MARKET OUTCOMES AND DYNAMIC PATENT BUYOUTS

Alberto Galasso
Matthew Mitchell
Gabor Virag

Market Outcomes and Dynamic Patent Buyouts
Alberto Galasso, Matthew Mitchell, and Gabor Virag
NBER Working Paper No. 20197
June 2014
JEL No. O30,O34

ABSTRACT

Patents are a useful but imperfect reward for innovation. In sectors like pharmaceuticals, where monopoly distortions seem particularly severe, there is growing international political pressure to identify alternatives to patents that could lower prices. Innovation prizes and other non-patent rewards are becoming more prevalent in government's innovation policy, and are also widely implemented by private philanthropists. In this paper we develop a model in which a patent buyout is effective, using information from market outcomes as a guide to the payment amount. We allow for the fact that sales may be manipulable by the innovator in search of the buyout payment, and show that in a wide variety of cases the optimal policy in our model still involves some form of patent buyout. The buyout uses two key pieces of information: market outcomes observed during the patent's life, and the competitive outcome after the patent is bought out. We show that such dynamic market information can be effective at determining both marginal and total willingness to pay of consumers in many important cases, and therefore can generate the right innovation incentives in our model.

Alberto Galasso
Joseph L. Rotman
School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
Canada
and NBER
alberto.galasso@rotman.utoronto.ca

Matthew Mitchell
Joseph L. Rotman
School of Management
University of Toronto
105 St. George St.
Toronto , ON M5S 3E6
Canada
Matthew.Mitchell@Rotman.Utoronto.Ca

Gabor Virag
Joseph L. Rotman
School of Management
University of Toronto
105 St. George St.
Toronto , ON M5S 3E6
Canada
Gabor.Virag@rotman.utoronto.ca

1 Introduction

Innovation is the main engine of economic growth, and the consensus among economists, beginning with Arrow (1962), is that the positive externalities from R&D imply under-investment relative to the socially optimal level. For example, a recent study by Bloom et al. (2013) estimates that the gross social rate of return to R&D substantially exceeds the private return, with the socially optimal R&D level being over twice as high as the currently observed R&D expenditure. A central policy question, therefore, is how one can best devise a mechanism that encourages innovation. This paper contributes to the recent literature focusing on designing prizes that infer demand from various market signals, and use that information to design a reward at least partially based on a cash prize. We develop a model in which social welfare is best served with a hybrid system that incorporates market power but then offers prizes in the form of a buyout of patent rights over time.

Mitigating patent rights as a way to improve welfare fits with current concern over the patent system. Concerns about the potentially deleterious effects of patents have been voiced in both academic and policy debates. In a prominent book, Jaffe and Lerner (2004) conclude that in the U.S. patents have become "*sand rather than lubricants in the wheels of American progress.*" Some have even gone so far as to recommend more draconian reductions in permissible patenting (Boldrin and Levine, 2008; Bessen and Maskin, 2009).

The buyout system replaces some of the rents that are obtained through monopoly rights with a prize. If the planner cannot discern the quality of innovations, pure prizes are difficult to implement because the value of the prize cannot be tied to the surplus generated by the innovation as the demand is unknown by the policy maker. Even if the planner could get good information about the number of units sold at a given price, for instance by observing units sold under perfect competition, this is insufficient to construct the inframarginal values of consumers, which is essential to estimating the full value of an innovation. In our model, the hybrid mechanisms can generate greater innovation incentives and increase social welfare if the policy maker has less than perfect information about market demand.

Our approach addresses the issue of needing to estimate inframarginal values. In order

to accomplish this we stress the dynamic approach to innovation rewards, since one point on the demand curve will generally be insufficient for reconstructing demand. In contrast with the previous literature, we assume that the policy maker can learn over time about market conditions by observing price and quantity realizations that arise from the choice of the innovator and the underlying demand function. As information about the market demand is revealed, the reward mechanism that maximizes social welfare may change according to the revealed information. Eventually the policy-maker can resort to allowing perfect competition, which generates additional information about the demand for the innovation.

We consider the possibility that innovations may generate a variety of different demand curves, and that the market signals coming from the sales of those products may be manipulable by the innovator. In all but the least-manipulable extensions of our model, the optimal policy begins with market power for the innovator, and gradually moves toward competitive pricing as information is generated by the experience of the innovation. We show that even in the *most* manipulable extension of the model, where the true price that gave rise to the observed sales can be completely obscured by the innovator, the optimal mechanism involves some reward through a contingent prize near the end of the period in which the innovator is rewarded. In our model the optimal policy is a sort of hybrid between a patent and a prize in the sense that it rewards innovators through prices above marginal cost initially, but then moves toward a reward that is focused on a cash prize and prices closer to, or reaching, marginal cost.

The results are of interest both to the design of intellectual property policy and private philanthropic organizations who use resources to reward innovators through prizes. From the policy maker's perspective, the difficulty introduced by limited information about cost and demand conditions provides a foundation for why policy makers may wish to grant temporal monopoly rights (i.e. patents) to innovators. It has long been argued that rewards through monopoly profits (as guaranteed by patents) are tightly connected to the surplus generated by the product, and therefore provide the appropriate incentives for an innovator who knows the demand much better than the policy maker. This idea was originally formulated by Mill (1848) who wrote that patents are an effective reward *"because the reward conferred by it depends upon the invention's being found useful, and the greater the usefulness, the greater the reward,"* and is at the heart of mechanism design approaches to patent policy like the one in Scotchmer (1999).

2

In the sense that the optimal policies in our model are a hybrid between a patent and a prize, this logic remains in effect; the planner simply tries to mitigate the social cost of market power through buyouts in some cases. The buyout is facilitated by information from data in much the same way that the impact of a merger on consumer surplus is assessed through estimation of an econometric model.

Our results are also of particular interest for philanthropists who have entered the business of rewarding innovators. Qualcomm and Nokia currently offer multi-million dollar prizes for the development of affordable devices that can recognize and measure personal health information. Similarly, the Gates Foundation has offered an innovation award to immunize children in the poorest parts of the world. In recent academic and policy debates, it has been recommended to link prize rewards to specific market outcomes. For example, the Center for Global Development advised that philanthropists willing to sponsor the development of a malaria vaccine pay the innovator 14 dollar for each of the first 200 million treatments sold at 1 dollar to the recipients (Glennerster, et al. 2006).[1] Our results suggest one approach to this philanthropy: use resources to buyout patents that have a track record of success.

We show the sense in which the details of this buyout approach are related to the philanthropists' and policy makers' concerns about the veracity of the market signals that are observed. In some cases, such as pharmaceuticals, quantity may be relatively well measured, but prices may be more opaque and companies have an incentive to manipulate their prices in order to obtain higher reimbursements through public funding.[2]

One can alternatively view our model as describing some policy choices for a regulator, antitrust or otherwise, who faces firms with monopoly granted through IP. The FDA already is involved in the administration of ex post rights for pharmaceuticals through the orange book program and the rights granted therein. Similarly, the Australian government offers

---

[1] Similar ideas have appeared in AgResult, an initiative launched by the governments of Australia, Canada, Italy, the United Kingdom, the United States, the Bill and Melinda Gates Foundation and the World Bank to mitigate R&D underinvestment in tropical agriculture. A key feature of the initiative is to focus on incentive schemes that link payments to demonstrated results.

[2] For example, in March 2001 the State of Wisconsin reached a $4.2 million settlement agreement with Merck, Schering and Warrick Pharmaceuticals in litigation charging the companies with defrauding the Wisconsin Medicaid Program. Wisconsin alleged that the pharmaceutical manufacturers manipulated wholesale prices information, knowing that Medicaid would rely on these prices to determine Medicaid reimbursement.

co-payments to selected drugs to mitigate monopoly distortions. More generally, Hovenkamp (2004) describes the sense in which antitrust policy might respond to growing IP protection.

We develop a model with discrete time and infinite horizon where the planner commits to a reward structure that depends on the history of prices and quantities realizations observed over time. The planner's problem in designing an appropriate prize is observing total benefit of the innovation. As in Kremer (1998), this requires information about the quality of the innovation; Weyl and Tirole (2012) point out that this problem is magnified by the need to discern the willingness to pay of non-marginal consumers. Our mechanism attacks both issues. Our first result is that, in the absence of demand manipulation, the first best can be approached arbitrarily closely in a large set of demand functions that includes those typically used in the industrial organization literature. We also show that the result generalizes to introducing noise in the demand and to allowing the demand to shift following a stochastic Markov process as in Battaglini (2005). An implication of this result is that, in the absence of demand manipulation, innovation can be efficiently rewarded without patents. This is because information can be extracted from a competitive market in which the product is sold at marginal cost. The policy maker can generate price variation taxing the firms and shifting their marginal costs of production. This market outcome information can then be exploited by the policy maker to implement the first best. Because the value of an innovation comes from both marginal and inframarginal consumers, monopoly power without price discrimination may deliver far less reward than the value of the innovation. In Kremer (1998), this is addressed by paying a fixed proportion above the monopoly value of the innovation. Here the planner directly addresses this issue by learning about inframarginal customers, and then can pay a reward that acts as if the monopolist had the opportunity to price discriminate.

Our findings also indicate that policy-makers may design an innovation reward systems exploiting structural demand estimation techniques. Typically, structural demand studies identify the primitives of a model from local price variation and exploit the estimated parameters to conduct out-of-sample welfare analysis (Figure 1 case A). In our contest, the policy-maker can request the innovator to generate price variation that will be used to identify the underlying demand curve of the technology and to compute a patent buy-out transfer that compensates the innovator for the surplus generated. By keeping the price variation concentrated around

4

the marginal cost of production, the policy maker can limit the loss of surplus associated with learning to a minimum (Figure 1 case B).
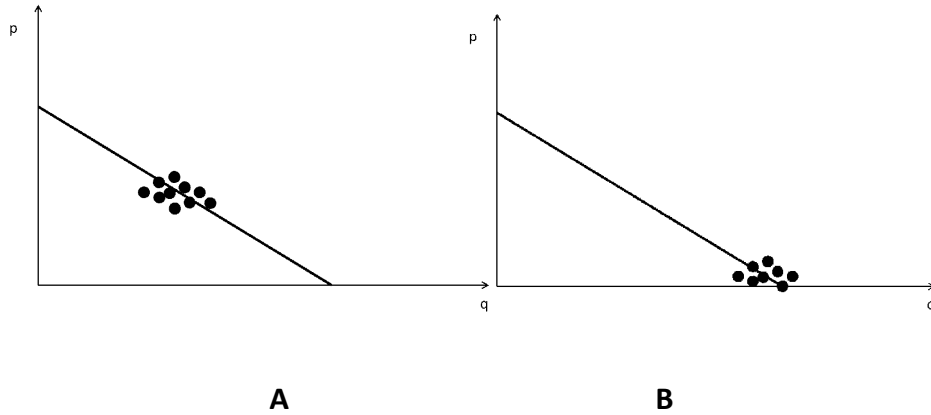

Figure 1: Market Outcomes and Demand Identification

We then investigate the case in which the innovator can manipulate demand. In keeping with the pharmaceutical price manipulation example, we assume that quantity is observable, while price may not be. We show that it is crucial to distinguish between the case in which demand manipulation is possible after the buyout takes place and the case in which post-buyout demand is non-manipulable. We show that pre-buyout manipulation, even if costless, may be ignored as long as manipulation after buyout it is not possible. This is because the planner can generate price variation after the buyout to learn the demand and to punish the innovator in the case of manipulation. This implies that market outcomes are relevant even after the buyout, because they are useful to detect and avoid manipulations.

The case in which the planner cannot generate price variation after the buyout is more complicated. We consider the case in which after the buyout the patent is sold in a competitive market and neither the planner nor the innovator can manipulate this outcome. We show that in this case, as long as pre-buyout manipulation is costly, the planner can construct a buyout scheme that generates the same R&D incentives as a patent and increases total welfare. Intuitively, the planner can induce the innovator to reveal the true monopoly profits by requiring a stream of pre-buyout outcomes that are too costly to manipulate.

Finally, we characterize the optimal mechanism in a model where price manipulation

is costless for the innovator. We show that even in this case the optimal mechanism differs substantially from a patent. In this extension of the model it is optimal for the planner to induce the innovator to produce quantities that are above the monopoly level and the output is larger for innovations generating lower surplus.

The paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the baseline model. Section 4 examines the optimal policy in a model without demand manipulation. Section 5 introduces costly demand manipulation. Section 6 studies the optimal mechanism in a model with costless demand manipulation. Section 7 summarizes and concludes. All the proofs are in Appendix 1.

2   Related Literature

This paper is connected to various strands of the literature on the economics of innovation. In an influential paper, Kremer (1998) suggests a buy-out mechanism, which is linked to an auction to incentivize research and maximize welfare. The role of the auction is to reveal information to the planner about the private value generated by the innovation. Innovation incentives are maximized because the planner would pay for the patent the private value times a fixed markup that compensates for the difference between social and private surplus. Consumer welfare is also maximized because the innovation would be placed in the public domain once acquired by the planner. An important assumption underlying the buyout scheme suggested by Kremer is that the competitors of the innovator know the value (and the cost) of the innovation and are willing to take part in the auction. In our model, we depart from this assumption and assume that only the innovator knows how valuable an innovation is.

