Incentivizing innovation: Adding to the toolkit[1]

Michael Kremer[2]
Harvard University
Brookings Institution
Center for Global Development
National Bureau of Economic Research

Heidi Williams[3]
Harvard University

This draft: 19 March 2009

Abstract: The expectation of receiving Intellectual Property Rights (IPR) such as patents creates incentives for research, but implementation imposes static efficiency losses and other costs. In this paper, we discuss recent proposals and present a case for incremental experimentation with other mechanisms for rewarding innovation, with an eye towards testing and refining these mechanisms. Some of these other mechanisms aim to promote innovation as well as access to new technologies conditional on their development. Voluntary mechanisms will involve lower risks of undermining expectations that research will be rewarded than mandatory mechanisms. Prizes, such as those recently offered by the X-Prize Foundation, have been successful in spurring research, but typically for demonstration projects rather than for innovations capable of being used at scale. To the extent that it is desirable to use reward systems to directly spur the creation of products for widespread use, in many contexts the design of prizes could be usefully extended by conditioning rewards on a market test, as in the recent $1.5 billion pilot Advance Market Commitment (AMC) for a pneumococcus vaccine. Experimentation with this and other AMCs can likely inform the design of other mechanisms incorporating more complex *ex post* reward triggers, such as the proposed Medical Innovation Prize Fund. We discuss AMC design issues in more detail, focusing on design differences between AMCs for technologically closer products (such as a pneumococcus vaccine) and technologically more distant products (such as a malaria, tuberculosis, or HIV vaccine). With time, these and other mechanisms could be added to the toolkit for encouraging innovation.

---

[1] Portions of this article draw on a previous paper written by the authors for the German Marshall Fund of the United States ("Promoting Innovation to Solve Global Challenges: Opportunities for R&D in Agriculture, Climate Change, and Health," available at http://www.gmfus.org); those portions of this paper are adapted with permission from the German Marshall Fund of the United States. Michael Kremer served on a Center for Global Development (CGD) working group on Advance Market Commitments, funded by the Bill & Melinda Gates Foundation, as well as on the Economic Expert Group for the pilot Advance Market Commitment for a pneumococcus vaccine. He is very grateful for discussions with members of both groups on the economic issues surrounding Advance Market Commitments which informed the substance of this paper – particularly Owen Barder, Jonathan Levin, Ruth Levine, and Christopher Snyder.
[2] Gates Professor of Developing Societies, Department of Economics, Harvard University; mkremer@fas.harvard.edu.
[3] PhD student, Department of Economics, Harvard University; hlwill@fas.harvard.edu.

I. Introduction

The expectation of receiving Intellectual Property Rights (IPR) such as patents creates incentives for research and development (R&D). However, patents impose static efficiency losses and other costs. In this paper, we present a case for incremental experimentation with other mechanisms for rewarding innovation, with an eye towards testing and refining these mechanisms.

We first discuss some of the trade-offs arising under patent systems, and briefly sketch the key ideas behind three other mechanisms proposed as ways to encourage innovation while also addressing some of the concerns arising under patent systems – prizes, Advance Market Commitments, and the proposed Medical Innovation Prize Fund (Section II). We discuss how some of these other mechanisms are motivated by a desire to provide incentives for innovation but also promote access to new technologies conditional on their development (Section III). We then review the distinction between voluntary and mandatory mechanisms for rewarding innovation (Section IV). Because the current system of innovation relies on firms anticipating that they will receive IPR, we argue that voluntary mechanisms will involve lower risks of undermining expectations that research will be rewarded than mandatory mechanisms. Prizes, such as those recently offered by the X-Prize Foundation, are a natural starting point. Historically, however, most prizes have been offered for demonstration projects, not for innovations capable of being used at scale. We describe (in Section V) various triggers for reward payments that can be used, and argue that in many contexts the design of prizes could be usefully extended by conditioning rewards on a market test. Advance Market Commitments (AMCs), such as the recent $1.5 billion pilot AMC for a

pneumococcus vaccine, rely on one type of market test. Experimentation with this pilot and other AMCs can likely inform the design of other mechanisms incorporating more complex *ex post* reward triggers, such as the proposed Medical Innovation Prize Fund.

We then focus on AMC design issues in more detail (in Section VI), discussing how design details differ between technologically closer products (such as a pneumococcus vaccine) and technologically more distant products (such as a malaria, tuberculosis, or HIV vaccine). Section VII concludes.

It is worth noting that it took time for institutions such as the patent system or the peer review process to evolve into their current forms: these institutions, which today we think of as integral to supporting our systems of innovation, required time as well as trial-and-error to develop. Since the first US Patent Act was put in place in 1790, rules have developed on what is allowed to be patented, who is allowed to file patents, for how long patents should be held, *etc*. Likewise, the peer review system has made tremendous progress over time. Weller (2001) discusses how prior to World War II, editors frequently made all decisions themselves with only informal advice from colleagues, and that only recently has the paradigmatic 'editor plus two referees' system become widespread (Rowland 2002). Work by individuals such as Vannevar Bush, who lobbied for the evaluation of scientific research by scientists, not government officials, led to the establishment of the modern system of federally-supported peer-review institutions for decision-making on federal funding for scientific research in the United States.

