary constraints caused by the 'failure' of the first APCs.

[284] K4:14.

12. PRIVATE FINANCE

Under an APC, all research costs are first covered by private finance, and then much later (on average) repaid from public taxes. Of course, most of those investing do not get repaid anything. R&D under an APC is a gamble with a large 'prize' but low odds of winning for any individual firm. The APC has to be set large enough to cover *all* expected privately-financed costs and *all* expected capital costs aggregated over *all* of those engaged in vaccine research. This is much the same logic as that behind the blockbuster model, and is the reasoning behind the calculations of Tufts of the typical cost of developing a drug[285].

The notion is that if financial markets are efficient, risk can be spread, with investors holding well-diversified portfolios and bearing little or no idiosyncratic risk. Some financers and researchers, however, have to bear risk to be motivated. As Kremer argues, an APC would "provide a strong financial incentive for researchers to focus on developing a marketable vaccine, rather than pursing other goals, like publishing academic articles," and the "more sensitive the research expenditure to the rewards for a successful project; the better the value-increasing programs perform." The exact mechanism for this however still needs some spelling out. We already saw how difficult it actually is to tailor rewards to required research expenditures to get a given outcome. It is not clear that the sensitivity of research expenditure to the rewards for a successful project is anywhere near as high as Kremer presumes, and therefore the performance of APC programs anywhere near as good in comparisons with alternatives.

Mechanisms other than APCs would also face many of the problems highlighted in this chapter.

12.1. The Financial Market Difficulties of Vaccine Research

The fundamental problem in financing vaccine research, as indeed with any research, is the usual one of the 'separation of ownership and control' of those firms engaged in research. This creates two co-existing and somewhat conflicting problems. Firstly, managers/scientists have a preference to invest in things that benefit them (a larger firm size, nicer offices, more staff under their control, higher pay, prestige projects, etc.) But, secondly, being risk averse, and certainly more risk averse than shareholders, they wish to avoid risky R&D.

Normally, leveraging would be useful to mitigate the first problem, but it is of limited use in the case of R&D-intensive firms. The knowledge asset created by R&D investments is intangible, often contains a lot of 'know-how', is partly, if not largely, embedded in human capital, and is often very specific to the firm. With banks and debt-holders reluctant to invest where there is no physical asset to secure loans[286] (and given that the sunk costs associated with R&D investments are higher than for

[285] DiMasi et al, ibid.

[286] Williamson, O.E. (1988) refers to 're-deployable' assets (those whose value is almost the same in alternatives to current use) as more suited to governance based on debt.

ordinary firms) capital structure is therefore less leveraged (i.e. less debt-based) than average. Servicing debt also requires a stable cash flow. Often R&D must be sustained at a certain stable level to be productive and it would make R&D even more expensive if it had to compete with this cash flow requirement[287]. Again this tends to reduce the use of debt finance for R&D, and again this may raise the cost of capital if alternatives are more expensive (depending on the tax treatment of debt versus equity, etc.). Furthermore, if bankruptcy is a possibility, managers may avoid variance-increasing R&D projects that shareholders want, leading to fewer long-term projects, and this too mitigates the use of debt. So, the apparent solution to the first problem that would seem to suggest reducing free cash flow would simply force the use of high cost external finance – which makes R&D more expensive[288] The optimal solution is to somehow increase the long-term incentives of managers rather than reduce free cash flow[289].

This leads to the conclusion that that part of vaccine research that is privately financed will be largely based on equity forms of finance. But this leads to a new set of problems.

Financial markets face asymmetric information problems too

One reason many companies do not do certain kinds of research is not because of the lack of an end market *per se* (the World Bank or the 'Global Fund for Health' would ultimately be the end market in this case), but because it is hard to communicate to equity-based markets the value of research and hence to raise the finance for it. Problems with asymmetric information and moral hazard create an extra gap between the private rate of return and the cost of capital when the innovator-investor and financier are different. Firms therefore do not invest in innovations that would pass the private returns hurdle[290]. The 'lemons premium' is higher for R&D than for ordinary investment because the difficulty of separating good from bad projects when projects are long-term R&D investments is much greater than with short-term low-risk projects[291]. The asymmetric information problems is made worse by the fact that many firms are also reluctant to release information to financial markets, afraid of revealing information to competitors. This reduces the quality of information signals that financial markets need to base investment decisions on[292]. In worse case scenarios the problem bites so severely that projects disappear altogether. This is particularly aggravated by the long gestation periods of pharmaceutical projects and is especially but for projects that would actually require information revelation and sharing (such as HIV, malaria, and TB research). There is a tendency in the APC literature to talk in the mantra of 'efficient financial markets' where none of these

[287] It is worth exploring how a stable cash flow generated by an R&D Treaty might help stabilise flows into R&D, though it would, no, doubt, also need some consideration of any agency problems also created.

[288] There is also good empirical evidence that limiting cash flow in R&D intensive firms is less desirable as a method to reduce the agency costs of the first problem.

[289] Incidentally, for the uninitiated, the last two paragraphs encapsulate why pharmaceutical R&D takes place in equity-based firms, older firms with already established cash flow records, or newer firms with access to venture capital... but certainly not debt-backed or bank-financed firms.

[290] K 1:5 see Blanc, (1999). Kremer mentions (K1:5) that in private correspondence with Jon Horton, GSK, Horton remarks that firms "like to see a return on investment by the end of year 3."

[291] Leland and Pyle (1977)

[292] Bhattacharya and Ritter (1983), and Anton and Yao (1998).

difficulties arise (Kremer just ignores it all; finance is simply a vale behind which real economic activity takes place). This is where a rôle for venture capital, VC, comes in.

The use of venture capital

Sometimes the arm's-length market-based financial systems of the US and UK are contrasted with the bank-centred capital markets of Continental Europe and Japan. VC is a combination of the good bits of both. It gives the strong incentives for manager-entrepreneurs of the stock-market and the monitoring of the bank-based system. The optimal form of the VC contract is actually a complex debt-equity hybrid[293]; more like debt when the firm does badly, but more like equity when it does well (since this is incentive compatible).

The VC solution to the financing of vaccines has its limits however. VC tends to concentrate on few sectors at a time and also tends to make investments of a minimal size that may be too large for some start-ups and smaller ventures. VCs also require a thick, active market in small and new stocks (NASDAQ and EASDAQ for example) to provide an exit strategy for early-stage investors, so they can move on to new projects, and to enable successful entrepreneurs to regain control of their firms (and to give them incentives to start up in the first place). VC also tends to be pro-cyclical (though it is hard to disentangle the direction of causation). Empirically, even though there is a great deal of entry to the VC industry, returns in the industry are still high, suggesting a high required rate of return.

The financial side of an APC does not automatically solve these particular problems. And, we will see, the modelling of APCs has largely ignored many of them.

12.2. The Financial Market Bias of APCs Towards Large Pharmaceutical Firms and away from Not-for-Profit, Biotechs, and Developing Country Research

The APC would not, for the above financial market reasons, be the most appropriate mechanism to encourage many small biotechs or companies from developing countries to join the push for a vaccine. We know from econometric evidence that small, new, innovative firms that have not yet had time to establish cash flow, experience difficulty accessing capital and hence face high capital costs. The evidence on costs of capital for large firms is more mixed. Venture capital plays an important rôle in filling this funding gap, but it is incomplete, especially in countries where public equity markets are limited. One side-effect of an over-reliance on APCs, with their supposed reliance *totally* on private financial markets[294] is that the distribution of firms working on vaccine research is biased in the direction of large pharmaceutical firms, that are already wealthy or already profitable, have good cash flows already, and exist in developed economies with strong equity-based financial markets – and away from smaller firms *ceteris paribus*, and especially those based in economies with weak or non-existent equity-based financial markets (indeed this is one of Kremer's declared intents). To the extent that the truly innovative research and the

[293] Aghion, P. and Bolton, P. (1992), Dewatripont, M, and Tirole, J. (1994).