Wright (1983) and Shavell and Van Ypersele (2001) provide a comparison of prizes and patents as mechanisms to incentivize innovation in a static framework. Scotchmer (1999) studies the optimal mechanism to reward innovation when the planner offers a menu of patents that differ in length and application fee. She shows that if market outcomes are not observed, then in the presence of asymmetric information on the cost and benefit of research, patent renewal mechanisms are optimal in the sense that in her model every incentive compatible and individually rational direct revelation mechanism can be implemented with a renewal mechanism. Cornelli and Schankerman (1999) characterize the optimal innovation mechanism in a model

with moral hazard and adverse selection where innovators have unobservable productivity parameters. As in Scotchmer (1999), the planner offers the innovator a menu of patents that differ in length and application fee. They show that in thier model the optimal patent scheme is typically differentiated and can be implemented through menu of patent renewals.

Hopenhayn and Mitchell (2001) and Hopenhayn, Llobet and Mitchell (2006) study the optimal patent design in a model where innovation is cumulative and each discovery is a building block to future innovations. Hopenhayn and Mitchell (2001) consider the case in which the quality of the idea is private information and there are two generations of the technology. They show that in their model to maximize innovation incentives, patents must vary in breadth, i.e. the policy maker needs to vary the set of products that at any given time may be prevented by the patent holder. Hopenhayn, Llobet and Mitchell (2006) study a dynamic model with multiple cumulative innovations and private information about the quality of ideas and R&D investments. They show that in such an environment the optimal mechanism is a patent buyout scheme where the innovator commits to a price ceiling at which he sells his rights to a future inventor.

We are aware of only two studies which consider observable market outcomes. The first one is Weyl and Tirole (2012) that studies the optimal reward structure in the presence of multidimensional heterogeneity and non-manipulable market outcomes. In a static model, they show that the optimal policy requires some market power but not full monopoly profits. Such a policy is similar to Mitchell and Moro (2006), who study a planner who trades off deadweight loss against over-transferring to a group that "loses" from elimination of the distortion generating deadweight loss. Our setup differs from Weyl and Tirole (2012) because we introduce dynamics and allow the innovator to manipulate market outcomes.

The second paper is by Chari, Golosov and Tsyvinski (2012) who compare prizes and patents when the planner can observe market signals over time. They develop a dynamic framework where both the innovator and his product market competitors (but not the planner) know the value of the innovation that is represented by a unidimensional parameter. Their main finding is that in their model patents are necessary if the innovator can manipulate market signals. Our model departs from their setting in a number of dimensions. First, we assume that only the innovator knows the value of his technology and we do not require the presence

of informed competitors. Second, we allow for multidimensional heterogeneity in innovation quality. Because in a multidimensional setting observing one market outcome is not enough to learn the entire demand curve, in our model the planner faces a nontrivial learning problem even under full (i.e. non-manipulated) observation of the market signals. This implies that, differently from Chari, Golosov and Tsyvinski (2012), in our model the planner finds it optimal to use information acquired over time in a truly dynamic way. Third, we do not restrict the planner to use either patents or prizes and we consider a large set of reward structures that depend on the quantity and prices practiced by the innovator. In particular, we allow for patents of different "strength" in which the price charged by the innovator differs from both the monopoly and the competitive prices.

3   The Model

Time is discrete and the horizon is infinite. Each innovation is characterized by an ex-ante cost of creating the innovation, $c \in C \subset \mathbb{R}_+$, and a demand function $q_t = D(p_t, \theta)$ that is twice continuously differentiable in the price variable. Let $\overline{p} > 0$ be the minimum price at which $D(\overline{p}, \theta) = 0$. We assume that $\theta \in \Theta$, a compact subset of $\mathbb{R}^N$, and that D is continuous in θ. To ensure the concavity of the static profit function, we assume that $D_p(p, \theta) < 0$ and $D_p(p, \theta) + p D_{pp}(p, \theta) < 0$ for each $p \geq 0$. The known marginal cost of production is normalized to zero. Demand and cost parameters, θ and c are private information for the innovator and are distributed according to a smooth probability density function $\psi(\theta, c)$ that is known by the planner.

We assume that the planner observes perfectly the quantities in each period but the innovator can manipulate the price observed by the planner. Specifically, he can make the planner observe $\widehat{p}_t$ by sustaining a cost $\phi(\widehat{p}_t, p_t)$ with $\phi(p_t, p_t) = 0$. Most of our analysis will focus on two polar cases: (i) no manipulation where

$$\phi(\widehat{p}_t, p_t) = \infty \text{ if } \widehat{p}_t \neq p_t$$

and (ii) costless manipulation where

$$\phi(\widehat{p}_t, p_t) = 0 \text{ for all } \widehat{p}_t.$$

8

Let us indicate with $h_t \in H_t$ the public history at time t, that can be defined recursively as $h_t = \{h_{t-1}, r_t\}$ where $r_t = (q_t, \widehat{p}_t)$ is the information revealed in period t and $h_0 = \phi$. Thus $H_t \in \mathbb{R}^{2t}_+$, the set of public histories at time t is the Cartesian product (t times) of the set of observable price quantity pairs.

Following the literature, we assume that the planner has full commitment power to any announced policy.[3] The planner designs a reward schedule that in each period transfers to the innovator a sum, $g_t(h_t)$, that depends on the history $h_t \in H_t$. The planner has also the option to set up a non-manipulable irreversible competitive market in period $T + 1$. The switching time may depend on the history, and can be infinite (i.e. switching to competition may never occur).

A strategy of the innovator is a sequence of pair of prices $(\widehat{p}_t, p_t)$ for each period t that satisfies the constraint that prices are set to zero after switching has occurred. Formally, a strategy for the innovator is an infinite sequence of price pairs $((\widehat{p}_t, p_t))_{t=1,2,3,\dots}$. Let $\alpha \in A$ denote any such generic strategy and A denote the set of all possible strategies.[4] The function $T(\alpha)$ captures the time period in which the planner's policy calls for a switch to the competitive and non-manipulable regime. This time is deterministic from the perspective of the innovator since it depends only on his strategy. In Appendix 2 we provide a recursive definition of $T(\alpha)$ given the planner's policy function.

Then the innovator's maximization problem is

$$\max_{\alpha \in A} \sum_{t=1}^{T(\alpha)} \delta^{t-1} \left(p_t D\left(p_t, \theta\right) + g_t(h_t) - \phi\left(\widehat{p}_t, p_t\right)\right) + \sum_{t=T(\alpha)+1}^{\infty} \delta^{t-1} g_t(h_t). \tag{1}$$

To simplify the notation, we leave the relationship between the switching time and the strategy of the innovator implicit and indicate $T(\alpha)$ as T in the remainder of the paper.

Let us indicate the optimal revealed and actual price for period t with $\widehat{p}_t^*(\theta)$ and $p_t^*(\theta)$ and with $h_t^*(\theta)$ the public history revealed by this optimal equilibrium play. Investment in innovation takes place if the net present value of the profits of the innovator (1) exceeds c. Let

---

[3] This is usuall justified by reputational concerns of the government. We discuss the case of limited commitment at the end of Section 4.1.

[4] Note, that the innovator's strategy is formed upon observing the planner's switching policy, and that the planner's switching policy affects the set of strategies for the innovator.

us indicate with $\Theta^*(c)$ the set of types for which this condition is satisfied.

The social surplus (net of manipulation costs) in the product market if the planner chooses functions $\{g_t\}_{t=1,2,3,\ldots}$ and T is equal to:

$$W(\theta) = \sum_{t=1}^{T} \delta^{t-1} \left[S(p_t^*(\theta), \theta) - \phi(\widehat{p}_t^*(\theta), p_t^*(\theta)) \right] + \sum_{t=T+1}^{\infty} \delta^{t-1} S(0, \theta)$$

with

$$S(p_t, \theta) = p_t D(p_t, \theta) + \int_{p_t}^{\infty} D(z, \theta)\, dz.$$

The social planner chooses functions g_t and T (taking the optimal strategy of the innovator as given) to maximize the expected total social welfare created by the innovation:

$$\max_{g_t, T} \int_c \int_{\theta \in \Theta^*(c)} [W(\theta) - c]\, \psi(\theta, c)\, dc\, d\theta.$$

The *first best* can now be defined formally: in the first best it holds that $p_t = 0$ for all $t \geq 1$, the innovator does not distort the observed price $(\widehat{p}_t = p_t)$, and the innovation is developed if and only if

$$c \leq \sum_{t=1}^{\infty} \delta^{t-1} S(0, \theta).$$

The first best can be easily implemented by the planner if θ is known. To do so, the planner transfers the entire surplus to the innovator if he observes the competitive quantity and punishes the innovator if a different quantity is observed (i.e. $g_t = S(0, \theta)$ if $q_t = D(0, \theta)$ and $g_t = -\infty$ if $q_t \neq D(0, \theta)$).

The functions g_t and T allow the planner to implement a number of different reward mechanisms. We provide some examples below.

Patents

When $g_t(h_t) = 0$ and $T = \varsigma$ the planner offers a ς-period patents that generates innovation incentives through product market profits. The setting also accommodates the payment of renewal fees. For example we can introduce a fee, f, to be paid at time $T_1 < \varsigma$, with expiration of the patent in the absence of payment:

$$g_t(h_t) = \begin{cases} -f & \text{if } t = T_1 \text{ and } \widehat{p}_{T_1} > 0 \\ 0 & \text{else} \end{cases}$$

$$T = \begin{cases} \varsigma & \text{if } \widehat{p}_{T_1} > 0 \\ T_1 & \text{else} \end{cases}.$$

10

Buyouts

The following specification

$$g_t(h_t) = \begin{cases} 0 & \text{if} \quad t < \varsigma \\ K & \text{if} \quad t = \varsigma \end{cases}$$

$$T = \varsigma$$

captures a simple buy-out scheme in which the planner commits to buy the patent after ς periods at a pre-specified amount K. The setting also allows to implement more general buyout mechanisms where transfer price K and acquisition time T may depend on observed market outcomes.

4 Optimal Mechanism in the Absence of Demand Manipulation

In this section we characterize the optimal mechanism when the government can dictate prices and the innovator cannot manipulate demand, i.e. $\phi(\widehat{p}_t, p_t) = \infty$ if $\widehat{p}_t \neq p_t$.

To develop the intuition, let us consider a simple setting where the demand is linear $q_t = \theta_1 - \theta_2 p_t$. In this simple environment the planner can identify the intercept of the demand by inducing a price equal to zero in the first period so that $q_1 = \theta_1$. In the second period he can induce $p_2 = \varepsilon > 0$ and identify θ_2 by inverting $q_2 = q_1 - \theta_2 \varepsilon$. This means that it takes only two periods for the planner to learn the demand function and the surplus generated by the innovation. Notice that the planner can set ε arbitrarily close to zero and minimize the deadweight loss generated by above marginal cost pricing. If the entire surplus generated by the innovation is transferred to the innovator, innovation incentives are set at the first best level.[5]

The above example suggests that transfers that depend on market outcomes can be powerful mechanisms to incentivize innovation. In the example the planner finds it optimal to use market information in a truly dynamic way that allows him to approximate the complete information (first best) solution. In particular, by conditioning rewards on quantities and prices, the planner can obtain the information required to trace-out the demand curve. Once

---

[5] A transfer that approximates the first best is $g_1(h_1) = 0$ for all h_1; $g_2(h_2) = S(\varepsilon) + S(0)/\delta$ if $h_2 = \{q_1, 0, q_2, \varepsilon\}$ and $g_2(h_2) = -\infty$ otherwise; $g_t(h_t) = S(0)$ for $t > 2$ if $h_t = \{q_1, 0, q_2, \varepsilon, q_3, 0, .., q_t, 0\}$ and $g_t(h_t) = -\infty$ otherwise.

the demand is known, the surplus generated by the innovation is transferred to the inventor to maximize his innovation incentives. In the linear case, the demand can be learned by observing only two data points: the quantity sold at marginal cost and the quantity sold at any strictly positive price. Exploiting this feature of the demand, the planner will learn the demand by inducing the innovator to sell at an arbitrarily small price. This makes the deadweight loss negligible and allows the planner to approximate the first best solution.

The result obtained in the simple linear setting suggests that in a dynamic model the planner can substantially improve welfare and innovation incentives relative to patent systems or other static multidimensional screening mechanisms as the one characterized by Weyl and Tirole (2012). We now turn to the question of how general the result is. We start with the definition of an analytic demand function.

Definition 1 *(Judd, 1998) A demand function $D(\cdot, \theta)$ is analytic on X if and only if for every $p \in X$ there is an r, and a sequence c_k such that whenever $|z - p| < r$:*

$$D(z, \theta) = \sum_{k=0}^{\infty} c_k (z - p)^k.$$

We generalize the result obtained for linear demands to analytic functions.

Proposition 1 *If $D(\cdot, \theta)$ is analytic on $[0, \overline{p}] \subset \mathbb{R}$, then the first best can be approached arbitrarily closely.*

Our proof builds on Aghion et al (1991) who show in the context of an uniformed decision maker that when a payoff function is analytic the approximate derivative at a single point can be used to estimate the global behavior of the function. We show that the demand function can be approximated by collecting price and quantity observations over a small neighborhood around a single price. These observations are used to approximate the derivatives of $D(\cdot, \theta)$ around that price and to learn about the global behavior of $D(\cdot, \theta)$.

