

Market Outcomes and Dynamic Patent Buyouts¹

Alberto Galasso
University of Toronto

Matthew Mitchell
University of Toronto

Gabor Virag
University of Toronto

August 23, 2013

¹We thank seminar participants at the University of Toronto. We are grateful for financial support from the University of Toronto Michael Lee-Chin Family Institute for Corporate Citizenship.

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 describe situations 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 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.

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 question in the economics of innovation, therefore, is how one can plan a mechanism that solves this under-investment problem in a manner that is least costly to society. 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.

The leading method of rewarding innovation, the patent system, is under fire. There are growing concerns voiced in both academic and policy debates about the potentially deleterious effects of patents in the modern innovation landscape. 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).

Potential alternatives exist. It is well known that if the policy maker observes the market demand for a new product, then the first-best can be implemented by simply setting a prize that rewards high quality research in proportion to the surplus obtained by the consumers. In reality, however, the policy maker often lacks information about the cost of and the demand for innovative products. Under limited information, 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. The main objective of this paper is to explore alternatives to patents and to identify mechanisms that generate greater innovation incentives and increase social welfare if the policy maker has less than perfect information about market demand.

Our paper stresses the dynamic approach to innovation rewards. In contrast with the previous literature, we assume that the policy maker can learn over time about market condi-

tions 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. 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 environments we study, 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 environments, the optimal mechanism involves some reward through a contingent prize near the end of the period in which the innovator is rewarded. The optimal policy in many cases is a sort of hybrid between a patent and a prize, using prices above marginal cost initially, but then moving 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). In the sense that the optimal policies we study are a hybrid between a patent and a prize, this logic remains in effect; the planner simply tries to limit the damage that the patent inevitably causes by eventually using a cash reward as a sort of patent buyout.

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 infor-

mation. 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, Kermer and Williams, 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 copayments 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.

¹Similar 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.

²For 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.

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

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

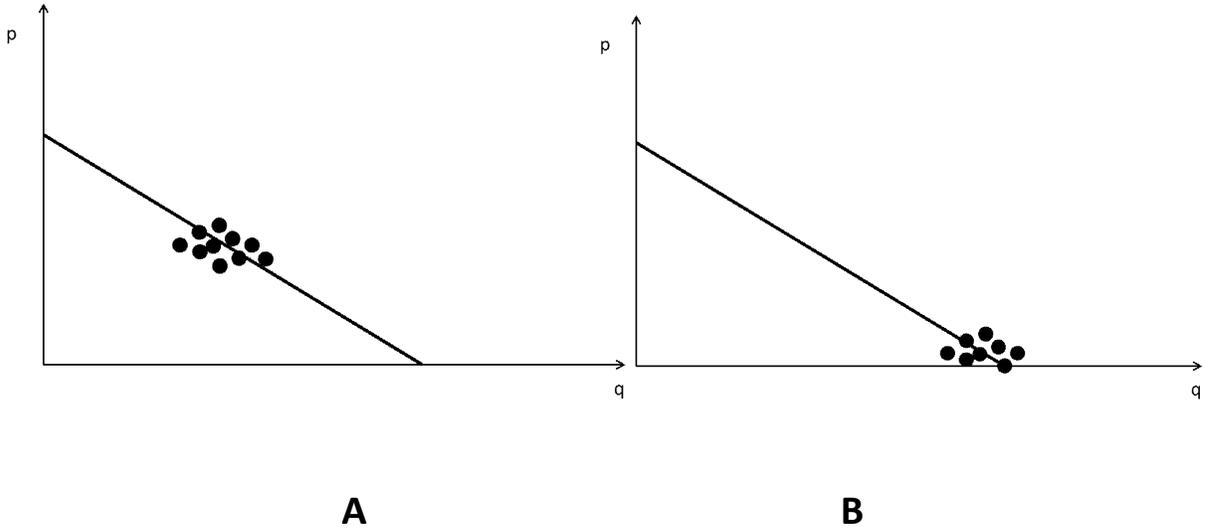


Figure 1: Market Outcomes and Demand Identification

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 when price manipulation is costless for the innovator. We show that even in this case the optimal mechanism differs substantially from a patent. 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 the absence of demand manipulation. Section 5 introduces costly demand manipulation. Section 6 studies the opti-

mal mechanism in the presence of costless demand manipulation. Section 7 summarizes and concludes. Most of the proofs are in the Appendix.

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 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 to the auction. In our model, we depart from this assumption and assume that only the innovator knows how valuable an innovation is.