We argue in this paper for an incremental approach to experimentation with other mechanisms for rewarding innovation, and offer one proposal for how best to move ahead. The issues involved with designing mechanisms for rewarding innovation are

complex, and on a first attempt any given idea may likely not work well. But the potential payoffs to adding new mechanisms to our toolkit for encouraging innovation are immense, and thoughtful experimentation with several potential mechanisms would be valuable.

II. Patents and other mechanisms for rewarding innovation

One broad class of mechanisms for encouraging innovation commits to reward successful products. The historical record provides several sources of evidence supporting the idea underlying such mechanisms, which is that innovation responds to incentives (Griliches 1957, Schmookler 1966, Hayami and Ruttan 1971, Acemoglu and Linn 2004, Finkelstein 2004, Brunt *et al.* 2008).

Intellectual Property Rights (IPR) systems such as patents are one such mechanism for rewarding innovation. The core idea of a patent is to provide incentives for innovation by allowing developers of new products a period of market exclusivity. A positive feature of this institution is that the rewards provided to firms under the patent system correspond at least in a rough sense to how desirable the product is to consumers, thus providing firms incentives to focus on producing products with characteristics valued by consumers. Note that this design also implies that the patent system is in some sense fragile as it is dependent on firms' expectations about future rewards.

Patents involve a trade-off between the benefit of providing dynamic incentives for innovation and the cost of providing these incentives in a way that imposes static distortions – that is, because patents make goods more expensive to consumers, at the

margin some goods will not be used even when their social value exceeds the cost of production.

Other trade-offs arise with patents as well. For example, patents may potentially deter downstream innovations in contexts where innovation is cumulative, in the sense that many innovators will build on prior developments and discoveries. Similar issues can arise in contexts where there are complementarities across innovations in a broader sense; for more discussion see, for example, Merges and Nelson (1990), Scotchmer (1991), Heller and Eisenberg (1998), Murray and Stern (2007), and Bessen and Maskin (forthcoming).

In recent years, several authors have argued that the patent system is in some sense a highly uncertain reward structure. Jaffe and Lerner (2004), for example, argue that firms engaged in R&D under the patent system face substantial risks of as-yet unknown or untested patents being asserted against them, claims which must be subject to decisions made by the judicial system.[4]

In part reflecting such concerns, in recent years there has been a resurgence of interest in both alternatives and supplements to the current system of IPR, motivated by the idea that alternative types of mechanisms for rewarding innovation may better mitigate some of the trade-offs arising under the patent system. In this section, we briefly sketch the key ideas behind each of three other mechanisms which have recently attracted attention – prizes (II.1), Advance Market Commitments (II.2), and the proposed Medical Innovation Prize Fund (II.3).

[4] Lemley and Shapiro (2005) argue that in this sense patents do not confer upon their owners a right to market exclusivity, but rather confer a right to *try* to exclude others by asserting the patent in court.

II.1. Prizes

Prizes are typically designed as voluntary mechanisms which lay out a set of technical specifications in advance and have a committee which decides *ex post* whether those technical specifications have been met. A series of such prizes called X-Prizes has recently been set out by the X-Prize Foundation. For example, the first X-Prize offered a $10 million dollar prize for the first non-government organization to launch a reusable manned spacecraft into space twice within two weeks. This prize was awarded in 2004 to a team led by aircraft designer Burt Rutan and financed by Microsoft co-founder Paul Allen. Similar subsequent X-prizes were later announced – including the Archon X-Prize in 2006, the Automotive X-Prize in 2006, and the Google Lunar X-Prize in 2007.

II.2. Advance Market Commitments

Advance Market Commitments (AMCs) are another type of voluntary mechanism. We here focus on the application of AMCs to the case of vaccines for neglected diseases concentrated in poor countries, because policy discussions of AMCs have largely focused on that application. For additional discussions of AMCs, see Kremer (2001a, b), Kremer and Glennerster (2004), and Barder *et al*. (2005).

Kremer and Glennerster (2004) argue that despite the enormous potential health benefits of new vaccines appropriate for use in poor countries, there are too few incentives in place for private firms to pursue research and development (R&D) on these vaccines. Moreover, conditional on vaccine development market failures prevent quick adoption and diffusion of these vaccines to poor countries – the historical record suggests there is typically a ten to fifteen year lag between the introduction of vaccines in rich

countries and their widespread use in poor countries. As we discuss more below, AMCs for early-stage vaccines – such as malaria – attempt to address both issues, whereas AMCs for later-stage vaccines – such as pneumococcus, the focus of the pilot AMC mentioned in the introduction – are relatively more focused on the latter issue.