By choosing a smaller and smaller neighborhood around $p = 0$, the planner minimizes the welfare losses associated with learning and increases the accuracy of the estimates of the derivatives of $D(\cdot, \theta)$. In the proof we show that exploiting a step-wise analytic continuation technique the planner can approach arbitrarily closely the first best even if $D(\cdot, \theta)$ can be expanded in a power series locally but not globally.

Proposition 1 substantially generalizes the result for linear demands. Polynomials, exponentials, logarithms, power functions and a number of other demand functions that are typically used in applied theory are analytic functions. Fox and Ghandi (2011) show how analyticity of the market demand is a property of various well know demand models used for structural estimation as the linear random coefficients model, the almost ideal demand system of Deaton and Muellbauer (1980) and the mixed logit of Berry, Levinsohn and Pakes (1995).

4.1 Implementation

The above result suggests that variation in prices and quantities may provide useful information for a planner who aims to maximize welfare by providing innovation incentives and minimizing distortions in the product market. For a large class of demand functions, we have shown that a policy maker can learn the surplus generated by the innovation and minimize market distortions by generating a price variation that is close to the marginal cost of production. In our model, this allows the planner to implement an outcome arbitrarily close to the first best.

The most intuitive way to generate this price variation is by awarding the innovator a patent that confers him the exclusive right to sell the product and to commit to a *patent buyout* scheme whose reward depends on the observed market outcomes. In other words, the planner can dictate to the patentee a price path and commit to buy-out the patent if the innovator follows the path with a reward that depends on the quantities sold. The computation of the reward resembles structural estimation studies which typically estimate the primitives of a model from local price variation and exploit these estimates for out-of-sample welfare analysis. An implication of our result is that policy-makers may affect innovation incentives by designing reward systems that exploit these techniques.

In the context of the malaria vaccine, the Center for Global Development proposes to reward the innovator with a prize if 200 million treatments are sold at 1 dollar to the recipients. The suggested prize is 2.8 billion (14 dollars per treatment). A possible concern with such scheme is that vaccine development may inefficiently not take place if such reward is too small compared to the social welfare generated by the vaccine. An implication of Proposition 1 is that this prize scheme can be improved by requesting the successful innovator to sell the 200 million treatments at different prices, even if the overall price level remains close to the 1

dollar benchmark. This is because the market outcomes generated by such price variation will allow the sponsor to obtain an estimate of the product market-surplus generated by the new vaccine. Such estimate will provide useful guidance in determining the reward and avoiding under-payment (or overpayment) for the innovation.

But buyouts are not the only way to implement the first best. An alternative approach is to start from a perfectly competitive market in which the product is sold at marginal cost. The price variation can then be generated by the planner taxing the firms and shifting their marginal costs of production. In such model, the information generated in this way will be the same as the one generated by the buyout scheme and can be exploited by the planner to implement the first best. An implication of this alternative implementation method is that market power is not essential to solve the asymmetric information problem between the policy maker and the innovator. In other words, in the context of our model for a large class of demand functions the socially optimal innovation level can be reached through minor perturbations of a competitive market.

The mechanism proposed in this Section assumes that the government is able to commit to truthfully revealing the observed price-quantity pair, and all the market participants agree with the revelation. Suppose instead that the government can freely manipulate the observed price, and thus can decide how much the innovator is paid. This case is similar to the case where the innovator can costlessly manipulate the price signal in the sense that the price variable becomes non-contractible, since it is not verifiable in front of the court.[6] The analysis of price manipulation by the innovator in Section 6 is therefore also applicable to the case where the planner has limited commitment.

4.2 Extensions and discussion

In Appendix 2 we describe a number of extensions of our baseline model.

Shifts in the demand. Following Battaglini (2005) we assume that the demand has

---

[6]This approach is the one followed by the incomplete contracts literature, see Tirole (2003) for a discussion. A caveat that is pointed out in the literature is that if both the agent and the principal can observe a variable, then it is possible to enforce a contract by requiring both agents to report the value of that variable. In case of disagreement both the principal and the agent can be punished, which allows truthtelling to be an equilibrium outcome. If that caveat is accepted, then limited commitment on the part of the government does not prevent the planner from implementing the first best.

two states high (H) and low (L) with $D_H(p,\theta) \geq D_L(p,\theta)$ for each p and that transition between states follows a Markov process. We show that also in this model, if demand functions are analytic the planner can maximize innovation incentives by approximating the first best outcome. To understand the intuition for the proof consider the case of linear demand. The planner can identify the intercepts of the two demand functions by dictating a price equal to zero and maintaining it until two different quantities are observed. Then he will set $p = \varepsilon$ until two different quantities are observed. With two observations along each demand line the planner learns the demand and welfare functions. By setting ε arbitrarily close to zero, the dead weight loss generated by above marginal cost pricing is minimized and the first best is approached arbitrarily closely. An interesting feature of this result is that the optimal incentive scheme is non-stationary and has unbounded memory even if the demand shifts follow a Markov process and the relevant economic environment has a memory of only one period.

Demand is observed with error. Second, we consider the case in which the demand is observed with error and assume that $q_t = D(p_t, \theta) + \varepsilon_t$ where ε_t is a mean zero i.i.d. noise over the support $[-\overline{\varepsilon}, \overline{\varepsilon}]$. Even in this case, analyticity of the demand function is sufficient to approach the first best arbitrarily closely. In the linear demand case, the planner can use the following two step scheme. In the first stage the planner induces the firm to charge $p = 0$ and obtains a sample of N quantities for this price. Then he sets a price equal to ε and obtains another sample of N quantities. The weak law of large numbers guarantees that, for N large enough, the sample averages are unbiased estimates of $D(0, \theta)$ and $D(\varepsilon, \theta)$ and therefore the demand parameters can be learned by the planner. While the law of large numbers guarantees that the estimate is unbiased, the variance of the estimate depends on the price variation and is smaller when the variation is larger. This is not an issue in our setting because we assumed that the planner and the innovator are both risk neutral. Even in the case of risk aversion, the variance can be made arbitrarily small by letting the sample size, N, be very large.[7]

Distortionary taxation. Our baseline model assumes that there is no loss incurred by society when raising revenue to buy the patent out. This is a typical assumption in the eco-

---

[7]In practice, the policy maker may not be able to collect a very large dataset. This introduces a trade-off between price distortion and precision of out-of-sample estimates. For a discussion of additional challenges faced in structural demand estimation see Chintagunta and Nair (2011).

nomics of innovation literature (Chari et al., 2012; Weyl and Tirole, 2012). The assumption is justifiable in our setting where buy-outs may be conducted by philanthropic foundations and may not be associated with distortionary taxation. In our third extension we show that Proposition 1 is robust to dropping this assumption. We extend our model considering the case in which the government finances transfer T at a cost $(1 + \kappa)T$ where $\kappa \geq 0$ denotes the cost of public funds due to the deadweight loss associated with taxation (as in Laffont and Tirole, 1993). The planner faces a trade-off between two types of welfare distortions: the cost of raising money through public taxation, κ, and the surplus losses due to market power. We show that in this case the first best involves a positive price, and that market power does not prevent the planner to approximate the welfare maximizing outcome arbitrarily closely. To see the intuition of this result, consider a mechanism design approach in which the innovator reports to the planner a type, $\widetilde{\theta}$, and the planner indicates a path of market outcomes. If the type is reported truthfully, the first best will be approximated by the planner by choosing market outcomes arbitrarily close to the welfare maximizing outcome for that type. Truthful revelation will occur because the planner can exploit market outcomes to learn the analytic demand and punish the innovator if the reported type is not consistent with the identified demand.

Demand growth. As we show it in Appendix 2, Proposition 1 may not generalize if demand grows over time. In such setting, when the demand does not grow too quickly, the planner will be able to approximate each level of demand and approach the first best. Nevertheless, the planner may not have enough time to learn the various demand levels when growth is fast and the first best may not be approached.

In general, Proposition 1 does not hold in settings where structural demand identification is not feasible. When the demand is not analytic, local price variation cannot be exploited to identify to estimate the global behavior of the demand function. In the same way, when demand grows very fast it may be unfeasible to collect the price quantity observations necessary to identify demand. These issues are typical in structural modeling, where it is assumed that the structural parameters identified through local data variation can be used to perform counterfactuals or policy simulations (Reiss and Wolak, 2007). From this perspective, Proposition 1 does not require extra assumptions to those typically imposed in structural industrial organization studies.

16

5   Demand Manipulation

The analysis in Section 4 focused on the case of no demand manipulation. In this Section we consider the case in which the innovator can manipulate the market outcomes.

5.1   Buyouts and Price Variation

In the general model described in Section 3, the innovator can affect the market outcomes and manipulate market signals received by the planner up to period T but not after T. Our model also assumes a constant competitive market outcome from $T + 1$. A natural interpretation of this assumption is that the patent is acquired by the planner at T, so in the following we will refer to T as the buyout time.[8]

For a moment, let us depart from that model and assume that the planner (but not the innovator) can affect market outcomes after T. In this setting the first best can be approximated as in the case in which manipulation is not possible. This is the case both if manipulation is costly and if it is costless. To see this, consider the case in which the demand is linear. Then the planner can acquire the patent in the first period, sell the innovation at $p_1 = \varepsilon$ and $p_2 = 0$ and reward the innovator in the second period. In other words, the planner can appropriate the patent, generate the market outcomes required to learn the surplus generated by the innovation and then compensate the innovator. Alternatively, the planner can induce the innovator to generate the market outcomes necessary to learn the surplus and use additional post-buyout market outcomes to detect demand manipulations. For example, the patentee can be required to sell at $p_1 = \varepsilon$ and $p_2 = 0$ in the first two periods. The planner can then acquire the patent and practice $p_3 = \varepsilon$ and $p_4 = 0$ in the third and fourth periods. If the outcomes generated by the innovator coincide with those generated by the planner, the innovator will be rewarded with a transfer that approximates the surplus generated. If there are differences between market outcomes generated by the innovator and those generated by the planner, the innovator receives no transfer.

The basic insight is that pre-buyout manipulation, even if costless, can be avoided as long

---

[8]Noticed that in the previous Section we ignored $T(h_T)$ and focused on $g(h_t)$. This is because, in the absence of price manipulation the planner can generate a competitive outcome using only $g(h_t)$ by punish the innovator if $p_t \neq 0$.

17

as manipulation after buyout is not possible and the planner can generate price variation after buyout to identify the demand and detect manipulation. Therefore, for manipulation to distort away from the first best, it has to be the case that either (i) manipulation by the innovator is feasible both before and after the buyout or (ii) the ability of the planner to generate price variation after the buyout is limited. In the next Section we study case (ii) from above.

5.2 Post-Buyout Competitive Outcome

We now consider the case in which after the buyout time T the innovation is sold in a competitive market and that neither the innovator nor the planner can affect (manipulate) this outcome. The quantity of product sold can be perfectly observed by the planner but the price and hence the revenue can be distorted by the innovator, as described in Section 3. This may arise, for example, when the innovator awards secret discounts to his consumers.

To provide a micro-foundation of the manipulation cost $\phi(\widehat{p}_t, p_t)$ we assume that the innovator can convince the planner that he is selling at $\widehat{p} > p$ by sustaining a cost equal to $\widehat{\phi}((\widehat{p} - p) D(p, \theta))$ with $\widehat{\phi}$ being twice differentiable, and $\widehat{\phi} > 0$, $\widehat{\phi}'' \geq 0$. Intuitively, the planner observes sales equal to $\widehat{p}D(p, \theta)$ whereas the true revenue is equal to $pD(p, \theta)$ and $(\widehat{p} - p) D(p, \theta)$ are fake revenues undermined by secret price discounts. A simple justification of a positive manipulation cost is that the secret discounts offered are wasteful, that is they cost more to the innovator to offer than they are worth for the consumers. Alternatively, there may be a difference between the cost of external and internal financing. As argued by Aghion and Tirole (1994), for innovative firms this difference arises naturally because of the informational asymmetries involving new products and technologies. In this case, to convince the planner that sale revenue is equal to $\widehat{p}D(p, \theta)$ the innovator will have to borrow $(\widehat{p} - p) D(p, \theta)$ sustaining a cost of $\widehat{\phi}((\widehat{p} - p) D(p, \theta))$.[9] A simple functional specification for the manipulation cost is $i (\widehat{p} - p) D(p, \theta)$, if $i > 0$ there is a positive cost of manipulating sales.

Proposition 2 *A patent of length T is Pareto dominated by a patent buyout scheme that depends on market outcomes.*

---

[9]Another microfundation of the cost $\widehat{\phi}$ is that with some probability the planner will detect the manipulation and the innovator will pay a fine that depends on the fake proceeds.

The proposition shows that in our model for any patent of length T the policy maker can design a buyout scheme that generates greater welfare than the patent. The planner commits to buy out the patent at a price that depends on the market outcomes observed during the first $\widehat{T} < T$ periods. The buyout time $\widehat{T}$ is chosen to allow the planner to learn about the value of the innovation and to remove the incentives of the innovator to manipulate sales. At this optimal time the marginal cost of manipulating sales for $\widehat{T}$ periods is equal to the marginal benefit of obtaining extra buyout reward.

