



ELSEVIER

Contents lists available at ScienceDirect

## International Journal of Industrial Organization

[www.elsevier.com/locate/INDOR](http://www.elsevier.com/locate/INDOR)



### Market outcomes and dynamic patent buyouts<sup>☆</sup>



Alberto Galasso<sup>a,b</sup>, Matthew Mitchell<sup>a</sup>, Gabor Virag<sup>a,\*</sup>

<sup>a</sup> Rotman School of Management, University of Toronto, Toronto, ON, Canada

<sup>b</sup> NBER, United States

#### ARTICLE INFO

##### Article history:

Received 19 October 2015

Revised 13 June 2016

Accepted 20 June 2016

Available online 11 July 2016

##### JEL Classification:

D82

D83

O31

O38

##### Keywords:

Innovation incentives

Patents

Buyouts

Market information

#### 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 new reward mechanisms which complement the patent system and reduce 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

<sup>☆</sup> We are grateful to Luis Cabral, Richard Gilbert, Mark Schankerman, Ralph Winter, Hugo Hopenhayn, Glen Weyl, Gerard Llobet, Chun-Hui Miao, Emanuele Tarantino and Emmanuel Dechenaux for comments on an earlier draft of the paper. We thank seminar participants at the University of Georgia, Hitotsubashi University, University of Toronto, the International Industrial Organization conference, the Stony Brook workshop on patent licensing, the EARIE conference, the Simposio Nacional de Microeconomía of Colombia, and the REER conference at the Georgia Institute of Technology. We are grateful for financial support from the University of Toronto Michael Lee-Chin Family Institute for Corporate Citizenship.

\* Corresponding author.

E-mail address: [gabor.virag@utoronto.ca](mailto:gabor.virag@utoronto.ca) (G. Virag).

pay of consumers in many important cases, and therefore can generate the right innovation incentives.

© 2016 Elsevier B.V. All rights reserved.

---

## 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 more than twice as high as the currently observed R&D expenditure. A central policy question, therefore, is how one can best devise mechanisms that encourage innovation.

The patent system is one the main instruments used by governments to increase R&D incentives. Recently, increased attention has been paid to alternative reward mechanisms, which complement the patent system and can preserve innovation incentives, especially for breakthrough technologies that generate a large improvement in social welfare. McKinsey estimates that the total funds available from innovation prizes have more than tripled over the last decade to surpass \$375 million with a large number of philanthropists entering the business of rewarding innovators ([McKinsey, 2009](#)). For example, Qualcomm and Nokia have offered multimillion-dollar prizes for the development of affordable devices that can recognize and measure personal health information. Similarly, the Gates Foundation has offered an innovation award to those who help immunize children in the poorest regions in the world, and the X PRIZE Foundation offered a \$10 million Ansari Prize for a private space vehicle that can launch a reusable manned spacecraft into space twice within 2 weeks ([Murray et al., 2012](#)). At the same time, government interest in innovation prizes has also increased substantially. In the United States, President Obama's Strategy for American Innovation strongly encouraged the use of innovation prizes, and the America Competes Reauthorization Act of 2010 provided all federal agencies with the power to offer innovation prizes ([Williams, 2012](#)).

Despite this growing trend, little theoretical work on the design of innovation prizes has been conducted. This paper contributes to the recent literature focused on designing prizes that infer demand from various market signals and uses that information to design a reward at least partially based on a cash prize. We study the problem of a philanthropic or government agency interested in rewarding a breakthrough innovation with limited information on the research cost and the social welfare generated by the new technology. Following the mechanism design literature, we refer to such an agency as the "social planner." We show that, in a wide variety of environments, social welfare can be improved by prizes in the form of a buyout of patent rights over time. The buyout system replaces some of the rents obtained through monopoly rights with a prize.

If the planner cannot discern the quality of the innovations, pure prizes are difficult to implement because the value of the prize cannot be tied to the surplus generated by the innovation, as the demand is unknown by the planner. Worse still, the market signals might be manipulated by an innovator, if the innovator knew that a prize was tied to the market outcomes. Even if the planner could obtain precise information about the number of units sold at a given price (for instance, by observing the number of units sold under perfect competition), this information is insufficient to construct the use values of inframarginal consumers, which is essential to estimating the full value of an innovation. Our approach addresses the need to estimate the surplus of inframarginal consumers. In order to accomplish this, we stress a dynamic approach to innovation rewards, as one point on the demand curve will generally be insufficient for reconstructing demand. In contrast to the previous literature, we assume that the planner can learn over time about the market conditions by observing price and quantity realizations that arise from the choice of the innovator and the underlying demand function. As information about the market demand is revealed, the reward mechanism that maximizes social welfare may change according to the revealed information. Eventually, the planner can resort to allowing perfect competition, which generates additional information about the demand for the innovation.

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, where the true price that gave rise to the observed sales can be completely obscured by the innovator, the optimal mechanism involves some reward through a contingent prize near the end of the period in which the innovator is rewarded. The optimal mechanism is a hybrid between a patent and a prize in the sense that it rewards innovators through prices above the marginal cost initially but then moves toward a reward that is focused on a cash prize and prices closer to, or reaching, the marginal cost.

Our results are directly relevant to the rising number of philanthropists who have entered the business of rewarding innovators. Recent proposals have considered linking 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 should pay the innovator \$14 for each of the first 200 million treatments sold for \$1 to the recipients (Glennerster et al., 2006). Our results suggest one approach to this philanthropy: Use resources to buy out patents that have a track record of success.<sup>1</sup>

Our results might also be of interest to policy makers. From the policy maker's perspective, our findings provide guidance to government agencies looking for tools that complement patents and can spur innovation while minimizing product market

---

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

distortions. We show that an effective mechanism is a patent buyout scheme whose reward depends on the observed market outcomes. The computation of the reward resembles structural estimation studies that typically estimate the primitives of a model from local price variation and exploit these estimates for out-of-sample welfare analysis. In other words, the buyout is facilitated by information from data in much the same way that the impact of a merger on consumer surplus is assessed by estimation of an econometric model. An important feature of the mechanism is that it does not require any change in the functioning of the current patent system. Policy makers simply have the option to complement patents with buyout schemes that depend on market outcomes. From this perspective, we believe that this policy tool is well suited for selected high-value technologies for which the welfare impact of free access is expected to exceed the cost of the public funds associated with the buyout.

Our model can also provide insights for a regulator, antitrust or otherwise, that faces firms that have monopoly granted through intellectual property (IP). For example, the Australian government offers copayments for selected drugs to mitigate monopoly distortions. Similarly, the FDA is involved in the administration of ex post rights for pharmaceuticals through the Orange Book program and the rights granted therein. 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 an 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 when designing an appropriate prize is observing the total benefits of the innovation. As in [Kremer \(1998\)](#), observing the total benefits 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 non-marginal consumers' willingness to pay. Our mechanism addresses both issues. Our first result is that, in the absence of demand manipulation, the first best can be approached arbitrarily closely in a large set of demand functions that includes those typically used in the industrial organization literature.

Discussing a number of extensions of the baseline model, we argue that the assumptions required to reach the first best are those typically imposed in structural industrial organization studies that identify the primitives of a model from local price variations and exploit the estimated parameters to conduct out-of-sample welfare analysis ([Fig. 1](#) case A). In our context, the planner can request that the innovator generate price variations that will be used to identify the underlying demand curve of the technology and compute a patent buyout transfer that compensates the innovator for the surplus generated. By keeping the price variations concentrated around the marginal cost of production, the planner can limit the loss of surplus associated with learning to a minimum ([Fig. 1](#) case B).