The basic design of an AMC is similar for early- and late-stage vaccines. Under an AMC, one or more credible sponsors (that is, sponsors that are financially solvent and thought to be unlikely to renege on a commitment) legally commit – in advance of product development and licensure – to underwrite a guaranteed price for a maximum number of pre-defined purchases of the vaccine. Vaccines are eligible if a committee deems they fulfill a set of technical specifications laid out in advance, detailing vaccine characteristics such as the maximum number of needed doses, the required level of efficacy in a certain population, *etc.* Sponsors guarantee to provide a top-up payment (say, $15 per treatment) conditional on poor countries expressing demand for a given product and them (or other qualified purchasers, on their behalf) paying a low, relatively affordable price (say, $1 per treatment). The higher, guaranteed price provides an economic return for developers of the product, and in exchange these developers agree to a cap on the long-run price that they charge for the product (or agree to license the technology to other manufacturers). If no suitable product is developed, no AMC payments would be made.

A number of governments – Italy, the United Kingdom, Canada, Norway, and Russia – together with the Bill and Melinda Gates Foundation recently announced a $1.5 billion pilot AMC for a pneumococcus vaccine suitable for children in the developing world. UK Prime Minister Gordon Brown has suggested that this be the first in a series

of AMCs to encourage the development of vaccines against diseases affecting the developing world.

II.3. Medical Innovation Prize Funds

The Medical Innovation Prize Fund proposal, such as that described by Love (2005), is for a mandatory mechanism.[5] Rather than rewarding new products through market exclusivity, generic companies would be allowed to freely compete and developers of new products would be financially rewarded through payments from the Medical Innovation Prize Fund.

The size of the fund in Love (2005) is proposed to be fixed at 0.5 percent of US GDP, although the size of the fund could be larger or smaller depending on the specific proposal. Love (2005) notes that some of the fund could be earmarked for priority projects, such as neglected diseases and orphan drugs. The payments to participating firms are proposed to be paid out over ten years and based on the incremental health benefits of new products.

A version of this proposal was introduced in the US House of Representatives in 2005 by Representative Bernard Sanders of Vermont, in HR 417 – The Medical Innovation Prize Act of 2005.

III. Promoting innovation and access

One broad difference across the three types of mechanisms discussed above is that prizes have historically tended to (although need not) focus on providing incentives for

[5] Aidan Hollis, Thomas Pogge, and others have advocated a Health Impact Fund proposal, which is similar to a voluntary version of the proposed Medical Innovation Prize Fund discussed above.

innovation and not on promoting access to new technologies conditional on their development, whereas both the proposed Medical Innovation Prize Fund and the AMC proposal aim to decouple these goals and to promote both.[6]

Access can be promoted through a variety of mechanisms – including patent buyouts, in which a patent is purchased and placed in the public domain[7]; requirements for patents to be placed in the public domain; or requirements that a product be sold at a particular price. Note that options which place patents in the public domain would also address the issue of potentially deterred cumulative or otherwise complementary innovations discussed above, whereas pricing rules would not.

The placement of patents in the public domain, allowing for generic competition, may often be an attractive way to promote access, but this option will not always be feasible. In the case of vaccines, for example, because production technologies are very specialized the original firm may effectively be the only feasible producer even if the patent for the vaccine is placed in the public domain. Thus, while the proposed Medical Innovation Prize Fund addresses access through promoting generic competition, the

---

[6] While we focus on the three types of innovation mechanisms discussed above, it is worth noting that recent years have seen a number of other proposals for new mechanisms to address the issue of access within the current IPR system. Lanjouw (2001), for example, proposes that pharmaceutical companies be required to choose whether to protect their drug patent in rich countries or poor countries, but not in both. For diseases affecting both rich and poor countries, companies would presumably choose to protect their patent in rich countries, thus allowing low-cost generic copies to be sold in poor countries. For neglected diseases which only have markets in poor countries, companies would presumably choose patent protection in poor countries. Terry Fisher has proposed a requirement where each pharmaceutical firm must achieve an aggregate ratio (disability adjusted life years saved, or DALYs, divided by gross revenues) above some threshold, arguing that companies could comply with such a requirement by increasing R&D for neglected diseases or by expanding access to existing drugs in poor countries. He proposes allowing DALY vouchers to be bought and sold across companies. Other proposals to encourage innovation on medical technologies for neglected diseases include the priority review voucher system advocated by Henry Grabowski and others.

[7] Kremer (1998) proposes a price mechanism for patent buyouts in which the private value of patents is determined using an auction.

AMC proposal for vaccines instead addresses access in the short-run through the topped-up payments, and in the long-run through the capped price.