In the linear case the optimal buyout time $\widehat{T}$ is pinned down by the formula

$$\frac{\delta^{\widehat{T}} - \delta^T}{1 - \delta^{\widehat{T}}} = i \tag{2}$$

that indicates how patent buyout takes place sooner as i gets larger. This result is reminiscent of Chari et al (2012) who consider patents and prizes, and show that in their model shorter patents are more likely to be optimal when manipulation costs are higher, but longer patents need to be used when manipulation costs are lower. In the next Section, we show that even with costless manipulation of the price signals (when an infinitely lived patent is implied by (2)), one can do better by considering mechanisms that are different from both prizes and patents.[10]

With additional assumptions on the relationship between surplus and monopoly profits, innovation incentives can be increased even more. Take for example the setting of Weyl and Tirole (2012) with $D(p, \theta) = \sigma Q(\frac{p}{m})$ where $\theta = (\sigma, m)$, m is the monopoly price, σ is the quantity sold at marginal cost price and $Q()$ is a function known to the planner. In their setting there is proportionality between monopoly profits $m\sigma Q(1)$ and surplus at zero price $m\sigma S(0)$. By inducing truthful revelation of monopoly profits, the buyout allows the planner to back out the surplus and to transfer the entire surplus to the innovator from period $\widehat{T} + 1$. The innovator will obtain the monopoly profits before the buyout and the entire consumer surplus for the post-buyout period. In this way consumers enjoy greater surplus than the case of a T-period patent and the innovator has greater innovation incentives. In particular, the

---

[10]It is important to note that there are two important differences between the setup of Chari et al (2012) and ours. First, we allow heterogeneous innovation costs. Second, Chari et al (2012) rule out positive transfers by allowing the innovator to produce a fake (and useless) "innovation".

outcome resembles the first best after the buyout, because there is marginal cost pricing and all the surplus is transferred to the innovator.[11]

One may speculate that when $\widehat{\phi}(\widehat{p}, p) = 0$ patents cannot be improved upon. This is not the case, as the next proposition shows.[12]

Proposition 3 *When $\widehat{\phi}(\widehat{p}, p) = 0$ there is a per unit subsidy level τ that Pareto dominates patents that last forever.*

Proposition 3 shows that in our model even when price manipulation is costless, the planner can improve upon patents by exploiting the observed quantities. In the proof we show that a small quantity subsidy increases product market surplus by reducing the market price and increasing firm's profits. We also show that for τ small enough, such positive welfare effect dominates any loss generated by entry of inefficient innovators induced by the subsidy.

Overall, Propositions 2 and 3 show that in our model for a broad class of demand functions patents are not the optimal mechanism to incentivize innovation when the planner can observe market outcomes, even when the innovator may substantially manipulate sales. In the next Section, in a simplified environment, we characterize the optimal mechanism.

6   Optimal Mechanism with Costless Manipulation

In this Section we study the optimal incentive system in which the quantity produced is observable by the planner, but the innovator can manipulate the price costlessly, so the price will not be contracted on. This assumption captures a situation where the innovator can offer secret price discounts to buyers at no cost (other than lowering revenues). As in the previous

---

[11]In the linear specification, if we interpret $i > 0$ as the difference between the cost of external and internal financing the planner can reduce manipulation incentives even more by combining the buyout of the patent with the requirement to purchase a bond. Specifically, the planner can request the innovator to purchase a bond that costs $\gamma \widehat{p} D(p, \theta)$, pays no interest and expires after T^B periods. If $p D(p, \theta)$ is the only revenue available to the innovator, he will have to borrow $\gamma \widehat{p} D(p, \theta) - p D(p, \theta)$ for T^B periods at a cost of

$$i \left(\gamma \widehat{p} D(p, \theta) - p D(p, \theta) \right) \frac{1 - \delta^{T^B}}{1 - \delta}.$$

This extra manipulation cost generated by the bond allows to accelerate the buyout time and therefore increases consumer welfare.

[12]It is also possible to show that patents of finite lengths can also be improved upon by a simple per unit subsidy mechanism but this result is somewhat more tangential to what we discuss below.

Section, we assume that after the buyout the innovation is sold in a competitive market and that neither the innovator nor the planner can affect (manipulate) this outcome.

We will study the problem with a mechanism design approach in which the innovator reports to the planner a type, $\widehat{\theta}$, and the planner requires that in period t the innovator produces a specific quantity, $q_t(\widehat{\theta})$, and receives a payment $\tau_t(\widehat{\theta})$. To simplify the analysis we focus on the linear demand case[13] $D(p) = \theta_1 - \theta_2 p$.

First, we show that there is no loss of generality in assuming that the planner knows the intercept θ_1.[14] More precisely, we can approximate the welfare of an auxiliary problem where the planner knows θ_1 from the outset arbitrarily closely. This is an upper bound because the planner cannot do better than in the hypothetical case where he observed θ_1 at the beginning.

Lemma 1 *The planner can approximate the welfare that can be induced under full information about θ_1 arbitrarily close.*

This result is quite intuitive: the planner can perfectly learn the demand intercept when the market becomes perfectly competitive and punish the innovator if θ_1 was not reported truthfully. Exploiting this Lemma, we focus on the linear demand case with known intercept (normalized to 1) and unknown slope that for simplicity we rewrite as $\theta_2 = 1/2\theta$. The demand is therefore

$$q = 1 - \frac{p}{2\theta}$$

and larger θ are associated with steeper demand curves and larger consumer surplus. Notice that the monopoly quantity is independent of θ and it is equal to $q^M = 1/2$.

6.1 Static Mechanisms

We first study a static setting where the profits are realized only for one period after the innovator reports his type. Let $p(\widehat{\theta}, \theta) = 2\theta(1 - q(\widehat{\theta}))$ be the price at which the innovator can sell quantity $q(\widehat{\theta})$ if the actual demand is characterized by θ. The profits from reporting $\widehat{\theta}$ when

---

[13]Most of our results would continue to hold if we had a more general demand function of the form $D = \theta - g(p)$ where g is an unknown convex function, and θ is an unknown parameter.

[14]We do not need the innovator to report his cost, c, because in our setting, as in Scotchmer (1999), the innovator's compensation cannot depend on the true c since he cannot be punished for lying about c.

the type is θ (gross of innovation costs) are:

$$U(\widehat{\theta}, \theta) = \tau(\widehat{\theta}) + p(\widehat{\theta}, \theta)q(\widehat{\theta})$$

$$= \tau(\widehat{\theta}) + 2\theta(1 - q(\widehat{\theta}))q(\widehat{\theta}).$$

Letting $V(\theta) = U(\theta, \theta) - c$ denote the rent under truth-telling, the envelope theorem implies that

$$V'(\theta) = \frac{\partial}{\partial \theta} U(\widehat{\theta}, \theta) \mid_{\widehat{\theta} = \theta} =$$

$$= q(\theta) \frac{\partial}{\partial \theta} p(\widehat{\theta}, \theta) \mid_{\widehat{\theta} = \theta} = 2q(\theta)(1 - q(\theta)). \tag{3}$$

The above condition (3) is a first order condition. The following result states a necessary and sufficient condition for implementability:

Lemma 2 *A schedule $q(\theta)$ can be implemented if and only if q is weakly decreasing in θ.*

Lemma 2 shows that the optimal mechanism requires the quantity sold to be decreasing in θ. Therefore, as the surplus created by the innovation increases, the quantity produced is reduced. The intuition for this result is the following. The planner exploits market power to induce truthful revelation and screen consumers' willingness to pay. When θ is large consumers are willing to pay high prices for the product and the innovator is likely to prefer market power to lump-sum transfers. Conversely, when θ is low consumers are price sensitive and market power would not be attractive to the innovator.

We are ready to formulate the planner's problem. First, note that the total surplus when q is implemented for an innovator with type θ is $W(q, \theta) = \int_0^q 2\theta(1 - x)dx = \theta(2q - q^2)$. Let $\widehat{c}(\theta)$ be the highest cost innovator who enters (endogenously determined by the mechanism by $V(\theta) = 0$). Then the objective function can be written as

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \int_0^{\widehat{c}(\theta)} \psi(c, \theta)(W(q(\theta), \theta) - c)dcd\theta.$$

The planner's problem is

$$\max_{q(\theta)} \Pi$$

$$\text{s.t.} \quad \widehat{c}'(\theta) = V'(\theta) = 2q(\theta)(1 - q(\theta)), \text{ and } q'(\theta) \leq 0 \ \forall \theta \in [\underline{\theta}, \overline{\theta}].$$

22

The main challenges are twofold: first, the monotonicity constraint on q; second, the fact that the state variable $\widehat{c}(\theta)$ has free initial *and* end conditions, a combination that is uncommon for standard dynamic optimization problems. To obtain a solution to this problem, let us assume uniform independent distributions for c and θ on $[0, 1]$ and $[\underline{\theta}, 1]$ for some $\underline{\theta} > 0$. Then the problem simplifies to (ignoring a few constants):

$$\max_{q(\theta)} \int_{\underline{\theta}}^{1} [\theta(2q(\theta) - q^2(\theta))\widehat{c}(\theta) - \frac{\widehat{c}^2(\theta)}{2}]d\theta$$
$$\text{s.t.} \quad \widehat{c}'(\theta) = 2q(\theta)(1 - q(\theta)), \text{ and } q'(\theta) \leq 0 \ \forall \theta \in [\underline{\theta}, 1].$$

Optimal static mechanism

In the next proposition we characterize the optimal quantity schedule in the presence of costless price manipulation.

Proposition 4 *In the optimal static mechanism, there exists $\overline{\theta} \in (\underline{\theta}, 1)$ such that it holds that q is strictly decreasing on interval $[\underline{\theta}, \overline{\theta}]$ and then constant on $[\overline{\theta}, 1]$. Moreover, $q(\theta) \geq 2/3 > 1/2 = q^{monopoly}$ for all θ and $q(\underline{\theta}) = 1 = q^{first \ best}$.*

To gain intuition for this result, our starting point is Lemma 2, which implies that the quantity schedule needs to be weakly decreasing to be incentive compatible. In the proof of the Proposition, we show that in the relaxed problem where the monotonicity constraint on q is ignored, the optimal solution is such that $q(\underline{\theta}) = q(1) = 1$.[15] Given this, it is not surprising that when one reintroduces the monotonicity constraint on q it is still true that $q(\underline{\theta}) = 1$. It is also not surprising that now $q(1) < 1$, because $q(1) = 1$ and the monotonicity constraint would imply that $q(\theta) = 1$ for all θ, that is all possibility for screening would be given up.

Proposition 4 shows that in our model the optimal mechanism differs substantially from a patent system even if the innovator can manipulate price signals costlessly. The optimal quantity schedule has three important characteristics. First, the quantity produced varies across types. This is a fundamental difference with the patent system that implements only the monopoly quantity that in our setting is constant across types. Second, the quantity produced by each type is above the monopoly quantity. Thus, despite costless price manipulation,

---

[15] The reason is that otherwise function q could be increased uniformly by the same amount ε, and the value of $\widehat{c}(\underline{\theta})$ adjusted so that an increase in welfare is induced without violating any incentive constraints.

information on the quantity produced allows the planner to reward the innovation generating less distortions than a traditional patent system. Finally, the optimal quantity is strictly decreasing in θ for low values of θ and constant for high surplus innovations as depicted in Figure 2.
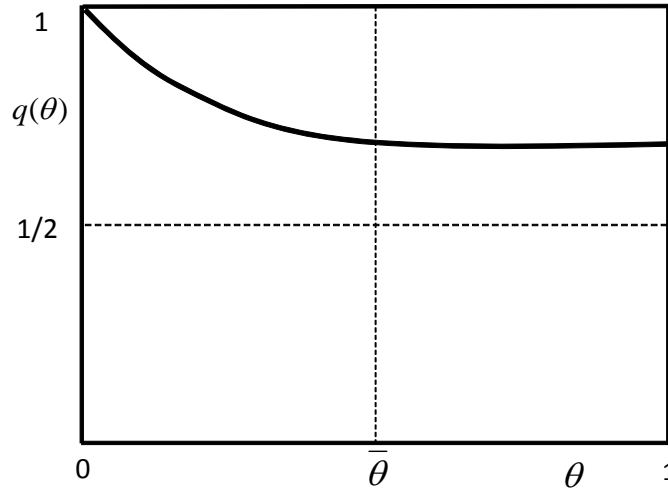

Figure 2: Optimal quantity schedule with costless price manipulation

The intuition behind this result is that the planner's welfare maximization involves a trade-off between a 'consumer welfare' effect and a 'screening' effect. When quantities decrease with θ, the planner can use market power to screen consumers' willingness to pay. Nevertheless, maximization of consumer surplus implies that larger quantities should be offered for innovation with larger θ since the impact on welfare of an increase in q is greater the greater is θ. For low values of θ, the 'screening effect' dominates and the planner exploits market power to screen willingness to pay. This is intuitive since for low θ it is crucial for the planner to avoid excess entry of low value innovators. As θ increases, the innovations have larger impact on consumer surplus and the planner has lower incentives to distort the market for screening purposes. For θ large enough, the 'consumer welfare' effect dominates and the planner implements a quantity schedule that is constant in θ. The idea that market power can be exploited to screen willingness to pay is similar to the logic in Weyl and Tirole (2012).[16]

---

[16]They restrict their attention to Cobb-Douglas reward policies (in our setting this restriction would generate a constant level of q across types). They show that q decreases with the variance of the type distribution. In our setting, we show that even with a fixed type distribution, the planner may use different quantities to screen

6.2 Optimal Dynamic Mechanism and Discussion

Having characterized the optimal quantity schedule in the static setting, we now consider the dynamic problem where the planner can choose a path $(q_t(\theta), \tau_t(\theta))$ for every $t \geq 0$.[17] Our main result shows that repeating the same quantity over time for all types θ is optimal.