We then investigate the case in which the innovator can manipulate demand. 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




Fig. 1. Market outcomes and demand identification.

obtain higher reimbursements through public funding.<sup>2</sup> In keeping with the pharmaceutical price manipulation example, we assume that the quantity is observable, while the price may not be. We show that distinguishing between the case in which demand manipulation is possible after the buyout takes place and the case in which the post-buyout demand is non-manipulable is crucial. We show that pre-buyout manipulation, even if it is costless, may be ignored as long as manipulation after the buyout is not possible. This is because the planner can generate price variation after the buyout to learn the demand and to punish the innovator in the case of manipulation. This implies that market outcomes are relevant even after the buyout, because they are useful for detecting and avoiding 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 the 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 a larger output for the innovations that generate lower surplus.

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

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 optimal mechanism in the presence of costless demand manipulation. [Section 7](#) summarizes and concludes. All the proofs are in the [Appendix A](#). Additional results and extensions appear in the online 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 buyout mechanism, which is linked to an auction to incentivize research and maximize welfare. The role of the auction is to reveal information to the planner about the private value generated by the innovation. Innovation incentives are maximized because the planner would pay for the patent the private value times a fixed markup that compensates for the difference between social and private surplus. Consumer welfare is also maximized because once the innovation is acquired by the planner it would be placed in the public domain. 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 participate in the auction. In our model, we depart from this assumption and assume that only the innovator knows how valuable an innovation is at the aggregate level. Therefore, the planner needs to design a mechanism that aggregates the information contained in consumers' individual valuations.

[Wright \(1983\)](#) and [Shavell and Ypersele \(2001\)](#) provide a comparison of prizes and patents as mechanisms that incentivize innovation in a static framework. [Scotchmer \(1999\)](#) studies the optimal mechanism to reward innovation when the planner offers a menu of patents that differ in length and application fee. She shows that if the market outcomes are not observed, then in the presence of asymmetric information about the cost and benefit of the 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 where innovators have unobservable productivity parameters. As in [Scotchmer \(1999\)](#), in [Cornelli and Schankerman \(1999\)](#), the planner offers the innovator a menu of patents that differ in length and application fee. [Cornelli and Schankerman \(1999\)](#) show that the optimal patent scheme is typically differentiated and can be implemented through menu of patent renewals.<sup>3</sup>

[Hovenkamp \(2004\)](#) and [Hopenhayn et al. \(2006\)](#) study the optimal patent design when innovation is cumulative and each discovery is a building block to future innovations.

<sup>3</sup> [Gans and King \(2007\)](#) extend the innovative environment to include timing as an important choice. They demonstrate that a finitely lived, but broad, patent can be socially desirable.

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

Acemoglu and Akcigit (2012) develop a dynamic framework with cumulative innovation and show that full patent protection is not optimal, whereas state-dependent property rights are preferable. Akcigit and Liu (2014) study the optimal mechanism for rewarding research investments when multiple firms compete with private information about the value of innovation.

We are aware of only a few studies that consider observable market outcomes. The first one is Weyl and Tirole (2012) who study the optimal reward structure in the presence of multidimensional heterogeneity and non-manipulable 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 Weyl and Tirole (2012) because we introduce dynamics and allow the innovator to manipulate the market outcomes.

The second paper is by Chari et al. (2012) who compare prizes and patents when the planner can observe market signals over time. They develop a dynamic framework in which the innovator and his product market competitors (but not the planner) know the value of the innovation that is represented by a unidimensional parameter. Their main finding is that patents are necessary if the innovator can manipulate the market signals. Our model departs from their setting in a number of dimensions. First, we assume that only the innovator knows the value of his technology, and we do not require the presence of informed competitors. Second, we allow for multidimensional heterogeneity in innovation quality. Because in a multidimensional setting observing one market outcome is not enough to learn the entire demand curve, in our model the planner faces a nontrivial learning problem even under full (i.e. non-manipulated) observation of the market signals. This implies that, differently from Chari et al. (2012), in our model the planner finds it optimal to use information acquired over time in a truly dynamic way. Third, we do not restrict the planner to use either patents or prizes, and we consider a large set of reward structures that depend on the quantity and prices practiced by the innovator. In particular, we allow for patents of different “strengths” in which the price charged by the innovator differs from both the monopoly and the competitive prices.

Finally, Brynjolfsson and Zhang (2007), in a policy paper, propose a “statistical couponing” mechanism for assessing the value of digital goods (e.g., software or music

files). Their idea is to identify consumer willingness to pay by exploiting coupons for a small (but representative) sample of consumers. The planner infers the total market valuation from the behavior of this sample of consumers and reward the innovator appropriately while keeping the price at the marginal cost level. Our results show that a similar outcome can be reached in the absence of coupons, by exploiting dynamic market information. Our analysis also stresses that innovators have strong incentives to manipulate the coupon market to obtain larger rewards.

### 3. The model

The time is discrete, and the horizon is infinite. Each innovation is characterized by variables  $c$  and  $\theta$ . These two parameters determine the cost and the value, respectively, of a particular innovation. First, the ex-ante cost of creating the innovation is  $c \in C \subset \mathbb{R}_+$ . Second, the demand function  $q = D(p, \theta)$  depends on the demand parameter  $\theta$ . We assume that  $\theta \in \Theta$ , a compact subset of  $\mathbb{R}^N$ , and that  $D$  is continuous in  $\theta$ . The demand and cost parameters,  $\theta$  and  $c$ , are private information for the innovator and are distributed according to a smooth probability density function  $\psi(\theta, c)$  that is common knowledge among the planner and the innovator. We make the regularity assumption that  $D$  is twice continuously differentiable in  $p$ . To ensure the concavity of the static profit function in  $p$ , we assume that  $D_p(p, \theta) < 0$  and  $D_p(p, \theta) + pD_{pp}(p, \theta) < 0$  for each  $p \geq 0$ . Let  $\bar{p} > 0$  be the minimum price at which  $D(\bar{p}, \theta) = 0$ . The marginal cost of production is normalized to zero.<sup>4</sup>

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

$$\phi = \infty \text{ if } \hat{p}_t \neq p_t$$

and (ii) costless manipulation where

$$\phi = 0 \text{ for all } \hat{p}_t.$$

Section 5 provides examples and a micro-foundation for the manipulation cost  $\phi(\hat{p}_t, p_t, \theta)$ .

Let us indicate with  $h_t \in H_t$  the public history at time  $t$ , which can be defined recursively as  $h_t = (h_{t-1}, r_t)$  where  $r_t = (q_t, \hat{p}_t)$  is the information revealed in period  $t$  and  $h_0 = \emptyset$ . 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 maximizes the expected total discounted surplus, with discount factor  $\delta$  between the periods. Following the literature, we assume that the planner has full

<sup>4</sup> The parameter  $\theta$ , by capturing the global shape of the demand function can be interpreted as a proxy for the perceived quality of the innovation. Our model abstracts from consumption externalities.



commitment power to any announced policy.<sup>5</sup> 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 also has 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). Our preferred interpretation of this reward schedule is that the innovator owns a patent up to period  $T$ , and that at period  $T + 1$ , the innovator’s patent rights are revoked and the innovation is placed in the open domain. When  $g_t(h_t) > 0$  and  $T$  is finite, the policy implies that the planner is paying the innovator to remove patent protection at time  $T$ ; that is, the planner buys the patent out at time  $T$ .

The innovator’s strategy is a sequence of a pair of prices  $(\hat{p}_t, p_t)$  for each period  $t$  that satisfies the constraint that prices are set to zero after switching has occurred. Let  $\alpha \in A$  denote any such strategy and  $A$  denote the set of all possible strategies.<sup>6</sup> 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 the time depends only on his strategy. In the online Appendix, we provide a recursive definition of  $T(\alpha)$  given the planner’s policy function.

Then the innovator’s maximization problem upon pursuing the innovation is

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

To simplify the notation, we leave the relationship between the switching time and the innovator’s strategy implicit and indicate  $T(\alpha)$  as  $T$  in the remainder of the paper.<sup>7</sup>

Let us indicate the optimal revealed and actual prices for period  $t$  with  $\hat{p}_t^*(\theta)$  and  $p_t^*(\theta)$  and the public history revealed by this optimal equilibrium play with  $h_t^*(\theta)$ . The innovation takes place if the net present value of the innovator’s profits (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:

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

with

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

<sup>5</sup> This is usually justified by reputational concerns of the planner. We further discuss the case of limited commitment at the end of Section 4.2.