[294] We saw that this was not actually the case. To the extent that they do *not* rely totally on private finance, many of the public-finance biases they were supposed to solve, including crowding out, creep back in.

cheapest research may be taking place in the smaller firms, and that it might be useful to have firms in developing countries (other than divisions of large pharmaceutical firms) also involved in vaccine research, this would slow vaccine research *ceteris paribus*, and require a larger APC price to encourage research in these innovative firms[295], and would lead to the conclusion that it may be socially more beneficial to have alternative front-end incentives to such companies or access to funds from an R&D Treaty, in order to reduce *their* capital costs.

Problems with finance at the start of projects

Indeed, the effect may be much worse if the technology probability distributions are as argued here and not as Kremer envisages. Uncertainty is often greatest at the start of a research programme. This leads to the options-based character of many projects. But it also leads to greater problems in acquiring access to capital. An *exclusively* privately-based funding mechanism based on equity finance may reduce the number of firms working on the early parts of vaccine research projects to just the few largest firms who have the internal capital to do so (clearly this also aggravates the auction problem mentioned above since there would be too few firms to get an optimal price out of the mechanism even at early stages in the auction, never mind in the late stages as discussed above). This, naturally, reduces the number of leads being followed, with potentially deleterious effects to the speed of vaccine discovery, if the number of leads being followed is an important determinant of the speed of vaccine discovery. By assuming stationary technology (the same average-probability-distribution-repeated-every-period) Kremer envisages firms as acquiring finance always relative to the whole path of development, with no problems in acquiring finance for the initial activity. In reality early projects with small probability of great success in the future are worth continuing even if they do not pass the expected-rate-of-return test and have problems acquiring access to finance[296]. Problems with early stage finance are being treated in far too off-hand a fashion in Kremer.

Bias against firms who need other government help

In addition, the APC is supposed to reward *private additional* finance to the vaccine effort, otherwise it simply crowds out other publicly-funded vaccine investment and the public pays twice. Kremer's claim that publicly-funded research should also chase the APC we showed to be patently wrong and actually encourages public waste. It was pointed out above that this might even require those who seek to claim the APC to prove that they did achieve the breakthrough with genuinely private finance and not with-tax-breaks or other public funding. Combining this fact with the fact that many innovative and small firms have problems attracting private finance and may therefore have to rely on such publicly-funded initiatives, we again find that the APC is biased against them (they would also find it harder to 'hide' this proportion of their funding for vaccine research since they would have fewer other programs to bury costs in).

Both here and in many other sections of this paper we have seen the way the APC is biased in its impact towards large pharmaceutical firms rather than towards not-for-

[295] The *ceteris paribus* is inserted since clearly the presence of the APC *will* increase the interest of venture capital firms in small biotechs, but comparing an increase in funding on APCs with an identical increase on alternatives, the argument being made goes through. This *ceteris paribus* reasoning carries into the following paragraphs too.

[296] Indeed Scherer (1998) shows that the distribution of returns to a project can be Paretian where variance does not exist, so that standard risk adjustments do not work.

profits and biotechs[297]. In the end it is an empirical issue as to where we expect the most innovative vaccine research to come from. Kremer claims that it is in the large pharmaceutical firms and makes this central to his model. But there is good evidence that the most innovative vaccine research goes on in the not-for-profit and biotechs. And there is also good evidence that encouraging research in developing countries and by locals, is a way to keep costs down, to spread knowledge (because of technology transfer, etc.), and to ensure that research at clinical stages is run in ways that produces products more appropriate for developing countries. Why – in the face of this evidence – use a mechanism that is deliberately biased towards those already with good cash flow?

The paradox of targeting those who had cash flow before

Paradoxically, large pharmaceutical firms may be being targeted precisely because of their *previous* good cash flow rather than because, at the margin, the most innovative vaccine research *possibilities* reside within them. This requires a counterfactual that is different from the world we exist in. A good way to think about this would be to consider what the impact of an equal reduction in cash flow problems for large pharmaceutical firms and biotechs/not-for-profits might look like. To the extent that reducing the cash flow problems of biotechs/not-for-profits has a differentially greater impact on vaccine research, *they* should be being pursued with methods that ease their cash flow problems the most. To the extent that Kremer is wrong, then the lower the marginal impact of spending a dollar on an APC will be compared to alternatives, and the more costly the APC will be as a method to stimulate vaccine research relative to alternatives.

12.3. Problems with the (Hidden) Financial Part of the APC Model: Stock Options to Incentivise Vaccine Research

Kremer argues that non-APC methods "place administrators in the position to judge what scientific avenues should be pursued, rather than scientists with a self-interested stake in the success of their work."[298] And it is further claimed that "Public funding is generally given to institutions that are already well-known in a particular field, which creates little incentive for *new players* to invest funds in a different idea or approach. In comparison, a pull program such as a purchase contract or tax credit for purchases is open, and encourages innovation form *any participant*"[299] (Italics added). But this is only to the extent that 'novel ideas' can find financial backers. The phrase "any participant" should be rephrased to "any participant with the resources".

The APC simulations so far presented to governments make heavy, but often implicit, use of the notion of stock options as a way to self-interestedly reward scientists who work on vaccine research and to enable "new players" and "any participant" to take part in the search for a vaccine. With stock options disciplining researchers, wasteful, poorly-targeted, R&D never takes place. Publicly-funded research lacks this disciplining device. Unfortunately, the technological assumption of Kremer rules out

[297] Again this is *ceteris paribus* thinking. It is all based on a dollar for dollar comparison of the impacts of alternative finance schemes for paying these firms to do research. Alternatives to the APC could be made to more directly target biotechs, not-for-profit, developing country firms, etc. for the same dollar expenditure as an APC.
[298] K2:10.
[299] K2:10.

many of the problems that stock options standardly have difficulty dealing with, so that very little can be said about their operation under an actual rather than under a hypothetical vaccine APC.

Stock options – and rewards linked to stock markets generally – have strengths but also many weaknesses as an incentive mechanism to encourage high quality vaccine research:

1) Options struggle to allocate efficiently in environments based on secrecy. We found that secrecy appeared as an important feature in holding together a non-idealised APC dealing with complicated technology.

2) Nor can they cope well with technology that generates information for later stages of the technological process but that does not yield to easily identifiable and defendable property rights, or even if it does, where patents would reveal too much to competitors. Private financial markets suboptimally invest in such information discovery.

This again points to APCs as being targeted at only the late stage parts of vaccine research – and the difficulty that this causes if it leads to companies keeping otherwise valuable public information secret.

3) If complicated technology is to be dealt with via APCs, the use of stock options is another force pushing in the direction of stronger IP.

4) Options would not reward private research with public good/collaborative aspects to it, since the option payment could not be based on the full value of returns to the research.

Asymmetric information problems ruled out in the technology anyway

5) In Kremer, it is not that stock options are presumed to solve asymmetric information problems perfectly every time; such problems are ruled out in the technology at the start. Even though the modern economics literature is replete with evidence of the difficulties that financial players face, is it right to rule out *a priori* such problems in all calculations of the APC? And is this reasonable given that such problems are ruled in very heavily in all alternative methods?

For example, it is claimed that "biotech and pharmaceutical firms could decide whether it is likely that, for example, work on one HIV clade will be effective in helping prevent other clades and the government will not have to make these scientific judgements."[300] The asymmetric information problem does not disappear however. If these are to be "new players" or "any participant" then financial backers will have to work out the value of the research strategy in order to work out whether to fund them.