Shavell and Van Ypersele (2001) provide a comparison of prizes and patents as mechanisms to incentivize innovation. In a static framework, they show that an optional reward system in which an innovator can choose between patents and rewards dominate patents. Chari, Golosov and Tsyvinski (2012) compare prizes and patents when the planner can observe market signals. Most of their analysis is under the assumption that both the innovator and his competitors know the value of the innovation as in Kremer (1998). Their main finding is that patents are necessary if the innovator can manipulate market signals. In our model 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. Weyl and Tirole (2012) study the optimal reward structure in the presence of multidimensional heterogeneity and observable market outcomes. In a static framework, 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 these models because we introduce dynamics and allow the innovator to

manipulate market outcomes.

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 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 when 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 the optimal patent scheme is typically differentiated and can be implemented through menu of patent renewals.

Hopenhayn and Mitchell (2001) and Hopenhayn, Llobett and Mitchell (2006) study the optimal patent design when innovation is cumulative and each discovery is a building block to future innovations. Hopenhayn and Mitchell (2001) consider the case in which idea quality is private information and there are two technology generations. They show that 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, Llobett and Mitchell (2006) study a dynamic framework with multiple cumulative innovations and private information about the quality of ideas and R&D investments. They show that in such environment the optimal mechanism is a patent buyout scheme where the innovator commits to a price ceiling at which he sell his rights to a future inventor.

3 The Baseline Model

Time is discrete and the horizon is infinite. Each innovation is characterized by an ex-ante cost of creating an innovation, $c \in C \subset \mathbb{R}_+$, and a demand function $q_t = D(p_t, \theta)$ with $\theta \in \Theta$ compact subset of \mathbb{R}^N . To ensure the concavity of the static profit function we assume that $D_p(p, \theta) < 0$ and $D_p(p, \theta) + pD_{pp}(p, \theta) < 0$ for each $p \geq 0$. We assume that the marginal cost of production is known and normalize it to zero. We assume that θ and c are private information for the innovator and are distributed according to some smooth probability density function $\psi(\theta, c)$ 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

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

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.

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). We refer to the Appendix for a formalization of the planner's policy.

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 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.³ 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 the Appendix we provide a recursive definition of $T(\alpha)$ given the planner's policy function.

Then the innovator's maximization problem is

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

³Note, 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.

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 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 $\{g_t\}_{t=1,2,3,\dots}$ and T is equal to:

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

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:

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

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 $q_t \neq D(0, \theta)$).

Innovation Rewards

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

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:

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

Buyouts

The following specification

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

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. $\phi(\hat{p}_t, p_t) = \infty$ if $\hat{p}_t \neq p_t$.

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.⁴

The above example suggests that transfers that depend on market outcomes can be powerful mechanisms to incentivize innovation. 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 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 environment the planner can substantially improve welfare and innovation incentives relative to patent systems or static multidimensional screening mechanisms as 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 $D(p, \theta)$ 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 $|z - p| < r$:*

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

The next result generalizes the result obtained for linear demands to analytic demand functions.

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

⁴A 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 $h_t = \{q_1, 0, q_2, \varepsilon, q_3, 0, \dots, q_t, 0\}$ and $g_t(h_t) = -\infty$ otherwise.

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

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

Notice that the coefficients of the polynomial can be estimated by charging $n^2 + 1$ distinct prices close to p^* . By setting p^* arbitrarily close to zero this is equivalent to approximating the demand function by distorting the first best welfare only very slightly. As n gets large the approximation improves. If $r < \bar{p}$ 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 $l = \bar{p}/n$ and take n large enough such that $l < r$. We can approximate $D(\cdot, \theta)$ in the interval $[p^*, p^* + l]$ by

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

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 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. ■

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 experimenting around $p = 0$ and choosing a very small price interval for experimentation, the planner minimizes the welfare cost of learning and approximates the first best complete information outcome.

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).

Structural demand estimation studies typically estimate the primitives of a model from local price variation and exploit these estimates for out-of-sample welfare analysis. Our result suggests that policy-makers may affect innovation incentives by designing reward systems that exploit these techniques. Specifically, the policy-maker can request the innovator to generate some price variation and then use this variation to construct a patent buy-out transfer that compensates the innovator for the surplus generated by the innovation. By keeping the price variation concentrated around the marginal cost of production, the policy-maker can limit the loss of surplus associated with learning to a minimum.

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. 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.

But this is 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. 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, *for a large class of demand functions the socially optimal innovation level can be reached through minor perturbations of a competitive market.*

Robustness

We explore robustness of our result extending the setting in two ways. First, we study shifts in the demand. Second, we look at noisy demand.

