DISCUSSION PAPER SERIES

DP16826

Acquisitions, innovation, and the entrenchment of monopoly

Vincenzo Denicolò and Michele Polo

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

Acquisitions, innovation, and the entrenchment of monopoly

Vincenzo Denicolò and Michele Polo

Discussion Paper DP16826
Published 20 December 2021
Submitted 16 December 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Acquisitions, innovation, and the entrenchment of monopoly

Abstract

We analyze a dynamic model of repeated innovation where inventors may be acquired by an incumbent or else challenge its leadership. In the short run, acquisitions always spur innovation because of the invention-for-buyout effect. In the long-run, however, acquisitions may stifle innovation because of a countervailing effect, the entrenchment of monopoly. The entrenchment-of-monopoly effect arises when the incumbent's dominance depends on its past activity levels and thus is reinforced by repeated acquisitions over time. We show that if the entrenchment-of-monopoly effect is sufficiently strong, forward-looking policymakers should prohibit acquisitions in the anticipation of their long-run negative impact on innovation. This argument provides a new theory of harm that can be used to block acquisitions that might otherwise go unchallenged.

JEL Classification: L10, L40

Keywords: Acquisitions, Innovation, market power, Antitrust Policy

Vincenzo Denicolò - vincenzo.denicolo@unibo.it
University of Bologna and CEPR

Michele Polo - michele.polo@unibocconi.it
Bocconi University

Acquisitions, innovation and the entrenchment of monopoly*

Vincenzo Denicolò[§] and Michele Polo[§§]
[§] University of Bologna and CEPR
[§§] Bocconi University, GREEN and IGIER

December 16, 2021

Abstract

We analyze a dynamic model of repeated innovation where inventors may be acquired by an incumbent or else challenge its leadership. In the short run, acquisitions always spur innovation because of the invention-for-buyout effect. In the long-run, however, acquisitions may stifle innovation because of a countervailing effect, the entrenchment of monopoly. The entrenchment-of-monopoly effect arises when the incumbent's dominance depends on its past activity levels and thus is reinforced by repeated acquisitions over time. We show that if the entrenchment-of-monopoly effect is sufficiently strong, forward-looking policymakers should prohibit acquisitions in the anticipation of their long-run negative impact on innovation. This argument provides a new theory of harm that can be used to block acquisitions that might otherwise go unchallenged.

Keywords: Acquisitions; Innovation; Market power; Invention-for buyout; Entrenchment-of-monopoly

---

*E-mail addresses: vincenzo.denicolo@unibo.it, michele.polo@unibocconi.it. We thank, without implicating, Igor Letina, Massimo Motta, Fausto Panunzi, Patrick Rey, Thom Schmitz, Piercarlo Zanchettin and participants at the 2021 CRESSE Conference for useful comments and discussions.

1 Introduction

Technology giants often acquire innovative start-ups with a high potential for growth.[1] Many of these acquisitions today escape antitrust scrutiny, or are cleared by the agencies, because merger control focuses on the size of the firms at the time of the acquisition, and target firms are often acquired while they still are small. In the last decade, however, this long-established approach has increasingly been called into question. Critics argue that it is ill suited to innovative industries, where the acquisition of small entrants may impede the mechanisms of Schumpeterian competition (i.e., the replacement of market leaders by new innovators), thereby stifling innovation.[2] Advocates of the permissive policy counter that the prospect of being acquired increases the incentives to innovate of small enterprises that do not possess the assets required to effectively bring their innovations to the market – a mechanism that has come to be known as the *invention-for-buyout* effect.[3]

The entrenchment of monopoly. To shed light on this debate, in this paper we propose a dynamic model of repeated innovation and acquisition derived from Segal and Whinston (2007). In the model, acquisitions have both pro- and anti-competitive effects. The pro-competitive effects are created by the invention-for-buyout mechanism mentioned above. The anti-competitive effects derive from a different mechanism, which we refer to as *entrenchment of monopoly.*

The entrenchment-of-monopoly effect arises when the market power of the current incumbent depends on its past activity levels. This dependence may be due to various reasons, such as for instance intertemporal network externalities, dynamic economies of scale, exclusive access to more and better data, switching costs, and the like. All these factors imply that by increasing the size of the incumbent, acquisitions may strengthen its market dominance. This makes it more difficult for future inventors to enter the industry, reducing their incentives to innovate. This is true even if future inventors are in turn acquired, because the entrenchment of monopoly worsens their outside options and therefore reduces the share of innovative rents they can obtain in the bargaining with the incumbent over the acquisition price.

Results. Our main result is that the competitive effects of acquisitions depend on the time horizon of the assessment. In the short run, acquisitions increase the incentive to innovate because of the invention-for-buyout effect. In the long run, however, acquisitions may reduce both the rate of innovation and consumers' surplus provided that the entrenchment-of-monopoly effect is strong enough. In other words, the invention-for-buyout effect dominates in the short run, but the entrenchment-of-monopoly effect may prevail in the long run.

We also show that the optimal policy can be state dependent: it may be optimal

[1]Recent examples include Facebook's acquisition of WhatsApp and Istagram, Google's acquisition of Youtube and Waze, and Microsoft's acquisition of Linkedin. These prominent cases are just the tip of the iceberg. Focusing only on the "big five", Motta and Peitz (2021) report 42 acquisitions by Amazon, 33 by Apple, 21 by Facebook, 48 by Google, and 53 by Microsoft in the 2015-2020 period.

[2]See, among the several policy reports prepared recently, Cremer et al. (2019), Furman et al. (2019), Scott Morton et al. (2019).

[3]The effect is so named after Rasmusen's (1988) entry-for-buyout.

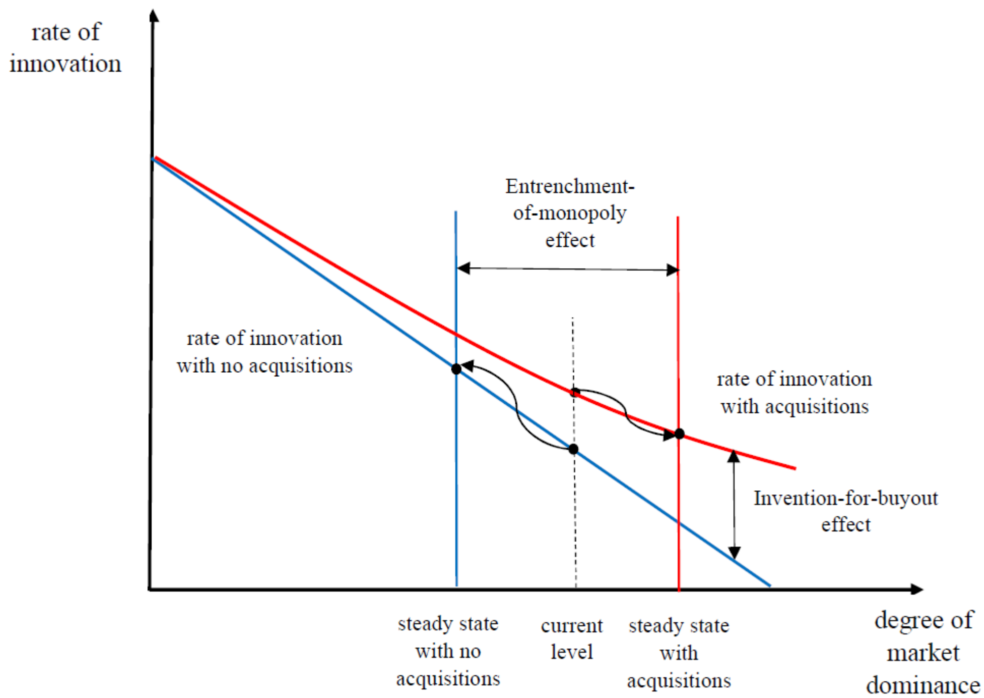

Figure 1: The decreasing curves represent the equilibrium rate of innovation when acquisitions are permitted (red) or prohibited (blue). The distance between the two curves measures the invention-for-buyout effect. The vertical lines instead represent the long-run level of market dominance, which is higher when acquisitions are pemitted because of the entrenchment-of-monopoly effect. The curley arrows represent the process of convergence to the steady states, starting from the current level of market dominance.

to permit acquisitions as long as market dominance is weak and prohibit them once repeated acquisitions have made it too strong.

Mechanism. The logic behind these results is illustrated in Figure 1, where the case that acquisitions are permitted is depicted in red, and the case that they are prohibited in blue. The figure displays three key properties of the model's equilibrium, which our conclusions rest on:

- both with and without acquisitions, the rate of innovation decreases with the degree of market dominance, as stronger dominance reduces innovators' ability to appropriate the returns from their innovations;
- for any given level of market dominance, the rate of innovation is higher when acquisitions are permitted, reflecting the invention-for-buyout effect;
- the long-run degree of market dominance is higher if acquisitions are permitted than if they are prohibited, reflecting the entrenchment-of-monopoly effect.

The implication is that prohibiting acquisitions may increase the long-run rate of innovation, as shown in the figure.

Policy implications. Our results have noteworthy implications for policy. On methodological grounds, they imply that acquisitions should not be assessed one by one, in isolation from each other. This myopic approach, which in our model would produce a lenient policy, is generally sub-optimal. Forward-looking policymakers should instead consider the cumulative dynamic effects of alternative policy rules.[4]

On substantive grounds, our analysis provides a theory of harm that can be used to block acquisitions that might otherwise go unchallenged. In particular, prohibiting acquisitions benefits consumers if the social discount factor is sufficiently high and the entrenchment-of-monopoly effect is sufficiently strong. The analysis also clarifies the role of other factors, such as for instance the inventors' bargaining power *vis-a-vis* the incumbent, or the speed with which innovations are imitated.

Related literature. Although the risk of entrenchment of monopoly is often mentioned in the acquisition policy debate,[5] to the best of our knowledge this is the first paper that provides a formal analysis of this possibility.[6] Previous research on the impact of acquisitions on innovation[7] has either focused on static models of isolated innovations, or else posited that the degree of market dominance is time invariant. In these settings, the entrenchment-of-monopoly effect cannot arise.

Models of isolated innovations provide the simplest analytical setting where the invention-for-buyout effect can be demonstrated: see, for instance, Mason and Weeds (2013), Phillips and Zhdanov (2013), and Letina et al. (2020).[8] Static models have also been used to uncover various adverse effects of acquisitions. In an important contribution, Cunningham et al. (2021) have shown, both theoretically and empirically, the profitability of "killer acquisitions." Acquisitions are killer when the acquirer suppresses one or more research projects initiated by the target firm in order to prevent the cannibalization of its own market. In a similar vein, Kamepalli et al. (2020) have suggested the possibility of a "kill zone," where entrants, whose innovations would challenge the incumbent's dominance, are discouraged by the threat of incumbent's aggressive reaction. Our analysis abstracts from these effects.

Static models have also been used to show that acquisitions can affect not only

[4] This marks an important difference with Nocke and Whinston (2010), where a myopic merger policy based on consumer surplus as a welfare criterion is optimal also in a dynamic setting where a series of mergers may be proposed over time. The reason for this difference is that the entrenchment-of-monopoly effect makes our model more intrinsically dynamic: acquisitions may affect not only the set of active firms, as in Nocke and Whinston (2010), but also future demand, the degree of market dominance, and the incentives to innovate.

[5] See, for instance, Scott Morton et al. (2019) and Bryan and Hovenkamp (2020).

[6] In a recent paper, however, Fons-Rosen et al. (2021) have developed an endogenous growth model with acquisitions where the incumbent is the technological leader and may lead by a margin of variable size. The size of the incumbent's lead determines the magnitude of Arrow's replacement effect. An increase in the incumbent's lead therefore reduces the incentives to innovate, with an effect similar to our entrenchment of monopoly. The model, however, cannot be solved analytically and is simulated numerically.

[7] There is also an extensive literature on the impact of mergers on innovation: see Bourreau et al. (2021) for an excellent synthesis. Unlike the case of acquisition of start-ups, this literature studies the case in which mergers take place before investments in R&D are chosen.

[8] There can also be further positive effects: for example, acquisitions may relax the inventors' financial constraints, as in Fumagalli et al. (2021).

the rate but also the direction of technological progress. In particular, acquisitions can impact the diversity of research projects (Letina et al., 2020), whether innovators target substitutes or complements of the incumbent's product (Shelegia and Motta, 2021; Dijk, Moraga-Gonzàlez and Motchenkova, 2021), or whether they target the product of the market leader or of the follower (Bryan and Hovenkamp, 2020).

This paper, in contrast, belongs in the strand of the literature that analyzes antitrust policy in dynamic models of repeated innovation. The pioneering contribution here is Segal and Whinston (2007). Below, we shall discuss at length the differences with their model; for now, suffices it to say that Segal and Whinston (2007) do not consider acquisitions and assume that the degree of market dominance does not change over time.