Proposition 5 *It is optimal for the planner to set a policy where $q_t(\theta)$ is constant in time for any θ, that is to adopt the optimal static mechanism.*

In our model a constant mechanism (over time) is optimal because of the desirable features of quantity (and price) smoothing over time. This is due to the fact that total surplus is concave in the quantity (and price), so inducing a temporal variation in quantities (as patents do) introduces extra distortion in the product market without improving innovation incentives. This finding resembles the result of Gilbert and Shapiro (1990).[18]

Proposition 5 confirms that in our model the optimal mechanism *differs from a patent system* even if manipulation (of prices) is costless. Welfare is maximized with the innovator selling a quantity that is above the monopolistic quantity until the buyout occurs, unlike the (optimal) patent system described by Scotchmer (1995). This result is related to our earlier finding (Proposition 3), which shows that a small quantity subsidy always improves welfare.

Notice the apparent tension between Proposition 5 and Lemma 1. Proposition 5 requires the planner to implement a constant quantity over time whereas Lemma 1 requires the planner to move to the competitive outcome for at least one period in order to learn the intercept of the demand function. This tension identifies a key trade-off. On one hand, the planner would like to smooth market outcomes over time to increase welfare. On the other hand, the planner would like to generate variation of market outcomes to learn the underlying demand parameters. In the linear context, this tension leads to a mechanism that resembles a buyout

for different types.

[17]Since no new information is revealed to the agent (the innovator), it is without loss of generality to concentrate on mechanisms where the agent reports his type only at the outset.

[18]They conclude that the optimal patent policy calls for infinitely lived patents when patent breadth is increasingly costly in terms of deadweight loss. In our setting, lowering the quantity produced can be thought of as an increase in patent breadth because a lower quantity reduces consumer surplus and increases the profits of the innovator.

where the patent is bought out after a long time (as long as possible) has elapsed.[19]

Proposition 5 also highlights the fact that learning from market signals over time may be substituted by an initial screening process where the innovator self reports his type. The literature on dynamic mechanism design cautions us that this result (no learning is optimal until the buyout) is only true because our agent (the innovator) has strictly superior information over the planner, and this advantage is maintained over time. However, in a large number of applications this may be a realistic assumption. In such applications, the optimal mechanism does not utilize learning on the part of planner, rather it relies on a single report of the innovator at the outset. Such a policy can be implemented by offering a menu of R&D subsidies and per unit quantity subsidies.

It is beyond the scope of our work to characterize the optimal mechanism in a general framework of dynamic market signals, but a few characteristics of our proposed mechanism appear to be robust. First, prices need not be set at the extremes of monopoly pricing (i.e. full patent protection) or fully competitive pricing. Second, buyout itself can be viewed in terms of its ability to generate information, an important aspect that has been overlooked by the previous literature. Relatedly, observed demand information *after* the patent buyout can be used to incentivize innovation.

7 Conclusions

In this paper we have examined the problem of a social planner aiming to maximize consumer welfare and innovation incentives while observing prices and quantities practiced by the innovator over time. We have shown that in our model information about market outcomes may allow the planner to generate more welfare than a traditional patent system through patent buyouts.

There are a number of historical experiences in which governments bought patents out. The most famous example of patent buyout took place in July 1839 when the French government purchased the patent for the Daguerreotype photography process. The inventor, Luis

---

[19]In reality, there may be legal or political reasons why the buyout cannot be delayed indefinitely. For example, it may happen that the product becomes obsolete, and in this case the planner may not be able to commit to a buyout that may not seem to promote consumer welfare ex-post.

Jacques Daguerre, was not able to find buyers for the process, but obtained the support of a politician that convinced the government to acquire the patent and put the rights in the public domain. Within a short period of time the process spread around the country to become the technology standard in photography (Kremer, 2001). In recent academic and policy debates pharmaceutical patent buyouts have been suggested as a strategy to improve health in low income countries. For example, Banerjee et al. (2010) propose that a Health Impact Fund compensate drug manufacturers that sell in low income countries at marginal cost. They suggest that the compensation to a given manufacturer would depend on use of the drug and evidence of realized health benefits.[20]

Our paper provides two main insights into the design and application of such buyout schemes. First, the planner may find it beneficial to collect market data before the buyout and use them to estimate the surplus generated by the innovation. In practice, surplus may be estimated through structural econometric models that allow policy makers to estimate the primitives of consumer preferences and to generate out of sample predictions (Cho and Rust, 2008). Such estimates can provide useful guidance in the determination the buyout compensation for the innovator. Second, the planner should consider the welfare cost associated with collecting price-quantity observations. As long as local variation in market outcomes can be exploited to learn about the global properties of the demand, prices close to marginal costs minimize the loss in consumer surplus.

References

[1] Arrow, Kenneth (1962), "Economic Welfare and the Allocation of Resources for Inventions," in Richard Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* (Princeton: Princeton University Press): 609–625

[2] Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien (1991),"Optimal Learning by Experimentation," *Review of Economic Studies* 55: 621-654

---

[20] A similar policy proposal is described in Guell and Fischbaum (1995).

[3] Aghion, Philippe and Jean Tirole (1994),"The Management of Innovation," *Quarterly Journal of Economics* 109: 1185-1209

[4] Banerjee, Amitava, Aidan Hollis and Thomas Pogge (2010), "The Health Impact Fund: incentives for improving access to medicines," *Lancet* 375:166-169

[5] Battaglini, Marco (2005), "Long-Term Contracting with Markovian Consumers,"*American Economic Review*, 95: 637-658

[6] Berry, Steven, James Levinshohn and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica* 63: 841-890

[7] Boldrin, Michele and David Levine (2008), "Perfectly Competitive Innovation", *Journal of Monetary Economics* 55: 435–453

[8] Bessen, James and Eric Maskin (2009), "Sequential innovation, patents, and imitation", *RAND Journal of Economics* 40: 611-635

[9] Bloom, Nick, Mark Schankerman and John Van Reenen (2013), "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica* 81: 1347–1393

[10] Chintagunta, Pradeep and Harikesh Nair (2011), "Discrete-Choice Models of Consumer Demand in Marketing," *Marketing Science* 30: 977–996

[11] Chari V. V., Mikhail Golosov and Aleh Tsyvinski (2012), "Prizes and Patents: Using Market Signals to Provide Incentives for Innovations", *Journal of Economic Theory* 147: 781-801

[12] Cho, Sungjin and John Rust (2008), "Is econometrics useful for private policy making? A case study of replacement policy at an auto rental company," *Journal of Econometrics* 145: 243–257

[13] Cornelli, Francesca and Mark Schankerman (1999), "Patent Renewals and R&D Incentives," *RAND Journal of Economics* 30: 197-213

[14] Deaton, Angus and John Muellbauer (1980), "An almost ideal demand system," *American Economic Review* 70: 312-326

[15] Fox, Jeremy and Amit Gandhi (2011), "Identifying Demand with Multidimensional Unobservables," NBER working paper 17557

[16] Galasso, Alberto and Mihkel Tombak (2014), "Switching to Green: the Timing of Socially Responsible Innovation," *Journal of Economics and Management Strategy*, forthcoming

[17] Glennerster, Rachel, Michael Kremer and Heidi Williams (2006), "Creating Markets for Vaccines," *Innovations: Technology Governance Globalization* 1: 67-79

[18] Gilbert, Richard and Carl Shapiro (1990), "Optimal patent length and breadth," *RAND Journal of Economics* 21:106-112

[19] Guell, Robert and Marvin Fischbaum (1995), "Toward Allocative Efficiency in the Prescription Drug Industry," *Millbank Quarterly* 73: 213-230

[20] Hellwig, Martin (2009), "A Maximum Principle for Control Problems with Monotonicity Constraints," Max Planck Institute working paper

[21] Hopenhayn, Hugo, Gerard Llobet and Matthew Mitchell (2006), "Rewarding Sequential Innovations: Prizes, Patents, and Buyouts," *Journal of Political Economy* 114: 1041-1068

[22] Hopenhayn Hugo and Matthew Mitchell (2001), "Innovation variety and patent breadth," *RAND Journal of Economics* 32: 152-166

[23] Hovenkamp, Herbert (2004), "United States Antitrust Policy in an Age of IP Expansion," University of Iowa Legal Studies Research Paper No. 04-03

[24] Judd, Kenneth (1998), Numerical Methods in Economics, MIT Press

[25] Jaffe, Adam B. 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

[26] Kremer, Michael (1998), "Patent Buyouts: A Mechanism for Encouraging Innovation," *Quarterly Journal of Economics* 113: 1137-1167

[27] Laffont, Jean-Jacques and Jean Tirole (1993), *A Theory of Incentives in Procurement and Regulation,* MIT Press

[28] Mastroianni, Giuseppe and Gradimir Milovanovic (2008), "Interpolation Processes: Basic Theory and Applications,"Springer-Verlag Berlin

[29] Mitchell, Matthew and Andrea Moro (2006), "Persistent distortionary policies with asymmetric information," *American Economic Review* 96: 387-393

[30] Reiss, Peter and Frank Wolak (2007), "Structural Econometric Modeling: Rationales and Examples from Industrial Organization," Handbook of Econometrics, Vol. 6A, Elsevier BV

[31] Scotchmer, Suzanne (1999), "On the Optimality of the Patent Renewal System", *RAND Journal of Economics* 30: 181-196

[32] Shavell, Steven and Tanguy Van Ypersele (2001), "Rewards versus Intellectual Property Rights," *Journal of Law and Economics* 44: 525-547

[33] Tirole, Jean (2003), "Incomplete Contracts: Where do We Stand", *Econometrica* 67: 741-781

[34] Weyl, Glenn and Jean Tirole (2012), "Market Power Screens Willingness to Pay," *Quarterly Journal of Economics* 127: 1971-2003

[35] Wright, Brian (1983), "The Economics of Invention Incentives: Patents, Prizes and Research Contracts," *American Economic Review* 73: 691-707

Appendix 1: Proofs

Proof of Proposition 1

First, note that if $r \geq \overline{p}$ then $D(\cdot, \theta)$ can be expanded globally on $[0, \overline{p}]$ and we can construct a global estimate of the demand function given approximate knowledge of the function $D(\cdot, \theta)$ around the point $(0, D(0, \theta))$. The global estimate is obtained with the following polynomial:

$$\sum_{i=0}^{k} \alpha_i p^i,$$

where α_i is an appropriate estimate of the i'th derivative of D with respect to p at $p = 0$ divided by $i!$ to use Taylor's formula. Notice that the coefficients of the polynomial can be estimated by charging $k + 1$ distinct prices close to 0.

The basis for this is that as k gets large, the approximation of the derivatives improves and thus our estimate of D approaches the true value of D arbitrarily close.[21] To formalize this, suppose that we have taken a sample of $k + 1$ observations such that the price was always below some $\widehat{x} > 0$. The error term (in absolute value) for the estimate of the i'th derivative can be bounded by $\max_{\theta \in \Theta, i, x \in [0, \widehat{x}]} | D^{(i+1)}(x, \theta) x^{i+1} / (i + 1)! | \leq K \widehat{x}^{i+1}$, which can be made arbitrarily small (in absolute value) if $\widehat{x}$ is small (Mastroianni and Milovanovic, 2008). Here we used the fact that there exists a $K > 0$ such that $\max_{x \in [0, \widehat{x}]} | D^{(i+1)}(x, \theta) / (i + 1)! | < K$ for all $i = 1, 2, ...$ and $\theta \in \Theta$. To establish that this is indeed true, note that by D being analytic there exists $\widetilde{K}(\theta)$ such that $\max_{x \in [0, \widehat{x}]} | D^{(i+1)}(x, \theta) / (i + 1)! | < \widetilde{K}(\theta)$ for all $i = 1, 2, ...$ Moreover, Weierstrass's theorem implies that there exists K such $\widetilde{K}(\theta) \leq K$ for $\theta \in \Theta$, because Θ is compact and $\widetilde{K}$ is a continuous function of θ because all the derivatives of D are continuous in θ by assumption.[22]

If $r < \overline{p}$ then $D(\cdot, \theta)$ can only be expanded locally and approximation by polynomial is valid only in intervals around p^* of size less than r. To estimate the demand in this case we apply an analytic continuation technique as in Aghion et al (1991). Let us define $l = \overline{p}/n$ and

---

[21] If the derivatives at 0 can be estimated with a known error ε, then the total error at $p > 0$ is less than $\varepsilon(p + p^2 + ...) = \frac{\varepsilon}{1-p}$ if $p < 1$. If the choking price cannot be bounded away from 1 (from above), then this procedure does not suffice, and local expansion is needed similarly to what is suggested below for the case where $r < \overline{p}$.