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 $D_H(p, \theta) > 0$. 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 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

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

where $V(D | h_t, \theta, D_t)$ 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 $\mu_H V(D | h_0, \theta, D_H) + \mu_L V(D | h_0, \theta, D_L) \geq c$ and the total social welfare created by the innovation is

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

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

Proposition 2 *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 $p^* = 0$, its analyticity can be exploited to learn their global behavior by following the procedure in the proof of Proposition 1. By choosing an experimentation interval arbitrarily close to $p^* = 0$ and n arbitrarily large the demand is approximated arbitrarily closely at an arbitrarily low welfare cost. ■

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.⁵

Proposition 2 implies 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 (as in Battaglini, 2005).

Noise

Our baseline 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:

$$q_t = D(p_t, \theta) + \varepsilon_t, \tag{2}$$

where ε_t is a mean zero i.i.d. noise over the support $[-\bar{\varepsilon}, \bar{\varepsilon}]$. 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.

⁵From the linear example it is also clear that we do not need that $D_H \geq D_L$. For analytic demand functions, the result holds as long as $D_H(p, \theta) \neq D_L(p, \theta)$ for $p \in (0, \varepsilon)$ with ε arbitrarily close to zero.

Proposition 3 *If D is 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 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. ■

To understand the intuition for the proof consider the case of linear demand. The planner can use the following two step scheme to obtain the required estimate for (θ_1, θ_2) . 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 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.⁶

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 $\tau + 1$

⁶In 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).

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.⁷

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.

5 Demand Manipulation

The previous analysis was conducted under the assumption of no demand manipulation. In this section we consider the case in which the innovator can manipulate the market outcomes.

Buyouts and Price Variation

In our baseline model the innovator can affect the market outcomes and manipulate market signals received by the planner up to period T but not after T . Our baseline 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 the baseline 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

⁷In 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.

⁸Noticed 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 $p_t \neq 0$.

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 as manipulation after buyout it 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 first best, it has either to be the case that (i) manipulation is feasible both before and after the buyout or that (ii) the ability of the planner to generate price variation after the buyout must be limited. In the next section, building on our baseline model, we study an environment in which neither the planner nor the innovator can manipulate the competitive outcome.

Post-Buyout Competitive Outcome

Following our baseline model, 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. To model the (pre-buyout) manipulations, we assume that 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. This may arise, for example, when the innovator awards secret discounts to his consumers.

We provide a micro-foundation of the manipulation cost $\phi(\hat{p}_t, p_t)$ and assume that the innovator can convince the planner that he is selling at $\hat{p} > p$ by sustaining a cost that depends on $(\hat{p} - p) D(p, \theta)$ and that we indicate as $\hat{\phi}((\hat{p} - p) D(p, \theta))$ with $\hat{\phi}' > 0$ and $\hat{\phi}'' \geq 0$. Intuitively, the planner observes sales equal to $\hat{p}D(p, \theta)$ whereas the true revenue is equal to $pD(p, \theta)$ and $(\hat{p} - p) D(p, \theta)$ are fake revenues undermined by secret price discounts. A simple justification

of a positive manipulation cost is if 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 $\hat{p}D(p, \theta)$ the innovator will have to borrow $(\hat{p} - p)D(p, \theta)$ sustaining a cost of $\hat{\phi}((\hat{p} - p)D(p, \theta))$.⁹ A simple functional specification for the manipulation cost is $i(\hat{p} - p)D(p, \theta)$, if $i > 0$ there is a positive cost of manipulating sales.

Proposition 4 *For each patent of length T there is a patent buyout scheme that Pareto dominates the patent.*

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

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

Now consider setting \hat{T} such that

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

so that the marginal benefit of manipulation when $\hat{p} = p$ is exactly equal to the marginal cost.¹⁰ Setting $\hat{p} = p$ is then optimal for the innovator because the first order condition holds by construction and the objective function is concave in \hat{p} . This removes the innovator's incentive to manipulate. Maximizing the payoff respect to p (with $\hat{p} = p$) gives:

$$\begin{aligned} & (1 - \delta^{\hat{T}}) \left[pD'(p, \theta) + D'(p, \theta) - \hat{\phi}'(0)\hat{p}D' + \hat{\phi}'(0)(pD'(p, \theta) + D'(p, \theta)) \right] + (\delta^{\hat{T}} - \delta^T)\hat{p}D'(p, \theta) \\ & = (1 - \delta^T)(pD'(p, \theta) + D'(p, \theta)) \end{aligned}$$

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

¹⁰ If \hat{T} is not an integer set it equal to the smallest integer for which $\delta^{\hat{T}} - \delta^T < (1 - \delta^{\hat{T}})\hat{\phi}'(0)$.