This latter assumption is also made by Cabral (2018, 2021). Cabral distinguishes between incremental and radical innovations. For incremental innovations, the invention-for-buyout effect implies that acquisitions spur innovation. Radical innovations are different, though. For them, the invention-for-buyout effect is mute as these innovations would not be transferred to the incumbent anyway. Still, acquisitions are not neutral because innovators may choose which type of innovation to target. When acquisitions are permitted, incremental innovations may therefore crowd-out the radical ones. Clearly, this crowding-out mechanism, which may reduce the overall rate of innovation, is different from the entrenchment of monopoly.

Katz (2021) also focuses on the invention-for-buyout mechanism. He notes that acquisitions raise the entrant's payoff; however, the incentive to innovate is determined by the rate with which the payoff increases with the size of the innovation, and this in principle may either increase or decrease. This remark applies also to our model, where however acquisitions increase both the level and the slope of the inventor's profit.

While we can solve our model in closed form and derive our results analytically, other papers have resorted to numerical analysis to study a richer industry dynamics. However, existing computational dynamic models do not feature the entrenchment-of-monopoly effect. The results they produce are therefore driven, essentially, by the invention-for-buyout effect.[9]

Structure of the paper. In the next section, we outline a tractable model of repeated innovation and acquisitions. Section 3 derives the model's equilibrium when acquisitions are permitted. Section 4 examines the effects of acquisitions on the rate of innovation. Section 5 and 6 analyze the optimal acquisition policy when antitrust authorities adopt non-contingent and state-contingent policy rules, respectively. Section 7 discusses the robustness of our findings. Section 8 summarizes and concludes the paper. Proofs are collected in an Appendix.

[9] In particular, Hollenbeck (2019) focuses on the trade-off between the static allocative effects of acquisitions, which are always negative, and the dynamic effects *via* the level of innovation, which in his model are always positive due to the invention-for-buyout effect. As a result, he finds that acquisitions are welfare-reducing in the short run but can be welfare-increasing in the long run. Mermelstein et al. (2020) consider a model where entry is inefficient because of economies of scale in production and investment. The invention-for-buyout effect implies that acquisitions facilitate entry, but the fact that entry is inefficient implies that a restrictive policy may be optimal.

2 The model

In this section, we propose a tractable model of repeated innovation, adapted from Segal and Whinston (2007). The model is tailored to industries where the ability to innovate is diffused, so that it is unlikely that the same firm may innovate repeatedly, and that the successful innovator may be identified *ex ante*.

This focus implies three main differences from Segal and Whinston (2007). First, we assume that in each period a new firm, randomly drawn from a large set of potential innovators, obtains an innovation.[10] The inventor enters the market and in the absence of acquisitions stays active for two periods. In the first period of its life-cycle, the inventor is the technological leader and competes with an incumbent that has some other sort of competitive advantage. In the second period, it becomes the new incumbent and competes with the next inventor, which has meanwhile entered the market. In the subsequent period, it is absorbed by a competitive fringe. Therefore, in the absence of acquisitions the incumbent's market power is temporary.

Second, we consider the possibility that incumbents may acquire the innovative entrants. Since inventors cannot be identified *ex ante*, acquisitions may take place only *ex post*, when the invention has already materialized. The acquisition has two consequences: the incumbent gains control over the entrant's innovative technology, and it also "rejuvenates," postponing exit by one period. Therefore, the incumbent can remain active forever if it systematically acquires the new inventors.

Third, while in Segal and Whinston (2007) the incumbent's advantage is time-invariant, we allow it to change over time as a function of the industry's past history. In particular, incumbent's dominance is strengthened by acquisitions, which can therefore lead to the entrenchment of market power.

These ingredients are embedded in a fully specified model of the industry, whose other components are kept as simple as possible in order to obtain closed-form solutions. After developing the analysis, we shall discuss the effect of relaxing some of these assumptions.

2.1 Demand

A vertically differentiated product is demanded by a mass of homogeneous consumers (normalized to one), who may purchase either 0 or 1 units. The net utility from purchasing one unit of a product of quality q^i at price p^i is

$$U^i = q^i - p^i, \tag{1}$$

where the willingness to pay for quality is also normalized to one. The utility of not purchasing is normalized to 0. These normalizations do not entail any loss of generality.

[10] Segal and Whinston's analysis mainly focuses on the case where the same two firms alternate in leading.

2.2 Innovation, entry and market structure

In each period $t = 1, 2, ..$, one outsider, randomly drawn out of a number of potential innovators, gets an idea for improving the existing technology q_{t-1}. This outsider then becomes the period-t inventor by developing its idea into an innovation, i.e., a product of quality $q_t > q_{t-1}$.[11] The inventor chooses the innovation size $\Delta_t = q_t - q_{t-1}$ so as to maximize its profits. The cost of raising quality by Δ_t is independent of the current level q_{t-1} and is quadratic in Δ_t. With another innocuous normalization, we can write:

$$C(\Delta_t) = \frac{1}{2}\Delta_t^2. \tag{2}$$

After developing its invention, inventor q_t enters the market. In the absence of acquisitions, in period t inventor q_t is the technological leader but faces competition from the incumbent (i.e., inventor q_{t-1}). In period $t+1$, inventor q_t becomes the new incumbent and competes with inventor q_{t+1}.

As time passes, inventions can be imitated by a competitive fringe. We assume that the innovation is used exclusively by the inventor for two periods and can henceforth be imitated freely. (For example, the invention might be protected by a patent that lasts for two periods.) As a consequence, in the absence of acquisitions inventor q_{t-2} is absorbed by the competitive fringe in period t. Thus, in each period t there are three types of firms: an entrant (E), which supplies a product of quality $q_t^E = q_t$, an incumbent (I) with quality $q_t^I = q_{t-1}$, and a competitive fringe (F) with $q_t^F = q_{t-2}$.[12]

The unit production cost is independent of quality and is normalized to 0.[13]

2.3 Market dominance

While being technological laggards, incumbents have been active for longer and thus may have acquired a competitive advantage of some other sort. This competitive advantage may be due, for instance, to intertemporal network externalities, dynamic economies of scale, exclusive access to more and better data, or switching costs. Conversely, entrants may face various entry hurdles; for example, some consumers may be unwilling to try new products, or may not be aware of their existence.

To capture these factors in a tractable way, we assume that demand is not entirely

[11] With an abuse of notation, we denote by q_t both the quality and the identity of inventor.

[12] One can allow for the possibility that imitation is faster, e.g., because intellectual property protection is imperfect. For example, continuing to assume that invention q_t is fully protected in period t and can be imitated freely in period $t+2$, the innovation could be imitated partially in period $t+1$. In this case, the competitive fringe's quality would be $q_t^F = q_{t-2} + \vartheta(q_{t-1} - q_{t-2})$, where parameter ϑ is an index of the speed of imitation, or an inverse index of the strength of intellectual property protection. When $\vartheta = 1$, the competitive fringe imitates the innovation in just one period, whereas the baseline case where imitation takes two periods is re-obtained for $\vartheta = 0$. One can easily verify that the only change in our formulas is that the discount factor δ must be replaced by $\delta(1 - \vartheta)$. Therefore, the parameter δ captures not only the private rate of time preference but also the strength of intellectual property protection.

[13] This normalization, like the previous ones, does not involve any loss of generality.

contestable:[14] in each period, a fraction μ_t of consumers are "captive" and cannot purchase from the new entrant; the remaining $1 - \mu_t$ consumers, on the other hand, are "free" and can purchase also the new product.[15] Specifically, we assume that the captive consumers in period $t + 1$ are a fraction of those who in the previous period purchased from the q_t entrant. Thus, the number of captive consumers evolves over time according to the following equation:

$$\mu_{t+1}^{NA} = \kappa(1 - \mu_t)x_t^E, \tag{3}$$

where x_t^E is the fraction of free consumers who purchase from the entrant, $\kappa \in [0, 1]$ is the fraction of such consumers who are fidelized and turned into captive, and the superscript NA stands for "no acquisition." This assumption captures the notion that the incumbent's advantage is related to its past activity level.

The assumption may be interpreted literally, or as a metaphor for various possible sources of the incumbent's competitive advantage. The literal interpretation may be justified as follows. Suppose that consumers face a cost of switching to the newest product, because they must learn how to use it, or must conduct a search to become aware of its existence. These learning costs are heterogeneous. For a fraction κ of consumers, they are sufficiently high that it is worth paying them only to move ahead by two quality levels. As a consequence, a high-cost consumer who purchased the state-of the art product q_t in period t would not be willing to switch to product q_{t+1} in period $t + 1$. In other words, such a consumer would be captive. For the remaining fraction $1-\kappa$ of consumers, on the other hand, learning costs are negligible. Therefore, these consumers can always purchase all products supplied, including the newest one.[16]

In this framework, in period t the q_t entrant can serve only the $1 - \mu_t$ consumers who are free. Suppose it serves a fraction x_t^E of them. The consumers who purchase the state-of-the-art product in period t are then $(1 - \mu_t)x_t^E$. A fraction κ of these consumers are high cost, and therefore in the next period, when the q_t entrant becomes the new incumbent, they will form the incumbent's captive consumer base. This explains equation (3).

Intuitively, learning costs create inertia in firms' market shares, translating the entrant's sales in the first period of its life-cycle into a captive consumer base that it can exploit in the next period. Parameter κ measures the size of this effect and therefore plays a crucial role in our model.

[14] This assumption is quite common in the literature on exclusionary conduct: see, for instance, Ide and Montero (2020), Oertel and Schmutzler (2021), and the literature cited therein.

[15] However, captive consumers can purchase from the fringe. This limits the extent to which they can be exploited by the incumbent and ensures the stationarity of the model.

[16] Consumers need not be permanently high or low cost: in fact, the parameter κ may be interpreted as the probability that a consumer is high cost in a given period. Any level of correlation across periods is consistent with our formulation.

2.4 Acquisitions

If acquisitions are permitted, the incumbent may acquire the inventor after the latter has entered the market. Let P_t denote the acquisition price paid by the incumbent.[17] For simplicity, we assume that one of the two firms makes a take-it-or-leave-it offer to the other. (With a slight abuse of terminology, we shall refer to this stage of the game as the "bargaining process."[18]) This implies that bargaining is efficient, so acquisitions will take place whenever they are jointly profitable. We denote by α the probability that the entrant is the proponent and the incumbent is the receiver; with probability $1-\alpha$, the roles of the firms are reversed. Thus, α is the share of the bargaining surplus obtained on average by the entrant – a measure of its bargaining power.

The entity resulting from the acquisition is denoted by M. The entrant provides to firm M its new technology q_t, while the incumbent brings in the old technology q_{t-1} and its captive consumer base. However, we assume that after the acquisition the new entity is able to remove the factors that would otherwise prevent captive consumers from purchasing the new product q_t.[19] For example, consumers do not have to bear any learning costs as long as they purchase from the same firm. Therefore, with acquisitions the dynamics of μ_t becomes :

$$\mu_{t+1}^{A} = \kappa x_t^M, \tag{4}$$

where x_t^M denotes the fraction of consumers served by firm M, and superscript A stands for "acquisition." The merged entity can serve more consumers than the entrant, so it can build a larger captive consumer base for the next period. This is how acquisitions increase the merged entity's market dominance in our model. We shall refer to this as the entrenchment-of-monopoly effect of acquisitions.

2.5 Timing

We consider an infinite horizon game in discrete time. Each period t is divided into three stages. In the first stage (*ex ante*), the inventor is randomly selected and chooses the innovation size, Δ_t. In the second stage (*interim*), the incumbent and the entrant bargain over the acquisition price and, if an agreement is reached, the acquisition occurs. In the third stage (*ex post*), firms compete in prices. This sequence of events is repeated in every period $t = 1, 2, ...$ The second stage is absent if acquisitions are prohibited.

Firms are risk neutral and maximize intertemporal profits, where future values are discounted by the common discount factor $\delta < 1$. Total discounted profits as of time t are denoted by Π_t^i, and current profits are denoted by π_t^i, with $i \in \{E, I, M\}$.

[17] The assumption that the incumbent acquires the entrant, and not the other way around, is just an accounting convention. Nothing would change if the roles of the firms in the acquisition process were reversed.

[18] We adopt a strategic approach to the bargaining process to avoid mixing notions from cooperative and non-cooperative game theory. However, many different bargaining solutions would lead to the same expected outcome as our non-cooperative assumptions.

[19] See Kamepalli et al. (2020) for a similar assumption.

We analyze the Markov perfect equilibria of the game. Under our assumptions, it is clear that at the beginning of each period t, the payoff-relevant variables are μ_t, q_{t-1} and q_{t-2}. At the *interim* stage, i.e. after the entrant has chosen the size of the new innovation Δ_t, the set of payoff-relevant variables includes also q_t.