[22] Note, that directly we only assumed that D itself is continuous in θ. However, if D is analytic, then continuity of D in θ implies that all the c_k coefficients are continuous in θ, which implies that all derivatives of D are also continuous in θ.

take n large enough such that $l < r$. We can approximate $D(\cdot, \theta)$ in the interval $[p^*, p^* + l]$ by setting $k = n^2$ and calculating

$$\sum_{i=0}^{n^2} \alpha_i (p - p^*)^i,$$

and approximate the first $(n^2 + 1) - n$ derivatives of $D(\cdot, \theta)$ by the first $(n^2 + 1) - n$ derivatives of the polynomial. Next, let $\langle \beta_i | 0 \leq i \leq n^2 - n \rangle$ be the values of these derivatives at $x^* + l$. We can now approximate $D(\cdot, \theta)$ in the interval $[p^* + l, p^* + 2l]$ by

$$\sum_{i=0}^{n^2 - n} \beta_i (p - p^* - l)^i$$

and approximate the first $(n^2 + 1) - 2n$ derivatives of $D(\cdot, \theta)$ by by the first $(n^2 + 1) - 2n$ derivatives of the polynomial. Proceeding this way one reaches $\bar{p}$ after at most n steps and similarly proceeding leftward one can estimate $D(\cdot, \theta)$ up to zero. Also in this case by choosing p^* arbitrarily small and n arbitrarily large the demand is approximated arbitrarily closely at a very low welfare cost.

Proof of Proposition 2

Consider the following mechanism. The innovator is awarded a patent for $\widehat{T} \leq T$ periods as long as the same prices and quantities $(\widehat{p}, D(p, \theta))$ are observed by the planner for the entire patent duration $\widehat{T}$. After $\widehat{T}$ periods the patent is acquired by the planner that will pay the innovator $\widehat{p}D(p, \theta)$ per period for the remaining $\widehat{T} - T$ periods and the innovation is sold at marginal cost. With $\widehat{p} \geq p$ the payoff of the innovator is

$$\frac{1 - \delta^{\widehat{T}}}{1 - \delta} \left[pD(p, \theta) - \widehat{\phi} \left((\widehat{p} - p)D(p, \theta) \right) \right] + \frac{\delta^{\widehat{T}} - \delta^T}{1 - \delta} \widehat{p} D(p, \theta).$$

Now consider setting $\widehat{T}$ such that

$$\delta^{\widehat{T}} - \delta^T = (1 - \delta^{\widehat{T}}) \widehat{\phi}'(0)$$

so that the marginal benefit of manipulation when $\widehat{p} = p$ is exactly equal to the marginal cost.[23] Setting $\widehat{p} = p$ is then optimal for the innovator because the first order condition

---

[23]If $\widehat{T}$ is not an integer set it equal to the smallest integer for which $\delta^{\widehat{T}} - \delta^T < (1 - \delta^{\widehat{T}})\phi'(0)$.

holds by construction and the objective function is concave in $\widehat{p}$. This removes the innovator's incentive to manipulate. Maximizing the payoff respect to p (with $\widehat{p} = p$) gives:

$$(1 - \delta^{\widehat{T}}) \left[pD'(p, \theta) + D'(p, \theta) - \widehat{\phi}'(0)\widehat{p}D' + \widehat{\phi}'(0)(pD'(p, \theta) + D'(p, \theta)) \right] + (\delta^{\widehat{T}} - \delta^T)\widehat{p}D'(p, \theta)$$

$$= (1 - \delta^T)(pD'(p, \theta) + D'(p, \theta))$$

so the innovator will truthfully report the monopolistic profits. The profits of the innovator will be the same as with a patent of infinite length but consumers will be better off.

Proof of Proposition 3

When a per unit subsidy is awarded there are two main changes in total welfare. First, the set of types who enter becomes larger as the profit of the innovator increases. Second, for a fixed type who enters even without a subsidy, the total surplus on the market changes as prices go down due to the subsidy. Both effects increase welfare when τ is small as we show below.

The following argument shows that there is a small enough per unit subsidy $\tau > 0$ such that for *any* specific value of θ social welfare is larger than in the absence of any subsidies ($\tau = 0$) . To save notation, we do not explicitly indicate that the optimal price is a function of θ, and not only of τ.

The profits for the patentee in the presence of a quantity subsidy are equal to $(p + \tau)D(p, \theta)$ where τ is the per unit subsidy. The first order and second order conditions are:

$$(p + \tau)D'(p, \theta) + D(p, \theta) = 0$$

$$2D'(p, \theta) + (p + \tau)D''(p, \theta) \leq 0.$$

Let us indicate with $p(\tau)$ the optimal price charged by the monopolist. Now we exploit the FOC and the implicit function theorem to obtain

$$\frac{dp(\tau)}{d\tau} = -\frac{D'(p, \theta)}{2D'(p, \theta) + (p + \tau)D''(p, \theta)} < 0$$

because $D'(p, \theta) < 0$ and the second order condition is satisfied. Profits of the firm when optimally charging price $p(\tau)$ can be written as $\pi(\tau) = R(\tau, p(\tau)) = (p(\tau) + \tau)D(p(\tau), \theta)$. The envelope theorem implies that

$$\pi'(\tau) = \frac{dR}{d\tau} = D(p(\tau), \theta) > 0,$$

33

so innovation incentives become larger as τ increases. Next, for a given θ the product market surplus S (net of subsidies) is equal to

$$S(\tau) = p(\tau)D(p(\tau), \theta) + \int_{p(\tau)}^{\infty} D(z, \theta)dz$$

and thus

$$
\begin{aligned}
S'(\tau) &= D(p(\tau), \theta)\frac{dp(\tau)}{d\tau} + p(\tau)D'(p(\tau), \theta)\frac{dp(\tau)}{d\tau} - D(p(\tau), \theta)\frac{dp(\tau)}{d\tau} \\
&= p(\tau)D'(p(\tau), \theta)\frac{dp(\tau)}{d\tau} > 0.
\end{aligned}
$$

Total welfare can be written as $W(\tau) = \int_{\underline{c}}^{\pi(\tau)} (S(\tau) - x)\psi(\theta, x)dx$. Thus for τ close to zero we obtain

$$W'(\tau) = \int_{\underline{c}}^{\pi(\tau)} S'(\tau)\psi(\theta, x)dx + \pi'(\tau)(S(\tau) - \pi(\tau))\psi(\theta, \pi(\tau)) > 0,$$

because

$$S(\tau) - \pi(\tau) = p(\tau)D(p(\tau), \theta) + \int_{p(\tau)}^{\infty} D(z, \theta)dz - (p + \tau)D(p, \theta) = \int_{p(\tau)}^{\infty} D(z, \theta)dz - \tau D(p, \theta) > 0$$

for τ close to zero. Take any $\tau > 0$ such that $\int_{p(\tau, \theta)}^{\infty} D(z, \theta)dz - \tau D(p(\tau, \theta), \theta) > 0$ for all θ.[24] By the above, any such subsidy level τ increases total welfare for all θ. In other words the same level τ is applicable to all θ.

Proof of Lemma 1

Take the hypothetical problem where the planner observes θ_1 so the innovator needs to report only θ_2. As we show it in the next Section, the optimal mechanism prescribes a quantity $q_t(\theta_2) = q^*(\theta_2)$ that is constant in time (t). Now, take our original problem where the planner does not observe θ_1 at the outset, and suppose that the planner provides a buyout at time T

---

[24] When τ goes to zero the difference $\int_{p(\tau, \theta)}^{\infty} D(z, \theta)dz - \tau D(p(\tau, \theta), \theta)$ is strictly positive for every θ. Therefore, as long as $D(p(0, \theta), \theta)$ is bounded below by a positive uniform bound for all θ, then there is a τ that works uniformly for all θ. If such a uniform bound is not available, then the proof goes through with a few straightforward modifications.

and sets the quantities produced before time T equal to the $q^*(\theta_2)$.[25] After the buyout, when the market becomes perfectly competitive the intercept will be observed by the planner. At that stage the innovator can be punished if the quantity sold at marginal cost, θ_1, differs from the report of the innovator $\widehat{\theta}_1$. By making the punishment large enough the innovator has no incentive to misreport. Moreover, letting T become arbitrarily large the welfare induced by this mechanism approximates the welfare under full information about θ_1.

Proof of Lemma 2

First, let us write up the incentive conditions $U(\widehat{\theta}, \theta) \leq U(\theta, \theta)$ and $U(\theta, \widehat{\theta}) \leq U(\widehat{\theta}, \widehat{\theta})$. Adding these constraints up and substituting $p(\widehat{\theta}, \theta) = 2\theta(1 - q(\widehat{\theta}))$ we obtain

$$2\theta \left[(1 - q(\widehat{\theta}))q(\widehat{\theta}) - (1 - q(\theta))q(\theta) \right] \leq 2\widehat{\theta} \left[(1 - q(\widehat{\theta}))q(\widehat{\theta}) - (1 - q(\theta))q(\theta) \right]$$

Because quantities are higher than the monopoly quantities $(1/2)$ then q has to be decreasing in θ. On the other hand, if q is decreasing in θ, then by choosing an appropriate transfer schedule τ the quantity schedule can be implemented.

Proof of Proposition 4
Part 1: Solution of the relaxed problem

To develop intuition for the optimal static mechanism as characterized in Proposition 4, we simplify the problem by looking at the optimal control problem ignoring the $q'(\theta) \leq 0$ constraint first. To obtain a solution continuos in θ, we follow Hellwig (2009) and specify the following Hamiltonian:

$$H = \lambda(\theta)2q(\theta)(1 - q(\theta)) + [\theta(2q(\theta) - q^2(\theta))\widehat{c}(\theta) - \frac{\widehat{c}^2(\theta)}{2}].$$

The state variable $\widehat{c}$ has neither an initial nor an end condition, which makes it different from other optimal control problems. The first order condition for the control variable is

$$0 = \frac{\partial H}{\partial q} = \lambda(\theta)2(1 - 2q(\theta)) + 2\theta(1 - q(\theta))\widehat{c}(\theta), \ \forall \theta. \tag{4}$$

The other co-state equation is

$$-\lambda'(\theta) = \frac{\partial H}{\partial \widehat{c}} = \theta(2q(\theta) - q^2(\theta)) - \widehat{c}(\theta). \tag{5}$$

---

[25] By standard arguments, there is a payment schedule τ that makes this quantity schedule incentive compatible.

Moreover, Hellwig (2009) shows that in this class of problems:

$$\lambda(\underline{\theta}) = \lambda(1) = 0. \tag{6}$$

The above conditions lead to the following result.

Lemma 3 $q(\underline{\theta}) = q(1) = 1.$

Proof. From (4) we obtain that

$$q(\theta) = \frac{\lambda(\theta) + \widehat{c}(\theta)\theta}{2\lambda(\theta) + \widehat{c}(\theta)\theta} \tag{7}$$

that is equal to 1 when $\theta = 1$ and when $\theta = \underline{\theta}$. $\blacksquare$

This result shows that in the relaxed problem there is efficient production both for the innovations that create the largest surplus and for those that create the smallest surplus. One may conjecture that the solution of the relaxed problem is then a prize and all innovations are produced without market distortions. The next proposition shows that this is not the case, and that the optimal quantity schedule is non-monotonic.

Lemma 4 *There exists a* $\overline{\theta}$ *such that* $q(\overline{\theta}) < 1$ *and* $q'(\overline{\theta}) = 0$. *Moreover* $q' \leq 0$ *for* $\theta \in [\underline{\theta}, \overline{\theta})$ *and* $q' > 0$ *for* $\theta \in (\overline{\theta}, 1]$.