In truth, given this asymmetric information problem, most APC benefit goes to "institutions that are already well-known" rather then "new players". Indeed we have

[300] K8:3.

found that Kremer is keen to get such institutions, the big developed-economy pharmaceutical firm, taking the leading rôle.

Financial players are presumed to see too much

6) Even then, financial players might need a great deal more information than just this one firm could provide to make a socially optimal decision (even if we ignore the secrecy and asymmetric information problem related to this one firm for a moment). To work out the value of this firm's one project, the financer really needs to know the value of other projects being undertaken by other firms. Kremer presumes that financial players work on the basis of probability distributions that are common knowledge and financial markets that are perfect everywhere. But if they experience asymmetric information with respect to this one firm, it is illogical for them (and us) to presume perfect, non-asymmetric information elsewhere. We cannot escape the fact that to make their decision financial backers need to form assessments over information elsewhere in the system.

Some lessons on compensation schemes generally

7) There is a standard argument in the economics literature regarding the trade-off between risk (requiring a fixed component of reward) and incentives (requiring a variable component of reward). Depending on the noisiness of the relationship between outcome and effort, technology, length of contracts, degree of monitoring, shape of indifference curves, etc, there is an *optimal* trade-off. The more unknown and risky the R&D process (HIV vaccine research for example) the larger the fixed component of compensation. In reality, scientists need a component (possibly, large) of insurance. A model that assumes that it is *all* incentive and no insurance is misleading.

8) Besides, as the variable component of any compensation scheme rises, the scheme detracts more risk averse scientists and attracts increasing numbers of less risk averse or even risk loving scientists. It is not *a priori* clear that such scientists are desired, especially if excessively risky research strategies are encouraged ('gambling for resurrection' for example). At the same time, some projects do need particularly risk-loving scientists in order for them to be carried out; one could imagine some option-based discrimination across projects with intent to match scientists to projects according to their degree of risk tolerance.

9) Nevertheless, scientists vary in their degree of risk aversion, so that no general options scheme will ever fully maximise incentives as the Kremer calculations presume.

10) And as the science base expands and more scientists are taken on, one might imagine that the average degree of risk aversion would in fact naturally rise, leading to lower and lower 'incentive' components on average. If a major push for vaccines and drugs for neglected diseases is made, this general equilibrium consideration matters even more.

11) An *individual* pharmaceutical project is especially risky. Most trials-based projects fail, especially at early stages. It is tautological to define the 'good' research projects as those that eventually succeed. Many failing projects were *ex ante* 'good' projects, worth pursuing. Indeed the level of knowledge of what would and would not

succeed is often poor (after all, if it were not so, many more failures would be avoided). Many of the risks have nothing to do with the efforts of a specific scientist. Those who work on such projects and who are paid via options are not able to diversify risks in the same way as investors. There is nothing in such settings to suggest that risk averse scientists can be motivated by 'options' to take part in such risky trials activities.

Quite likely the notion is that those financing, leading, or monitoring projects should face compensation packages forcing *them* to take risks in order to incentivise *them* towards picking good quality projects, while many of the scientists working on projects should not. But this still leaves the pharmaceutical firms themselves facing plenty of internal principal-agent problems and the need to create incentive mechanism to motivate effort and risk taking within the firm.

12) The more unknown and risky the R&D process – this is often the way it is with vaccines – then the larger the fixed component of compensation will have to be for scientists, with less and less of an options-based element anyway. This reduces the 'incentives' argument of Kremer. A project that is very risky but that yields a very useful result when it works may not generate enough options-based payment schemes to attract scientists to work on it (this is even more so if the useful outcome has public good aspects to it and hence the scientists will not be able to internalise the outcome exclusively to themselves). And why should scientists be forced to face these risks anyway?

13) The 'success', in terms of usefulness of results produced, is often very hard to quantify. From a social welfare perspective it is sometimes useful to discover what does *not* work. For some drugs the result is purely serendipitous. If success is only revealed over time after development of the product (usually the case with vaccines), options will not efficiently price this.

14) If the project is likely to take many years, as is typical of vaccine research, it is not clear that scientists will get the full payout from their efforts. To efficiently enforce 'effort', the value of share options needs to be linked *not* to the time of vaccine discovery, but to the value added by individual researchers. Many (if not most) results are intermediate. If scientists are being paid via options, the equity market needs to be good at pricing this *expectation of discovery* at *all* points during the process of discovery so that researchers can get rewarded for the *value they add*.

We saw above the way non-stationary technology creates the need for links between stages. This is paralleled in the world of the researchers. An early 'successful' activity – if it enhances the probability of success of later stages – should be reflected in share prices and option values, so that those who did the good research in the earlier period should see reward via their options. To the extent that this does not work (maybe because of asymmetric information) incentives are weakened.

The inefficiencies of stock market booms and busts

15) Booms and busts in the stock market force those paid on such options-based incentive schemes to face instability in their own welfare that has nothing to do with their effort and investment choices. This is inefficient. It adds to the required return to those being paid by options and it distorts investment decisions. In the late 1990s, IT,

communications, and biotechnology saw great inflows of capital, but when the bubble collapsed, capital dried up. If the market perceives that the chance of discovery has receded (say because a financial bubble has burst and access to finance has dried up) they downgrade the stocks, hence the option values. This all feeds into a higher needed APC price *ex ante* to compensate researchers for this risk.

There is of course an argument that an 'inefficient' bubble, if it creates finance where finance would otherwise have been very constrained, may correct one distortion with another, and is therefore not entirely bad. But this still interferes with options-based incentive schemes since bubbles often do not discriminate across good and bad projects, and they continue to force those paid on options to face extra risk.

16) There are issues of team efforts versus individual efforts. If the stock price is not conditioned on a scientist's exact act, but on the choices of others in a 'team' or indeed a multitude of teams (often this is the way with vaccine research) the individual scientist may not have the strong individual incentives implied by the simple theory. This is especially so if the choice of research strategy imposes a great deal of risk on the individual researcher but where the gains are spread over many.

Credit where credit is due – and also not where it is not due

17) Reward should not be given to or taken away from privately-paid scientists for something that they had no affect upon. If vaccine technology improves because of publicly-funded research, the options mechanism should not reward scientists at later stages who did nothing to improve the chances of success. Similarly, scientists should be protected from deteriorations in vaccine technological possibilities caused by others failing to perform their part earlier in the chain – including publicly-funded science. This makes options a less efficient instrument for paying scientists. It also aggravates the interplay between public and private stages of technology discussed above.

18) In analogy to the coordination needed between the non-APC-funded parts of the research process and the APC-funded parts of the process, those not being rewarded in ways linked to the APC nevertheless need to commit to actions that feed into the rewards of players who *are* paid according to the APC, otherwise the APC funded players have to be compensated for this uncertainty (meaning that again the APC has to be set higher). Coordination is a problem, and the risk of failed coordination has to be priced in to options too.

The worry about discretion

19) If a flexible APC program is in place (it was argued above that this was much more likely than a fixed program, though even the fixed program turned out to have many discretionary elements) then players have to work out how the discretionary choices of those administering the program will impinge on share prices, hence their options. This has several consequences. First, there is a clear conflict of interest with those from the industry populating the boards administering the programs, with a potential for regulatory capture and intent to manipulate the program to alter option payoffs. Second, given the way discretion and rules impinge on share values and options, if there are litigious aspects to those working in the industry, this may mitigate against 'regulators' taking actions that might *ex post* reduce share prices and lower options values; when the regulator gets it right, the regulator gains nothing,

when they get it wrong, the regulator suffers much, so, on average, they hold back from doing anything that risks pushing share prices down. Third, there are multiple layers of administrators/regulators of the program, and many layers of committees, and this imposes even more risk on those paid via options.