To simplify the presentation, we adopt the following tie-breaking rule: when a consumer or a firm is indifferent among different actions, it chooses the one that maximizes aggregate profits.[20]

3 The acquisition game

In this section, we find the equilibrium under the assumption that acquisitions are always permitted. Since the merged entity can replicate any behavior of the entrant and the incumbent, acquisitions are weakly profitable in our model. (In fact, we shall presently show that they are strictly profitable.) This implies that acquisitions will always take place in equilibrium.

To ensure perfectness of the equilibrium, we start from the pricing subgames and then proceed to the bargaining over the acquisition price and the choice of innovation size.

3.1 Pricing subgames

Post-acquisition. We begin from the pricing subgame that is actually played on the equilibrium path, i.e., the one starting after the incumbent and the entrant have merged.

The competitive constraint in this case comes only from the fringe, which supplies the best freely available quality, $q_t^F = q_{t-2}$, and prices it at cost, $p_t^F = 0$. The equilibrium strategy of the merged entity is given by the following lemma.

Lemma 1 *The merged entity supplies only one product of quality $q_t^M = q_t$. It serves all consumers ($x_t^M = 1$) at price $p_t^M = \Delta_t + \Delta_{t-1}$, reaping a profit of*

$$\pi_t^M = \Delta_t + \Delta_{t-1}. \tag{5}$$

The intuition is simple. Even if the competitive fringe does not sell any output in equilibrium, it exerts a competitive pressure by providing an outside option to consumers. The merged entity must then undercut the fringe in utility space, charging a price equal to the value of the quality differential.

The presence of the competitive fringe prevents prices and profits from increasing over time in spite of the fact that the quality level grows unboundedly. From an economic point of view, this guarantees that all benefits from technological progress eventually accrue to consumers; from an analytical point of view, it guarantees the stationarity of the equilibrium.

[20] This assumption captures the idea that the stronger firm could slighty reduce the price to break the indifference.

Two further remarks are in order. First, the merged entity always uses the state-of-the-art technology q_t, so there are no "killer acquisitions" in our model. Second, in equilibrium the merged entity serves all consumers and therefore creates a base of κ captive customers for the next period.

No acquisition. Next, we characterize the price equilibrium that arises, out of the equilibrium path, if the incumbent does not acquire the entrant.[21] In this case, there are two active firms beyond the competitive fringe. In the baseline specification of the model, we assume that these two firms choose their prices sequentially, with the incumbent acting as the price leader.

This assumption serves two purposes. First, it implies that, for any given state of the technology, acquisitions do not reduce consumer surplus, which is always pinned down by the outside option provided by the fringe. For our purposes, this is a conservative property that biases the analysis against the prohibition of acquisitions. Second, the assumption guarantees the existence of a pure strategy equilibrium. With simultaneous moves, in contrast, a pure-strategy pricing equilibrium generally fails to exist. Intuitively, the existence of captive consumers is similar to a capacity constraint, as the entrant cannot supply more than $(1-\mu_t)$ units. Later we shall also consider the alternative timings.

Lemma 2 *If the incumbent acts as a price leader, it serves all captive consumers and the entrant serves all free consumers ($x_t^E = 1$). The incumbent prices at $p_t^I = \Delta_{t-1}$ and obtains a profit of*

$$\pi_t^I(\mu_t) = \mu_t \Delta_{t-1}. \tag{6}$$

The entrant's equilibrium price is $p_t^E = \Delta_t + \Delta_{t-1}$, so the profit it earns in the first period of its life cycle is

$$\pi_t^E(\mu_t) = (1-\mu_t)\left(\Delta_t + \Delta_{t-1}\right). \tag{7}$$

When the incumbent acts as a price leader, both the incumbent and the entrant slightly undercut the competitive fringe in utility space, and the entrant also slightly undercuts the incumbent. As a result, consumers obtain the same net utility from any firm they could buy from.[22] The incumbent has no incentive to compete more aggressively for the free consumers because it anticipates that it would be outpriced by the entrant.

Implications. From Lemmas 1 and 2, it follows immediately:

Corollary 1 *Acquisitions are always strictly profitable.*

Acquisitions are profitable for two reasons. From a static viewpoint, they facilitate the diffusion of the innovation: the state-of-the-art product is sold not only

[21] The same equilibrium arises also on the equilibrium path if acquisitions are prohibited, as shown in footnote 24 below.

[22] Equilibrium outputs follow from our tie-breaking rule. As noted, that rule reflects the notion that the more efficient firm could outprice the rival by offering tiny price discounts that rivals cannot afford.

to the free consumers but also to the captive ones, and the extra surplus is reaped by firm M. From a dynamic viewpoint, acquisitions increase the fraction of captive consumers that the merged entity can exploit in the next period.

Next consider the impact of acquisitions on consumer surplus. Both with and without acquisitions, consumers obtain exactly the surplus guaranteed to them by the fringe:

$$CS_t = q_{t-2}. \tag{8}$$

From the viewpoint of consumers, therefore, acquisitions matter only to the extent that they affect innovation, and hence their future surplus.

Finally, Lemmas 1 and 2 imply that the dynamics of the share of captive consumers reduces to

$$\mu_{t+1}^{NA} = \kappa(1 - \mu_t) \tag{9}$$

with no acquisitions, whereas with acquisitions we simply have

$$\mu_{t+1}^{A} = \kappa. \tag{10}$$

3.2 The acquisition price

Proceeding with the backward induction, consider next the bargaining over the acquisition price.

Firms are forward looking and consider all future consequences of their current choices. These consequences are anticipated correctly and are embedded into the firms' value functions. In a Markov perfect equilibrium, these value functions can depend only on the payoff relevant variables. From the previous analysis, it appears that profits depend only on the quality differentials Δs and not on the quality levels qs. Thus, the period-t payoff-relevant variables are $\{\mu_t, \Delta_{t-1}\}$ at the *ex ante* stage and $\{\mu_t, \Delta_{t-1}, \Delta_t\}$ at the *interim* stage. Accordingly, denote by $V_t^i(\mu_t, \Delta_{t-1})$ the firms' value functions at the *ex ante* stage, and by $v_t^i(\mu_t, \Delta_{t-1}, \Delta_t)$ the *interim* value functions, for $i \in \{E, I, M\}$.

Since entrants are systematically acquired, the entrant's value function (gross of the innovation cost) must coincide with the acquisition price. Furthermore, the value functions must obey the following conditions (to simplify the notation, we shall suppress the dependence of the *interim* value functions on the relevant variables when this does not create confusion):

$$v_t^M = \pi_t^M + \delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t) \tag{11}$$

$$v_t^E = (1-\alpha)\left[\pi_t^E + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)\right] + \alpha\left(v_t^M - \pi_t^I\right) \tag{12}$$

$$v_t^I = v_t^M - v_t^E. \tag{13}$$

Equation (11) says that the merged entity obtains profits π_t^M in period t and then becomes the new incumbent with μ_{t+1}^A captive consumers, which gives a continuation value of $\delta V_{t+1}^{I}(\mu_{t+1}^{A}, \Delta_t)$. According to (12), the acquisition price (which as said coincides with the entrant's value function) equals the entrant's disagreement payoff plus a fraction α of the bargaining surplus. The entrant's disagreement payoff is equal to the entrant's current profit if it resists being acquired, π_t^E, plus the continuation value, $\delta V_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t)$. The "one-shot deviation principle" implies that the continuation value must be calculated on the expectation that even if there was no acquisition in period t, entrant q_t will nevertheless acquire entrant q_{t+1} in period $t+1$. However, an entrant that is not acquired will have only μ_{t+1}^{NA} captive consumers in the next period. As for the period-t incumbent, its disagreement payoff is simply π_t^I, as in the absence of an agreement it would exit the market in the next period. The bargaining surplus is therefore $v_t^M - \left[\pi_t^E + \delta V_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t) + \pi_t^I\right]$, whence condition (12) follows. The final condition says that the value of being the incumbent must be equal to the value of the merged entity minus the acquisition price. In other words, the acquisition does not change the sum of the firms' values because the extra-profits created by the merger are already included in the forward-looking valuation of the firms.

The system of equilibrium conditions (11)-(13) cannot be solved for the *interim* value functions yet, because it involves also the *ex ante* value functions $V_{t+1}^{I}(\mu_{t+1}^{A}, \Delta_t)$ and $V_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t)$, which depend on Δ_t and, implicitly, also on the expected future values $\Delta_{t+1}^{\exp}$, $\Delta_{t+2}^{\exp}$ etc. To proceed, we must therefore consider the optimal choice of the innovation size.

3.3 The innovation size

The equilibrium innovation size must satisfy the following condition:

$$\Delta_t^A(\mu_t, \Delta_{t-1}) = \arg\max \left[v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) - \frac{1}{2}\Delta_t^2\right]. \tag{14}$$

In a Markov perfect equilibrium, this optimal choice of Δ_t is anticipated by all players. This provides a link between the *ex ante* and *interim* value functions:

$$V_t^i(\mu_t, \Delta_{t-1}) = v_t^i\left[\mu_t, \Delta_{t-1}, \Delta_t^A(\mu_t, \Delta_{t-1})\right] \qquad \text{for } i \in \{E, I, M\}. \tag{15}$$

This completes the set of conditions that must simultaneously hold in equilibrium.

3.4 Equilibrium

It is easy to see that the set of Markov perfect equilibria coincides with the set of solutions to the system of equilibrium conditions (11)-(15), given the profit functions (5), (6) and (7).

The solution can be calculated explicitly thanks to a key simplifying property of the model: the profit functions π_t^i are additively separable in Δ_t and Δ_{t-1}. This separability implies that the marginal value of increasing the innovation size, $\frac{\partial v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)}{\partial \Delta_t}$, is independent of Δ_{t-1}, and so is the optimal value of Δ_t. This in turn implies that in spite of the forward-looking nature of system (11)-(15), Δ_t does not affect the expected future values $\Delta_{t+1}^{\exp}$, $\Delta_{t+2}^{\exp}$, ...

Separability allows for a two-stage solution procedure. Instead of solving for Δ_t^A and $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ simultaneously, in the first stage one calculates the derivative $\frac{\partial v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)}{\partial \Delta_t}$ and finds the equilibrium innovation size $\Delta_t^A(\mu_t)$. With the equilibrium function $\Delta_t^A(\mu_t)$ at hand, in the second stage one can find the entire value function $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ by a guess-and-verify method. This solution procedure is detailed in the proof of Lemma 3. It yields:

Lemma 3 *In the baseline model, the equilibrium innovation size is*

$$\Delta_t^A(\mu_t) = (1 + \delta\kappa)\left[1 - (1 - \alpha)\mu_t\right]. \tag{16}$$

The ex ante value functions are

$$V_t^E(\mu_t, \Delta_{t-1}) = \phi_0 + \phi_1\mu_t + \phi_2\mu_t^2 + (1 - \mu_t)\Delta_{t-1} \tag{17}$$
$$V_t^I(\mu_t, \Delta_{t-1}) = \varphi_0 + \varphi_1\mu_t + \varphi_2\mu_t^2 + \Delta_{t-1} \tag{18}$$

The coefficients ϕ_n and φ_n, for $n = 0, 1, 2$, are reported in the appendix, which also verifies that $V_t^I(\mu_t, \Delta_{t-1})$ increases with μ_t. Given $V_t^E(\mu_t, \Delta_{t-1})$ and $V_t^I(\mu_t, \Delta_{t-1})$, one can easily recover $V_t^M(\mu_t, \Delta_{t-1})$ and the *interim* value functions v^i from conditions (11)-(15).

4 Acquisitions and innovation

In this section, we analyze the impact of acquisitions on innovation. We show that prohibiting acquisitions always reduces the equilibrium innovation size in the short run but can increase it in the long run if the entrenchment-of-monopoly effect is sufficiently large.

4.1 Benchmark: no acquisitions

To proceed, we determine the innovation size when acquisitions are prohibited and thus never occur. In this case, the entrant's payoff is

$$\Pi_t^{E,NA} = \pi_t^E(\mu_t) + \delta\pi_{t+1}^I(\mu_{t+1}^{NA}), \tag{19}$$

where the profit functions are the same as in Lemma 2.[23] The equilibrium innovation size with no acquisitions then is $\Delta_t^{NA}(\mu_t) = \arg\max\left[\Pi_t^{E,NA} - \frac{1}{2}\Delta_t^2\right]$. Simple calculations lead to the following:

[23] In fact, this is not obvious, as firms are forward looking, and the entrant's continuation value is different with and without acquisitions. In the former case, the continuation value is $\delta V^I(\mu_{t+1}^A, \Delta_t)$; in the latter, it is $\delta\pi_{t+1}^I(\mu_{t+1}^{NA})$. However, the proof of Lemma 2 shows that all that matters is that the continuation value is non-decreasing in μ_{t+1}, which is true in both cases.