IV. Mandatory versus voluntary institutions

Public policies to provide incentives for innovation can either be mandatory – thus being alternative to the IPR system that is currently in place – or be voluntary – thus supplementing the IPR system that is currently in place. The Medical Innovation Prize Fund proposal is a mandatory mechanism, whereas both X-Prizes and Advance Market Commitments (AMCs) are voluntary mechanisms.

Mandatory programs such as the proposed Medical Innovation Prize Fund would be alternative to the current IPR system. Whether the incentives provided by such an alternative system would be higher or lower than the level of incentives provided by the current IPR system would be a function of the design of the prize fund proposal. Voluntary programs such as X-Prizes and AMCs, on the other hand, would supplement the current IPR system, and thus at least weakly increase the total available incentives for R&D - since if the price in a voluntary program was set low enough such that firms would realize lower revenue if they chose to participate than they would realize if they chose not to participate, presumably firms would select out of participating in the voluntary mechanism.

Because, as argued above, the current system of IPR depends on firms' expectations of future rewards, experimenting with voluntary mechanisms involves lower risks than mandatory mechanisms. If an experiment with a voluntary mechanism shows promise, it can be refined and applied in a broader range of settings; if it fails, the

voluntary mechanism can either be revised or abandoned. In contrast, if an experiment with a mandatory mechanism fails, it may shake the confidence of R&D investors, who may be concerned that IPR will disappear and that no adequate alternative incentives will take its place to reward them for their investments. Mandatory mechanisms for encouraging innovation cannot be costlessly turned on and off because of the dynamic element inherent in any market in which firms make long-term R&D investments.

V. Triggers for reward payments[8]

Ideally, mechanisms for rewarding innovation would credibly commit to reward appropriate innovations, while not committing sponsors to pay for innovations which end up not being useful or desirable. The design of triggers for reward payments needs to balance these objectives.

In this section we discuss three potential elements of a system for triggering reward payments: fulfillment of technical specifications set *ex ante* (Section V.1), *ex post* discretion (Section V.2), and measures of *ex post* use, willingness-to-pay, or impact (Section V.3). As we discuss below, most mechanisms will use a combination of two or three of these triggers.

V.1. *Ex ante* technical specifications

The relative weight put on fulfilling a pre-specified set of technical specifications varies across mechanisms. At one end of the continuum are mathematics prizes, such as the Wolfskehl prize which was established in 1908 to reward the first person to prove

[8] Many of the examples in this section are drawn from Kremer and Glennerster (2004).

Fermat's Last Theorem. For such mathematics prizes, sponsors can very clearly describe in advance what they are looking for. Putting relatively higher weight on fulfillment of technical specifications is likely also appropriate in contexts in which sponsors are primarily looking to spur a demonstration project as opposed to spurring the development of a product that will see widespread use. For example, a series of prizes established in 1959 by Henry Kremer sought to encourage innovation in human-powered flight through offering prizes for demonstration projects including the first human-powered aircraft to fly a figure eight.

If the aim is more ambitious, not just to reward proofs of mathematical theorems or demonstration projects but instead to reward applied innovations that will see widespread use, it will be difficult to rely as heavily on *ex ante* technical specifications. Moreover, in some cases – such as for the Post-It Note or the GUI (Graphical User Interface) technology – sponsor likely could not have described the product specifically enough in advance for this type of reward trigger to be useful. For many types of technologies, choosing to solely use fulfillment of technical specifications as a basis for reward payments (rather than combining this reward trigger with another type of mechanism) through attempting to write down completely exhaustive technical specifications *ex ante* might result in projects specifications that are either too tight (so that many useful products would fail to satisfy the specifications) or too loose (thus rewarding a product that is not useful or desirable to consumers).

V.2. *Ex post* discretion

When *ex ante* technical specifications are used as a reward trigger they will almost always need to be combined with some sort of committee to make an *ex post* decision about whether the technical specifications have been met. More generally, essentially any mechanism for rewarding innovation will involve some sort of *ex post* discretion.[9] The relevant questions are instead how much *ex post* discretion is allowed, and to whom *ex post* discretion is allocated.

The issue of how much *ex post* discretion to allow can be thought of along a continuum. For example, a committee given a relatively high amount of discretion is used to award the Nobel Peace Prize, whereas the committee which awards the Nobel Prize for Chemistry has more limited discretion given the bounds of the field within which the prize must be relevant. The committee for the Wolfskehl prize mentioned above had even less discretion.

Ex post discretion is usually, although not always, allocated to some type of committee – as in the Nobel Prize examples discussed above. One broad issue which can arise in the context of rewards based on *ex post* discretion by committee is that a committee may have incentives to reward based on different criteria *ex ante* relative to *ex post*. For example, *ex ante* the committee may want to reward innovation, but *ex post* the committee may prefer to reward the individuals who have made the most substantial contribution to pure science, or prefer to reward those individuals who might make the best use of the prize money going forward.