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. ■

The proposition shows that 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 $\hat{T} < T$ periods. The buyout time \hat{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 \hat{T} periods is equal to the marginal benefit of obtaining extra buyout reward.

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

$$\frac{\delta^{\hat{T}} - \delta^T}{1 - \delta^{\hat{T}}} = i \quad (3)$$

that indicates how patent buyout takes place sooner as i gets larger. This result is reminiscent of Chari et al (2012) who only consider patents and prizes, and show that 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. However, in the next Section we show that while this result is interesting, it has some limitations: Even with costless manipulation of the price signals (when an infinitely lived patent is implied by (3)), one can do better by considering mechanisms that are different from both prizes and patents.¹¹

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 $\hat{T} + 1$. The innovator will obtain the monopoly profits before the buyout and the entire consumer

¹¹It is important to note that there are several 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".

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 outcome resembles the first best after the buyout, because there is marginal cost pricing and all the surplus is transferred to the innovator.¹²

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

Proposition 5 *When $\hat{\phi}(\hat{p}, p) = 0$ there is a per unit subsidy level τ that Pareto dominates patents.*

Proof. See Appendix. ■

Intuitively, even when price manipulation is costless, the planner can improve upon patents by exploiting the observed quantities. Proposition 5 shows that in our setting a minor perturbation of the patent system with small quantity subsidies improves social welfare. Overall, Propositions 4 and 5 show that 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 the case of costless demand manipulation. When $i = 0$ 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

¹²In 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 $\gamma\hat{p}D(p, \theta)$, pays no interest and expires after T^B periods. If $pD(p, \theta)$ is the only revenue available to the innovator, he will have to borrow $\gamma\hat{p}D(p, \theta) - pD(p, \theta)$ for T^B periods at a cost of

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

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

(other than lowering revenues). As in the previous 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, $\hat{\theta}$, and the planner requires the innovator to produce a specific quantity, $q(\hat{\theta})$, and to receive payment $\tau(\hat{\theta})$. To simplify the analysis we focus on the linear demand case $D(q) = \theta_1 - \theta_2 p$. In the next lemma we show that there is no loss of generality in assuming that the planner knows the intercept θ_1 .¹³

Lemma 1 *The innovator will truthfully report the intercept.*

Proof. After a 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 $\hat{\theta}_1$. ■

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 $q^M = 1/2$.

Static Mechanism

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

¹³We do not allow 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:

$$\begin{aligned} U(\hat{\theta}, \theta) &= \tau(\hat{\theta}) + p(\hat{\theta}, \theta)q(\hat{\theta}) \\ &= \tau(\hat{\theta}) + 2\theta(1 - q(\hat{\theta}))q(\hat{\theta}). \end{aligned}$$

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

$$\begin{aligned} V'(\theta) &= \frac{\partial}{\partial \theta} U(\hat{\theta}, \theta) \Big|_{\hat{\theta}=\theta} \\ &= q(\theta) \frac{\partial}{\partial \theta} p(\hat{\theta}, \theta) \Big|_{\hat{\theta}=\theta} = 2q(\theta)(1 - q(\theta)). \end{aligned} \quad (4)$$

The above condition (4) is a first order condition. The following result states a necessary and sufficient condition for implementability:

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

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

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

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. ■

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

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

The planner's problem is

$$\begin{aligned} & \max_{q(\theta)} \Pi \\ \text{s.t. } & \widehat{c}'(\theta) = V'(\theta) = 2q(\theta)(1 - q(\theta)), \text{ and } q'(\theta) \leq 0 \ \forall \theta \in [\underline{\theta}, \bar{\theta}]. \end{aligned}$$

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):

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

Solution for the relaxed problem

To develop some intuition for the optimal mechanism we simplify the problem by looking at the optimal control problem ignoring the $q'(\theta) \leq 0$ constraint first. To obtain a solution continuous 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

$$0 = \frac{\partial H}{\partial q} = \lambda(\theta)2(1 - 2q(\theta)) + 2\theta(1 - q(\theta))\widehat{c}(\theta), \ \forall \theta. \quad (5)$$

The other co-state equation is

$$-\lambda'(\theta) = \frac{\partial H}{\partial \widehat{c}} = \theta(2q(\theta) - q^2(\theta)) - \widehat{c}(\theta). \quad (6)$$

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

$$\lambda(\underline{\theta}) = \lambda(1) = 0. \quad (7)$$

The above conditions lead to the following result.