Options distort truth too

20) Options can distort truth and incentives (as can be seen from the experience of heavy options usage in the late 1990s stock market). Given that many options are written for short periods of time, players have an incentive to hide bad information and to make out that things are going a lot better than they truly are in order to boost the firm's stock market value, attract finance for their particular projects, and/or to enhance their options' values. And they concentrate on actions that reveal lots of positive information, even if other actions with less-easy-to-communicate information might have been more optimal.

Information becomes a far more tricky concept than the original APC models presume, since information has a dual rôle. It is crucial to making the APC mechanism efficient, but it is a commodity the production of which affects the range of strategies open to firms and their payoffs. We can't automatically presume that the production and revelation of information in response to the latter set of incentives will be sufficient to perform the first function. All the above problems add to the incentive to distort information flows – spoiling other aspects of the APC story.

Scientists face an asymmetric information problem

21) Kremer claims that with pay linked to stock options, scientists join companies where they "believe the scientific prospects are good." But this suggests that we should explore how it is that scientists discover such information, such that they do not simply find themselves on one side of another asymmetric information problem. And it suggests further incentives for firms to 'spin' prospects of research projects to attract scientists as well as finance.

22) Many of these problems interact with other problems discussed in previous sections. For example, incentives are created to distort or hide information in order to strategically deter entry so as to boost share prices, and to engage in rent-seeking (by trying to manipulate the APC price to extract more of the social surplus) rather than in genuinely innovative research – to go for short term advantage over more long-term objectives, etc.

23) Most choices of strategy are about choosing over risky actions. To be efficient, the share price has to directly connect to a player's act and efficiently reflect the consequences of that act. Options are less useful in a complicated situation like vaccine research where the connection is not close and where valuation mistakes may be made.

24) Scientists in large pharmaceutical companies are offered shares that relate to the 'overall' performance of the company, and not to this particular part of the company's activities. This generates perverse incentives, especially with the generation and use of information.

25) Kremer concedes that over-optimism goes on in pharmaceutical and biotech firms but claims that this over-optimism is corrected by investors requiring a higher hurdle rate for projects before they approve them. Normally, this would be visualised as feeding into higher overall cost of projects and be interpreted as a bad thing. In studies, hurdle rates are often typically 20% or more. The 4% discount rate of Kremer looks even less likely to be correct and it is not clear why he uses it and then relies on high hurdle rates of 20% or so elsewhere to drive his logic.

26) There is still a winners curse problem, on top of adverse selection, for any company rewarding via options.

27) These finance-based difficulties as well as the patent-based difficulties mentioned earlier, suggest that APCs may also not be suitable if developments require a certain sequencing of events, possibly across many firms. This is, perhaps, another rôle for a central coordinator (such as the NIH or an open collaborative allocation mechanism).

The pressure is to create profit, not a low APC price

28) It is not at all clear that private financial markets would discipline firms not to 'chase' social surplus by manipulating the auction towards the top of the band of possible APC prices. The overriding principle of finance is profit maximisation, and, once much of the simple APC structure of Kremer is changed, it would almost certainly be privately more profitable for financial markets if pharmaceutical companies were allowed to 'chase' the social surplus by seeking ways to push the APC price higher.

An auction would not discover the price consistent with the supposed 'optimal' level of R&D, but the level consistent with *profit maximisation*. Intuitively, it is just as profitable (at the margin the return is the same) for a financial institution to hold stock in a pharmaceutical firm with some monopoly power expending resources rent-seeking, as it is to hold stock in competitive firms not expending resources rent-seeking – given that the APC price has to be adjusted up in the former case to compensate for the rent-seeking (and this is fully understood *ex ante*). Collectively, society is worse off, but there is no private way out of this prisoners' dilemma. We are simply lead again to the realisation of just how important the unrealistic assumption of perfect competition is to Kremer's results.

29) Like much of the analysis above, all this suggests that APCs – this time from the angle of finance – are being thought as solutions to relatively uncomplicated scientific problems. And, as above, it suggests that the suggested breadth of usage of APCs is exaggerated. Anything less than 100% financial market efficiency is not modelled, or even discussed, and many knotty financial market problems are simple assumed away.

This is simply not good enough.

The financial side of the APC model needs a great deal more elucidation. Unfortunately, like so many other aspects of the Kremer approach, the inclusion of financial options seems to be for convenience rather than the result of a carefully analysed thought process. Their inclusion allows those promoting APCs to vaguely allude to *some* mechanism that would supposedly enforce the desired efficient

solution without having to spell it out, while meanwhile castigating alternatives for not having such a mechanism.

What does it do to the veracity and trustworthiness of the apparent 'superiority' of the APC to discover that yet another layer of potential imperfection has been cleansed entirely from the APC model before subjecting it to comparison with alternatives?

12.4. The 'Replacement Effect'-Financial-Market Interplay

We saw earlier that the development of cheap, one-off, vaccines for conditions like HIV/AIDS will replace profitable, expected, long-term treatment programs, thus generating less of an incentive to develop them in the first place. Total (expected, discounted) *industry* profits are lower if such vaccines are developed[301]. As pointed out before, this is not intended to cast aspersions. It is an effect that is being forced on pharmaceutical firms through the natural workings of financial markets – as well as being a function of the structure of the pharmaceutical industry and the nature of IPR[302]. If equity markets (and we just explored why the large pharmaceutical firms are almost entirely equity-based) correctly price all future expected discounted profit flows, then those firms working on projects that risk replacing profitable programs (profitable in the expected sense, which may be an important sense for a growing market like HIV/AIDS), will experience a depressing influence on their equity valuations, and this will increase their capital costs generally – not just for this research project but for others too[303]. This leads to them requiring an even higher rate of return on projects. The figures are not inconsequential. Even at the currently much lower prices than a few years ago (one can imagine how the equations must have looked then) the costs of the drugs alone for life-time treatment of HIV/AIDS, generates a cost of nearly $1,200 per DALY saved in developing countries[304] compared to probably a few dollars per DALY saved for a vaccine. If there is already a 'lack of a market' for HIV/AIDS vaccines, this simply reinforces this problem.

The fewer the firms that are already being relied on for *both* treatments and vaccines, the larger the 'replacement effect' and the lower the incentives to invest in vaccine R&D. Conversely, the more competitive the pharmaceutical industry then the stronger the incentive for firms to work on vaccine R&D since success would replace the treatments of *other* companies. The 'replacement effect' is also stronger the more able are incumbents, through tight IPR, to restrict access to information that might undermine their competitive positions.

The problems of an aggregate condition

[301] Kremer hints at something similar going on in the TB drugs market. K1:2.

[302] The issues are certainly controversial, but that should not prevent us from tackling them. If it turns out that 'replacement effects' are part of the problem in raising finance in certain vaccine markets such as HIV/AIDS, then better policy will result from considering rather than from ignoring such effects – as the following section will hope to show.

[303] Notice that it does not have to be 'actual' replacement; risk of replacement is sufficient.

[304] Based on approximately $430 per year of drug costs (K10). The author has no up-to-date (2005) figure for this based on $120-$140 per year drug costs, and would welcome a correct updated calculation (rather than improvising an approximate calculation). K2:25 lists a host of programs where the extension of vaccine coverage would cost just a few dollars per DALY, as low as $3-$4 for measles, which accounts for one in seven of the DALYs in Table 1 above.