Lemma 4 *If acquisitions are always prohibited, the equilibrium level of innovation is*

$$\Delta_t^{NA}(\mu_t) = (1 + \delta\kappa)(1 - \mu_t). \tag{20}$$

4.2 Innovation and market dominance

As noted, a key property of our model is the negative impact of market dominance on innovation. Formally, we have:

Proposition 1 *The equilibrium level of innovation is a decreasing function of the degree of market dominance both when acquisitions are prohibited and when they are permitted.*

Proposition 1 says that incumbent's market dominance μ_t always exerts a negative effect on the entrant's incentives to innovate, irrespective of whether acquisitions are prohibited or not. To understand why this is so, consider first the case with no acquisitions. Entrant q_t's marginal benefit from increasing its innovation size is equal to the discounted sum of the derivatives of the profit functions (6) and (7) with respect to Δ_t. Inspection of the formulas reveals that these derivatives are given by the number of free consumers $1 - \mu_t$ in the first period of the entrant's life-cycle, and by the number of its captive consumers $\mu_{t+1}^{NA} = k(1 - \mu_t)$ in the second period. Both decrease with the degree of market dominance of the current incumbent, μ_t.

In the case with acquisitions, the effect is similar. The q_t entrant's outside option when bargaining on the acquisition price includes the profit $\Pi_t^E(\mu_t)$ that it would obtain if it resisted being acquired. A higher $\Pi_t^E(\mu_t)$ therefore allows the entrant to capture a larger fraction of the innovation rents. But we have just seen that the marginal impact of Δ_t on $\Pi_t^E(\mu_t)$ decreases with μ_t. This is therefore true also of the acquisition price. This implies that a higher μ_t reduces the entrant's incentives to invest in innovation, even when it expects to be acquired.

4.3 The short run

Comparing (16) and (20) one immediately obtains:

Proposition 2 *Provided that $\alpha > 0$, for any given t and $\mu_t > 0$ prohibiting acquisitions reduces the equilibrium innovation size in the current period:*

$$\Delta_t^A(\mu_t) > \Delta_t^{NA}(\mu_t). \tag{21}$$

Proposition 2 reflects the invention-for-buyout effect. Intuitively, the innovation is more valuable in the hands of the incumbent, which can supply the state-of-the-art product not only to the free consumers but also to the captive ones. By transferring the new technology to the incumbent, acquisitions therefore create a surplus, a share of which, when $\alpha > 0$, is obtained by the inventors. The prospect of being acquired thus increases the value of the innovation to forward-looking inventors, and hence their incentives to innovate.[24] The greater the entrant's bargaining power α, the stronger the invention-for-buyout effect.

4.4 The long run

However, acquisitions affect also the dynamics of μ_t. Starting from an arbitrary μ_0, μ_t will eventually converge towards its steady state level $\bar{\mu}$, which is

$$\bar{\mu}^A = \kappa \tag{22}$$

if acquisitions always occur, and

$$\bar{\mu}^{NA} = \frac{\kappa}{1+\kappa} \tag{23}$$

if they never occur. Clearly, $\bar{\mu}^A > \bar{\mu}^{NA}$. The difference

$$\bar{\mu}^A - \bar{\mu}^{NA} = \frac{\kappa^2}{1+\kappa} \tag{24}$$

reflects the entrenchment of monopoly created by acquisitions. It increases with κ, which may therefore be viewed as the entrenchment-of-monopoly parameter.

In the long run, the total effect of acquisitions is the combination of the short-run effect and the effect of the change in $\bar{\mu}$. If acquisitions are always prohibited, the long-run level of innovation is

$$\Delta^{NA}(\bar{\mu}^{NA}) = \frac{1+\delta\kappa}{1+\kappa}. \tag{25}$$

If acquisitions are always permitted, on the other hand, the level of innovation is

$$\Delta^A(\bar{\mu}^A) = (1+\delta\kappa)\left[1-(1-\alpha)\kappa\right]. \tag{26}$$

Comparing (25) and (26), it appears that if the entrenchment-of-monopoly effect is sufficiently strong, the positive short-run impact of acquisitions on innovation may be reversed in the long run.

Proposition 3 *In the long run, prohibiting acquisitions increases the equilibrium innovation size if*

$$\kappa > \frac{\alpha}{1-\alpha}. \tag{27}$$

Intuitively, the long-run impact of acquisitions is the sum of two components, namely, the difference between Δ_t^A and Δ_t^{NA} for any given μ_t, and the difference

[24] Yet, this is not a foregone conclusion. The incentive to innovate is not determined by the impact of acquisitions on the inventor's profit level, but by the marginal profitability of the innovation size. Proposition 1 guarantees that in our model the marginal and total effects go hand in hand. See Katz (2021) for a model where this property does not necessarily hold.

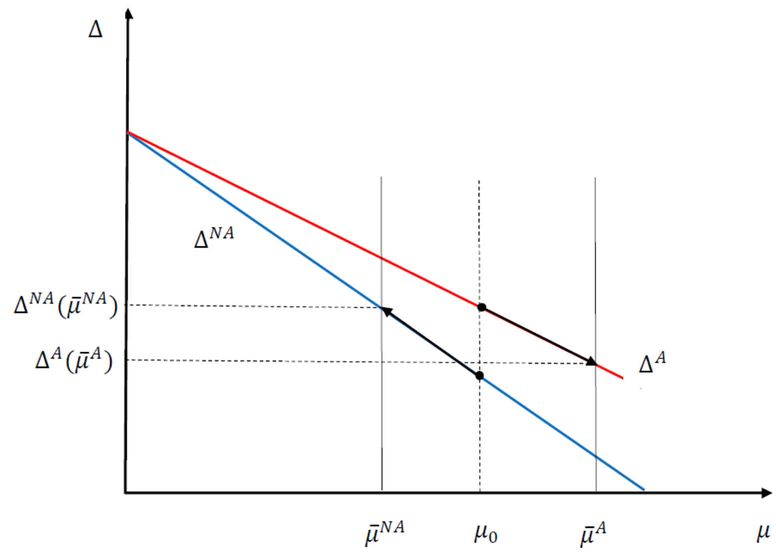

Figure 2: Equilibrium innovation size as a function of μ. For any given μ_0, Δ is higher with acquisitions. If acquisitions are prohibited, however, the incumbent's competitive advantage μ_t will decrease in the subsequent periods, whereas if acquisitions are permitted it will increase. The two vertical lines are the steady state values of μ in the two cases.

between $\bar{\mu}^A$ and $\bar{\mu}^{NA}$:

$$\Delta_t^A(\bar{\mu}^A) - \Delta_t^{NA}(\bar{\mu}^{NA}) = \left[\Delta_t^A(\bar{\mu}^A) - \Delta_t^{NA}(\bar{\mu}^A)\right] + \left[\Delta_t^{NA}(\bar{\mu}^A) - \Delta_t^{NA}(\bar{\mu}^{NA})\right]. \quad (28)$$

The first component reflects the invention-for-buyot effect and is positive by Proposition 2. The second component reflects the entrenchment of monopoly effect and is negative by Proposition 1. Condition (27) essentially determines when the second component prevails over the first one. Intuitively, the condition says that the entrenchment-of-monopoly parameter κ must be large and the entrant's bargaining power α, which determines the size of the invention-for-buyout effect, must be small.

Two properties of the baseline model are worth noting here. First, prohibiting acquisition may spur innovation only if $\alpha < \frac{1}{2}$. Second, the long-run effects of acquisitions do not depend on the discount factor δ (and hence on the speed of imitation, or the strength of intellectual property protection). These conclusions however no longer holds in simple extensions of the baseline model, as we shall see below.

Figure 2 illustrates the difference between the short-run and the long-run effects. The two downward sloping lines represent the equilibrium levels of innovation with and without acquisitions, as a function of μ_t. For any given μ_t, the level of innovation is higher with acquisitions. However, the level of innovation is, in both cases, a decreasing function of the degree of market dominance μ_t (Proposition 1), and in the long run market dominance is higher with acquisitions. Under condition (27), the move along the curve outweighs the gap between the curves.

4.5 Transitory dynamics

Our model is tractable enough to allow explicit calculation of the equilibrium dynamics of the innovation size Δ_{t+n}, starting from an arbitrary μ_t. When acquisitions are permitted, μ_t jumps immediately at its steady state level $\bar{\mu}^A = \kappa$, and so does the level of innovation:

$$\Delta_t^A = (1 + \delta\kappa) [1 - (1 - \alpha)\kappa] . \tag{29}$$

When, on the contrary, acquisitions are prohibited, we have

$$\mu_{t+n}^{NA} = \frac{\kappa}{1+\kappa} + \left(\mu_t - \frac{\kappa}{1+\kappa}\right)(-\kappa)^n \tag{30}$$

and

$$\Delta_{t+n}^{NA} = \frac{1+\delta\kappa}{1+\kappa} - (1 + \delta\kappa)\left(\mu_t - \frac{\kappa}{1+\kappa}\right)(-\kappa)^n . \tag{31}$$

Figure 3 illustrates the dynamics of Δ_t starting from $\mu_t = \kappa (= \bar{\mu}^A)$. Consider a shift from a lenient policy (the red curve) to a restrictive one (the blue curve) at time $t = 0$. In period 0, the level of innovation jumps down, as the invention-for-buyout effect vanishes. When acquisitions are prohibited, however, the share of captive consumers μ_{t+n} shrinks, reducing the entrenchment of the incumbent's monopoly. This has a positive effect on the entrant's innovative effort, which increases over time. In the counterfactual where acquisitions are permitted, on the other hand, μ_{t+n} remains at its steady state level $\bar{\mu}^A$. The figure represents the case where condition (27) holds. In this case, at some point in time the innovation size without acquisitions overcomes the counterfactual where acquisitions continue to be permitted.

5 Acquisition policy

In this section, we analyze the optimal antitrust policy, assuming that antitrust authorities adopt consumer surplus as a welfare criterion and discount future values by the social discount factor δ_S.[25] We assume that acquisitions are either always approved or never approved. The choice is made once and for all in period t.

Consider a generic period t. In view of (8), the policymaker's objective function is

$$\sum_{n=0}^{\infty} CS_{t+n}\delta_S^n = \frac{1}{1-\delta_S}\left(q_{t-2} + \delta_S q_{t-1} + \delta_S^2 \sum_{s=0}^{\infty} \delta_S^n \Delta_{t+n}\right) . \tag{32}$$

The first two terms inside brackets are pre-determined, so the policymaker's objective

[25] The social discount factor δ_S is generally greater than the private discount factor δ because benevolent policymakers ought to be more patient than private firms, and because δ may reflect not only the private rate of time preference but also the speed of imitation, as discussed in footnote 10 above. However, our formulas would continue to hold even if $\delta_S < \delta$.

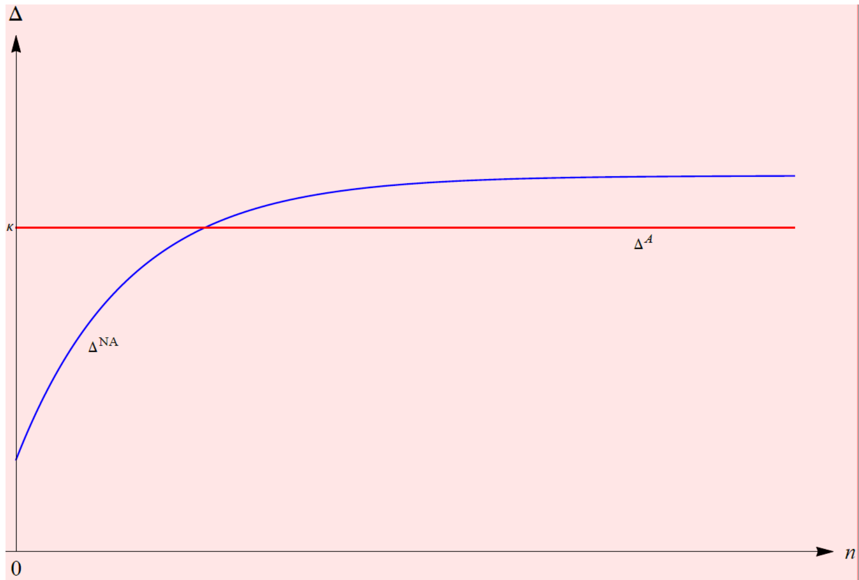

Figure 3: The dynamics of equilibrium innovation size when acquisitions are permitted and prohibited. The figure depicts a continuous time approximation of the discrete dynamics, which eliminates the oscillations that may be exhibited by the discrete dynamics. The picture has been drawn for $\kappa = \frac{3}{4}$, $\delta = \frac{19}{20}$, $\alpha = 0.35$ and $\mu_t = \kappa$.

function in period t effectively reduces to

$$W_t = \sum_{n=0}^{\infty} \delta_S^n \Delta_{t+n}. \tag{33}$$

Thus, social welfare comparisons boil down to the comparison of the discounted sum of current and future innovation sizes.