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

Proof. From (5) we obtain that

$$q(\theta) = \frac{\lambda(\theta) + \widehat{c}(\theta)\theta}{2\lambda(\theta) + \widehat{c}(\theta)\theta} \quad (8)$$

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

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. Then next proposition shows that this is not the case, and that the optimal quantity schedule is non-monotonic.

Proposition 6 *There exists a $\bar{\theta}$ such that $q(\bar{\theta}) < 1$ and $q'(\bar{\theta}) = 0$. Moreover $q' \leq 0$ for $\theta \in [\underline{\theta}, \bar{\theta}]$ and $q' > 0$ for $\theta \in (\bar{\theta}, 1]$.*

Proof. See Appendix. ■

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.

Solution of the original problem

We now reintroduce the monotonicity constraint $q'(\theta) \leq 0$. The following result shows that there is efficient production for the lowest innovation type, since for such a type there is no incentive to misreport in general.

Lemma 4 *It holds that $q(\underline{\theta}) = 1$ for the solution of the original problem.*

Proof. See Appendix. ■

Building on the previous lemma, in the next proposition we characterize the optimal quantity schedule in the presence of costless price manipulation.

Proposition 7 *In the solution for the original problem there exists $\bar{\theta} \in (\underline{\theta}, 1)$ such that it holds that q is strictly decreasing on interval $[\underline{\theta}, \bar{\theta}]$ and then constant on $[\bar{\theta}, 1]$. Moreover, $q(\theta) \geq 2/3$ for all θ .*

Proof. See Appendix. ■

Proposition 7 shows that 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, 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.

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

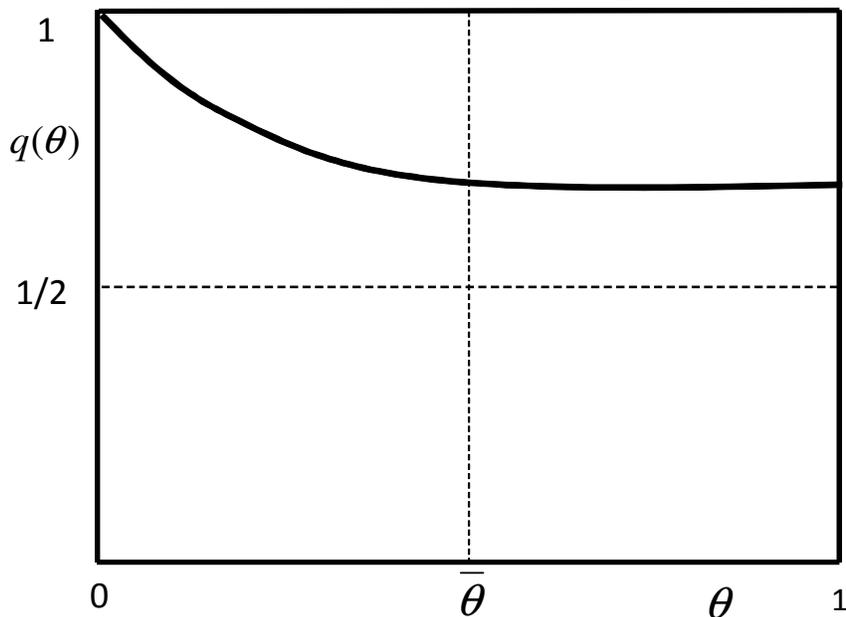


Figure 2: Optimal quantity schedule with costless price manipulation

willingness to pay is similar to the logic in Weyl and Tirole (2012). They restrict their attention to Cobb-Douglas reward policies that in our setting would generate a constant level of q across types, and 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 for different types.

Dynamic Mechanism

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 $t \geq 0$.¹⁴ Our main result shows that repeating the same quantity over time for all types θ is optimal.

Proposition 8 *It is optimal for the planner to set a policy where $q_t(\theta)$ is constant in time for any θ .*

Proof. Take any (potentially non-constant) path q_t, τ_t . The proof establishes that the same entry function \hat{c} can be induced by an appropriate policy that is constant over time.

¹⁴Since 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.

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 onetime up-front transfer is without loss of generality as the innovator only cares about the present value of the transfers. The utility from reporting $\hat{\theta}$ when the type is θ is

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

By construction, $p_t(\hat{\theta}, \theta) = 2\theta(1 - q_t(\hat{\theta}))$, and thus $U(\hat{\theta}, \theta) = \tau(\hat{\theta}) + 2\theta \sum_{t=0}^{\infty} \delta^t q_t(\hat{\theta})(1 - q_t(\hat{\theta}))$.