This is also complicated by the fact that the 'replacement effect' is an *aggregate* condition. Clearly, if the expenditure on HIV/AIDS treatments in Sub-Saharan Africa is already pitifully low, then vaccine developers might not expect much of a 'replacement effect' there. However, the HIV/AIDS treatment market also includes potentially very profitable segments, and the effect on *these* segments from vaccines developed for the poor segments works against private incentives to research towards vaccines for the poorer, low 'replacement', segments. This is much the same logic as that found at work in anti-retroviral drugs markets, where firms are very unwilling to price-discriminate (normally the profitable thing to do) by setting very low prices in very poor markets for fear that this will alert consumers in much richer markets to the potentially extremely low marginal costs of the drugs, risking agitation for prices to be set much lower there[305]. Given the one-off nature of vaccines, and the very low prices that could ever be expected from them in very poor countries, the effect need only be tiny.

'Replacement effects' might even be at work for vaccines that do not obviously compete with treatment programs – such as vaccines for diseases that affect mostly only the poor and for which there is low current treatment – if cheap only-once-ever-used drugs (costing cents or a few dollars at most) weaken pricing power in profitable treatment markets[306]. This weakening only has to be tiny, maybe even fractions of a percentage, given the size and duration of the latter market compared to the former (all compounded by the fact that the latter market refers to multiple periods of future sales of treatments whereas the former refers to one-off sales), and that the prices in the former could never be very high at all. And the effect is strengthened further if there is any expectation that any resources being made available might otherwise go to treatments in the poor markets.

Reinforcing factors

There are three further financial mechanisms reinforcing this problem:

1) If the current system relies on 'small' firms (entrants, biotechs, not-for-profits, etc.) to work on vaccines to achieve this 'replacement', such entrants will need access to sources of finance[307]. If these firms are much more credit-constrained than large incumbents – as we have just seen that they are – then their cost of researching vaccines is much higher and profitability much lower. Their ability to do the 'replacement' is much weakened as a result.

2) In addition, biotechs usually have to sell the promising discoveries they make onto large pharmaceutical firms since they lack access themselves to the large amounts of capital needed to take projects right the way through to an end product (and this may be especially so for something like a HIV/AIDS vaccine). Even if biotechs are marginal, competitive, players and might not suffer from the 'replacement effect'

[305] Scherer and Watal 2002 contains a diagram showing the weak correlation found between price and country-level income for 15 antiretroviral drugs (They also point out that the empirical evidence is complicated by import duties, local tariffs, price controls, taxes and wholesale profits, etc.).

[306] Kremer points out that one of the advantages of the APC is that it enables firms not to have to be transparent about what it actually costs to manufacture drugs, for fear of these effects. Though, we also found that they have to reveal a great deal of information to those running the mechanism.

[307] Observe that overall profits to all companies would be lower after replacement, illustrating again the low incentives to do such activities by any player other than a purely marginal player.

themselves, the need to turn to large pharmaceutical firms at late stages, feeds the 'replacement affect' onto them. Biotechs in turn find it more difficult to raise the finance to do early stage vaccine work since financial investors know that they will face less of a market for the results of such projects because of the 'replacement effect' of the buyers, and because of the risk that buyers will not be so interested in sinking heavy investments themselves to bring a project to completion.

3) Currently, not-for-profit firms and 'not-profitable' biotech firms can only take advantage of tax-breaks to the extent that they can be bought out by much larger pharmaceutical companies to 'cash in' on the value of the tax-break (the smaller firms amass all their unused tax-breaks as an asset reflected in their equity valuations until taken over). This is unfortunate given that more than 50% of current vaccine research takes place in biotechs. That their research needs to boost their share valuations in ways that appeal to large pharmaceutical firms, gives another route for the 'replacement effect' to enter. A mechanism that is less reliant on this feature may enable a greater number of firms to exist in equilibrium and a lower impact of the 'replacement effect'[308].

Incidentally, given the way the APC is designed to create additional private finance, and incentives additional to tax-breaks, it would supposedly have to find some way to exclude the value of the tax-breaks of biotechs when it was being allocated (at least that is the assumption running through the APC cost-effectiveness calculations).

The APC, since it is differentially more targeted at large pharmaceutical firms over small biotechs and not-for-profits firms, makes this problem worse where it exists[309]. It is also an ironic strategy to pitch towards large pharmaceutical firms, if the reason for low vaccine research is, in some cases, in part generated by a 'replacement effect' induced by an over-reliance on large pharmaceutical firms.

'Replacement effect' crowding out effect reduces APC cost-effectiveness

It may affect how we measure the cost-effectiveness of the APC if there are replacement effects in the system. There is what might be called a 'replacement effect' crowding-out effect working against the APC. The APC has to be set sufficiently high that the marginal positive return on vaccine research minus the marginal negative return caused by the 'replacement effect' produces an overall return that equals that on all other research projects that the firm engages in. And this crowding out effect is worse if the APC concentrates incentives even more in a few large pharmaceuticals firms and leads to a tightening of IPR in ways that make research more difficult and expensive for small firms[310].

'Replacement effect' crowding in effect boosts alternative

If there is a 'replacement effect', it is not clear why an APC would be preferred over alternative finance mechanisms that more directly tackle the 'replacement effect' – for example, mechanisms that feed finance more directly towards biotechs and not-for-

[308] This observation affects other features of APCs including the auction mechanism and other strategic behaviours that drive up the APC price.
[309] All of this section is under *ceteris paribus* assumptions, since clearly the APC could be set so high that these problems become insignificant.
[310] Observe that this refers to the 'crowding out' effect, not the overall effect, of a dollar of government finance.

profit firms, enabling them to take projects further without needing to rely on large pharmaceutical firms, and measures that generally create more of a competitive industry with ease of entry and greater numbers of firms, and an IPR system that better works to allow firms to freely acquire technology that might undermine those firms experiencing (and causing) a 'replacement effect'. If there is a 'replacement effect' at work, there is what might be termed a 'replacement effect' 'crowding-in' effect boosting the effectiveness of these alternatives[311].

It may be that this 'replacement effect' 'crowding-in' effect can even be boosted further. The flip-side to the notion that overall (expected, discounted) industry profits are lower if vaccines are developed in areas with large 'replacement effects', is that large institutions who might otherwise spend heavily on treatment programs, like the World Bank and the WHO, would be better off. That this fact does not automatically lead vaccine developers (and their financiers) to reason that it is in their interests to develop vaccines *even if* they replace treatments, is at least in some part down to the previous under-purchase and under-use of vaccines by such institutions[312]. It is sometimes claimed that the simple purchase of currently-available vaccines (and, indeed, acts that enable their usage) by these institutions has little effect on vaccine research incentives[313]. However, once the 'replacement effect' is recognised, the 'demonstration effect' of the purchase of current vaccines is stronger. Quite literally, the purchase of current vaccines in part unlocks the credit constraints (i.e. makes finance cheaper) of biotechs and not-for-profits, and others by 'demonstrating' that the 'replacement effect' is now weaker. This also indicates a 'demonstration effect' from investments into healthcare infrastructure too[314]. With a 'replacement effect' present, a stimulus package including expenditure on previous vaccines and on health infrastructure might have the added externality benefit of 'crowding in' some privately-financed vaccine R&D[315]. This stimulus package would be strengthened further if finance mechanisms were set to give differentially greater impact to biotechs, not-for-profits, and all those working on 'replacement' projects, rather than on those suffering from and, indeed, creating the 'replacement effect'.

Clearly, this would alter the APC cost-comparison figures too.

---

[311] Looked at another way, it is *cheaper* to use other modes of support targeted at small biotechs/not-for-profit, etc, since they do not have to contain this extra cost.