If acquisitions are always permitted, using (29) social welfare becomes:

$$W_t^A(\mu_t) = \frac{(1+\delta\kappa)\left[1-(1-\alpha)\delta_S\kappa\right]}{(1-\delta_S)} - (1-\alpha)(1+\delta\kappa)\,\mu_t. \tag{34}$$

If acquisitions are always prohibited, on the other hand, using (31) social welfare becomes:

$$W_t^{NA}(\mu_t) = \frac{1+\delta\kappa}{(1-\delta_S)(1+\delta_S\kappa)} - \frac{1+\delta\kappa}{1+\delta_S\kappa}\mu_t. \tag{35}$$

Comparing $W_t^{NA}(\mu_t)$ and $W_t^A(\mu_t)$, one obtains:

Proposition 4 *Prohibiting acquisitions increases social welfare if and only if*

$$\kappa > \frac{\alpha}{(1-\alpha)\,\delta_S}. \tag{36}$$

Even though in principle our formulation allows acquisition policy to depend on μ_t, it turns out that the optimal policy does not.

When $\delta_S \to 1$, condition (36) collapses to (27). In this limiting case, the weight of the transitory dynamics in the social welfare calculation becomes negligible, and the welfare comparison depends only on the steady state levels of innovation.

More generally, the effect of κ and α are the same as in Proposition 3, and for the same reasons. The only additional effect is that prohibiting acquisitions is more likely to be optimal, the higher the social discount factor δ_S. This makes intuitive sense: in our model prohibiting acquisitions is socially costly in the short run (Proposition 2) but may bring about long run benefits (Proposition 3). It is therefore natural that a restrictive policy may be optimal only if the policymaker is sufficient patient.

6 State-dependent policies

In this section, we inquiry into the possibility that acquisition policy may vary over time. In particular, if antitrust authorities can observe the state of the industry μ_t, they may condition acquisition policy on it.[26] It is therefore interesting to ask if it may be optimal to take advantage of this possibility.

We start with a useful preliminary result. So far, we have considered only the case where acquisitions are always permitted or always prohibited. Under a state-depenent policy, however, acquisitions may be permitted in certain periods and prohibited in others, depending on the current level of market dominance. A convenient simplifying property of the model is that the level of innovation in period t in fact depends only on whether acquisitions are permitted or prohibited in period t; it does not depend on whether acquisitions will be permitted or prohibited in subsequent periods.

Lemma 5 *If acquisitions are permitted in period t, then $\Delta_t(\mu_t) = \Delta_t^A(\mu_t)$ irrespective of acquisition policy in all subsequent periods. Likewise, if acquisitions are prohibited in period t, then $\Delta_t(\mu_t) = \Delta_t^{NA}(\mu_t)$ irrespective of acquisition policy in all subsequent periods.*

Firms are forward looking, and future acquisition policy affects the continuation values in the dynamic game. However, it does so by adding to v_t^E terms that do not depend on Δ_t and thus does not change the marginal profitability of R&D investment.[27]

6.1 Baseline model

In the baseline model, the possibility of conditioning acquisition policy on the degree of market dominance μ_t is in fact valueless.

[26] Strictly speaking, the state also includes Δ_{t-1}. However, this variable does not affect future innovation and affects consumer surplus in an additive, separable way. As such, it is evident that Δ_{t-1} cannot affect the optimal acquisition policy in period t.

[27] This simplifying property of the model rests on the separability of the profit functions in Δ_t and Δ_{t-1}. Therefore, Lemma 5 continues to hold also in the variant of the model considered in the second part of this section.

Proposition 5 *In the baseline model, the optimal acquisition policy does not depend on μ_t. Acquisitions should always be prohibited or always permitted, depending on whether condition (36) holds or not.*

This result rests on the fact that the profits functions π_t^i are linear in μ_t. This implies that $\Delta_t^A(\mu_t)$ and $\Delta_t^{NA}(\mu_t)$ are also linear in μ_t, which in turn implies that the sign of the difference $W_t^{NA}(\mu_t) - W_t^A(\mu_t)$ is independent of the initial condition μ_t.

6.2 The entrant as the price leader

To allow for the possibility that the optimal policy may be state dependent, we now consider a variant of the model where the entrant's profit π_t^E is a non-linear function of μ_t. Let us assume that if the acquisition does not take place, the entrant, rather than the incumbent, acts as the price leader. Reversing the order of moves in the pricing game changes the pricing equilibrium. Lemma 2 is now replaced by the following:[28]

Lemma 6 *If the entrant acts as a price leader, the incumbent serves all captive consumers and the entrant serves all free consumers ($x_t^E = 1$). The incumbent prices at $p_t^I = \Delta_{t-1}$ and obtains a profit of*

$$\pi_t^I(\mu_t) = \mu_t \Delta_{t-1}. \tag{37}$$

The entrant prices at $p_t^E = \Delta_t + \mu_t \Delta_{t-1}$, so the profit it earns in the first period of its life cycle is

$$\pi_t^E(\mu_t) = (1 - \mu_t)\left(\Delta_t + \mu_t \Delta_{t-1}\right). \tag{38}$$

The incumbent's price and profit are the same as in the baseline model. Differently from the baseline model, however, when the entrant acts as a price leader it cannot just undercut the rivals in utility space by setting $p_t^E = \Delta_t + \Delta_{t-1}$. If the entrant priced this way, the incumbent would now have an incentive to lower the price below Δ_{t-1} in order to capture the $(1 - \mu_t)$ free consumers. The entrant must therefore further reduce its own price down to the point where the incumbent's incentive to compete for the free consumers vanishes – a form of limit pricing.

If the acquisition takes place, the price equilibrium does not change. Compared to the new and more competitive benchmark, however, acquisitions now raise the equilibrium price paid by the free consumers, by an amount equal to $(1 - \mu_t)\Delta_{t-1}$. Thus, consumer surplus is $CS_t^A = q_{t-2}$ if the acquisition takes place in period t, as in the baseline model, but now it is

[28] Under the assumption that the dynamics of market dominance is given directly by (9) and (10), the case of simultaneous moves would lead to the same results as when the entrant acts as a price leader. The reason for this is that firms would price myopically, as the impact of acquisitions on future market dominance would not depend on current output levels. Now, in a static pricing game of simultaneous moves, each firm obtains the same payoff as it would if it acted as the price leader (as in Kreps and Scheinkman, 1983). Therefore, the entrant's profit would be (38), and the incumbent's profit would be (6), which coincides with (37). If one sticks to the original assumptions (3) and (4), however, the simultaneous-move pricing game becomes untractable.

$$CS_t^{NA} = q_{t-2} + (1-\mu_t)^2 \Delta_{t-1} \tag{39}$$

if the acquisition does not take place. Therefore, acquisitions have a direct, negative effect on consumer surplus for each given level of the technology. In addition, they have again a dynamic effect, *via* the equilibrium size of innovations.

Indeed, proceeding as in the baseline model one finds that when the entrant acts as a price leader and acquisitions are permitted, the equilibrium innovation size is:[29]

$$\Delta_t^A(\mu_t) = \chi_0 + \chi_1\mu_t + \chi_2\mu_t^2 \tag{40}$$

where:

$$\begin{aligned}
\chi_0 &= 1+\delta(1-\kappa+\kappa^2) \\
\chi_1 &= -(1-\alpha)\left(1-\delta\kappa+2\delta\kappa^2\right) \\
\chi_2 &= (1-\alpha)\,\delta\kappa^2 > 0.
\end{aligned}$$

Thus, the non-linearity of the profit function (38) translates into a non-linearity of the equilibrium innovation size function, $\Delta_t^A(\mu_t)$. The function $\Delta_t^A(\mu_t)$ is still decreasing in μ_t, provided that α is not too large, but now it is convex. When acquisitions are prohibited, on the other hand, the equilibrium does not change and is still given by condition (20).

The short-run effect of acquisitions on innovation (Proposition 2) is still positive. In fact, in this variant of the model the short-run effect is strictly positive even when $\alpha = 0$, $\mu_t = 0$, or both.[30] In the baseline model, in contrast, the invention-for-buyout effect vanishes for $\alpha = 0$ or $\mu_t = 0$.

The reason for this difference is as follows. Even if the q_t entrant has no bargaining power ($\alpha = 0$), its intertemporal payoff is higher when acquisitions are permitted because it will appropriate the extra-rents from the acquisition in period $t+1$ – when, having become the new incumbent, it will have all the bargaining power. The extra-rents generated by the period-$t+1$ acquisition now are $\mu_{t+1}\Delta_{t+1} + \left(1-\mu_{t+1}\right)^2\Delta_t$ and therefore increase with Δ_t. As a result, when acquisitions are permitted entrants now have an extra-incentive to invest in R&D even if $\alpha = 0$.[31]

The long-run effect of acquisitions on innovation (Proposition 3) is still the combination of the short-run effect and the change in the long-run degree of market dominance. If acquisitions are always prohibited, the long-run level of innovation

[29] The derivation of equation (40) may be found in the Appendix.

[30] To be precise, we have

$$\Delta_t^A - \Delta_t^{NA} = \delta\left[1-\kappa(1-\mu_t)\right]^2 + \alpha\mu_t\left\{1-\delta\kappa\left[1-2\kappa+(2-\delta)\kappa\mu_t\right]\right\}$$

whence the claims in the text immediately follow. Even if the difference $\Delta_t^A - \Delta_t^{NA}$ is now positive also when $\alpha = 0$ and $\mu_t = 0$, the gap is still an increasing function of both α and μ_t.

[31] In the baseline model, in contrast, the period-$t+1$ extra-rents are just $\mu_{t+1}\Delta_{t+1}$ and thus are independent of Δ_t. Thus, when $\alpha = 0$ acquisitions would increase the entrant's intertemporal payoff but would not affect the marginal profitability of R&D investment. This difference between the effect on the total and the marginal payoff illustrates a point made by Katz (2021).

does not change. If acquisitions are always permitted, on the other hand, the level of innovation now is

$$\Delta^A(\bar{\mu}^A) = 1 - (1-\alpha)\kappa + \delta\left[1 - \kappa\,(1-\kappa) - (1-\alpha)\kappa^2(1-\kappa)^2\right]. \qquad (41)$$

Proposition 3 changes as follows:

Proposition 3′. *In the long run, prohibiting acquisition increases the level of innovation if and only if*

$$\kappa > \tilde{\kappa}$$

where the critical threshold $\tilde{\kappa}$ is an increasing function of α and δ.

In this variant of the model, an increase in the discount factor δ makes it less likely that a prohibition of acquisition may increase the long-run level of innovation. The intuitive reason for this is that the extra-rents created by the period-$t+1$ acquisition now depend on Δ_t, as noted above. Therefore, the size of the invention-for-buyout effect now increases if future profits count more.

Since a higher discount factor δ captures also the possibility of slower imitation due to e.g. stronger intellectual property protection, as discussed in footnote 12, Proposition 3′ suggests that in our model acquisition policy and patent policy may be interconnected: when entrants are more strongly protected against imitation, acquisition policy should become more lenient; conversely, weaker patent protection calls for a stricter acquisition policy.

Continuing to assume that antitrust authorities maximize consumer welfare, discounted consumer surplus now is

$$\sum_{n=0}^{\infty}\delta_S^n CS_{t+n} = \frac{q_{t-2} + \delta_S q_{t-1} + \delta_S^2 \sum\limits_{s=0}^{\infty}\delta_S^n \Delta_{t+n}}{1-\delta_S} + \sum_{n=0}^{\infty}\mathbf{1}_{t,NA}\left(1-\mu_{t+n}\right)^2\delta_S^n\Delta_{t+n}, \quad (42)$$

where $\mathbf{1}_{t,NA}$ is an indicator function that is 1 if acquisitions are prohibited in period t and 0 if they are permitted. Compared to (29), the additional term on the right-hand side captures the static negative effect of acquisitions on consumer surplus. This static effect arises because firms would compete more aggressively when the entrant acts as the price leader, as discussed above.

For simplicity, however, we shall henceforth focus on the limiting case $\delta_S \to 1$. This allows us to abstract from this static effect, whose relative weight in the social welfare calculation becomes negligible as δ_S approaches 1. Intuitively, the static allocative effects of acquisitions are transitory, whereas the effects on the size of innovations are permanent. As $\delta_S \to 1$, welfare comparisons therefore rest uniquely on the impact of acquisitions on the long-run level of innovation.