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

$$\begin{aligned} V'(\theta) &= \frac{\partial}{\partial \theta} U(\hat{\theta}, \theta) \Big|_{\hat{\theta}=\theta} = \\ &= 2 \sum_{t=0}^{\infty} \delta^t q_t(\hat{\theta})(1 - q_t(\hat{\theta})). \end{aligned} \tag{9}$$

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 $\sum_{t=0}^{\infty} \delta^t q^*(\hat{\theta})(1 - q^*(\hat{\theta})) = \sum_{t=0}^{\infty} \delta^t q_t(\hat{\theta})(1 - q_t(\hat{\theta}))$. This will then guarantee that the payoffs of the innovator, and thus the entry function is preserved.¹⁵ 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 $\sum_{t=0}^{\infty} \delta^t \theta q^*(\theta)(2 - q^*(\theta))\hat{c}(\theta) \geq \sum_{t=0}^{\infty} \delta^t \theta q_t(\theta)(2 - q_t(\theta))\hat{c}(\theta)$ 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 . ■

Proposition 8 confirms that even in a dynamic setting with costless manipulation the optimal mechanism differs dramatically from a patent system. Our result shows that welfare is maximized with the innovator selling a constant quantity over time until the buyout occurs.

¹⁵The incentive conditions are not affected either, see (9).

In particular, it is not the case that the quantity sold is equal to the monopoly level for a period of time and then it switches to the competitive level as in the patent system described by Scotchmer (1995).

A constant mechanism is optimal because of the desirable features of quantity (and price) smoothing over time. Specifically, a constant quantity path increases total surplus on the product market while maintaining R&D incentives. Indeed, the proof of the Proposition finds a constant price (quantity) path, which provides the same profits to the innovator as a time varying path, but yields a higher consumer surplus. The key advantage of constant quantity mechanisms comes from the fact that the 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. This finding resembles the result of Gilbert and Shapiro (1990) who show how 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 has the same effect of an increase in patent breadth because consumer surplus is reduced and innovator profits increase.

Notice that there is a tension between Proposition 8 and Lemma 1. Proposition 8 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 associated with the ability to observe market outcomes. 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 where the patent is acquired after a long time (as long as possible) has elapsed.¹⁶ A key difference with a traditional patent buyout is that the planner restricts the innovator to produce a specific quantity before the patent is acquired. A possible way to implement the mechanism is to offer a menu of quantities and transfers. Each innovator will then sell his chosen quantity forever and receive a lump sum R&D subsidy from

¹⁶In 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.

the government.

The optimal reward structure involves two distinct but interrelated trade-offs that the social planner faces to incentivize innovation. First, there is a static trade-off associated with variation in market power “between innovations”. Market power allows the planner to screen consumers’ willingness to pay and thus subsidize entry of more valuable innovations but generates pricing distortions and therefore reduces consumer surplus. Second, there is a dynamic trade-off associated with variation of market power “within innovations” over time. This variation allows the planner to identify the demand parameters but it is costly for consumers that are better-off with a constant price.

Finally, Proposition 8 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 our result (no learning is optimal until the buyout) is only true because the agent (the innovator) has strictly superior information over the planner, and any learning over time is public, so this advantage is maintained over time. It is beyond the scope of the current paper to assess the role of screening versus learning in the optimal mechanism if dynamic market signals are available to the planner in a general framework.

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 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 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.¹⁷

Our paper provides two main insights into the design and application of such buyout schemes. First, the planner may find it beneficial to induce the innovator to generate price variation before the buyout. This is because the market outcomes generated before the buyout provide information that is useful to estimate the surplus generated by the innovation. Second, the planner may also find it beneficial to collect information on market outcomes after the buyout. Such post-buyout market outcomes not only provide information about the surplus generated but also help identifying demand manipulation that took place before the buyout. Our model also suggests that gains in welfare can be generated by keeping the price variation as close as possible to marginal costs. This occurs if local price variation is sufficient to learn about the global properties of the demand and the total surplus generated by the innovation that is typically the case for demand functions employed in modern industrial organization. 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).

¹⁷A similar policy proposal is described in Guell and Fischbaum (1995).