<sup>6</sup> Note that the innovator’s strategy is formed upon observing the planner’s switching policy, and that the planner’s switching policy affects the set of strategies for the innovator.

<sup>7</sup> We assume that the innovator cannot manipulate the market after time  $T$ . To prevent manipulation, the planner can implement a whistle-blowing system or exclude the innovator from producing after period  $T$  and only allow the competitive fringe to manufacture the good.

The social planner chooses functions  $g_t$  and  $T$  (taking the innovator’s optimal strategy as given) to maximize the expected total social welfare created by the innovation:

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

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 ( $\hat{p}_t = p_t$ ), 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  $\theta$  is known. To do so, the planner transfers the entire surplus to the innovator if he observes the competitive quantity, and the planner punishes the innovator if he observes a different quantity (i.e.,  $g_t = S(0, \theta)$  if  $q_t = D(0, \theta)$  and  $g_t = -\infty$  if  $q_t \neq D(0, \theta)$ ). In this case, one can interpret the punishment as part of the contract to which potential prize winners agree.

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

*Patents*

When  $g_t(h_t) = 0$  and  $T = \varsigma$ , the planner offers a  $\varsigma$ -period patent 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 the expiration of the patent in the absence of a payment:

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

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

*Simple buyout*

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 buyout scheme in which the planner commits to buy the patent after  $\varsigma$  periods at a pre-specified amount  $K$ . The setting also allows the implementation more complex buyout mechanisms where transfer price  $K$  and acquisition time  $T$  may depend on the observed market outcomes.

#### 4. Optimal mechanism in the absence of demand manipulation

In this section, we characterize the optimal mechanism when the government (the planner) can dictate prices and the innovator cannot manipulate demand; that is,  $\phi = \infty$  if  $\hat{p}_t \neq p_t$ . In this case, the planner can essentially dictate prices by requiring that the prize winner follow the prescribed price path. Existing prizes, such as the malaria prize offered by the Center for Global Development (described below), include a pricing requirement.

To develop an 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  $\theta_2$  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  $\varepsilon$  arbitrarily close to zero and minimize the deadweight loss generated by pricing above the marginal cost. If the entire surplus generated by the innovation is transferred to the innovator, the innovation incentives are set at the first best level.<sup>8</sup> This generates two benefits: a lower deadweight loss, and the prize can reward inframarginal values that should motivate innovation but are not included in simple monopoly pricing.

This example suggests that transfers that depend on market outcomes can be powerful mechanisms for incentivizing 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.

Compared to static multidimensional screening mechanisms, such as the one characterized by [Weyl and Tirole \(2012\)](#), it is natural that the dynamic model can do more, because it offers more instruments. The result obtained in the simple linear setting shows that this gain can be substantial. Our next result shows that this logic expands to a broad set of demand functions. We start with the definition of an analytic demand function.

<sup>8</sup> A transfer that approximates the first best is  $g_1(h_1) = 0$  for all  $h_1$ ;  $g_2(h_2) = S(\varepsilon) + S(0)/\delta$  if  $h_2 = \{q_1, 0, q_2, \varepsilon\}$  and  $g_2(h_2) = -\infty$  otherwise;  $g_t(h_t) = S(0)$  for  $t > 2$  if  $h_t = \{q_1, 0, q_2, \varepsilon, q_3, 0, \dots, q_t, 0\}$  and  $g_t(h_t) = -\infty$  otherwise.

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

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

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

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

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

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

[Proposition 1](#) substantially generalizes the result for linear demands. Polynomials, exponentials, logarithms, power functions and a number of other demand functions that are typically used in applied theory are analytic functions. [Fox and Gandhi \(2011\)](#) show how analyticity of the market demand is a property of various well-known 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 et al. \(1995\)](#).

[Proposition 1](#) does not require the planner to know the specific functional form of the demand; she needs to know only that it is an analytic function. For instance, if the planner knows that the demand is one of  $N$  functional forms (or a linear combination of the  $N$  forms), experimentation means not only learning the shape of a given functional form but also experimenting across the functional forms. Given that  $N$  can be arbitrarily large, this specification allows us to consider an arbitrarily rich set of possible demand functions.

One may also wonder about the role of the planner in an environment in which the demand is known but the marginal costs are not. This alternative problem is likely to be simpler for the planner, because it involves using variations in the market outcomes to identify the marginal value of the cost function instead of identifying both the marginal and the inframarginal willingness to pay as in our baseline model. As long as the planner

can set up a non-manipulable competitive market with a fringe that produces at the same (unknown) marginal cost as the innovator, then our approach remains valid. The planner can simply observe the competitive price and back out the marginal cost of production.

#### 4.1. Discussion

The prior result can be generalized in many ways and in some (but not all) cases still allow the planner to implement the first best. In the online Appendix, we describe these extensions in more detail.

##### 4.1.1. Non-stationary demand

*Demand shifts.* Following Battaglini (2005), we assume that the demand has two states, high ( $H$ ) and low ( $L$ ), with  $D_H(p, \theta) \geq D_L(p, \theta)$  for each  $p$  and that the transition between the states follows a Markov process. We show also that in this setting, if the demand functions are analytic the planner can maximize the innovation incentives by approximating the first best outcome. To understand the intuition for the proof, consider the case of linear demand. The planner can identify the intercepts of the two demand functions by dictating a price equal to zero and maintaining it until two different quantities are observed. Then he will set  $p = \varepsilon$  until two different quantities are observed. With two observations along each demand line, the planner learns the demand and welfare functions. By setting  $\varepsilon$  arbitrarily close to zero, the dead weight loss generated by pricing above the marginal cost is minimized, and the first best is approached arbitrarily closely. An interesting feature of this result is that the optimal incentive scheme is non-stationary and has unbounded memory even if the demand shifts follow a Markov process and the relevant economic environment has a memory of only one period.

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

##### 4.1.2. Public funds and menu costs

*Social costs of public funds and imperfect capital markets.* When the innovator does not have access to a frictionless capital market, he prefers to be paid as soon as possible. However, a simple buyout mechanism delays much of the payment until after the buyout occurs, which can pose a long delay because learning the entire demand curve takes time. This is not a problem if the planner can raise funds at no cost, that is, if taxation does not have any deadweight losses. In this case, the planner can easily reimburse the innovator even before the buyout. Intuitively, the function  $g_t$  allows the planner to act as a capital market removing any friction the innovator may face in raising funds. In other words, the planner can spread the reward to the innovator over time by adjusting the amount paid depending on the market outcomes observed in the past.

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

This analysis also implies that even if there is a cost of raising public funds, the planner can still bring the innovator's payments forward in time, if reimbursing the innovator as soon as possible is socially beneficial. In particular, the planner can still implement the welfare maximum and pay out the innovator arbitrarily early.

*Menu costs.* In our model, the only cost of price variations is their impact on consumer welfare.<sup>9</sup> In the presence of menu costs, the planner's problem becomes substantially more complex. In fact, introducing a cost of changing the price leads to a trade-off between the marginal welfare benefit from observing an additional data point along the demand curve and the cost of changing the price.