6.3 Acquisition cycles

For simplicity, we shall henceforth focus on a class of simple policy rules, where the policy-maker permits acquisitions as long as $\mu_t < \hat{\mu}$ and prohibits them when $\mu_t \geq \hat{\mu}$

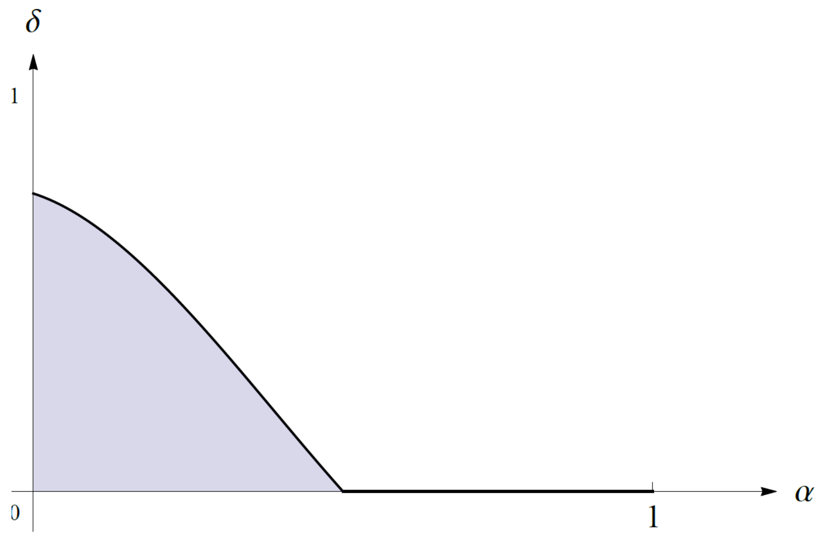

Figure 4: The region of parameter values where a state-dependent policy is optimal when κ is close to $\tilde{\kappa}$.

for some cut-off value $\hat{\mu}$.[32] Plainly, such a policy may produce cycles where the industry oscillates between periods where market dominance is low, and acquisitions are permitted, and periods where market dominance is high, and acquisitions are prohibited.

In this simple version of the model, these cycles cannot have a period longer than 2. (Later we shall consider an extension where cycles can be longer.) To see why, let us focus on the case where $\kappa(1-\kappa) < \hat{\mu} < \kappa$; otherwise, cycles are degenerate and the industry settles to states where acquisitions are either always prohibited or always permitted. Suppose that initially μ is below the threshold $\hat{\mu}$, so that acquisitions are permitted. As a result of the first acquisition, the degree of market dominance μ jumps up to κ, crossing the threshold $\hat{\mu}$. At this point, acquisitions are prohibited, and μ jumps down to $\kappa(1-\kappa)$, where a new cycle starts.

We are interested in ascertaining whether such period-2 cycles may be optimal and, if so, when. The following proposition says that a state-dependent policy may indeed be optimal when the policy-maker is almost indifferent between always permitting or always prohibiting acquisitions. This is true, in particular, when α and δ are not too large.

Proposition 6 *If α and δ are not too large, there exists a non empty neighborhood of $\tilde{\kappa}$ such that when κ lies in that interval, the optimal acquisition policy entails a period-2 cycle.*

The analytical characterization of the region where a state contingent policy is optimal in a neighborhood of $\kappa = \tilde{\kappa}$ is unmanageable, but the region can be identified by numerical methods and is depicted in Figure 4.

[32] We believe that the optimal policy belongs to this class but have not been able to prove this conjecture.

Any value of $\hat{\mu}$ in the interval $[\kappa(1-\kappa), \kappa]$ leads to the same limit cycle, so the asymptotic dynamics of the industry does not depend on the exact value of $\hat{\mu}$ within that interval. Nevertheless, the choice of $\hat{\mu}$ requires some caution. The reason for this is that when $\hat{\mu}$ is close to κ, the merged firm might choose to restrain its behavior in order to retain the possibility of engaging in further acquisitions in the future. To this end, the merged firm should ration its demand so as to prevent μ_t from crossing the threshold $\hat{\mu}$.[33] In particular, when $\mu_t = \hat{\mu} < \kappa$ the merged firm should serve only a fraction $x^M = \frac{\hat{\mu}}{\kappa}$ of its potential demand, so that $\mu_{t+1} = \kappa \frac{\hat{\mu}}{\kappa} = \hat{\mu}$. When $\hat{\mu}$ is close to κ, the profit foregone by contracting output in this way would be small, and the strategy would therefore be profitable.

One may wonder whether it may be optimal to induce such self-restraining behavior on the part of the merged firm. In fact, the answer is no.[34] Therefore, the policy-maker must be careful not to set $\hat{\mu}$ too high. But if it sets $\hat{\mu}$ just above $\kappa(1-\kappa)$, it is easy to see that the merged entity would definitely serve its entire demand. The reason for this is that to avoid crossing the threshold, the merged entity would have to serve only $1-\kappa$ consumers. If it did so, however, its profit would be lower than the aggregate profit of the incumbent and the entrant in case of no acquisition. In other words, acquisitions would no longer be profitable, and retaining the right to engage in further acquisitions would be valueless. Therefore, when $\hat{\mu}$ is just above $\kappa(1-\kappa)$ there will be no rationing, and the industry will oscillate between the high- and low-dominance states.

7 Robustness

Our results rests on the three key properties of the model illustrated in Figure 1. In addition, however, we have made several ancillary assumptions for reasons of tractability. These assumptions can be relaxed without changing our main results, but at the cost of complicating the analysis. In this section, we briefly report on the analysis of some of such extensions.

7.1 Dynamics of market dominance

In the baseline model, when acquisitions are permitted there is no transitory dynamics: μ_t jumps immediately to the steady state level $\bar{\mu}^A = \kappa$. This property follows from the assumption that the fidelization rate κ does not depend on the consumer's past purchasing history. However, one can argue that consumers who repeatedly purchase from the same firm tend to become more loyal over time.

To account for this possibility, we have analyzed an extension where firm M fidelizes a fraction κ of the free and a fraction ξ, possibly greater than κ, of the captive consumers that it serves. As a result, in the next period it will have

$$\mu_{t+1}^A = \kappa(1-\mu_t) + \xi\mu_t \tag{43}$$

[33] It cannot increase its price because of the competitive pressure from the fringe.

[34] This claim is proved in the proof of Proposition 6 in the Appendix.

captive consumers.

In the learning-cost interpretation of the model, this assumption may be justified as follows. As noted, the baseline model is obtained when a consumer is high cost with probability κ and low cost with the complementary probability $1-\kappa$, irrespective of what he did in the past. One may instead assume that these probabilities depend on the consumer's past purchasing history. In particular, it seems reasonable to assume that a consumer who purchased from the merged firm, and therefore did not have to sustain any learning cost, may have somewhat lost his ability to learn or search. As a result, in the next period he may be high-cost with a probability $\xi > \kappa$.

Under this assumption, the steady state value of μ_t^A becomes:

$$\bar{\mu}^A = \frac{\kappa}{1+\kappa-\xi}. \tag{44}$$

Since the size of the entrenchment-of-monopoly effect is given by the difference

$$\bar{\mu}^A - \bar{\mu}^{NA} = \frac{\kappa\xi}{(1+\kappa)(1+\kappa-\xi)}, \tag{45}$$

it appears that in this variant of the model, the entrenchment effect increases with ξ but may decrease with κ. Therefore, the entrenchment-of-monopoly parameter is now ξ, whereas a higher κ may actually soften the entrenchment effect.

The model can be solved following the same procedure as in the baseline case.[35] Prohibiting acquisitions now increases the long-run innovation level if

$$\xi > \frac{\alpha(1+\kappa)}{1+\delta\kappa-\alpha\delta(1+\kappa)}, \tag{46}$$

and it increases social welfare if

$$\xi > \frac{\alpha(1+\delta_S\kappa)}{(1-\alpha)\delta_S\delta\kappa+(\delta_S-\alpha\delta)}. \tag{47}$$

The intuition is, again, that the entrenchment-of-monopoly effect must be sufficiently strong. The impact of parameters α, δ and δ_S is qualitatively the same as in the baseline model. However, a higher κ now reduces the likelihood that prohibiting acquisitions may be optimal, for the reasons explained above.

When acquisitions are permitted, the degree of market dominance now converges to the steady state $\bar{\mu}^A$ gradually. As a result, under a state-dependent policy with a cut-off value $\hat{\mu}$, there can be cycles of period greater than 2. To be precise, the industry converges to a limit cycle of period ℓ when

$$\mu^H(\ell) < \hat{\mu} \leq \mu^H(\ell+1), \tag{48}$$

[35] Details for this extension may be found in the Appendix.

where

$$\mu^{H}(\ell) = \frac{\kappa \sum_{n=0}^{\ell} (\xi - \kappa)^{n}}{1 + \kappa (\xi - \kappa)^{\ell}}.^{36} \tag{49}$$

Thus, a state-dependent policy with cut-off $\hat{\mu}$ will generate cycle of period $\ell(\hat{\mu})$ implicitly given by the above conditions.[37]

Such state-dependent policies may be optimal for some parameter values when the entrant acts as a price leader.[38] The optimal cut-off $\hat{\mu}$ decreases with ξ, meaning that the optimal policy becomes more restrictive. On the other hand, the optimal cut-off increases, and hence policy becomes more lenient, as α, δ and κ increase. These results are in line with the comparative statics of the optimal state independent policy.

7.2 Product market competition

Our baseline model makes very specific assumptions about product market competition. These assumptions guarantee that the profit functions π_t^M, π_t^I and π_t^E are increasing in the size of innovations, and that π_t^E is decreasing in μ_t. These properties are essential to the model's results and cannot be dispensed with.

In addition, however, the baseline model's assumptions guarantee that the profit functions are stationary and are additively separable in Δ_t and Δ_{t-1}. These properties allow us to obtain a closed-form solution, but are not essential for our results.

Suppose, for instance, that the utility function is

$$U^i = \theta q^i - p^i, \tag{50}$$

where the willingness to pay for quality θ is uniformly distributed between 0 and 1. In this case, period-t profits would depend on the quality levels q_t, q_{t-1} and q_{t-2} and not only on the quality gaps Δ_t and Δ_{t-1}; furthermore, the profit functions would not be additively separable in the quality levels.

Without separability and stationarity, a solution of the complete dynamic model seems out of reach. However, one could consider a reduced-form, two-period version of the model. Solving a two-period model is a relatively standard exercise. The analysis shows that acquisitions spur innovation in the first period but may impede innovation in the second period, if the entrenchment of monopoly effect is sufficiently

[36] Limit cycles are degenerate, and the policy is effectively state independent, if

$$\hat{\mu} \le \mu^{H}(0) = \bar{\mu}^{NA},$$

in which case acquisitions are always prohibited, and if

$$\hat{\mu} > \lim_{\ell \to \infty} \mu^{H}(\ell) = \bar{\mu}^{A},$$

in which case acquisitions are always permitted.

[37] This relationship is illustrated in Figure 5 in the Appendix.

[38] When the incumbent acts as a price leader, on the other hand, Proposition 5 continues to hold.

strong.[39] In this reduced-form model, the first period represents, in a stylized way, the short run, the second period the long run. Therefore, we re-obtain the same qualitative results as in the fully dynamic model. The intuition is also the same.

7.3 Competition in research

The baseline model assumes that in each period, only one outsider may innovate. Therefore, the outsider will behave as a research monopolist, maximizing the difference $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) - C(\Delta_t)$.

With competition in research, the level of innovation will generally be higher. Specifically, suppose that if two or more outsiders innovate, they may then enter the market by paying a positive but arbitrarily small entry cost. In this case, as soon as there are two potential innovators, the equilibrium innovation size will be determined by a zero-profit condition

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) - C(\Delta_t) = 0. \tag{51}$$

The reason for this is as follows. Firstly, in equilibrium only one firm will make a positive R&D investment. This follows from the fact that, given the entry cost, only one outsider will enter the market, i.e., the one with the largest innovation size. Anticipating this, only one firm will invest in R&D. Secondly, if the entrant chose a value of Δ_t that entails a positive profit, i.e. $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) > C(\Delta_t)$, a competitor would have a profitable deviation by developing an innovation of greater size. Therefore, the zero-profit condition (51) must hold.

In this variant of the model, the equilibrium rate of innovation is therefore determined by the level of the acquisition price, $P = v^E$, rather than by its rate of increase with the size of the innovation. This implies that one cannot apply the two-stage solution procedure mentioned above. The model can only be solved numerically. However, the key properties of the model continue to hold, and so do our main conclusions.

7.4 Sources of market dominance

The baseline model uses a specific interpretation of the notion of market dominance, namely, the fraction of demand that is not contestable. However, our insights apply also to other interpretations.

For example, market dominance can be modeled as a cost advantage enjoyed by the incumbent, as in Stein (1997). Specifically, Stein assumes that the incumbent's cost advantage increases with the length of its incumbency. In this subsection, we adapt Stein's assumption to our framework.