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
- [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] Glennerster, Rachel, Michael Kremer and Heidi Williams (2006), "Creating Markets for Vaccines," *Innovations: Technology Governance Globalization* 1: 67-79
- [17] Gilbert, Richard and Carl Shapiro (1990), "Optimal patent length and breadth," *RAND Journal of Economics* 21:106-112
- [18] Guell, Robert and Marvin Fischbaum (1995), "Toward Allocative Efficiency in the Prescription Drug Industry," *Millbank Quarterly* 73: 213-230
- [19] Hellwig, Martin (2009), "A Maximum Principle for Control Problems with Monotonicity Constraints," Max Planck Institute working paper
- [20] Hopenhayn, Hugo, Gerard Llobet and Matthew Mitchell (2006), "Rewarding Sequential Innovations: Prizes, Patents, and Buyouts," *Journal of Political Economy* 114: 1041-1068
- [21] Hopenhayn Hugo and Matthew Mitchell (2001), "Innovation variety and patent breadth," *RAND Journal of Economics* 32, 152-166
- [22] Hovenkamp, Herbert (2004), "United States Antitrust Policy in an Age of IP Expansion," U Iowa Legal Studies Research Paper No. 04-03
- [23] Judd, Kenneth (1998), *Numerical Methods in Economics*, MIT Press

- [24] 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
- [25] Kremer, Michael (1998), "Patent Buyouts: A Mechanism for Encouraging Innovation," *Quarterly Journal of Economics* 113: 1137-1167
- [26] Mitchell, Matthew and Andrea Moro (2006), "Persistent distortionary policies with asymmetric information," *American Economic Review* 96: 387-393
- [27] Scotchmer, Suzanne (1999), "On the Optimality of the Patent Renewal System", *RAND Journal of Economics* 30: 181-196
- [28] Shavell, Steven and Tanguy Van Ypersele (2001), "Rewards versus Intellectual Property Rights," *Journal of Law and Economics* 44: 525-47
- [29] Weyl, Glenn and Jean Tirole (2012), "Market Power Screens Willingness to Pay," *Quarterly Journal of Economics* 127: 1971-2003

Appendix

Formalization of Switching Time

The primitive of the planner's buyout policy is a function $\tau_t : H_t \rightarrow \{0, 1\}$ 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:

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

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

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

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 $\tau = (\tau_0, \tau_1, \tau_2, \dots)$ 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 $\tau_T(h_T(\alpha_T)) = 1$.¹⁸

Proof of Proposition 5

The following argument does not depend on any specific value of θ , but we keep θ in the argument for D to emphasize the generality of our approach.¹⁹ 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.

¹⁸Note, 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.

¹⁹On the other hand, to save notation we do not explicitly indicate that the optimal price is a function of θ , and not only of τ .

The first order and second order conditions are:

$$\begin{aligned}(p + \tau)D'(p, \theta) + D(p, \theta) &= 0 \\ 2D'(p, \theta) + (p + \tau)D''(p, \theta) &\leq 0.\end{aligned}$$

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

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

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

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

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

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

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 . 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 .²⁰

Proof of Proposition 6

Differentiating (5) 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 (6) and also using the formula for \widehat{c}' yields

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

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

$$\begin{aligned} & (\widehat{c}(\theta) - \theta(2q - q^2))(1 - 2q(\theta)) + (1 - q(\theta))\widehat{c}(\theta) + \theta(1 - q(\theta))2q(1 - q) \\ &= \widehat{c}(2 - 3q) + \theta q[2(1 - q)^2 - (2 - q)(1 - 2q)] \\ &= \widehat{c}(2 - 3q) + \theta q^2. \end{aligned} \tag{10}$$

From (8) 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 $\bar{\theta}$ such that there exists a $\bar{\theta} \in (\underline{\theta}, 1)$ such that $q'(\bar{\theta}) = 0$ and $q''(\bar{\theta}) > 0$. Now assume that there exists some $\tilde{\theta} > \bar{\theta}$ for which $q'(\tilde{\theta}) < 0$. This means that there exists a $\theta' \in (\bar{\theta}, \tilde{\theta})$ 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 $\tilde{\theta}$ and implies that $q' > 0$ for each $\theta > \bar{\theta}$.

Proof of Lemma 4

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 $\tilde{q}(\theta) = 1$, and for other values of θ we maintain the original candidate optimum. We

²⁰When τ goes to zero the difference $\int_{p(\tau, \theta)}^{\infty} f(z, \theta) dz - \tau f(p(\tau, \theta), \theta)$ is strictly positive for every θ . Therefore, as long as $f(p(0, \theta), \theta)$ is bounded above by a uniform bound for all θ , then there is a τ that works uniformly for all θ .

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 $\widehat{c}(\theta) = \widehat{c}(\underline{\theta} + \varepsilon) - \int_{\underline{\theta}}^{\underline{\theta} + \varepsilon} 2q(x)(1 - q(x))dx$. But then $\widetilde{c}(\theta) - \widehat{c}(\theta) = \int_{\underline{\theta}}^{\underline{\theta} + \varepsilon} 2q(x)(1 - q(x))dx$ which goes to zero when ε goes to zero. Therefore, the component of the change in welfare that results from changing the entry function 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 of the lemma.