[312] This indicates that part of the problem may refer to the lack of healthcare infrastructure, and again emphasises one of the arguments of this paper that the APC price would need to be set higher as much on account of the lack of infrastructure as on account of the 'lack of a market'. The hepatitis B vaccine and the Hib vaccine discussed above are cases in point. After 13 years of being largely unavailable, even though the hepatitis B vaccine is supposedly now generally available, 40% of children in Sub-Saharan Africa still do not receive it. After 11 years of being largely unavailable, Hib vaccine usage even when supposedly generally available is heavily skewed towards rich countries, with only tiny percentages of coverage in poor countries. Millions of children do not get a yellow fever vaccine costing cents to manufacture.

[313] K7:46.

[314] The unwinding of the 'replacement effect' boosts the marginal impact of investment in infrastructure.

[315] It is not clear what the size of the effect might be, and the effect will be reduced somewhat by the fact that investment on vaccine R&D could well be a ten year plus program, followed by returns over a further ten years, with an average time to repayment of maybe fifteen years. And developers may still worry about the commitment of large institutions to such programs.

12.5. APCs for Vaccines Would be Mostly Capital Costs

No assessment has been made in the calculations so far presented of the likely private capital cost component of APCs. This should reflect all the usual risks, but also all the multiple new risks forced on pharmaceutical firms by the APC mechanism itself[316]. This would have to be priced into an APC. Failure to price it *all* in would slow vaccine research intensity or even halt it altogether.

Kremer[317] uses a 4% private real discount rate. He claims that "the cost of capital may be lower for the government than for pharmaceutical firms, but the difference is not that large." This is extremely wishful thinking. The required rates of return in the VC industry are high. Kremer must therefore be presuming a very large rôle for large pharmaceutical firms and not for the sort of companies that rely on VC. But the problem then is that the large pharmaceutical firms claim much higher costs of capital than Kremer presumes. The recent Tufts calculations of the average cost to develop a new drug derives an 11% real cost of capital, a nominal rate in the region of 15%. Many in the industry claim that this is an underestimate. Yet others claim it is too high[318]. The costs of a mechanism that achieves its maximal effect when it targets all of its incentives towards large pharmaceutical firms cannot be modelled on the basis of capital costs that are *not* derived from the capital costs of that industry. Given the lengths of periods at issue, even one or two percent more discounting has significant impact on the overall costs of R&D. Being out by a factor of nearly three is a very significant issue.

What are the levels of capital costs in a typical APC?

In addition, it is not clear how many extra percent of capital costs would need be added to take account of all the new layers of risk to capital that would form part and parcel of any realistic APC. And these new risks, being non-idiosyncratic, would not be easily diversifiable. As a very ball-park, totally unscientific figure, one could easily imagine the 15% nominal required rate of return becoming more like 20% or quite likely higher. Given the lengths of periods of compounding, this would have major impact on the overall costs of vaccine research and the needed APC price. Kremer's calculations[319] use the figure that it will take about ten years to develop a vaccine (though it is not clear where this figure comes from), with costs recouped over ten years via the APC (though, if less than perfect vaccines are allowed, this is unclear). So, relevant sales are 10-20 years off, with an average of 15 years of compounding of private capital costs. Assuming 2% inflation he adjusts the $250m figure to what it would be 'worth' in future years. The problem is that if the Tufts real figure of 11% is used, plus extra adjustments are included to account for the risks emanating from the APC mechanism itself, this severely underestimates capital costs of a typical APC, and sends the APC price much higher compared to that derived from the discounting method used by Kremer.

So far, no figures have been presented revealing the breakdown of the proportion of capital costs and out of pocket costs in a typical APC. The Tufts study of the cost of

[316] Neither has there been any assessment of how all the manipulation of information might feed into these costs.
[317] K4:7.
[318] For an assessment, see Farlow (2005) 'New Estimates of Drug Development Costs: An Evaluation'.
[319] K4:49.

bringing a new drug to market[320], calculated that of the $802m (in 2000 dollars) overall total cost, the out-of-pocket costs came to US$403m and the capital costs to $399m[321]. With all the layers of extra risk generated by the APC mechanism itself, the use of an APC to stimulate the process of vaccine R&D would easily generate the situation that the *majority* of the total cost of development of a new vaccine via an APC would be the costs of capital.

It is not clear that the difference between the use of finance underlying the APC and the use of other forms of finance would be small. This is a figure that needs to be derived rather than asserted. Neither is it clear that the capital costs involved using APCs would not be *much* higher than the capital costs of using an R&D Treaty or something like the UK's International Financial Facility plus open collaborative research.

[320] DiMasi, J, et al (2003).
[321] K7:32 refers to the earlier DiMasi study (1991), recognising that roughly half is capital costs.

13. CONCLUSIONS

It is claimed that "Our modelling exercise suggests that this approach... is the most cost-effective way to substantially increase research and development. This is because it is a cost-effective way to increase R&D and because it addresses access."[322] However, although the figures produced in favour of the APC look sophisticated, they are only as good as the assumptions and data chosen at the start.

We have seen that the assumptions have all been chosen deliberately to favour the case for the APC. In particular all of the potential problems of APCs have been screened from the modelling so that the cost-effectiveness figures revealed are always drawn from the lower bound of possibilities. And, then, anyway, the calibration of the model has been done on the basis of what even Kremer describes as "rule of thumb" figures provided by large pharmaceutical firms, while pointing out that "it would be a mistake to attach even a moderate degree of precision to these estimates."[323] Such figures incorporate many of the distortions currently present in pharmaceutical markets generally – in a sense conditioning them away – so that comparisons with alternative mechanisms that might be able to remove some of these distortions, are further biased. On a more reasonable modelling of the way vaccine research takes place, and of the way an APC would work, it will cost a great deal more to stimulate vaccine research via an APC than suggested. In particular, most of the costs of an APC would be just to cover the capital costs of large pharmaceutical firms.

All mechanisms for generating incentives for vaccine R&D involve distortions and imperfections. It really is not sensible to adjudicate between options by setting up an extremely distorted version of all but the favoured choice and then pretending to get a fair and realistic comparison. Hopefully by now the reader will have realised that the APC is not as straightforward as a superficial first look might have lead one to believe, and that there are sound theoretical reasons for doubting that a perfect APC could every exist. It is this realistic, imperfect, APC against which realistic, imperfect, alternatives should be judged. Once the many layers of imperfections and extra costs are factored in, and many of the biases against alternatives removed, it is very likely that APCs would actually cost a great deal more than some of the alternative mechanisms.

13.1. Alternative Incentive Mechanisms and the Rôle of Information

It is useful to view the alternative finance methods as different incentive mechanisms, based on different, but potentially overlapping, sets of information, with timing of information revelation an important factor in efficiency. The cost of various methods of vaccine finance is related to the way they handle information.

The fundamental stance of APC proponents is that payments based on *ex ante* information create much weaker incentives to develop vaccines than payments based on '*ex post* ' information, and that therefore the former route leads to more costly

[322] K2
[323] K4:17. Indeed any calibration exercise is very sensitive to the *initial* conditions.

drugs. As Kremer puts it: "information asymmetries between funders and researchers may hamper programs which fund researchers in advance." However this result is based on very simple APCs where the 'path' of technology has been so severely restricted that most information problems have been removed anyway. In other words, these models do not show how these information problems are handled. They just presume they get handled and can thus be just ignored. And, much more damming, time and time again we find that it is based on a contradictory attitude towards asymmetric information and the timing of the revelation of information: After so many sections analysing the potential distortions of an APC mechanism, we realise that to set the terms of an APC in order to avoid these distortions would itself require very good quality assessments over information *in advance*. Kremer's calculations are all based on the naïve assumption that this asymmetry of information simply disappears with an end-loaded program. But many of the decisions about the end-loaded program *require front-loaded information*.