To this end, we assume that all consumers are free, so the cost advantage is the only source of market power of the incumbent. Consumers' utility function is still given by (1).[40] We now assume that firms have a production cost of c. However,

[39] Details are available from the authors upon request.

[40] Under these assumptions, the pricing equilibrium no longer depends on the order of moves.

acquisition entail cost synergies. Thus, an incumbent that has been active for more than two periods, thanks to its past acquisitions, may have a cost lower than c.

To simplify the analysis, we assume that the cost falls to zero after one acquisition. This implies that in each period t, the industry can be in one of two states: state 0, where an acquisition occurred in period $t-1$, and thus the incumbent benefits from the cost synergy, and state 1, where no acquisition occurred in period $t-1$, and thus all firms have cost c. The cost c is taken to be lower than the quality differential Δ.

Generally speaking, cost synergies may have both pro- and anti-competitive effects. The pro-competitive effects are obvious; the anti-competitive effects derive, in our model, from the fact that a stronger incumbent reduces the entrant's ability to appropriate the value of its innovation. To further simplify the analysis, we focus only on the anti-competitive effect. To this end, we assume that if in state 1 the incumbent acquires the q_t entrant, its cost of producing the state of the art product q_t is still c, and only the cost of producing product q_{t-1} falls to zero. This means that the acquisition produces synergies that can be exploited only off the equilibrium path, in case the next entrant is not acquired.

Under these assumptions, in the absence of acquisitions the incumbent has no advantage of any sort. Thus, it will price at cost and make zero profit: $\pi_t^I = 0$. The entrant undercuts the incumbent and makes a profit of $\pi_t^E = \Delta_t$. Therefore, with an R&D cost function $C(\Delta_t) = \frac{1}{2}\Delta_t$, the equilibrium innovation size is simply $\Delta_t^{NA} = 1$ with monopoly in research, and $\Delta_t^{NA} = 2$ with competition in research (see the previous subsection).

Consider next the case with acquisitions. The model's solution can be obtained by calculating the values $v_t^M(0)$, $v_t^E(0)$, $V_t^I(0)$ and $V_t^I(1)$, with obvious notation.[41] These values are pinned down by the following conditions:

$$v_t^M(0) = \Delta_t + \Delta_{t-1} + \delta V_{t+1}^I(0) \quad (52)$$

$$v_t^E(0) = (1-\alpha)\left[\Delta_t - c + \delta V_{t+1}^I(1)\right] + \alpha v_t^M(0) \quad (53)$$

$$V_t^I(0) = (1-\alpha)\left\{\Delta_t + c + \delta\left[V_{t+1}^I(0) - V_{t+1}^I(1)\right]\right\} \quad (54)$$

$$(55)$$

$$V_t^I(1) = (1-\alpha)\left\{\Delta_t + \delta\left[V_{t+1}^I(0) - V_{t+1}^I(1)\right]\right\} \quad (56)$$

Condition (52) says that the value of the merged firm is its current profit $\pi_t^M = \Delta_t + \Delta_{t-1}$ plus the continuation value, which is the discounted value of being the incumbent in state 0. Condition (53) says that the entrant's payoff is a weighted combination of its disagreement payoff and the bargaining surplus, with weights given by its bargaining power α. The disagreement payoff is the sum of the current profit, which is $\pi_t^E = \Delta_t - c$ as an entrant that resists being acquired would be competing with a more efficient incumbent, plus the value of becoming a standard

[41] Clearly, $v_t^M(1)$=$v_t^M(0)$. As for $v_t^E(0)$, it may be easily calculated given the other values but is not necessary to derive the model equilibrium.

incumbent. The joint payoff in case of agreement, on the other hand, is $v_t^M(0)$. (The disagreement payoff of the current incumbent is nil, as this incumbent would be absorbed by the competitive fringe in case of disagreement.) Conditions (54) and (55) determine the value of being an incumbent in state 0 and 1, respectively. In both states, the incumbent makes no profit in case of disagreement but obtains a share $(1-\alpha)$ of the bargaining surplus. In state 0, the incumbent has a cost of zero when competing with the new entrant. Therefore, the bargaining surplus created by a period-t acquisition is the increase in current profit, $\pi_t^M - \pi_t^E = \Delta_{t-1} + c$, plus the increase in the continuation value, $\delta\left[V_{t+1}^I(0) - V_{t+1}^I(1)\right]$. In state 1, on the other hand, the bargaining surplus is lower, as the increase in current profit is just Δ_{t-1}.[42]

Solving the system (52)-(55) one gets

$$v_t^E(0) = \alpha\Delta_{t-1} + \left[1 + (1-\alpha)\delta\right]\Delta_{t-1} - (1-\alpha)(1-\delta)\left[1+(1-\alpha)\delta\right]c \qquad (57)$$

Under monopoly in research, acquisitions always increase the level of innovation as

$$\Delta_t^A = 1 + (1-\alpha)\delta > \Delta_t^{NA} = 1. \qquad (58)$$

The reason for this is that the entrenchment-of-monopoly effect, the strength of which is now captured by parameter c, causes a parallel downward shift in $v_t^E(0)$ but does not affect the marginal value of increasing the innovations size. The latter is higher when acquisitions are permitted because of the invention-for-buyot effect. As a result, the equilibrium level of innovation is unambiguously higher when acquisitions are permitted.

Under competition in research, however, things are different. The equilibrium level of innovation is higher when acquisitions are prohibited provided that

$$c > \frac{2\alpha(1-\delta)+\delta}{(1-\alpha)(1-\delta)\left[1+(1-\alpha)\delta\right]}.$$[43]

Once again, the intuition is that the entrenchment-of-monopoly effect must be sufficiently strong to outweigh the invention-for-buyout effect. We believe that this conclusion applies to many other possible sources of market dominance.

8 Conclusion

We have analyzed a tractable model of repeated innovation, where incumbents may either compete with innovative entrants or else acquire them. In the model, acquisitions have both pro-competitive and anti-competitive effects. The pro-competitive

[42] The increase in the current profit is Δ_{t-1} rather than $\Delta_{t-1} + c$ because in state 1 the incumbent will have a production cost of c rather than 0.

[43] The critical value of c is lower than Δ provided that

$$\alpha < \frac{1+\sqrt{1-\delta}-\delta^2}{\delta(1-\delta)}.$$

effects are created by the invention-for-buyout mechanism: inventors earn more by transferring their innovations to the incumbent than by exploiting them directly, so their incentive to innovate is higher when such technology transfers are permitted. The anti-competitive effects, on the other hand, are created by the entrenchment of monopoly caused by acquisitions. When acquisitions are permitted, that is to say, incumbents come to enjoy a higher degree of market dominance, and this reduces the entrants' incentive to innovate.

We have shown that the invention-for-buyout effect always prevails in the short run but can be outweighed by the entrenchment-of-monopoly effect in the long run. As a result, if policymakers are sufficiently patient and the entrenchment-of-monopoly effect is sufficiently strong, prohibiting acquisitions may be the optimal policy. In some cases, the optimal policy may be state-dependent. In other words, it may be optimal to permit acquisitions as long as market dominance is weak and prohibit them once repeated acquisitions have made it too strong.

Throughout our analysis, we have assumed that antitrust authorities focus exclusively on consumer welfare. It is sometimes claimed that this narrow focus is responsible for the leniency of antitrust policy. However, we have shown that if antitrust authorities are forward looking and consider the cumulative dynamic effects of different policy rules, the use of consumer surplus as a welfare criterion may actually lead to a restrictive policy towards acquisitions.

Our results also imply that the small size of the target firm should not provide a safe harbour against antitrust scrutiny. The critical variable which policy should focus on is not the size of the target of the acquisition, but the degree of market dominance enjoyed by the incumbent. To the extent that in innovative industries the degree of market dominance correlates to the size of the incumbents, it is that size that should matter in the antitrust assessment of acquisitions.

We conclude by mentioning three possible extensions of the model, which we leave for future work. First, our model assumes that incumbents do not innovate. This assumption is common in dynamic models of repeated innovation,[44] but it is restrictive. Relaxing this assumption is necessary to analyze the impact of acquisitions on the incumbents' innovative efforts.

Second, our model assumes an exogenous rate of arrival of ideas – one per period. Thus, it determines the size of innovations but not their frequency. One could endogenize the frequency of innovations by assuming that it depends on the entrant's prospective profit, a variable that is endogenously determined. This would add another dimension along which acquisitions could impact technological progress.

Finally, our partial-equilibrium model could be embedded into a general equilibrium model of the economy, so as to analyze the possible macroeconomic consequences of acquisition policy, and in particular its impact on the rate of growth of the economy.

[44]The fact that incumbents do not invest in R&D is often an endogenous property of the model, which follows from the fact that incumbents have a lower incentive to innovate than outsiders because of Arrow's replacement effect.

References

[1] Bryan, K. A. and Hovenkamp, E. (2020a). Antitrust limits on startup acquisitions. Review of Industrial Organization.

[2] Bryan, K. A. and Hovenkamp, E. (2020b). Startup acquisitions, error costs, and antitrust policy. The University of Chicago Law Review, 87(2):331–356.

[3] Bourreau, M., Jullien, B., and Lefouili, Y. (2021). Mergers and demand-enhancing innovation.

[4] Cabral, L. (2018). Standing on the shoulders of dwarfs: Dominant firms and innovation incentives, CEPR DP 13115.

[5] Cabral, L. (2021). Merger policy in digital industries. Information Economics and Policy, 54, 100866.

[6] Cunningham, C., Ederer, F., and Ma, S. (2021). Killer acquisitions. Journal of Political Economy, 129(3), 649-702.

[7] Fons-Rosen, C., Roldan-Blanco, P., & Schmitz, T. (2021). The Aggregate Effects of Acquisitions on Innovation and Economic Growth. Available at SSRN 3785485.

[8] Fumagalli, C., Motta, M., and Tarantino, E. (2021). Shelving or developing? The acquisition of potential competitors under financial constraints. Mimeo.

[9] Hollenbeck, B. (2020). "Horizontal mergers and innovation in concentrated industries," Quantitative Marketing and Economics, 18(1): 1-37.

[10] Ide, E., and Montero, J. P. (2020). A Theory of Must-Haves, mimeo.

[11] Kamepalli, S. K., Rajan, R., and Zingales, L. (2020). Kill zone. Stigler Center for the Study of the Economy and the State.

[12] Katz, M. L. (2021). Big Tech mergers: Innovation, competition for the market, and the acquisition of emerging competitors. Information Economics and Policy, 54, 100883.

[13] Katz, M. L. (2020). Big-Tech Mergers: Schumpeterian Competition, Innovation, and the Treatment of Emerging Competitors.

[14] Kreps, D. M., and Scheinkman, J. A. (1983). Quantity precommitment and Bertrand competition yield Cournot outcomes. The Bell Journal of Economics, 326-337.

[15] Letina, I., Schmutzler, A., and Seibel, R. (2020). Killer acquisitions and beyond: policy effects on innovation strategies. University of Zurich, Department of Economics, Working Paper, (358).

[16] Mason, R., and Weeds, H. (2013). Merger policy, entry, and entrepreneurship. European Economic Review, 57, 23-38.

[17] Mermelstein, B., Nocke, V., Satterthwaite, M. A., & Whinston, M. D. (2020). Internal versus external growth in industries with scale economies: A computational model of optimal merger policy. Journal of Political Economy, 128(1), 301-341.

[18] Motta, M., and Peitz, M. (2021). Big tech mergers. Information Economics and Policy, 54, 100868.

[19] Nocke, V., and Whinston, M. D. (2010). Dynamic merger review. Journal of Political Economy, 118(6), 1200-1251.

[20] Oertel, C., and Schmutzler, A. (2021). Challenging the incumbent: Entry in markets with captive consumers and taste heterogeneity. University of Zurich, Department of Economics, Working Paper, (386).

[21] Phillips, G. M., and Zhdanov, A. (2013). R&D and the incentives from merger and acquisition activity. The Review of Financial Studies, 26(1), 34-78.

[22] Rasmusen, E. (1988). Entry for buyout. The Journal of Industrial Economics, 281-299.

[23] Scott-Morton, F., Bouvier, P., Ezrachi, A., Jullien, B., Katz, R., Kimmelman, G., Melamed, D., and Morgenstern, J. (2019). Committee for the Study of Digital Platforms. Report. Stigler Center for the Study of the Economy and the State.

[24] Segal, I., & Whinston, M. D. (2007). Antitrust in innovative industries. American Economic Review, 97(5), 1703-1730.

[25] Shelegia, S., and Motta, M. (2021). The "kill zone": Copying, acquisition and start-ups' direction of innovation.

[26] Spulber, D. F. (2013). How do competitive pressures affect incentives to innovate when there is a market for inventions?. Journal of Political Economy, 121(6), 1007-1054.