#### 4.1.3. Additional identification challenges

*Asymmetric production costs.* Our baseline model assumes that the marginal cost of production for the innovator is the same as the one for the competitive fringe that produces after the buyout. In the Appendix, we show that Proposition 1 extends to the case in which the innovator has a cost advantage. Interestingly, the optimal policy no longer involves a buyout. Instead, even when the planner has full knowledge of the demand curve, the innovator is selected to produce and sell the product at his own (lower) marginal cost. We also show that in the case in which the competitive fringe has a cost advantage over the innovator, Proposition 1 goes through if the following is assumed: (i) Either the production technology of the fringe is available to the innovator through

<sup>9</sup> This is the typical assumption in the innovation literature, and can be justified by the mixed empirical evidence on the significance of menu costs (Nakamura and Steinsson, 2008) especially after the diffusion of the internet and modern information and communication technologies (Brynjolfsson and Smith, 2000).

licensing or contract manufacturing, (ii) or it is possible for the planner to levy a tax after the buyout to learn the demand after the buyout. Under case (i), the planner requires the innovator to pick one firm in the competitive fringe as the exclusive licensee for the innovation and induces price variation close to the marginal cost of the manufacturer. This is a natural assumption in cases in which the innovator lacks the complementary assets required for efficient large-scale production. Under case (ii), the planner places the innovation in the competitive market immediately and generates price variation through taxation. If assumptions (i) and (ii) are both violated, [Proposition 1](#) does not hold, and the first-best cannot be achieved. In this case, learning the demand curve can occur only by experimenting at prices that are above the post-buyout production costs.

*Demand is observed with error.* We consider the case in which the demand is observed with an error and assume that  $q_t = D(p_t, \theta) + \varepsilon_t$  where  $\varepsilon_t$  is a mean zero i.i.d. noise over the support  $[-\bar{\varepsilon}, \bar{\varepsilon}]$ . Even in this case, analyticity of the demand function is sufficient to approach the first best arbitrarily closely. In the linear demand case, the planner can use the following two-step scheme. In the first stage, the planner induces the firm to charge  $p = 0$  and obtains a sample of  $N$  quantities for this price. Then he sets a price equal to  $\varepsilon$  and obtains another sample of  $N$  quantities. The weak law of large numbers guarantees that, for large enough  $N$ , 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. Although the law of large numbers guarantees that the estimate is unbiased, the variance of the estimate depends on the price variation and is smaller when the variation is larger. This is not an issue in our setting because we assumed that the planner and the innovator are both risk neutral. Even in the case of risk aversion, the variance can be made arbitrarily small by letting the sample size,  $N$ , be very large. For a given sample size, one can employ econometric techniques to estimate the residual uncertainty about the demand curve and the induced total surplus when the pricing is at marginal cost. Back-of-the-envelope calculations show that even for a relatively low sampling error, the variance in the estimate of higher-order derivatives of the demand function can be substantial when the sample size is not large enough. For a discussion of additional challenges faced in structural demand estimation, see [Chintagunta and Nair \(2011\)](#).

*Small sample.* [Proposition 1](#) assumes that the planner can collect an arbitrarily large sample of market outcomes. The use of large samples that assure asymptotic convergence is the norm in modern applied econometrics. Nonetheless, one way to study the robustness of our results is to assume instead that only a small finite sample of  $N$  market outcomes can be observed. Because our interpolation technique relies on large sample convergence, the small sample implies that the demand is approximated with an error, and thus, the first best is no longer achievable. A more interesting observation concerns the optimal prices set by the planner in this case. For small samples, the interpolation literature suggests that the demand will be learned more precisely when the data points are not all drawn from a small interval ([Mastroianni and Milovanovic, 2008](#)). In our context, this implies that the planner faces a trade-off. If he charges only prices close to the marginal cost, then the deadweight loss from experimentation is low, but the demand is not learned

precisely. If the planner sometimes charges higher prices, learning is accelerated at the cost of increasing current deadweight losses for consumers. Although a formal analysis is beyond the scope of this work (we do not report this extension in the Appendix), with small samples a clear trade-off exists between generating demand information to incentivize innovation and mandating low prices to reduce the deadweight loss. However, the more general lesson that learning the demand is useful for incentivizing innovations, and that such learning can occur by dynamically experimenting taking into account consumer welfare, remains valid with small samples.

In general, [Proposition 1](#) does not hold in settings where structural demand identification is not feasible. When the demand is not analytic, local price variations are not sufficient to estimate the global behavior of the demand function. In the same way, when the demand grows very fast, collecting enough price–quantity observations to identify demand may not be feasible. These issues are typical in structural modeling, where it is assumed that the structural parameters identified through local data variation can be used to perform counterfactuals or policy simulations [Reiss and Wolak \(2007\)](#). From this perspective, [Proposition 1](#) does not require assumptions other than those typically imposed in structural industrial organization studies. Without demand manipulation (studied in [Section 5](#)), then, in many contexts market signals can be used to construct a prize that improves outcomes by lowering the deadweight loss from the monopoly while at the same time rewarding the inframarginal value from the innovation.

When the first best can be obtained, the planner may need to generate a large number of observations by experimenting at different price points near the marginal cost. The rate of convergence to full learning, and thus to the first best surplus, depends on the demand parameters. It remains for future research to study the speed of this convergence and to identify additional properties of the estimator.<sup>10</sup>

#### 4.2. Implementation

[Proposition 1](#) 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 the 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 with a patent that confers on 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

<sup>10</sup> However, in the proof of [Proposition 3](#) we show that even if experimentation away from the marginal cost is desirable, the chosen price is always strictly below the monopoly price.



out the patent if the innovator follows the path with a reward that depends on the quantities sold. The computation of the reward resembles structural estimation studies that typically estimate the primitives of a model from local price variation and exploit these estimates for out-of-sample welfare analysis. An implication of our result is that policy makers may affect the innovation incentives by designing reward systems that exploit these techniques.

In the context of the malaria vaccine, the Center for Global Development proposes to reward the innovator with a prize if 200 million treatments are sold for \$1 to the recipients. The suggested prize is \$2.8 billion (\$14 per treatment). A possible concern with this scheme is that the development of the vaccine may not take place if the reward is too small compared to the social welfare generated by the vaccine. An implication of [Proposition 1](#) is that this prize scheme can be improved by requesting that the successful innovator sell the 200 million treatments at different prices, even if the overall price level remains close to the \$1 benchmark. This is because the market outcomes generated by this price variation will allow the sponsor to obtain an estimate of the product market surplus generated by the new vaccine. This estimate will provide useful guidance for determining the reward and avoiding under-payment (or overpayment) for the innovation.

However, buyouts are not the only way to implement the first best. An alternative approach is to start from a perfectly competitive market in which the product is sold at the marginal cost. The price variation can then be generated by the planner, perhaps interpreted as a government, taxing the firms and shifting their marginal costs of production. The information generated in this way will be the same as that 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.

By taking a mechanism design approach, our construction of policies by a “planner” takes a fairly agnostic view of whether the policy is implemented by a government or a philanthropist. To the extent that the inventor has property rights, something that sounds like a function of the government, such as taxing a competitive market, can be implemented by a philanthropist by varying the royalty rate on nearly-free licenses. The only variable determined by the government is an upper bound on the length of time the allocation can use as a reward. A philanthropist can relinquish rights but cannot extend them beyond the statutory patent length, whereas a government can choose any  $T$  it wishes.

The mechanism proposed in this section assumes that the planner is able to commit to truthfully revealing the observed price–quantity pair, and all the market participants agree with the revelation. Suppose instead that the government can freely manipulate the observed price and, thus, can decide how much the innovator is paid. This case is similar to the case in which the innovator can costlessly manipulate the price signal in the sense that the price variable becomes non-contractible, as it is not verifiable in front

of the court.<sup>11</sup> The analysis of the price manipulation by the innovator in [Section 5](#) is therefore also applicable to the case in which the planner has a limited commitment.

## 5. Demand manipulation

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

### 5.1. Buyouts and price variation

In the general model described in [Section 3](#), the innovator can affect the market outcomes and manipulate market signals received by the planner up to period  $T$  but not after  $T$ . Our model also assumes a constant competitive market outcome from  $T + 1$ . A natural interpretation of this assumption is that the patent is acquired by the planner at  $T$ , so in the following we will refer to  $T$  as the buyout time.<sup>12</sup>

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

The basic insight is that pre-buyout manipulation, even if costless, can be avoided as long as manipulation after buyout is not possible and the planner can generate price variation after buyout to identify the demand and detect manipulation. Therefore, for

<sup>11</sup> This approach is the one followed by the incomplete contracts literature, see [Tirole \(2003\)](#) for a discussion. A caveat that is pointed out in the literature is that if both the agent and the principal can observe a variable, then it is possible to enforce a contract by requiring both agents to report the value of that variable. In case of disagreement both the principal and the agent can be punished, which allows truth-telling to be an equilibrium outcome. If that caveat is accepted, then limited commitment on the part of the planner does not prevent the planner from implementing the first best.