Real-world APCs and more collaborative research methods would both involve monitoring and control issues. It would simply happen at different points in the process. Each creates a different set of potential distortions that need to be calculated and compared. And both struggle with the problem of motivating agents. Both involve institutions and 'committees', with their concomitant problems, and it is not clear, given the great number of such committees needed to make an APC work, that more open and collaborative approaches do not actually involve fewer.

There is a good case to be made for the notion that some forms of more open and more front-loaded mechanisms may be better ways to handle complicated technology compared to end-loaded mechanisms like APCs. In reality an open approach would be less dependent on *ex ante* expectations of information, would not have to form many of the elaborate calculations we have come across in the application of APCs, would probably be a great deal more informationally efficient, would avoid some of the strategic problems discussed, would not need to contain large 'option price elements', would economise on capital costs, and would not need to make elaborate adjustments *ex post*, all the time desperately trying to keep behaviour efficient while knowing full well that all kinds of new dynamic inconsistency problems are being created.

We also saw the importance of the interaction of the push non-APC with the pull APC. This needs to be explored much more. It may be difficult to contractually and practically synchronise and police optimal adjustments in the former in response to the changes that affect the value of the latter. This forces extra risks on to the latter, runs the risk of 'double-paying' for research, creates distortions in vaccine and drug research generally, and raises the global public costs of achieving vaccines via this route. Open collaborative routes might be a more adaptable, and result in lower costs.

One other interesting discovery on closer inspection of APCs, is that for all their claimed 'free market' credentials they end up relying on surprisingly many layers of institutions, committees, regulators and the like, and a high degree of efficient information processing and discretion by those running the mechanism, to *try* to make it operate efficiently. Yet the promoters of APCs make little public reference to how such institutions might work, the interplay of responsibilities of institutions at different levels, the manipulation of information, and the dangers of institutional failure or capture. Since, in essence, the criticisms of non-APC models are based on

the notion of institutional failure, it is rather surprising to find so little effort devoted to these issues in the case of APCs (explained perhaps in part by the way Kremer's technology removes institutions from the picture anyway). And it rather numbs the criticism that alternatives involve 'institutions'.

We also discovered the hugely important nature of capital costs in judging between finance mechanisms, and concluded that capital costs can only be *much* higher in APC-based models in response to the layers of uncertainties, risks, and potential inefficiencies of the mechanism. We went as far as to suggest that most of the cost of an APC might turn out to be capital costs. Kremer's simplifications deliberately exclude these extra capital costs.

13.2. An Ideologically-Driven Model?

Unfortunately, debate about the potential difficulties and disadvantages of APCs has been largely stifled. The promoters of APCs have played their part in this. By now – having read in this paper so many of the potential 'disadvantages' of APCs – it might come as a surprise to the reader to discover that in a comparison of some 19 push and pull policy instruments, while room is made for pages of disadvantages for other mechanisms (many of which are perfectly valid), Kremer has not even inserted a 'disadvantages' section for the APC[324]. Either this is because there truly aren't any potential disadvantages (and the last 80,000 words have been a complete waste of time), or it is simply a reflection of the fact that all the publicly-available[325] modelling of APCs so far has simply deliberately ignored them. Kremer easily picks up on any distortion created by non-APC mechanisms[326], but he repeatedly fails to discuss whole swathes of possible distortions and problems that arise in response to, and are created by, the APC mechanism itself. Why is this?

To bore the reader by repeating a quote from above, in his executive summary Kremer explains the power of his findings thus – and *only* thus: "The cost-effectiveness of government R&D is limited by the potential of crowding out private R&D, difficulties in picking winners among competing research projects, potential politicization of funding decisions, and difficulties in shutting down unpromising research projects" We saw that this is justified on highly-selective grounds. Since the argument in favour of APCs rests almost entirely on this one quote, it should at least worry policymakers enough to make them request that this is backed up with more evidence and that at least a few of the potential 'disadvantages' of the APC itself are considered.

One is left with the unfortunate conclusion that, although the figures produced in favour of the APC and against other mechanisms of finance for vaccines are made to

[324] K 2:2.

[325] This allows for the possibility that these things are being modelled outside of the public domain. But if so, this raises the issue of why this is being done in such a closed fashion given the serious public policy issue at hand.

[326] For example, he correctly observes that when giving tax-breaks, HIV researchers will claim credits for working on clades prevalent in developing countries but then work on clades prevalent in developed countries, and those working on malaria will claim for work on a malaria vaccine for developing countries but then work on a vaccine that would only be suitable for travellers.

look sophisticated, the methods used to derive them are crude and mostly ideologically driven.

Why is the standard of evidence set so much lower for mechanisms like the APC than for alternatives, like, for example, open collaborative research (for which the standard of evidence is set so high that we do not even have a public debate about it)? Could those arguing for more of a front-loaded mechanism, get away with so little of *its* mechanism clarified after six years of arguing in its favour? And would it not be expected that there would be a strong onus of proof on the promoters of more open collaborative research methods to prove (openly, in the public domain) that it would be safe against all possible distortions?

Several times supporters of APCs have argued that since the social value of vaccines is so very high, throwing a great deal of money at an APC is a perfectly reasonable way to proceed. It is hard to imagine that the supporters of any approach other than APCs would be able to get away with this logic. It is patently ironic (and hypocritical) to argue for the inefficiency of other mechanisms compared to the APC, not provide any convincing proof of this, and then to wastefully throw everything at the APC approach anyway. At the very least, throwing the entire social surplus of a vaccine at an R&D approach removes an important potential disciplining device on costs and the behaviour of players. At the worst it is dangerous, with serious ramifications for other vaccines and drugs for neglected diseases generally.

13.3 The Political Appeal of APCs – Even Though They are Expensive

APCs seem to cost nothing, yet they can make politicians or policymakers look as if they are doing something to tackle the vaccine problem. It doesn't require them to have anything as radical as a 'research strategy' or a 'set of priorities'. They abrogate all of that to the 'markets' and to 'big pharmaceutical' firms. It is perhaps somewhat ironic – given all the talk of government/political failure at the heart of many mechanisms – that this failure might itself happen at the level of choosing the APC mechanism itself, and even be encouraged by the divisors of the mechanism.

It isn't clear why the supporters of APCs should be quite so happy when governments and others agree to support the APC idea without asking many extremely basic questions. And, it is even more baffling, given all their concerns about failure, that the models and figures presented to governments should be so biased, with little interest in stimulating a debate about the potential disadvantages (indeed, according to the literature accompanying the proposal, there are none).

13.4. The Need for a Broader Debate, and Why IPR Issues Cannot be Avoided

Setting up APCs has some of the characteristics of a 'poison pill' about it[327]. Once various APCs are in place for vaccines of various diseases, any attempt to modify the system away from patents (towards more open collaborative science for example) triggers the poison-pill. In particular, even small movements towards a more open and collaborative system might (to the extent it is believed) reduce investment in the R&D that is being motivated by APCs – unless some commitment could be made that these particular vaccines (and any subsequent vaccines building on their discovery) could be allowed to operate under the old closed APC system (and this would need to be known and credible in advance at all points of the move to more open methods). This would increase the cost of deriving vaccines via APCs, and immediately cause pressure to be applied to reverse the movement towards more open approaches. The advance purchase commitments are therefore a neat extra pressure to entrench the patents system and prevent alternatives, including open collaborative research, from gaining favour.

Why would policymakers promote a system that would enforce even more strongly patents and a closed science approach, without analysing the ramifications of this, especially on research that benefits from a more open science approach? Since, contrary to the arguments of some APC promoters, the APC is not transparently more efficient than more open and collaborative research and certainly not the panacea suggested, there is no obvious reason for closing down the debate about alternatives like open collaborative research and an R&D Treaty, as some governments and lobbyists seem intent on doing. Indeed, since it is absolutely key to the efficient workings of the APC that tight IPR has *no* deleterious consequences for the cost of research, even those favouring APCs should be keen to explore the impact of tighter IPR before initiating any APCs.