Committees may also use *ex post* discretion to "raise the bar." As an example, consider a £20,000 prize offered in the 1700s by the British government. The prize

[9] For example, in the case of the patent system *ex post* discretion is essentially left in the hands of the legal system (judges and jurors).

aimed to spur the development of a solution to the "Longitude Problem" – that is, of allowing ships to determine their longitude while at sea. The Board of Longitude expected astronomers and mathematicians to develop a solution through celestial observations of the positions and motions of heavenly bodies, but in fact the solution was developed by a clockmaker named John Harrison. Harrison developed a timepiece that was sufficiently accurate to determine time at the port of departure even on rolling ships. It took twelve years of tests to prove the worth of Harrison's invention to the committee and to reward him with his prize. In her popular book on the subject, Sobel (1996) argued that these delays were unnecessary; others have argued that the Board of Longitude was justified in requiring these tests.

One way to address the potential concern of committees "raising the bar" *ex post* is to require that the committee award a certain amount of money within a given time frame. Such a requirement is often used in architectural contests, where a committee must choose a winner to award a given contract to by a specified deadline. In architectural contests, committees are relatively certain to receive a sufficient number of high quality entrants such that choosing the best entrant will likely not result in a poor outcome. Thus, although a payment would have to be made no matter what, the risk that the committee will have to award the contract to an undesirable proposal is low.

However, in other contexts this may be more of a concern. With very challenging technological goals, such as the development of an HIV vaccine, the probability that no firm would have a high quality vaccine available at a given deadline is likely much higher. Moreover, in markets with a small number of firms the firms could potentially collude to slow innovation.

The proposed Medical Innovation Prize Fund uses a version of this type of requirement, but for several reasons this type of requirement is less concerning in that context. For the proposed Medical Innovation Prize Fund, there is a commitment to award a certain amount of money each year, but the scope of coverage is great enough – covering virtually all pharmaceutical research – to smooth out variation in the arrival of eligible products, and to minimize opportunities for collusion across firms.

Both prizes and AMCs rely on a combination of *ex ante* technical specifications and *ex post* discretion by committee to trigger rewards, with the committee's role primarily being one of determining whether the technical specifications have been met.

V.3. Metrics of *ex post* use, willingness to pay, or impact

One issue with basing reward payments solely on technical specifications set *ex ante* is that products may be developed which in a strict sense meet the technical specification but for some reason are not desirable to consumers. The types of aviation prizes discussed above, for example, were primarily intended to provide incentives for demonstration projects – not for the production of commercially useable products. Although demonstration projects may be the explicit goal of some mechanisms, for mechanisms which aim to spur the development of products desirable to consumers it may often be useful to base reward payments at least in part on some measure of *ex post* valuation of the product by consumers.

Allocating reward payments based in part on a measure of *ex post* use is one way of leaving *ex post* discretion relatively more in the hands of consumers instead of in the

hands of a committee. This can limit the amount of discretion given to a committee, and can also help address concerns of time inconsistency problems or political capture.

Although basing reward payments in part on some measure of *ex post* use can be useful in many contexts, measuring *ex post* use may not always be straightforward. At least two potential costs could arise with relying on measures of *ex post* use: static costs of rent seeking, and dynamic losses from inappropriate incentives. Ideally the measure would be objective, and something that is difficult for participating firms to manipulate. In some markets a natural metric that is easy to measure will be available, and in some markets relying on this type of reward trigger may be more difficult or even infeasible.

The major design difference between prizes and AMCs is that AMCs base reward payments in part on a measure of *ex post* use whereas prizes typically do not. Specifically, under AMCs the reward to the company is not paid simply for developing a product that meets a set of technical specifications, but rather is tied to actual adoption and use of that product. The practical implementation of this relies on the fact that vaccines used in poor countries are largely purchased through a centralized system (that is, through UNICEF) through which use of vaccines can be tracked relatively easily. Basing AMC payments in part on this measure of *ex post* use provides incentives for companies to focus their R&D efforts on products that would actually be used, rather than focusing on producing a product which somehow fits a set of pre-decided technical specifications but is not a good fit with what developing countries need or want.

Medical Innovation Prize Funds propose basing reward payments on the measured incremental health benefits. In practice, measuring the incremental health benefits of many products may be difficult. AMCs use a metric of use as a reward

trigger, but a metric of *usefulness* requires measuring the benefit per person – which may be quite challenging in many contexts. AMCs for neglected vaccines are a relatively easy case because what one would theoretically want the reward value to correspond to is the total social value of the vaccine. For vaccines that did not previously exist, the total social value is equal to the number of users multiplied by the benefit per user, the latter of which can be thought of as efficacy of the vaccine multiplied by the expected burden of the disease to an individual. However, in most other cases a given product will be a substitute and/or a complement for some other products currently on the market, and may be effective for some patients but not others. Appropriately calculating the incremental health benefits of a new technology would require taking into account such potentially complex factors. Small-scale experimentation with various ways of valuing new products under such a mechanism would likely be valuable.