<sup>12</sup> Noticed that in the previous section we ignored  $T(h_T)$  and focused on  $g(h_t)$ . This is because, in the absence of price manipulation the planner can generate a competitive outcome using only  $g(h_t)$  by punishing the innovator if  $p_t \neq 0$ .

manipulation to distort away from the first best, it has to be the case that either (i) manipulation by the innovator is feasible both before and after the buyout or (ii) the ability of the planner to generate price variation after the buyout is limited. In the next section we study case (ii) from above.

### 5.2. Post-buyout competitive outcome

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

To provide a micro-foundation of the manipulation cost  $\phi$ , we assume that the innovator can convince the planner that he is selling at  $\hat{p} > p$  by sustaining a cost equal to  $\phi(\hat{p}, p, \theta) = \hat{\phi}((\hat{p} - p)D(p, \theta))$  with  $\hat{\phi}$  being twice differentiable, and  $\hat{\phi} > 0$ ,  $\hat{\phi}'' \geq 0$ . 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 that the secret discounts offered are wasteful, that is they cost more to the innovator to offer than they are worth for the consumers. Alternatively, there may be a difference between the cost of external and internal financing. As argued by Aghion and Tirole (1994), for innovative firms this difference arises naturally because of the informational asymmetries involving new products and technologies. In this case, to convince the planner that the 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  $\hat{\phi}((\hat{p} - p)D(p, \theta))$ .<sup>13</sup> 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 2.** *A patent of length  $T$  is Pareto dominated by a patent buyout scheme that depends on market outcomes.*

The proposition shows that for any patent of length  $T$  the planner (philanthropist) can design a buyout scheme that improves welfare. 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.

<sup>13</sup> Another microfoundation 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.

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

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

that indicates how patent buyout takes place sooner as  $i$  gets larger. This result is reminiscent of Chari et al. (2012) who consider patents and prizes, and show that shorter patents are more likely to be optimal when manipulation costs are higher, but longer patents need to be used when manipulation costs are lower. In the next section, we show that even with costless manipulation of the price signals (when an infinitely lived patent is implied by (2)), one can do better by considering mechanisms that are different from both prizes and patents.<sup>14</sup>

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,  $\sigma$  is the quantity sold at marginal cost price and  $Q(\cdot)$  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 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.<sup>15</sup>

One may speculate that when  $\phi(\hat{p}, p, \theta) = 0$  patents cannot be improved upon. This is not the case, as the next proposition shows; even with fully manipulable prices, the planner can improve on patents.<sup>16</sup>

**Proposition 3.** *When  $\phi(\hat{p}, p, \theta) = 0$  for all  $\hat{p}$ ,  $p$  and  $\theta$ , there is a per unit subsidy level  $\tau$  that Pareto dominates patents that last forever.*

<sup>14</sup> It is important to note that there are two important differences between the setup of Chari et al. (2012) and ours. First, we allow heterogeneous innovation costs. Second, Chari et al. (2012) rule out positive transfers by allowing the innovator to produce a fake (and useless) “innovation”.

<sup>15</sup> In the linear specification, if we interpret  $i > 0$  as the difference between the cost of external and internal financing the planner can reduce manipulation incentives even more by combining the buyout of the patent with the requirement to purchase a bond. Specifically, the planner can request the innovator to purchase a bond that costs  $\gamma\hat{p}D(p, \theta)$ , 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.

<sup>16</sup> It is also possible to show that patents of finite lengths can also be improved upon by a simple per unit subsidy mechanism but this result is somewhat more tangential to what we discuss below.

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

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

## 6. Optimal mechanism with costless manipulation

In this section we study the optimal incentive system in which the quantity produced is observable by the planner, but the innovator can manipulate the price costlessly, so the price will not be contracted on. This assumption captures a situation where the innovator can offer secret price discounts to buyers at no cost (other than lowering revenues). As in the previous 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 that in period  $t$  the innovator produces a specific quantity,  $q_t(\hat{\theta})$ , and receives a payment  $\tau_t(\hat{\theta})$ . To simplify the analysis we focus on the linear demand case

$$D(p) = \theta_1 - \theta_2 p.$$

While this demand function allows us to simplify substantially the exposition, our key results hold with a more general demand of the form  $D(p) = \theta_0 - \sum_{i=1}^K \theta_i p^{i-1}$  where  $K$  is known and the  $\theta_i \geq 0$  are unknown.<sup>17</sup>

First, we show that there is no loss of generality in assuming that the planner knows the intercept  $\theta_1$ .<sup>18</sup> More precisely, we can approximate the welfare of an auxiliary problem where the planner knows  $\theta_1$  from the outset arbitrarily closely. This is an upper bound because the planner cannot do better than in the hypothetical case where he observed  $\theta_1$  at the beginning.

**Lemma 1.** *The planner can approximate the welfare that can be induced under full information about  $\theta_1$  arbitrarily close.*

<sup>17</sup> In particular, all the steps in the proof of **Proposition 4** would hold as the argument readily generalizes to any case where first-order conditions are sufficient for optimality, for which it is sufficient if  $D$  is concave in  $p$ .

<sup>18</sup> We do not need the innovator to report his cost,  $c$ , because in our setting, as in **Scotchmer (1999)**, the innovator's compensation cannot depend on the true  $c$  since he cannot be punished for lying about  $c$ .

This result is quite intuitive: the planner can perfectly learn the demand intercept when the market becomes perfectly competitive and punish the innovator if  $\theta_1$  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  $\theta$  are associated with steeper demand curves and larger consumer surplus. Notice that the monopoly quantity is independent of  $\theta$  and it is equal to  $q^M = 1/2$ .

6.1. Static mechanisms

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

$$\begin{aligned} V'(\theta) &= \frac{\partial}{\partial \theta} U(\hat{\theta}, \theta) |_{\hat{\theta}=\theta} \\ &= q(\theta) \frac{\partial}{\partial \theta} p(\hat{\theta}, \theta) |_{\hat{\theta}=\theta} = 2q(\theta)(1 - q(\theta)). \end{aligned} \tag{3}$$

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

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

Lemma 2 shows that the optimal mechanism requires the quantity sold to be decreasing in  $\theta$ . 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  $\theta$  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  $\theta$  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  $\theta$  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  $\theta$  on  $[0, 1]$  and  $[\underline{\theta}, 1]$  for some  $\underline{\theta} > 0$ . Then the problem is equivalent to

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

*Optimal static mechanism*

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

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

To gain intuition for this result, our starting point is [Lemma 2](#), which implies that the quantity schedule needs to be weakly decreasing to be incentive compatible. In the proof of the Proposition, we show that in the relaxed problem where the monotonicity constraint on  $q$  is ignored, the optimal solution is such that  $q(\underline{\theta}) = q(1) = 1$ .<sup>19</sup> Given this, it is not surprising that when one reintroduces the monotonicity constraint on  $q$  it is still true that  $q(\underline{\theta}) = 1$ . It is also not surprising that now  $q(1) < 1$ , because  $q(1) = 1$  and the monotonicity constraint would imply that  $q(\theta) = 1$  for all  $\theta$ , that is all possibility for screening would be given up.

<sup>19</sup> The reason is that otherwise function  $q$  could be increased uniformly by the same amount  $\varepsilon$ , and the value of  $\widehat{c}(\underline{\theta})$  adjusted so that an increase in welfare is induced without violating any incentive constraints.