Proof. Differentiating (4) with respect to θ and dividing through by 2 yields

$$\lambda'(\theta)(1 - 2q(\theta)) - 2q'(\theta)\lambda(\theta) + (1 - q(\theta))\widehat{c}(\theta) - \theta q'(\theta)\widehat{c}(\theta) + \theta(1 - q(\theta))\widehat{c}'(\theta) = 0.$$

Substituting in from (5) and also using the formula for $\widehat{c}'$ yields

$$\left(\widehat{c}(\theta) - \theta(2q - q^2)\right)(1 - 2q(\theta)) + (1 - q(\theta))\widehat{c}(\theta) + \theta(1 - q(\theta))2q(1 - q) = q'(2\lambda + \theta\widehat{c}),$$

so the sign of q' is equal to the sign of

$$\left(\widehat{c}(\theta) - \theta(2q - q^2)\right)(1 - 2q(\theta)) + (1 - q(\theta))\widehat{c}(\theta) + \theta(1 - q(\theta))2q(1 - q)$$

$$= \widehat{c}(2 - 3q) + \theta q[2(1 - q)^2 - (2 - q)(1 - 2q)]$$

$$= \widehat{c}(2 - 3q) + \theta q^2. \tag{8}$$

From (7) it follows that for all $\theta \geq \underline{\theta}$ it holds that $q(\theta) \leq 1$, therefore $q'(\underline{\theta}) \leq 0$ holds because $q(\underline{\theta}) = 1$ by the previous Lemma. Because $q(\underline{\theta}) = q(1) = 1$, it means that there exists a $\overline{\theta} \in (\underline{\theta}, 1)$ such that $q'(\overline{\theta}) = 0$ and $q''(\overline{\theta}) > 0$. Now assume that there exists some $\widetilde{\theta} > \overline{\theta}$ for which $q'(\widetilde{\theta}) < 0$. This means that there exists a $\theta' \in \left(\overline{\theta}, \widetilde{\theta}\right)$ such that $q'(\theta') = 0$ and $q''(\theta') < 0$. Notice that $q'(\theta') = 0$ implies that $A'(\theta') = q(6q^2 - 9q + 4)$ that is strictly positive for any value of $q > 0$. This implies that if $q'(\theta') = 0$ then $q''(\theta') > 0$ that contradicts the existence of $\widetilde{\theta}$ and implies that $q' > 0$ for each $\theta > \overline{\theta}$. ∎

The intuition for this result is related to the fact that ignoring the monotonicity constraint on q is essentially equivalent to ignoring the global optimality conditions of the innovator (agent), just taking the first order conditions of his problem into account. Therefore, the relaxed problem still includes some aspects of the incentive constraints of the innovator to report truthfully. The result indicates that a non-constant quantity schedule can be used to screen the different types of the innovators and make sure that (first-order) innovation incentives reflect the underlying demand conditions. This feature will play a substantial role in the solution of the original problem.

Part 2: The optimal static mechanism

We now reintroduce the monotonicity constraint $q'(\theta) \leq 0$. We first show that there is efficient production for the lowest innovation type ($q(\underline{\theta}) = 1$), since for such a type there is no incentive to misreport in general. Suppose that $q(\underline{\theta}) = q^* < 1$. Then take a small deviation where for all $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$ the quantity is set at $\widetilde{q}(\theta) = 1$, and for other values of θ we maintain the original candidate optimum. We show that this increases welfare, and still satisfies all the constraints. First, it is obvious that the monotonicity constraint is still satisfied. Second, we keep $\widehat{c}(\theta)$ unchanged for all θ outside the interval. This means that for all $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$ it holds that the modified entry function $\widetilde{c}(\theta) = \widehat{c}(\underline{\theta} + \varepsilon)$ because $\widetilde{c}'(\theta) = 0$ for all $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$ as $\widetilde{q}(\theta) = 1$ for such values of θ. The original value of the entry cost is such that for all $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$ it holds that the $\widehat{c}(\theta) = \widehat{c}(\underline{\theta} + \varepsilon) - \int_\theta^{\underline{\theta}+\varepsilon} 2q(x)(1 - q(x))dx$. But then $\widetilde{c}(\theta) - \widehat{c}(\theta) = \int_\theta^{\underline{\theta}+\varepsilon} 2q(x)(1 - q(x))dx$ which goes to zero when ε goes to zero. Therefore, the component of the change in welfare that results from changing the entry unction for types in $[\underline{\theta}, \underline{\theta} + \varepsilon]$ is second order in ε. The gain in welfare that comes from the fact that quantities are increased is first order in ε. Therefore, for

a small enough ε this change is welfare improving. This concludes the proof that $q(\underline{\theta}) = 1$.

We know from above that $q(\underline{\theta}) = 1$ and that the entire solution must be constrained, since the relaxed problem has an optimal solution that violates the monotonicity constraint. Therefore, there exist $\theta', \overline{\theta}$ such that $1 \geq \theta' > \overline{\theta} > \underline{\theta}$ and the solution involves $q(\theta) = q^*$ for all $\theta \in [\overline{\theta}, \theta']$, and $q(\theta)$ is strictly decreasing on $[\underline{\theta}, \overline{\theta}]$.[26] We provide a proof by contradiction. Suppose that there exist exist $\theta'' < 1$ and $\theta''' > \theta''$ such that q is strictly decreasing on $[\theta'', \theta''']$, while $q(\theta) = q^*$ for all $\theta \in [\overline{\theta}, \theta'']$. We derive a contradiction for such a point θ'' to conclude our proof. To derive this contradiction we study an auxiliary problem. Take the solution for interval $[\underline{\theta}, \overline{\theta}]$ as given, and let us maximize the objective function $\int_{\overline{\theta}}^{1} [\theta(2q(\theta) - q^2(\theta))\widehat{c}(\theta) - \frac{\widehat{c}^2(\theta)}{2}] d\theta$ taking $q(\overline{\theta}), \widehat{c}(\overline{\theta})$ as given, and placing the further condition that

$$q(\theta) \leq q(\overline{\theta}) \text{ for all } \theta \geq \overline{\theta}. \tag{9}$$

We show that the solution of this problem is a constant path on interval $[\overline{\theta}, 1]$, and thus the required θ'', θ''' cannot exist. The Hamiltonian is unchanged as the extra constraint (9) is incorporated as a standard Kuhn-Tucker condition:

$$H = \lambda(\theta)2q(\theta)(1 - q(\theta)) + [\theta(2q(\theta) - q^2(\theta))\widehat{c}(\theta) - \frac{\widehat{c}^2(\theta)}{2}].$$

The binding monotonicity constraint on $[\overline{\theta}, \theta'']$ means that $\frac{\partial H}{\partial q} |_{q=q^*} \geq 0$ $\forall \theta \in [\overline{\theta}, \theta'']$, and in particular

$$\frac{\partial H}{\partial q} |_{q=q^*, \theta=\overline{\theta}} \geq 0. \tag{10}$$

The fact that the monotonicity constraint ceases to bind at θ'' means that

$$\frac{\partial H}{\partial q} |_{q=q^*, \theta=\theta''} = 0. \tag{11}$$

Using that $q(\theta) = q^*$ for all $\theta \in [\overline{\theta}, \theta'']$ we obtain that

$$\frac{\partial H}{\partial q} |_{q=q^*} = 2\lambda(\theta)(1 - 2q^*) + 2\theta\widehat{c}(\theta)(1 - q^*),$$

and thus

$$\frac{\partial^2 H}{\partial q \partial \theta} |_{q=q^*} = 2\lambda'(\theta)(1 - 2q^*) + 2(\theta\widehat{c}(\theta))'(1 - q^*).$$

---

[26] In other words, $\overline{\theta}$ is the lowest type where the monotonicity constraint binds in the solution of the original problem.

We know that

$$\lambda'(\theta) = \widehat{c}(\theta) - \theta(2q^* - (q^*)^2),$$

and

$$\widehat{c}'(\theta) = 2q^*(1 - q^*).$$

Therefore,

$$\frac{\partial^2 H}{\partial q \partial \theta}\Big|_{q=q^*} = 2(1 - 2q^*)\lambda'(\theta) + 2(1 - q^*)[\theta\widehat{c}'(\theta) + \widehat{c}(\theta)] =$$

$$= 2(1 - 2q^*)\left(\widehat{c}(\theta) - \theta(2q^* - (q^*)^2)\right) + 2(1 - q^*)\widehat{c}(\theta) +$$

$$+2(1 - q^*)\theta 2q^*(1 - q^*) =$$

$$= 2(\widehat{c}(2 - 3q^*) + \theta\,(q^*)^2). \tag{12}$$

Because the monotonicity constraint starts binding at $\theta = \overline{\theta}$, we can conclude two observations at that point. First, ignoring the monotonicity constraint there locally is valid, second in the relaxed problem $q'(\overline{\theta}) = 0$ holds[27]. Then the same argument as above (see (8)) implies that $\widehat{c}(\overline{\theta}))(2 - 3q^*) + \overline{\theta}\,(q^*)^2 = 0$. Therefore, $\frac{\partial^2 H}{\partial q \partial \theta}\Big|_{q=q^*, \theta=\overline{\theta}} = 0$ must hold by (12). Also, $\frac{\partial}{\partial \theta}\left(\frac{\partial^2 H}{\partial q \partial \theta}\Big|_{q=q^*}\right) = 2(\widehat{c}'(2 - 3q^*) + (q^*)^2) = 2(2q^*(1 - q^*)(2 - 3q^*) + (q^*)^2) = 2q^*(2(1 - q^*)(2 - 3q^*) + q^*) > 0$ for all relevant values of q^*. Therefore, together with $\frac{\partial^2 H}{\partial q \partial \theta}\Big|_{q=q^*, \theta=\overline{\theta}} = 0$ we obtain that for all $\theta \in (\overline{\theta}, \theta'']$

$$\frac{\partial^2 H}{\partial q \partial \theta}\Big|_{q=q^*} > 0. \tag{13}$$

But comparing (10), (11), and (13) yields a contradiction, which concludes our proof of the shape of q. Finally, $\widehat{c}(\overline{\theta}))(2 - 3q^*) + \overline{\theta}\,(q^*)^2 = 0$ implies that $q^* > 2/3$, which provides the last result.

Proof of Proposition 5

Take any (potentially non-constant) path q_t, τ_t. The proof establishes that the same entry function $\widehat{c}$ can be induced by an appropriate policy that is constant over time. Moreover, total welfare is higher under this policy as the sum of consumer and producer surplus is larger than under the original non constant policy. First, it is clear that a one-time up-front transfer is

---

[27] This is an instance of the smooth pasting condition at point $\overline{\theta}$ where the function switches from being strictly decreasing to being flat.

without loss of generality as the innovator only cares about the present value of the transfers. The utility from reporting $\widehat{\theta}$ when the type is θ is

$$U(\widehat{\theta}, \theta) = \tau(\widehat{\theta}) + \sum_{t=0}^{\infty} \delta^t p_t(\widehat{\theta}, \theta) q_t(\widehat{\theta}).$$

By construction, $p_t(\widehat{\theta}, \theta) = 2\theta(1 - q_t(\widehat{\theta}))$, and thus $U(\widehat{\theta}, \theta) = \tau(\widehat{\theta}) + 2\theta \sum_{t=0}^{\infty} \delta^t q_t(\widehat{\theta})(1 - q_t(\widehat{\theta}))$. Letting $V(\theta)$ denote the rent (under truth-telling), the envelope theorem implies that

$$V'(\theta) = \frac{\partial}{\partial \theta} U(\widehat{\theta}, \theta) \mid_{\widehat{\theta} = \theta} =$$

$$= 2 \sum_{t=0}^{\infty} \delta^t q_t(\widehat{\theta})(1 - q_t(\widehat{\theta})). \tag{14}$$

A similar argument as in Lemma 2 implies that incentive compatibility requires that $\sum_{t=0}^{\infty} \delta^t q_t(\theta)$ is decreasing in θ. Take a constant quantity scheme that satisfies $\sum_{t=0}^{\infty} \delta^t q^*(\widehat{\theta})(1 - q^*(\widehat{\theta})) = \sum_{t=0}^{\infty} \delta^t q_t(\widehat{\theta})(1 - q_t(\widehat{\theta}))$. This will then guarantee that the payoffs of the innovator, and thus the entry function is preserved.[28] It is then sufficient to prove that for any θ the realized total surplus is larger than the one under the original policy. That is, it is sufficient to show that for all θ it holds that $\sum_{t=0}^{\infty} \delta^t \theta q^*(\theta)(2 - q^*(\theta))\widehat{c}(\theta) \geq \sum_{t=0}^{\infty} \delta^t \theta q_t(\theta)(2 - q_t(\theta))\widehat{c}(\theta)$ or

$$\sum_{t=0}^{\infty} \delta^t q^*(\theta)(2 - q^*(\theta)) \geq \sum_{t=0}^{\infty} \delta^t q_t(\theta)(2 - q_t(\theta))$$

if $\sum_{t=0}^{\infty} \delta^t q^*(\theta)(1 - q^*(\theta)) = \sum_{t=0}^{\infty} \delta^t q_t(\theta)(1 - q_t(\theta))$. Using Jensen's inequality this follows if we show that $x(2 - x)$ is a concave transformation of $x(1 - x)$ restricting x to be on $[0.5, 1]$. Letting $y = x(2 - x)$ and $z = x(1 - x)$ it holds that $y = z + x$. So, it is sufficient to show that y is concave in z for which it is sufficient that x is concave in z. But this holds because z is a concave and decreasing function of x.

---

[28] The incentive conditions are not affected either, see (14).

Appendix 2: Additional Results

Formalization of Switching Time

The primitive of the planner's buyout policy is a function $\tau_t : H_t \to \{0, 1\}$ indicating whether the switch to a competitive market has occurred at or before time t given the history h_t. Let us define as $H_j(h_t)$ the set of histories at time $j > t$ following a history h_t. To interpret $\tau_t(h_t)$ as an irreversible switch to a competitive market we require that $\tau_t(h_t) = 1 \Rightarrow \tau_j(H_j(h_t)) = 1$ for each $j > t$.