It is worth noting that not all types of mechanism "failures" will offer opportunities for learning on how to better design future mechanisms. If a mechanism "succeeds" in the sense that the desired product is developed, is nearly impossible to know whether a better-designed mechanism would have resulted in a higher quality product being developed. If a mechanism "fails" in the sense that the desired product is not developed, it is nearly impossible to know whether no mechanism – no matter how well designed – would have successfully produced the product (say, because the scientific challenges are too great). If a low quality product is developed which ends up being able to collect reward funds, this type of failure would likely offer a learning opportunity for how to better design future mechanisms. Likewise if a mechanism is put in place and a high quality product is developed which ends up not being able to collect

reward funds. Although the specific reasons for failure of a given mechanism in any context may sometimes be obvious and sometimes be debatable, some learning will almost assuredly be possible with the failure of any given experimentation.

VI. Design issues for early- and late- stage Advance Market Commitments

We argued above that although prizes, such as those recently offered by the X-Prize Foundation, are a natural starting point for encouraging innovation, that in many contexts the design of prizes could be usefully extended by conditioning rewards on a market test. Because AMCs are a voluntary mechanism that incorporates one type of market test, experimentation with AMCs can likely inform the design of other mechanisms incorporating more complex *ex post* reward triggers, such as the proposed Medical Innovation Prize Fund. In this section, we discuss AMC design issues in more detail, focusing on how AMC design details differ between technologically closer products (such as a pneumococcus vaccine) and technologically more distant products (such as a malaria, tuberculosis, or HIV vaccine).

In an early policy report on AMCs by the Center for Global Development in Washington DC (Barder *et al.* 2005), it was argued that AMCs could likely be applied cost-effectively to both technologically closer products and to technologically more distant products. While the basic economic logic behind an AMC is clear, the general framework leaves substantial latitude in terms of specific program design. In particular, there are several substantive differences between how an AMC should be designed for technologically closer products relative to technologically more distant products, two of which we discuss below: the appropriate price provided to developers (Section VI.1), and

the appropriate role for demand guarantees (Section VI.2). For more discussion of these issues, see Kremer *et al.* (2008).

It is worth giving a brief background on pneumococcal diseases, which are the focus of the $1.5 billion pilot AMC. Although not as well-known as malaria or HIV, pneumococcal diseases kill more than 1.6 million people annually, including up to one million children under age five. In rich countries, child deaths from pneumococcus are rare, but in poor countries pneumococcus is a leading cause of child mortality. Pneumococcal vaccines for adults have existed for some time but it is important to protect children as well, both because of the high death toll among children and because children are important in spreading the disease. A pneumococcal vaccine which protects children against some strains of bacteria has been available in the US for several years. However, the cost per dose of pneumococcus vaccine in the US is approximately $65, far out of reach of poor countries. In addition, existing versions of the vaccine are optimized to cover the strains of the diseases common in the US, and do not provide protection against some key strains common in poor countries. Two pneumococcus vaccines were recently licensed (from different suppliers), after the announcement of the $1.5 billion pilot AMC although before the legal details of the contract were fully in place.

VI.1. Setting prices under an AMC

A first difference between designing AMCs for technologically closer and technologically more distant products arises in thinking about how vaccine prices under an AMC should be set. Somewhat ironically, many aspects of AMC design are likely easier for technologically more distant products. For a technologically more distant

vaccine like HIV, policy makers' goal is to design an AMC that will attract a socially efficient amount of research effort to search for the vaccine. Setting the price paid under an AMC in this case is not so much about guessing what firms' production costs will be in the future as it is a question of determining the price at which a new vaccine would be cost-effective relative to other health expenditures.

On the other hand, for a technologically closer vaccine like pneumococcus, much of the R&D has already been completed, and the challenge is primarily one of designing a long-term procurement contract that will incentivize a small number of specific firms that have the necessary expertise to construct the large-scale capacity needed to serve the world's poorest countries as well as the rich- and middle-income world. If policy makers knew how much it would take to get the one or two specific firms that can currently produce childhood pneumococcus vaccine or are likely to be able to do so in the near future to participate in an AMC, they would set the AMC price at that level but no higher. Unfortunately, knowing this price is difficult and there are risks associated with either setting the price too high or with setting the price too low. The risk of setting the price too high is that firms will profit too much from the AMC mechanism. The risk of setting the price too low is that the status quo will continue – that is, that firms won't build the capacity needed to serve poor country markets in the near term, and we will face the historically typical ten to fifteen year lag between the introduction of vaccines in rich countries and their widespread use in poor countries, potentially resulting in the loss of millions of lives of children in poor countries.