Fig. 2. Optimal quantity schedule with costless price manipulation.

Proposition 4 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  $\theta$  for low values of  $\theta$  and constant for high surplus innovations as depicted in Fig. 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  $\theta$ , 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  $\theta$  since the impact on welfare of an increase in  $q$  is greater the greater is  $\theta$ . For low values of  $\theta$ , the ‘screening effect’ dominates and the planner exploits market power to screen willingness to pay. This is intuitive since for low  $\theta$  it is crucial for the planner to avoid excess entry of low value innovators. As  $\theta$  increases, the innovations have larger impact on consumer surplus and the planner has lower incentives to distort the market for screening purposes. For  $\theta$  large enough, the ‘consumer welfare’ effect dominates and the planner implements a quantity schedule that is constant in  $\theta$ . The idea that market power can be exploited to screen willingness to pay is similar to the logic in Weyl and Tirole (2012).<sup>20</sup>

<sup>20</sup> They restrict their attention to Cobb-Douglas reward policies (in our setting this restriction would generate a constant level of  $q$  across types). They show that  $q$  decreases with the variance of the type distribution. In our setting, we show that even with a fixed type distribution, the planner may use different quantities to screen for different types.



## 6.2. Optimal dynamic mechanism and discussion

Having characterized the optimal quantity schedule in the static setting, we now consider the dynamic problem where the planner can choose a path  $(q_t(\theta), \tau_t(\theta))$  for every  $t \geq 0$ .<sup>21</sup> Our main result shows that repeating the same quantity over time for all types  $\theta$  is optimal.

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

A constant mechanism (over time) is optimal because of the desirable features of quantity (and price) smoothing over time. This is due to the fact that 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 without improving innovation incentives. This finding resembles the result of [Gilbert and Shapiro \(1990\)](#).<sup>22</sup>

[Proposition 5](#) confirms that the optimal mechanism *differs from a patent system* even if manipulation (of prices) is costless. Welfare is maximized with the innovator selling a quantity that is above the monopolistic quantity until the buyout occurs, unlike the (optimal) patent system described by [Scotchmer \(1999\)](#). This result is related to our earlier finding ([Proposition 5](#)), which shows that a small quantity subsidy always improves welfare.

Notice the apparent tension between [Proposition 5](#) and [Lemma 1](#). [Proposition 5](#) requires the planner to implement a constant quantity over time whereas [Lemma 1](#) requires the planner to move to the competitive outcome for at least one period in order to learn the intercept of the demand function. This tension identifies a key trade-off. On one hand, the planner would like to smooth market outcomes over time to increase welfare. On the other hand, the planner would like to generate variation of market outcomes to learn the underlying demand parameters. In the linear context, this tension leads to a mechanism that resembles a buyout where the patent is bought out after a long time (as long as possible) has elapsed.<sup>23</sup>

[Proposition 5](#) also highlights the fact that learning from market signals over time may be substituted by an initial screening process where the innovator self reports his type. The literature on dynamic mechanism design cautions us that this result (no learning is optimal until the buyout) is only true because our agent (the innovator) has strictly

<sup>21</sup> Since no new information is revealed to the agent (the innovator), it is without loss of generality to concentrate on mechanisms where the agent reports his type only at the outset.

<sup>22</sup> They conclude that the optimal patent policy calls for infinitely lived patents when patent breadth is increasingly costly in terms of deadweight loss. In our setting, lowering the quantity produced can be thought of as an increase in patent breadth because a lower quantity reduces consumer surplus and increases the profits of the innovator.

<sup>23</sup> In reality, there may be legal or political reasons why the buyout cannot be delayed indefinitely. For example, it may happen that the product becomes obsolete, and in this case the planner may not be able to commit to a buyout that may not seem to promote consumer welfare ex-post.

superior information over the planner, and this advantage is maintained over time.<sup>24</sup> However, in a large number of applications this may be a realistic assumption. In such applications, the optimal mechanism does not utilize learning on the part of planner, rather it relies on a single report of the innovator at the outset. Such a policy can be implemented by offering a menu of R&D subsidies and per unit quantity subsidies.

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

## 7. Conclusions

In this paper, we have examined the problem of a social planner aiming to maximize consumer welfare and innovation incentives while observing the 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.

Governments have patents out in a number of historical examples. 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 who persuaded the government to acquire the patent and put the rights in the public domain. Within a short period of time, the process spread throughout the country to become the technology standard in photography (Kremer, 1998). In recent academic and policy debates, pharmaceutical patent buyouts have been suggested as a strategy for improving 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 the marginal cost. They suggest that the compensation to a given manufacturer would depend on the use of the drug and evidence of realized health benefits.<sup>25</sup>

Our paper provides two main insights into the design and application of such buyout schemes. First, the planner may find it beneficial to collect market data before the buyout and use them to estimate the surplus generated by the innovation. In practice, the 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). These estimates can provide useful guidance in the

<sup>24</sup> Baron and Besanko (1984) show that if adverse selection parameters are perfectly correlated over time, then under full commitment the optimal policy is the repetition of the optimal static contract. Our result is not a direct consequence of theirs as we allow the agent (the innovator) to manipulate price signals costlessly.

<sup>25</sup> A similar policy proposal is described in Guell and Fischbaum (1995).

determination of the buyout compensation for the innovator. Second, the planner should consider the welfare cost associated with collecting price–quantity observations. As long as local variation in market outcomes can be exploited to learn about the global properties of the demand, prices close to the marginal costs minimize the loss in the consumer surplus.

### Appendix A. Proofs

**Proof of Proposition 1.** First, note that 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 the point  $(0, D(0, \theta))$ . The global estimate is obtained with the following polynomial:

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

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

The basis for this is that as  $k$  gets large, the approximation of the derivatives improves and thus our estimate of  $D$  approaches the true value of  $D$  arbitrarily close.<sup>26</sup> To formalize this, suppose that we have taken a sample of  $k + 1$  observations such that the price was always below some  $\hat{x} > 0$ . The error term (in absolute value) for the estimate of the  $i$ 'th derivative can be bounded by  $\max_{\theta \in \Theta, i, x \in [0, \hat{x}]} |D^{(i+1)}(x, \theta)x^{i+1}/(i + 1)!| \leq K\hat{x}^{i+1}$ , which can be made arbitrarily small (in absolute value) if  $\hat{x}$  is small (Mastroianni and Milovanovic, 2008). Here we used the fact that there exists a  $K > 0$  such that  $\max_{x \in [0, \hat{x}]} |D^{(i+1)}(x, \theta)/(i + 1)!| < K$  for all  $i = 1, 2, \dots$  and  $\theta \in \Theta$ . To establish that this is indeed true, note that by  $D$  being analytic there exists  $\tilde{K}(\theta)$  such that  $\max_{x \in [0, \hat{x}]} |D^{(i+1)}(x, \theta)/(i + 1)!| < \tilde{K}(\theta)$  for all  $i = 1, 2, \dots$ . Moreover, Weierstrass's theorem implies that there exists  $K$  such  $\tilde{K}(\theta) \leq K$  for  $\theta \in \Theta$ , because  $\Theta$  is compact and  $\tilde{K}$  is a continuous function of  $\theta$  because all the derivatives of  $D$  are continuous in  $\theta$  by assumption.<sup>27</sup>

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

<sup>26</sup> If the derivatives at 0 can be estimated with a known error  $\varepsilon$ , then the total error at  $p > 0$  is less than  $\varepsilon(p + p^2 + \dots) = \frac{\varepsilon p}{1-p}$  if  $p < 1$ . If the choking price cannot be bounded away from 1 (from above), then this procedure does not suffice, and local expansion is needed similarly to what is suggested below for the case where  $r < \bar{p}$ .

<sup>27</sup> Note, that directly we only assumed that  $D$  itself is continuous in  $\theta$ . However, if  $D$  is analytic, then continuity of  $D$  in  $\theta$  implies that all the  $c_k$  coefficients are continuous in  $\theta$ , which implies that all derivatives of  $D$  are also continuous in  $\theta$ .