We start by defining the set of admissible histories in each period $t \geq 1$. The set of admissible histories in period 1 consists of all positive price-quantity pairs if $\tau_0 = 0$ but the price is restricted to be equal to zero if $\tau_0 = 1$. Formally:

$$H_1 = \{x \in \mathbb{R}_+^2 : x = (q, \widehat{p}), q \in \mathbb{R}_+, \widehat{p} = 0 \text{ if } \tau_0 = 1\}.$$

An inductive step defines the set of admissible histories H_t for all $t \geq 2$

$$H_t = \{x \in \mathbb{R}_+^{2t} : x = (y, q, \widehat{p}), \ y \in H_{t-1}, q \in \mathbb{R}_+, \widehat{p} = 0 \text{ if } \tau_{t-1}(y) = 1\}.$$

We are ready to define the switching time T taking the planner's policy and the innovator's strategy as given. Given any strategy of the innovator $\alpha \in A$, let α_t denote the truncation of α up to period t. We indicate with $h_t(\alpha_t)$ the admissible public history generated by α_t. Taking the policy of the planner $\tau = (\tau_0, \tau_1, \tau_2, ...)$ as given, the switching time $T(\alpha)$ is defined as follows: $\tau_k(h_k(\alpha_k)) = 0$ for all $k \leq T - 1$ and $\tau_T(h_T(\alpha_T)) = 1$.[29]

Markov Shifts

We extend our setting and assume that the demand has two states. Let us indicate with $D_L(p, \theta)$ the quantity consumed in the low demand state and with $D_H(p, \theta)$ the quantity consumed in the high demand state. For simplicity, we assume that $D_H(p, \theta) \geq D_L(p, \theta)$ for each p and that the inequality is strict if $D_H(p, \theta) > 0$.[30] We follow Battaglini (2005) and denote with $\Pr(D_L | D_k) \in (0, 1)$ the probability that state L is reached if the demand is in

---

[29] Note, that function τ_T is defined only on histories such that switching has not occurred by period $T - 1$, but this is satisfied by assumption here.

[30] Proposition A1 holds as long as $D_H(p, \theta) \neq D_L(p, \theta)$ for $p \in (0, \varepsilon)$ with ε arbitrarily close to zero.

state k. At date zero the prior on the demand states are (μ_H, μ_L). In this extended setting the problem for the inventor is to choose

$$\max_{p_t} \tau_t(r(p_t), h_{t-1}) + \delta E\left[V\left(D\left|h_t, \theta, D_t\right.\right)\right]$$

where $V\left(D\left|h_t, \theta, D_t\right.\right)$ is the value function of an innovator type θ after public history h_t at the demand state D_t. Investment in innovation takes place if $\mu_H V\left(D\left|h_0, \theta, D_H\right.\right) + \mu_L V\left(D\left|h_0, \theta, D_L\right.\right) \geq c$ and the total social welfare created by the innovation is

$$\int_c \int_{\theta \in \Theta^*(c)} \left[\sum_{t=0}^{\infty} \sum_{i \in \{L,H\}} \delta^t S(D_t(p_t^*)) \Pr(D_t = D_i) - c\right] \psi(\theta, c) d\theta dc.$$

Also in this setting the planner can maximize innovation incentives by approximating the first best outcome.

Proposition A1 *If D_L and D_H are analytic the first best can be approached arbitrarily closely.*

Proof. As in the proof of Proposition 1 we approximate the demand functions by polynomials that are estimated by charging $n^2 + 1$ distinct prices close to $p^* = 0$. For the estimation we now need two different quantities for each of these prices. The smaller quantity observed at a price is used for the estimation of D_L and the larger one to estimate D_H. Once the two demand functions have been approximated around $p^* = 0$, analyticity can be exploited to learn their global behavior by following the procedure in the proof of Proposition 1. By choosing and experimentation interval arbitrarily close to $p^* = 0$ and n arbitrarily large the demands are approximated arbitrarily closely at an arbitrarily low welfare cost. ■

Noise

Our setting assumes that the planner can perfectly observe the demand. We can relax this assumption and consider the case in which the demand is observed with error. To analyze such a setting, we assume that:

$$q_t = D(p_t, \theta) + \varepsilon_t, \tag{15}$$

where ε_t is a mean zero i.i.d. noise over the support $[-\overline{\varepsilon}, \overline{\varepsilon}]$. In the next proposition we show that even in this case the planner can estimate the surplus generated by the innovation and transfer it to the innovator.

Proposition A2 *If D is analytic, the first best can be approached arbitrarily closely.*

Proof. As in the proof of Proposition 1 we approximate the demand function by a polynomial estimated by charging $n^2 + 1$ distinct prices close to $p^* = 0$. For the estimation we now need N different quantities for each of these prices. Once N quantities are observed at a price p, $N^{-1} \sum_{i=1}^{N} q(p)$ is used for the estimation of D. Because of the weak law of large numbers, the sample average converges in probability to $D(p, \theta)$. Once the demand function has been approximated around $p^* = 0$, its analyticity can be exploited to learn its global behavior exploiting the procedure illustrated in the proof of Proposition 1. By choosing and experimentation interval arbitrarily close to $p^* = 0$ and N, n arbitrarily large, the demand is approximated arbitrarily closely at an arbitrary low welfare cost. ∎

Social Cost of Public Funds

Following Laffont and Tirole (1993) and Galasso and Tombak (2014) we assume that the government finances transfer T at a cost $(1 + \kappa)T$ where $\kappa \geq 0$ represents the cost of public funds due to the deadweight loss associated with taxation. We start by characterizing the first best in the case in which the planner knows the demand parameter θ. Consider a constant price p and a transfer of T per period.[31] The per-period product market surplus net of the deadweight loss associated with taxation is equal to

$$S(p, \theta) = pD(p, \theta) + \int_p^\infty D(z, \theta)\, dz - \kappa T.$$

The innovator invests if and only if

$$\frac{pD(p, \theta) + T}{1 - \delta} \geq c.$$

Let us indicate with $\widehat{c}$ the marginal innovator (whose profits are zero) and let $\widehat{C} = \widehat{c}(1 - \delta)$. If the planner aims to induce entry of innovator $\widehat{c}$ the problem becomes

$$\max_{T, p \geq 0} \frac{1}{1 - \delta} \left(pD(p, \theta) + \int_p^\infty D(z, \theta)\, dz - \kappa T \right)$$
$$\text{such that} \quad pD(p, \theta) + T - \widehat{C} = 0$$

---

[31] We discuss the optimality of a constant price path in Section 6.2.

The corresponding Lagrangian is

$$L = \frac{1}{1-\delta} \left(pD\left(p,\theta\right) - \int_{\infty}^{p} D\left(z,\theta\right)dz - \kappa T \right) + \lambda \left(\widehat{C} - pD\left(p,\theta\right) - T \right),$$

and the first order conditions are[32]

$$\frac{1}{1-\delta}\left(D + pD' - D\right) - \lambda\left(D + pD'\right) = 0$$

$$-\frac{\kappa}{1-\delta} - \lambda = 0.$$

These conditions yield that the optimal price p^* satisfies

$$\frac{\varepsilon(p^*)}{1+\varepsilon(p^*)} = -\kappa \tag{16}$$

where $\varepsilon(p) = pD'/D$ is the (negative) elasticity of the demand. To induce an innovator with cost $\widehat{c}$ to enter, the planner faces a trade-off between two types of welfare distortions: the cost of raising money through public taxation and the surplus losses due to market power. In a simple linear setting where $p = 1 - q$ the condition implies an optimal price $p^* = \kappa/(2\kappa+1)$ that ranges from 0 (in the case of no cost of public funding) to the monopoly level (when $\kappa = \infty$).

Let us denote with $p^*(\kappa,\theta)$ the price that satisfies condition (16) and with $T(\theta,\kappa,\widehat{c}) = \widehat{c} - p^*D\left(p^*,\theta\right)$ the transfer that induces entry of innovators with $c \leq \widehat{c}$. The maximization problem of the planner is then equivalent to finding the optimal entry level $\widehat{c}$:

$$\max_{\widehat{c}} \int_0^{\widehat{c}} \frac{1}{1-\delta} \left(p^*(\kappa,\theta)D\left(p^*(\kappa,\theta),\theta\right) + \int_{p^*(\kappa,\theta)}^{\infty} D\left(z,\theta\right)dz - \kappa T(\theta,\kappa,\widehat{c}) - c \right) \psi(\theta,c)dc.$$

Let us denote with $\widehat{c}^*(\kappa,\theta)$ the solution to this maximization problem. From the values of $\widehat{c}^*(\kappa,\theta)$ and $p^*(\kappa,\theta)$ we obtain the optimal transfer $T^*(\theta,\kappa)$.[33]

The above results show that with known demand, the first best is reached with prices equal to $p^*(\kappa,\theta)$ quantity $D\left(p^*(\kappa,\theta),\theta\right)$ and transfer $T^*(\theta,\kappa)$. We now show that the first best can be approximated arbitrarily closely even if the demand, θ, is unknown. We attack the problem with a mechanism design approach in which the innovator reports to the planner a

---

[32] By using an equality version of the first order condition with respect to p, we assumed that $p > 0$ in the optimum. The positivity of p follows if $D'(0) > -\infty$ because then increasing the price slightly above zero yields a second order loss for the consumers but a first order gain in terms of need for pulic funds.

[33] In the linear case $\widehat{c}^* = (2\kappa + \kappa^2 + 1)/(8\kappa + 8\kappa^2 + 2)$ and $T^*(\theta,\kappa) = \left(3\kappa^2 + 4\kappa + 1\right)/2\left(2\kappa+1\right)^2$.

type, $\widehat{\theta}$, and the planner requires that in period t the innovator produces a specific market outcome $q_t(\widehat{\theta}, h_t)$, $p_t(\widehat{\theta}, h_t)$ and receives a payment $\tau_t(\widehat{\theta}, h_t)$.

Consider the following mechanism. The innovator reports $\widehat{\theta}$ and the planner requires N quantity observations arbitrarily close to $D\left(p^*(\kappa, \widehat{\theta}), \widehat{\theta}\right)$. Set N large enough so that the analytic demand can be identified. Each time a quantity is produced, the planner assesses whether the prices are consistent with the revealed demand function, i.e. the market outcome lies on the demand curve $D\left(., \widehat{\theta}\right)$. If the quantity produced is not the one requested by the planner or the price is not consistent with the demand, then the innovator is punished and τ_t is set to $-\infty$ ever after. If the quantity is the one requested by the planner, then the transfer is $T^*(\theta, \kappa)$. After the first N observations the demand has been identified, and the reward is equal to $T^*(\theta, \kappa)$ per-period ever after conditional on observing $D\left(p^*(\kappa, \widehat{\theta}), \widehat{\theta}\right)$. It is easy to see that the innovator has no incentive to report his type untruthfully and the first best is approximated arbitrarily closely.

Demand Growth

A natural assumption with new technologies is that demand grows over time. Suppose, for example that for τ periods the demand is $D_L(p, \theta)$ and it becomes $D_H(p, \theta)$ from period $\tau + 1$ with $D_H(p, \theta) > D_L(p, \theta)$. If the functions are polynomials:

$$D_\alpha(p, \theta) = \sum_{i=0}^{I} c_i^\alpha(\theta) p^i$$

with $\alpha \in \{L, H\}$ then, under the restriction that only one price-quantity observation can be obtained in each period, the amount of time required to identify the low state demand is increasing in the complexity of the demand.

This simple specification suggests that when the demand does not grow too quickly, the first best can be implemented since the planner can learn the parameters of the demand fast enough. In particular, when $\tau \geq I + 1$ the first best can be approached arbitrarily closely: it takes $I + 1$ distinct price-quantity observations to identify all the coefficients of the polynomial. By requiring the innovator to charge in each period a distinct p_t arbitrarily close to zero the welfare cost of learning is minimized.

Nevertheless, the planner may not have enough time to learn the demand when growth is

fast. Take for example the case in which the demand is linear and the planner can observe only one price-quantity combination for the low demand regime. In this case the planner cannot approach the first best and will have to reward the innovator for the surplus generated in the low demand state by granting a one period patent or by using the Weyl and Tirole (2012) mechanism for one period.

It is important to notice that when $\tau < I + 1$ it is not optimal to give a τ period patent and then learn costlessly the demand $D_H(p, \theta)$. This is because a patent that lasts τ periods generates a loss in consumers' surplus in each period. The planner can improve the overall welfare by granting a patent that lasts only for one period and observe the quantities and prices practiced by the innovator. For periods 2 to τ the planner can transfer an amount equal to the observed first period profits to the innovator under the requirement that the product is sold at marginal cost. In this case the innovation incentives are the same as with a τ periods patent but the loss in consumer surplus is substantially lower.

Demand identification may be problematic also when the demand starts at a high level and then suddenly drops or disappears. This may occur when a follow-on superior technology is developed. Also in this case, the planner may not be able to reach the first best if the high demand state does not last for a period of time long enough to identify the demand curve.[34]

This discussion suggests that it is crucial for the planner to collect market outcomes in a timely manner. Nonetheless, the restriction that only one price-quantity can be observed in each period can be relaxed if the planner can generate variation geographically. When sudden demand shifts are expected, the planner may prefer to collect market outcomes through geographic (cross-markets) price variation rather than intertemporal (within market) price variation.

---

[34]In this case the planner may actually use intellectual property protection to prevent the new innovator to sell the innovation until the surplus generated by the previous innovator is estimated. Nonetheless this delay would affect negatively consumers surplus. A more careful examination of how market outcomes may help designing patent protection in the presence of cumulative innovation is left to future research.