Conceptually, many factors are relevant in thinking about what price would induce firms to participate in the pneumococcus AMC mechanism. At first blush, a price

equal to the cost of production may seem reasonable, but data on production costs is proprietary. In addition, it is important to realize that firms face risks from participating and not participating in the AMC that are likely difficult to quantify. On one hand, firms may realize public relations benefits from selling a product which addresses a major health need of individuals in poor countries. On the other hand, firms also face potential public relations costs from participating in an AMC to the extent that price discrimination across different markets is politically difficult. Specifically, firms may fear that if they sell the vaccine at a low price in the poorest countries (like Mozambique), that governments in middle income countries like Brazil will demand lower prices as well. This could set up a trickle up of lower prices which could put a serious dent in firms' revenues from middle- and high-income country market sales that represent a different order of magnitude of potential revenue relative to sales under an AMC.

Such political economy concerns area likely very salient to firms. For example, after Senator Paula Hawkins of Florida asked a major vaccine manufacturer how it could justify charging nearly three times as much to the US government for vaccines as to foreign countries, US manufacturers stopped submitting bids to UNICEF to supply vaccines (Mitchell *et al.* 1993; US Congress, Senate 1982). When President Bill Clinton announced his plan to immunize all children against a standard list of diseases in 1993, he said, "I cannot believe that anyone seriously believes that America should manufacture vaccines for the world, *sell them cheaper in foreign countries*, and immunize fewer kids as a percentage of the population than any nation in this hemisphere but Bolivia and Haiti" (Mitchell *et al*. 1993, italics added). In the face of such statements, potential public relations risks facing firms seem real.

When weighing the risks from setting the price either too high or too low, it is also worth comparing the costs from a given price to the expected benefits. Compared to the $65 cost per dose of pneumococcus vaccine in the US, the currently planned long-run price of $3.50 is low. The $3.50 price is also low compared to the estimated value for the spending in terms of lives saved, as Sinha *et al.* (2007) estimated that at a price of $5 or lower the AMC would cost less than $100 per year of life saved. (For comparison, rich countries typically spend up to $50,000 or even $100,000 to save a statistical life year, and even in poor countries saving a life year at $100 is considered very cost-effective.[10])

VI.2. Role of demand guarantees

A second difference arises over whether donors should guarantee some portion of demand. A general principal of contracting or mechanism design is that whoever is best placed to affect a risk should, all else equal, bear that particular risk. For earlier stage products, firms still have opportunities to affect product characteristics and thus should bear more risk – implying demand guarantees would be less appropriate. For later stage products, the situation is quite different. Once a product has already reached the stage where pneumococcus vaccines currently are, the donor community has more opportunity to influence demand – implying demand guarantees may be beneficial.

For a technologically close product, like a pneumococcus vaccine, it is fairly clear what a product will look like, and the main problem is to incentivize capacity construction. Firms will be more inclined to build capacity if they know they will be able

[10] In the US (Neumann *et al*. 2000) and the UK (Towse 2002; Devlin and Parkin 2001), the cost-effectiveness threshold for medical interventions is estimated to be $50,000 to $100,000 per life-year saved. In poor countries, health interventions are generally considered extremely cost-effective if the cost per life-year saved is either less than $100 (World Bank 1993) or less than a country's annual per capita income (GAVI 2004).

to sell an amount which will utilize that capacity, and donors may thus be able to get away with a slightly lower price if they guarantee demand. On the other hand, it would not make sense to guarantee demand for a vaccine that is still very technologically distant, since otherwise a firm might wind up creating a vaccine that complies with a list of technical specifications, but that no countries would want, and donors might wind up having to buy the vaccine. For technologically distant products, donors to AMCs arguably should condition payments on countries being willing to use the product and some buyer being willing to make a modest co-payment (as proposed above), so as to create incentives for firms to develop vaccines developing countries will want. Once a particular product is developed, and the problem shifts to one of capacity construction, donors could then move into a phase in which they would guarantee a portion of demand. AMCs could also specifically be linked to capacity installation by firms.

VII. Conclusion

Technological progress is a key determinant of economic growth, and finding ways to improve institutions to encourage technological progress could potentially do more to encourage economic growth than virtually any other area of public policy. While the patent system offers one mechanism for rewarding innovation, it involves some important trade-offs. In this paper, we presented a case for incremental experimentation with other mechanisms for rewarding innovation, with an eye towards testing and refining new mechanisms to encourage R&D. The potential payoffs to adding new mechanisms to our toolkit for encouraging innovation are immense, and thoughtful experimentation with several potential mechanisms would be valuable.

References

Acemoglu, Daron and Joshua Linn (2004) "Market size in innovation: Theory and evidence from the pharmaceutical industry," *Quarterly Journal of Economics* 119(3): 1049-1090.

Barder, Owen, Michael Kremer, and Ruth Levine (2005) *Making Markets for Vaccines: Ideas to Action*, Center for Global Development: Washington DC.