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

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

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

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

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

**Proof of Proposition 2.** Consider the following mechanism. The innovator is awarded a patent for  $T$ . The planner offers to buy the patent out after  $\hat{T} \leq T$  periods as long as the same prices and quantities  $(\hat{p}, D(p, \theta))$  are observed by the planner for the entire 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  $T - \hat{T}$  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.<sup>28</sup> 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 with respect to  $p$  (with  $\hat{p} = p$ ) gives:

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

<sup>28</sup> If  $\hat{T}$  is not an integer set it equal to the smallest integer for which  $\delta^{\hat{T}} - \delta^T < (1 - \delta^{\hat{T}})\hat{\phi}'(0)$ .

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

so the innovator will truthfully report the monopolistic profits. The profits of the innovator will be the same as with a patent of length  $T$  and no buyout but consumers will be better off.

**Proof of Proposition 3.** When a per unit subsidy is awarded there are two main changes in total welfare. First, the set of types who enter becomes larger as the profit of the innovator increases. Second, for a fixed type who enters even without a subsidy, the total surplus on the market changes as prices go down due to the subsidy. Both effects increase welfare when  $\tau$  is small as we show below.

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

The profits for the patentee in the presence of a quantity subsidy are equal to  $(p + \tau)D(p, \theta)$  where  $\tau$  is the per unit subsidy. 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  $\tau$  increases. Next, for a given  $\theta$  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

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

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

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

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

because

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

for  $\tau$  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  $\theta$ .<sup>29</sup> By the above, any such subsidy level  $\tau$  increases total welfare for all  $\theta$ . In other words the same level  $\tau$  is applicable to all  $\theta$ .

**Proof of Lemma 1.** Take the hypothetical problem where the planner observes  $\theta_1$  so the innovator needs to report only  $\theta_2$ . As we show it in the next section, the optimal mechanism prescribes a quantity  $q_t(\theta_2) = q^*(\theta_2)$  that is constant in time ( $t$ ). Now, take our original problem where the planner does not observe  $\theta_1$  at the outset, and suppose that the planner provides a buyout at time  $T$  and sets the quantities produced before time  $T$  equal to the  $q^*(\theta_2)$ .<sup>30</sup> After the buyout, when the market becomes perfectly competitive the intercept will be observed by the planner. At that stage the innovator can be punished if the quantity sold at marginal cost,  $\theta_1$ , differs from the report of the innovator  $\hat{\theta}_1$ . By making the punishment large enough the innovator has no incentive to misreport. Moreover, letting  $T$  become arbitrarily large the welfare induced by this mechanism approximates the welfare under full information about  $\theta_1$ .

**Proof of Lemma 2.** 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

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

<sup>29</sup> When  $\tau$  goes to zero the difference  $\int_{p(\tau, \theta)}^{\infty} D(z, \theta)dz - \tau D(p(\tau, \theta), \theta)$  is strictly positive for every  $\theta$ . Therefore, as long as  $D(p(0, \theta), \theta)$  is bounded below by a positive uniform bound for all  $\theta$ , then there is a  $\tau$  that works uniformly for all  $\theta$ . If such a uniform bound is not available, then the proof goes through with a few straightforward modifications.

<sup>30</sup> By standard arguments, there is a payment schedule  $\tau$  that makes this quantity schedule incentive compatible.

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

**Proof of Proposition 4.**

**Part 1: Solution of the relaxed problem**

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

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

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

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

The other co-state equation is

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

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

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

The above conditions lead to the following result.

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

**Proof.** From (4) we obtain that

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

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

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

**Lemma 4.** *There exists a  $\bar{\theta}$  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.** Differentiating (4) with respect to  $\theta$  and dividing through by 2 yields

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

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

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

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

$$\begin{aligned} & (\hat{c}(\theta) - \theta(2q - q^2))(1 - 2q(\theta)) + (1 - q(\theta))\hat{c}(\theta) + \theta(1 - q(\theta))2q(1 - q) \\ &= \hat{c}(2 - 3q) + \theta q[2(1 - q)^2 - (2 - q)(1 - 2q)] \\ &= \hat{c}(2 - 3q) + \theta q^2. \end{aligned} \tag{8}$$

From (7) it follows that for all  $\theta \geq \underline{\theta}$  it holds that  $q(\theta) \leq 1$ , therefore  $q'(\underline{\theta}) \leq 0$  holds because  $q(\underline{\theta}) = 1$  by the previous Lemma. Because  $q(\underline{\theta}) = q(1) = 1$ , it means that there exists a  $\bar{\theta} \in (\underline{\theta}, 1)$  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  $\theta > \bar{\theta}$ .  $\square$

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

**Part 2: The optimal static mechanism**

We now reintroduce the monotonicity constraint  $q'(\theta) \leq 0$ . We first show that there is efficient production for the lowest innovation type ( $q(\underline{\theta}) = 1$ ), since for such a type there is no incentive to misreport in general. Suppose that  $q(\underline{\theta}) = q^* < 1$ . Then take a small deviation where for all  $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$  the quantity is set at  $\tilde{q}(\theta) = 1$ , and for other values of  $\theta$  we maintain the original candidate optimum. We show that this increases welfare, and still satisfies all the constraints. First, it is obvious that the monotonicity constraint is still satisfied. Second, we keep  $\hat{c}(\theta)$  unchanged for all  $\theta$  outside the interval.



This means that for all  $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$  it holds that the modified entry function  $\tilde{c}(\theta) = \widehat{c}(\underline{\theta} + \varepsilon)$  because  $\tilde{c}'(\theta) = 0$  for all  $\theta \in [\underline{\theta}, \underline{\theta} + \varepsilon]$  as  $\tilde{q}(\theta) = 1$  for such values of  $\theta$ . 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  $\tilde{c}(\theta) - \widehat{c}(\theta) = \int_{\underline{\theta}}^{\underline{\theta} + \varepsilon} 2q(x)(1 - q(x))dx$  which goes to zero when  $\varepsilon$  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  $\varepsilon$ . The gain in welfare that comes from the fact that quantities are increased is first order in  $\varepsilon$ . Therefore, for a small enough  $\varepsilon$  this change is welfare improving. This concludes the proof that  $q(\underline{\theta}) = 1$ .

We know from above that  $q(\underline{\theta}) = 1$  and that the entire solution must be constrained, since the relaxed problem has an optimal solution that violates the monotonicity constraint. Therefore, there exist  $\theta', \bar{\theta}$  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  $[\underline{\theta}, \bar{\theta}]$ .<sup>31</sup> 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  $\theta''$  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

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

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

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

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

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

The fact that the monotonicity constraint ceases to bind at  $\theta''$  means that

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

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

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

<sup>31</sup> In other words,  $\bar{\theta}$  is the lowest type where the monotonicity constraint binds in the solution of the original problem.

and thus

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

We know that

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

and

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

Therefore,

$$\begin{aligned} \frac{\partial^2 H}{\partial q \partial \theta} \Big|_{q=q^*} &= 2(1 - 2q^*)\lambda'(\theta) + 2(1 - q^*)[\theta \widehat{c}'(\theta) + \widehat{c}(\theta)] \\ &= 2(1 - 2q^*)\left(\widehat{c}(\theta) - \theta(2q^* - (q^*)^2)\right) + 2(1 - q^*)\widehat{c}(\theta) \\ &\quad + 2(1 - q^*)\theta 2q^*(1 - q^*) \\ &= 2(\widehat{c}(2 - 3q^*) + \theta(q^*)^2). \end{aligned} \tag{12}$$

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  $q'(\bar{\theta}) = 0$  holds.<sup>32</sup> Then the same argument as above (see (8)) implies that  $\widehat{c}(\bar{\theta})(2 - 3q^*) + \bar{\theta}(q^*)^2 = 0$ . Therefore,  $\frac{\partial^2 H}{\partial q \partial \theta} \Big|_{q=q^*, \theta=\bar{\theta}} = 0$  must hold by (12). Also,  $\frac{\partial}{\partial \theta} \left(\frac{\partial^2 H}{\partial q \partial \theta} \Big|_{q=q^*}\right) = 2(\widehat{c}'(2 - 3q^*) + (q^*)^2) = 2(2q^*(1 - q^*)(2 - 3q^*) + (q^*)^2) = 2q^*(2(1 - q^*)(2 - 3q^*) + q^*) > 0$  for all relevant values of  $q^*$ . Therefore, together with  $\frac{\partial^2 H}{\partial q \partial \theta} \Big|_{q=q^*, \theta=\bar{\theta}} = 0$  we obtain that for all  $\theta \in (\bar{\theta}, \theta'']$

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

But comparing (10), (11), and (13) 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.