Good policy never comes out of uncritically accepting one solution while uncritically ruling out all others. The WHO, WIPO, the World Bank and others need to be free to explore all methods. It makes a lot more sense to engage in debate before starting any grand experiment. Compared to the potential cost of mistakes, a broader debate would cost practically nothing.

Ultimately, APCs are just a reflection of our poor sense of priorities. Our interest in them, and our refusal to consider alternatives, ultimately stems from our unwillingness to put resources into neglected diseases and vaccines for the poor. It's time to broaden our attention to include other policy responses.

[327] A poison pill is a device written into a contract that only bites in certain states of the world. For example, a firm might write penalty terms into a contract that only bite when, say, the contract is terminated, or a manger may have terms in his/her contract that generate millions of dollars of payments when he/she is fired – so reducing the incentive of the firm to fire them.

BIBLIOGRAPHY

DRAFT VERSION. INCOMPLETE

Aghion, P. and Bolton, P. (1992) "An Incomplete Contracts Approach to Financial Contracting", Review of Economic Studies, 77, 338-402.

Anton and Yao (1998).

Bhattacharya and Ritter (1983).

Blanc (1999).

Blumenthal D et al. NEJM (1996) 334: 368-73.

Benowitz S. (1996) "Progress Impeded?" The Scientist; 10[7]:1.

Branstetter, L., and Sakakibara, M. (2000) "When Do Research Consortia Work Well and Why? Evidence from Japanese Panel Data", National Bureau of Economic Research Working Paper No. 7972.

Bunk, S., (1999) "Researchers feel threatened by disease gene patents." The Scientist; 14[20]:7;

David, P.A., Hall, B.H., and Toole, A.A. (2000) "Is Public R&D a Complement or a Substitute for Private R&D? A Review of the Econometric Evidence" Research Policy, 29, 497-530.

Dasgupta, P, Stiglitz, J. (1980) "Industrial Structure and the Nature of Innovative Activity". Economic Journal; 266[90];279;

Dasgupta, P, Stiglitz, J. (1980) "Uncertainty, Industrial Structure, and the Speed of R&D", Bell Journal of Economics; 1[11]:3;

Dewatripont, M, and Tirole, J. (1994) "A Theory of Debt and Equity: Diversity of Securities and Manager-Shareholder Congruence" Quarterly Journal of Economics, 109, 1207-1054.

Diamond (1998).

DiMasi, J. A. Ronald W. Hansen, Henry G. Grabowski, (2003) "The price of innovation: new estimates of drug development costs," Journal of Health Economics 22, 151-185.

Farlow, A.W.K., (2003) "Costs of Monopoly Pricing Under Patent Protection.", http://www.earthinstitute.columbia.edu/cgsd/accesstomedicines_papers.html

Freundlich, N. (1998) "Will increasingly aggressive licensing terms on research tool patents hurt basis research?" Signals June 4.

Gangi (1999).

General Accounting Office. (2003) "Technology Transfer: NIH-Private Sector Partnership in the Development of Taxol", General Accounting Office, Washington D.C.

The Global Alliance for TB Drug Development, (2001) "The Economics of TB Drug Development,"

Granstrand (1999).

B.H. Hall, (2002) Oxford Review of Economic Policy. Vol. 18. No. 1.

Hall, BH, and Zeidonis, R.H. (2001) "The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry, 1979-1995", RAND Journal of Economics, 32,(1), 101-128.

Heller M.A. and Eisenberg R.S. (1998) "Can patents deter innovations? The Anticommons in biomedical research", Science; 280 95364): 698-701;

Hirschleifer, J, Riley, J.G. (1979) "The analytics of uncertainty and information", Journal of Economic Literature; 1375 [17]:1404;

Jaffe (2002) Oxford Review of Economic Policy. Vol. 18. No. 1.

Jaffe, A. lerner, J, and Stern, S. (Eds) Innovation Policy and the Economy, Cambridge, MA, MIT Press.
Kettler, (2000).
Klemperer, P. (2002) "What Really Matters in Auction Design" Journal of Economic Perspectives, Vol 16, Number 1, 169-189.
Klette, T.J. Møen, J. and Griliches, Z. (2000) "Do Subsidies to Commercial R&D Reduce Market Failures? Microeconomic Evaluation Studies", Research Policy, 29, 471-496.
Kremer, M. (2002) Journal of Economics Perspectives, Vol 16, Number 4, Fall, 67-
Lach, S. (2000) "Do R&D Subsidies Stimulate or Displace Private R&D? Evidence from Israel", National Bureau of Economic Research Working Paper No. 7943.
Lanjouw and Cockburn (2001)
Leland and Pyle (1977)
Lerner (1995)
Lerner, J. (1999) "The government as Venture Capitalist: The Long-run Effects of the SBIR Program", Journal of Business, 72, 285-319.
Love, J. (2003) Evidence Regarding Research and Development Investments in Innovative and Non-Innovative Medicines", Consumer Project on Technology, Washington D.C.
Mansfield, E. et al. (1977) *The Production and Application of New Industrial Technology* (New York, NY: W.W Noton & Company.
Mercer Management Consultancy, (2000) "HIV Vaccine Industry Study October-December 1998" World Bank Task Force on Accelerating the Development of an HIV/AIDS Vaccine for Developing Countries", March. http://www.iaen.org/files.cgi/103_batson.pdf
Nadiri, M. I. "Innovations and Technological Spillovers", National Bureau of Economic Research, Cambridge, MA) Working Paper no. W4423, 1993,
Neumann et. al. (2000).
Nichols, (1994) "Scientific Management at Merck: An Interview with CFO Judy Lewent" Harvard Business Review,.
Rivette and Klein (2000).
Rosenberg S.A. (1996) "Secrecy in Medical Research", NEJM: 334; 392-4;
Rowberg, R.E., (2001) "Pharmaceutical Research and Development: A Description and Analysis of the Process", Congressional Research Service, Report for Congress,
Scherer (1998)
Scherer and Watal (2002)
Scotchmer, S. (1999) *Cumulative Innovation in Theory and Practice*, University of California, Berkeley, CA,
Scotchmer, S., (2004) "Innovation and Incentives", The MIT Press.
Shapiro, C. (2001) "Navigating the Patent Thicket: Cross licenses, Patent Pools and Standard Setting" in OX73
Solsted, J. (2004) Presentation at MSF conference, Malaysia, February.
Stiglizt, J. (2003) "The Roaring Nineties – Seeds of Destruction" Allen Lane, Penguin, p 251
Spivack, R.N. (2001) "The Economic Evaluation of Technological Change", Conference Proceedings of the Advanced Technology Program, Washington, DC, National Institute of Standards and Technology.
Tandon P. (1983) "Rivalry and the Excessive Allocation of Resources to Research", Bell Journal of Economics, 152[14].
Wallsten (2000). Regulation v23 n4 2000, "the R&D Boondoggle

Whitehead, P (1999) “Public sector vaccine procurement approaches: A discussion paper prepared for the Global Alliance for Vaccines and Immunization.”
Williamson, O.E. (1988) “Corporate Finance and Corporate Governance”, Journal of Finance, 43, 567-91,
World Bank, (1993) World Development Report (p 8, 64, 68).
1993 World Development Report (p 8, 64, 68).
Varian, H. R., (2002) “Tales of Manipulation and Design Flaws from the Crypt of Auction History”. New York Times; New York, N.Y.; Aug 1.