Proof of Proposition 7

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 $\theta', \bar{\theta}$ such that $1 \geq \theta' > \bar{\theta} > \underline{\theta}$ and the solution involves $q(\theta) = q^*$ for all $\theta \in [\bar{\theta}, \theta']$, and $q(\theta)$ is strictly decreasing on $[\underline{\theta}, \bar{\theta}]$.²¹ We provide a proof by contradiction. Suppose that there exist $\theta'' < 1$ and $\theta''' > \theta''$ such that q is strictly decreasing on $[\theta'', \theta''']$, while $q(\theta) = q^*$ for all $\theta \in [\bar{\theta}, \theta'']$. 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 $[\underline{\theta}, \bar{\theta}]$ as given, and let us maximize the objective function $\int_{\bar{\theta}}^1 [\theta(2q(\theta) - q^2(\theta))\widehat{c}(\theta) - \frac{\widehat{c}^2(\theta)}{2}]d\theta$ taking $q(\bar{\theta}), \widehat{c}(\bar{\theta})$ as given, and placing the further condition that

$$q(\theta) \leq q(\bar{\theta}) \text{ for all } \theta \geq \bar{\theta}. \quad (11)$$

We show that the solution of this problem is a constant path on interval $[\bar{\theta}, 1]$, and thus the required θ'', θ''' cannot exist. The Hamiltonian is unchanged as the extra constraint (11) is incorporated as a standard Kuhn-Tucker condition:

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

²¹In other words, $\bar{\theta}$ is the lowest type where the monotonicity constraint binds in the solution of the original problem.

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

$$\frac{\partial H}{\partial q} |_{q=q^*, \theta=\bar{\theta}} \geq 0. \quad (12)$$

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

$$\frac{\partial H}{\partial q} |_{q=q^*, \theta=\theta''} = 0. \quad (13)$$

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

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

and thus

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

We know that

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

and

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

Therefore,

$$\begin{aligned} \frac{\partial^2 H}{\partial q \partial \theta} |_{q=q^*} &= 2(1 - 2q^*)\lambda'(\theta) + 2(1 - q^*)[\theta\widehat{c}'(\theta) + \widehat{c}(\theta)] = \\ &= 2(1 - 2q^*)\left(\widehat{c}(\theta) - \theta(2q^* - (q^*)^2)\right) + 2(1 - q^*)\widehat{c}(\theta) + \\ &\quad + 2(1 - q^*)\theta 2q^*(1 - q^*) = \\ &= 2(\widehat{c}(2 - 3q^*) + \theta(q^*)^2). \end{aligned} \quad (14)$$

Because the monotonicity constraint starts binding at $\theta = \bar{\theta}$, we can conclude two observations at that point. First, ignoring the monotonicity constraint there locally is valid, second in the relaxed problem $q'(\bar{\theta}) = 0$ holds²². Then the same argument as above (see (10)) implies that $\widehat{c}(\bar{\theta})(2 - 3q^*) + \bar{\theta}(q^*)^2 = 0$. Therefore, $\frac{\partial^2 H}{\partial q \partial \theta} |_{q=q^*, \theta=\bar{\theta}} = 0$ must hold by (14). Also, $\frac{\partial}{\partial \theta} \left(\frac{\partial^2 H}{\partial q \partial \theta} |_{q=q^*} \right) = 2(\widehat{c}'(2 - 3q^*) + (q^*)^2) = 2(2q^*(1 - q^*)(2 - 3q^*) + (q^*)^2) = 2q^*(2(1 - q^*)(2 -$

²²This is an instance of the smooth pasting condition at point $\bar{\theta}$ where the function switches from being strictly decreasing to being flat.

$3q^*) + q^*) > 0$ for all relevant values of q^* . Therefore, together with $\frac{\partial^2 H}{\partial q \partial \theta} \big|_{q=q^*, \theta=\bar{\theta}} = 0$ we obtain that for all $\theta \in (\bar{\theta}, \theta'']$

$$\frac{\partial^2 H}{\partial q \partial \theta} \big|_{q=q^*} > 0. \quad (15)$$

But comparing (12), (13), and (15) yields a contradiction, which concludes our proof of the shape of q . Finally, $\widehat{c}(\bar{\theta})(2 - 3q^*) + \bar{\theta}(q^*)^2 = 0$ implies that $q^* > 2/3$, which provides the last result.