**Proof of Proposition 5.** Take any (potentially non-constant) path  $q_t, \tau_t$ . The proof establishes that the same entry function  $\widehat{c}$  can be induced by an appropriate policy that is constant over time. Moreover, total welfare is higher under this policy as the sum of consumer and producer surplus is larger than under the original non constant policy. First, it is clear that a one-time up-front transfer is without loss of generality as the innovator only cares about the present value of the transfers. The utility from reporting  $\widehat{\theta}$  when the type is  $\theta$  is

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

<sup>32</sup> This is an instance of the smooth pasting condition at point  $\bar{\theta}$  where the function switches from being strictly decreasing to being flat.

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

$$\begin{aligned} V'(\theta) &= \frac{\partial}{\partial \theta} U(\hat{\theta}, \theta) \Big|_{\hat{\theta}=\theta} \\ &= 2 \sum_{t=0}^{\infty} \delta^t q_t(\hat{\theta})(1 - q_t(\hat{\theta})). \end{aligned} \tag{14}$$

A similar argument as in Lemma 2 implies that incentive compatibility requires that  $\sum_{t=0}^{\infty} \delta^t q_t(\theta)$  is decreasing in  $\theta$ . 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.<sup>33</sup> It is then sufficient to prove that for any  $\theta$  the realized total surplus is larger than the one under the original policy. That is, it is sufficient to show that for all  $\theta$  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  $x$ .

### Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.ijindorg.2016.06.007](https://doi.org/10.1016/j.ijindorg.2016.06.007).

### References

Acemoglu, D., Akgigit, U., 2012. Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association* 10, 1–42.

Aghion, P., Bolton, P., Harris, C., Jullien, B., 1991. Optimal learning by experimentation. *Review of Economic Studies* 55, 621–654.

Aghion, P., Tirole, J., 1994. The management of innovation. *Quarterly Journal of Economics* 109, 1185–1209.

Akgigit, U., Liu, Q., 2014. The Role of Information in Innovation and Competition, Working Paper.

Arrow, K., 1962. Economic welfare and the allocation of resources for inventions. In: Nelson, R. (Ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press, Princeton, pp. 609–625.

<sup>33</sup> The incentive conditions are not affected either, see (14).

- Banerjee, A., Hollis, A., Pogge, T., 2010. The health impact fund: incentives for improving access to medicines. *Lancet* 375, 166–169.
- Baron, D., Besanko, D., 1984. Regulation and information in a continuing relationship. *Information Economics and Policy* 1, 267–302.
- Battaglini, M., 2005. Long-term contracting with Markovian consumers. *American Economic Review* 95, 637–658.
- Berry, S., Levinshohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. *Econometrica* 63, 841–890.
- Bloom, N., Schankerman, M., Reenen, J.V., 2013. Identifying technology spillovers and product market rivalry. *Econometrica* 81, 1347–1393.
- Brynjolfsson, E., Smith, M.D., 2000. Frictionless commerce? A comparison of internet and conventional retailers. *Management Science* 46, 563–585.
- Brynjolfsson, E., Zhang, M., 2007. Innovation incentives for information goods. *Innovation Policy and the Economy* 7, 99–123.
- Chari, V.V., Golosov, M., Tsyvinski, A., 2012. Prizes and patents: using market signals to provide incentives for innovations. *Journal of Economic Theory* 147, 781–801.
- Chintagunta, P., Nair, H., 2011. Discrete-choice models of consumer demand in marketing. *Marketing Science* 30, 977–996.
- Cho, S., Rust, J., 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.
- Cornelli, F., Schankerman, M., 1999. Patent renewals and R&D incentives. *RAND Journal of Economics* 30, 197–213.
- Deaton, A., Muellbauer, J., 1980. An almost ideal demand system. *American Economic Review* 70, 312–326.
- Fox, J., Gandhi, A., 2011. Identifying Demand with Multidimensional Unobservables. NBER Working Paper 17557.
- Gans, J., King, S., 2007. Patent length and the timing of innovative activity. *Journal of Industrial Economics* 55, 772–772.
- Gilbert, R., Shapiro, C., 1990. Optimal patent length and breadth. *RAND Journal of Economics* 21, 106–112.
- Glennester, R., Kremer, M., Williams, H., 2006. Creating markets for vaccines. *Innovations: Technology Governance Globalization* 1, 67–79.
- Guell, R., Fischbaum, M., 1995. Toward allocative efficiency in the prescription drug industry. *Millbank Quarterly* 73, 213–230.
- Hellwig, M., 2009. A Maximum Principle for Control Problems with Monotonicity Constraints. Max Planck Institute Working Paper.
- Hopenhayn, H., Llobet, G., Mitchell, M., 2006. Rewarding sequential innovations: prizes, patents, and buyouts. *Journal of Political Economy* 114, 1041–1068.
- Hovenkamp, H., 2004. United States Antitrust Policy in an Age of IP Expansion. University of Iowa Legal Studies Research Paper No. 04-03.
- Hopenhayn, H., Mitchell, M., 2001. Innovation variety and patent breadth. *RAND Journal of Economics* 32, 152–166.
- Judd, K., 1998. *Numerical Methods in Economics*. MIT Press.
- Kremer, M., 1998. Patent buyouts: a mechanism for encouraging innovation. *Quarterly Journal of Economics* 113, 1137–1167.
- Laffont, J.-J., Tirole, J., 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press.
- Mastroianni, G., Milovanovic, G., 2008. *Interpolation Processes: Basic Theory and Applications*. Springer-Verlag, Berlin.
- McKinsey, 2009. And the Winner is Capturing the Promise of Philanthropic Prizes. McKinsey and Co, Sydney, Australia.
- Mitchell, M., Moro, A., 2006. Persistent distortionary policies with asymmetric information. *American Economic Review* 96, 387–393.
- Murray, F., Stern, S., Campbell, G., MacCormack, A., 2012. Grand innovation prizes: a theoretical, normative and empirical evaluation. *Research Policy* 41, 1779–1792.
- Nakamura, E., Steinsson, J., 2008. Five facts about prices: a re-evaluation of menu cost models. *Quarterly Journal of Economics* 123, 1415–1464.
- Reiss, P., Wolak, F., 2007. Structural econometric modeling: rationales and examples from industrial organization. *Handbook of Econometrics*, Vol. 6A. Elsevier BV.

- Scotchmer, S., 1999. On the optimality of the patent renewal system. *RAND Journal of Economics* 30, 181–196.
- Shavell, S., Ypersele, T.V., 2001. Rewards versus intellectual property rights. *Journal of Law and Economics* 44, 525–547.
- Tirole, J., 2003. Incomplete contracts: where do we stand. *Econometrica* 67, 741–781.
- Weyl, G., Tirole, J., 2012. Market power screens willingness to pay. *Quarterly Journal of Economics* 127, 1971–2003.
- Williams, H., 2012. Innovation inducement prizes: connecting research to policy. *Journal of Policy Analysis and Management* 3, 1–25.
- Wright, B., 1983. The economics of invention incentives: patents, prizes and research contracts. *American Economic Review* 73, 691–707.