Bessen, James and Eric Maskin (forthcoming) "Sequential innovation, patents, and imitation," *RAND Journal of Economics.*

Brunt, Liam, Josh Lerner, and Tom Nicholas (2008) "Inducement prizes and innovation," mimeo, available: https://nber15.nber.org/c/2008/si2008/DAE/lerner.pdf.

Devlin, Nancy and David Parkin (2001) "Does NICE have a cost effectiveness threshold and what other factors influence its decisions? A Discrete Choice Analysis," mimeo, City University London.

Finkelstein, Amy (2004) "Static and dynamic effects of health policy: Evidence from the vaccine industry," *Quarterly Journal of Economics* 119(2): 527-564.

Global Alliance for Vaccines and Immunization (GAVI, 2004) "Health, immunization, and economic growth, Research briefing #2, Vaccines are cost-effective: A summary of recent research," author.

Griliches, Zvi (1957) "Hybrid corn: An exploration in the economics of technological change," *Econometrica* 25(4): 501-22.

Hayami, Yujiro and Vernon Ruttan (1971) *Agricultural Development: An International Perspective* Johns Hopkins University Press: Baltimore MD.

Heller, Michael and Rebecca Eisenberg (1998) "Can patents deter innovation? The anticommons in biomedical research," *Science* 280: 698-701.

Jaffe, Adam and Josh Lerner (2004) *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What To Do About It* Princeton University Press: Princeton NJ.

Kremer, Michael (1998) "Patent buy-outs: A mechanism for encouraging innovation," *Quarterly Journal of Economics* 113(4): 1137-1167.

Kremer, Michael (2001a) "Creating markets for new vaccines, Part I: Rationale," in Adam Jaffe, Josh Lerner, and Scott Stern (editors), *Innovation Policy and the Economy, Volume I* MIT Press: Cambridge MA.

Kremer, Michael (2001b) "Creating markets for new vaccines, Part II: Design issues," in Adam Jaffe, Josh Lerner, and Scott Stern (editors), *Innovation Policy and the Economy, Volume I* MIT Press: Cambridge MA.

Kremer, Michael and Rachel Glennerster (2004) *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*, Princeton University Press: Princeton NJ.

Kremer, Michael, Jonathan Levin, and Christopher Snyder (2008) "Designing Advance Market Commitments for new vaccines," mimeo, Harvard University.

Lanjouw, Jean (2001) "A patent policy proposal for global diseases," mimeo, Yale University.

Lemley, Mark and Carl Shapiro (2005) "Probabalistic patents," *Journal of Economic Perspectives* 19(2): 75-98.

Love, James (2005) "Two ideas regarding innovation and access," Presentation for the World Health Organization (WHO) Commission on Intellectual Property Rights, Innovation, and Public Health Open Forum.

Merges, Robert and Richard Nelson (1990) "On the complex economics of patent scope," *Columbia Law Review* 90: 839-916.

Mitchell, Violaine, Nalini Philipose, and Jay Sanford (1993) *The Children's Vaccine Initiative: Achieving the Vision*, National Academies Press: Washington DC.

Murray, Fiona and Scott Stern (2007) "Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis," *Journal of Economic Behavior and Organization* 63(4): 648-687.

Neumann, Peter, Eileen Sandberg, Chaim Bell, Patricia Stone, and Richard Chapman (2000) "Are pharmaceuticals cost-effective? A review of the evidence," *Health Affairs* 19(2): 92-109.

Rowland, Fytton (2002) "The peer review process," *Learned Publishing* 15: 247-258.

Schmookler, Jacob (1966) *Innovation and Economic Growth* Harvard University Press: Cambridge MA.

Scotchmer, Suzanne (1991) "Standing on the shoulders of giants: Cumulative research and the patent law," *Journal of Economic Perspectives* 5(1): 29-41.

Sinha, Anushua, Orin Levine, Maria Knoll, Farzana Muhib, and Tracy Lieu (2007) "Cost-effectiveness of pneumococcal conjugate vaccination in the prevention of child mortality: an international economic analysis," *Lancet* 369(9559): 389-396.

Sobel, Dava (1996) *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time* Fourth Estate: London.

Towse, Adrian (2002) “What is NICE's threshold? An external view,” Chapter 2 in Nancy Devlin and Adrian Towse (editors), *Cost Effectiveness Thresholds: Economic and Ethical Issues*, Kings Fund / Office of Health Economics: London.

US Congress, Senate (1982) “Hearing to review federal and state expenditures for the purchase of children’s vaccines,” Subcommittee on Investigations and General Oversight, Committee on Labor and Human Resources, 22 July.

Weller, Ann (2001) *Editorial Peer Review: Its Strengths and Weaknesses*, American Society for Information Science and Technology: Silver Springs MD.

World Bank (1993) *Disease Control Priorities in Developing Countries* Oxford University Press for the World Bank: Washington DC.