[27] Stein, J. C. (1997). Waves of creative destruction: Firm-specific learning-by-doing and the dynamics of innovation. The Review of Economic Studies, 64(2), 265-288.

Appendix

Proof of Lemma 1. The merged entity's objective function is $\pi_t^M + \delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$, where δ is the discount factor,

$$\pi_t^M = x_t^M p_t^M$$

is the current period profit, and $\delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$ is the continuation value, i.e., the discounted value of being the incumbent in the next period with $\mu_{t+1}^A = \kappa x_t^M$ captive consumers.

To begin with, assume that the merged entity prices myopically, i.e., $\delta = 0$. Since all consumers are effectively identical in this case, there is no incentive to price discriminate by supplying different quality levels. Thus, the dominant firm will supply only the highest quality, q_t. Since the competitive fringe guarantees to all consumers an outside option of $U_t^F = q_{t-2}$, the merged entity must match this utility level:

$$U_t^M = q_t^M - p_t^M = U_t^F,$$

with a tiny price discount to break the indifference, if necessary. Therefore, $p_t^M = \Delta_t + \Delta_{t-1}$. In this myopic equilibrium, $x_t^M = 1$.

Next suppose that $\delta > 0$. It is intuitive (and we shall confirm below) that the continuation value $V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$ is a non-decreasing function of μ_{t+1}^A, which is in turn a non-decreasing function of x_t^M. Therefore, a forward-looking firm would have an incentive to further reduce the price so as to increase x_t^M if possible. But since x_t^M is already equal to 1, the myopic price remains optimal also for a forward-looking firm. ■

Proof of Lemma 2. Plainly, all firms supply the highest quality level that they control: $q_t^E = q_t$, $q_t^I = q_{t-1}$, and $q_t^F = q_{t-2}$, and the fringe prices at marginal cost (i.e., 0). The incumbent and the entrant, on the other hand, price so as to maximize their respective profits:

$$\pi_t^I = \mu_t x_t^I p_t^I$$

and

$$\Pi_t^E = (1 - \mu_t) x_t^E p_t^E + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$$

where X_t^i denotes the mass of consumers who purchase from firm i in period t, and $V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$ is the value of being the incumbent in the next period with $\mu_{t+1}^{NA} = \kappa(1 - \mu_t)x_t^E$ captive consumers. The incumbent, which is due to exit in the next period, prices myopically. A forward-looking entrant, in contrast, must keep into account the impact of its current price on the number of captive consumers that it will inherit in the second period of its life-cycle, as this affects the profits that it will earn in its capacity as the new incumbent.

To begin with, however, suppose that the entrant prices myopically ($\delta = 0$). Given the behavior of the fringe, consider the entrant's best response to p_t^I. Free consumers choose to purchase from the entrant if $U_t^E > \max\{U_t^I, U_t^F\}$, that is, if

$p^E_t < \min\left\{p^I_t + \Delta_t, \Delta_t + \Delta_{t-1}\right\}$. Therefore, the entrant's best response is

$$p^E_t(p^I_t) = \begin{cases} p^I_t + \Delta_t & \text{if } p^I_t \leq \Delta_{t-1}, \\ \\ \Delta_t + \Delta_{t-1} & \text{if } p^I_t > \Delta_{t-1}. \end{cases}$$

Next, consider the incumbent's strategy as a price leader. The incumbent makes no sales if $p^I_t > \Delta_{t-1}$. On the other hand, it anticipates that if it reduces the price below Δ_{t-1}, it would always be undercut by the entrant and would therefore serve only the captive consumers anyway. Therefore, the incumbent must price exactly at Δ_{t-1} (with a tiny discount to break the captive consumers' indifference, if necessary). By doing so, it gets a profit of $\pi^I_t = \mu_t \Delta_{t-1}$. In response, the entrant prices at $p^E_t = \Delta_t + \Delta_{t-1}$ (again with with a tiny discount if necessary) and will serve all free consumers.

If $\delta > 0$, so that the entrant is forward looking, it would have a further incentive to reduce the price to increase x^E_t if that were possible, as the continuation value V^I_{t+1} is increasing in μ_{t+1}. However, x^E_t is already equal to 1, so the myopic price remains optimal also for a forward-looking firm. ■

Proof of Corollary 1. With no acquisition, the firms' aggregate payoff is

$$(1 - \mu_t)\left(\Delta_t + \Delta_{t-1}\right) + \mu_t \Delta_{t-1} + \delta V^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t).$$

If the incumbent acquires the entrant, in contrast, the aggregate payoff becomes

$$(\Delta_t + \Delta_{t-1}) + \delta V^I_{t+1}(\mu^A_{t+1}, \Delta_t).$$

The lemma then immediately follows by comparing the above expressions, keeping in mind that V^I_{t+1} increases with the fraction of captive consumers μ_{t+1}, and that $\mu^A_{t+1} = \kappa \geq \mu^{NA}_{t+1} = \kappa(1 - \mu_t)$. ■

Proof of Lemma 3. From the optimization problem (14) it appears that the equilibrium innovation size depends only on the derivative of $v^E_t(\mu_t, \Delta_{t-1}, \Delta_t)$ with respect to Δ_t, which is the marginal profitability of increasing the size of the innovation. To calculate the derivative, let us substitute (11) into (12), obtaining

$$\begin{aligned} v^E_t(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\alpha)\left[\pi^E_t + \delta V^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t)\right] + \alpha[\pi^M_t + \delta V^I_t(\mu^A_{t+1}, \Delta_t) - \pi^I_t] \\ &= (1-\alpha)\pi^E_t + \alpha(\pi^M_t - \pi^I_t) + \delta\left[(1-\alpha)V^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t) + \alpha V^I_t(\mu^A_{t+1}, \Delta_t)\right]. \end{aligned}$$

The on-path continuation value is

$$\begin{aligned} V^I_{t+1}(\mu^A_{t+1}, \Delta_t) &= v^I_{t+1}(\mu^A_{t+1}, \Delta_t, \Delta^{\exp}_{t+1}) \\ &= v^M_{t+1}(\mu^A_{t+1}, \Delta_t, \Delta^{\exp}_{t+1}) - v^E_{t+1}(\mu^A_{t+1}, \Delta_t, \Delta^{\exp}_{t+1}) \\ &= \pi^M_{t+1} + \delta V^I_{t+2}(\mu^{A,A}_{t+2}, \Delta^{\exp}_{t+1}) + \\ &\quad -(1-\alpha)\left[\pi^E_{t+1} + \delta V^I_{t+2}(\mu^{A,A}_{t+2}, \Delta^{\exp}_{t+1})] - \alpha[v^M_{t+1}(\mu^A_{t+1}, \Delta_t, \Delta^{\exp}_{t+1}) - \pi^I_{t+1}\right], \end{aligned}$$

where $\Delta_{t+1}^{\exp}$ is the anticipated innovation size in period $t+1$ and $\mu_{t+2}^{A,A}$ is the fraction of captive consumers in period $t+2$ if the entrant is acquired both in period t and in period $t+1$. Likewise, the off-path continuation value is

$$\begin{aligned} V_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t) &= v_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t, \Delta_{t+1}^{\exp}) \\ &= v_{t+1}^{M}(\mu_{t+1}^{NA}, \Delta_t, \Delta_{t+1}^{\exp}) - v_{t+1}^{E}(\mu_{t+1}^{NA}, \Delta_t, \Delta_{t+1}^{\exp}) \\ &= \pi_{t+1}^{M} + \delta V_{t+2}^{I}(\mu_{t+2}^{NA,A}, \Delta_{t+1}^{\exp}) + \\ &\quad -(1-\alpha)\left[\pi_{t+1}^{E} + \delta V_{t+2}^{I}(\mu_{t+2}^{NA,A}, \Delta_{t+1}^{\exp})] + \alpha[v_{t+1}^{M}(\mu_{t+1}^{A}, \Delta_t, \Delta_{t+1}^{\exp}) - \pi_{t+1}^{I}\right]. \end{aligned}$$

where $\mu_{t+2}^{NA,A}$ is the fraction of captive consumers in period $t+2$ if the entrant is not acquired in period t but is acquired in period $t+1$. (It follows from the one-shot deviation principle that this is indeed the relevant value of μ.)

Next, note that all current-period profit functions

$$\begin{aligned} \pi_t^{M} &= \Delta_t + \Delta_{t-1} \\ \pi_t^{I}(\mu_t) &= \mu_t \Delta_{t-1} \\ \pi_t^{E}(\mu_t) &= (1-\mu_t)(\Delta_t + \Delta_{t-1}) \end{aligned}$$

are additively separable in Δ_{t-1} and Δ_t, and that all other terms in the expression for $v_t^{E}(\mu_t, \Delta_{t-1}, \Delta_t)$ do not depend on Δ_{t-1}. This implies that $v_t^{E}(\mu_t, \Delta_{t-1}, \Delta_t)$ depends on Δ_{t-1} in an additively separable way and that, as a result, the optimal choice of Δ_t does not depend on Δ_{t-1}.

Since a similar argument applies to all subsequent periods, it follows that $\Delta_{t+1}^{\exp}$ does not depend on Δ_t, and the same is true of $\Delta_{t+2}^{\exp}$, $\Delta_{t+3}^{\exp}$ etc. These future values depend only on μ_t. In particular

$$\begin{aligned} V_{t+1}^{I}(\mu_{t+1}^{A}, \Delta_t) &= (1-\alpha)\left[\pi_{t+1}^{M}(\mu_{t+1}^{A}) - \pi_{t+1}^{E}(\mu_{t+1}^{A})\right] + \alpha\pi_{t+1}^{I}(\mu_{t+1}^{A}) + \\ &\quad + \text{terms that depend only on } \mu_t \end{aligned}$$

and

$$\begin{aligned} V_{t+1}^{I}(\mu_{t+1}^{NA}, \Delta_t) &= (1-\alpha)\left[\pi_{t+1}^{M}(\mu_{t+1}^{NA}) - \pi_{t+1}^{E}(\mu_{t+1}^{NA})\right] + \alpha\pi_{t+1}^{I}(\mu_{t+1}^{NA}) + \\ &\quad + \text{terms that depend only on } \mu_t \end{aligned}$$

Thus, we have

$$\begin{aligned} v_t^{E}(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\alpha)\pi_t^{E}(\mu_t) + \alpha\left[\pi_t^{M} - \pi_t^{I}(\mu_t)\right]) + \\ &\quad +\delta\,\{\, (1-\alpha)\left[\pi_{t+1}^{M} - \pi_{t+1}^{E}(\mu_{t+1}^{NA})\right] + \alpha\pi_{t+1}^{I}(\mu_{t+1}^{NA}) + \\ &\quad +\alpha\left[\pi_{t+1}^{M} - \pi_{t+1}^{E}(\mu_{t+1}^{A})\right] + \alpha\pi_{t+1}^{I}(\mu_{t+1}^{A})\,\} \\ &\quad + \text{terms that depend only on } \mu_t \end{aligned}$$

Collecting all terms that depend on Δ_{t-1} and Δ_t, we finally have

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) = (1-\mu_t)\Delta_{t-1} + + (1+\delta\kappa)\left[1-(1-\alpha)\mu_t\right]\Delta_t + \text{terms that depend only on } \mu_t \tag{59}$$

From this expression, (16) follows immediately, proving the first part of the lemma.

With the equilibrium innovation size at hand, we can now determine the equilibrium value function, and hence the acquisition price. To this end, we make a guess on the functional form of the value functions and find them by the method of undetermined coefficients. Since Δ_t^A is linear in μ_t and (59) shows that the expression for the value function $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ involves the product $\mu_t \times \Delta_t$, we conjecture that the *ex ante* value functions are polynomials of degree 2:

$$V_t^E(\mu_t, \Delta_{t-1}) = (1-\mu_t)\Delta_{t-1} + \phi_0 + \phi_1\mu_t + \phi_2\mu_t^2 \tag{60}$$

$$V_t^M(\mu_t, \Delta_{t-1}) = \Delta_{t-1} + \varphi_0 + \varphi_1\mu_t + \varphi_2\mu_t^2. \tag{61}$$

Given $V_t^E(\mu_t, \Delta_{t-1})$ and $V_t^M(\mu_t, \Delta_{t-1})$, we have $V_t^I(\mu_t, \Delta_{t-1}) = V_t^M(\mu_t, \Delta_{t-1}) - V_t^E(\mu_t, \Delta_{t-1})$. We then identify the coefficients $\phi_0, \phi_1, \phi_2, \varphi_0, \varphi_1$ and φ_2 by imposing the condition that (60)-(61) must be identically satisfied. Proceeding in this way, we get

$$\phi_0 = \frac{(2-\alpha)(1+\delta\kappa)^2\left\{1+(1-\alpha)\delta\left[1-\alpha(1-\kappa)\right]\kappa\right\}}{\left[1+(1-\alpha)\delta\right]\left[1-(1-\alpha)\delta\kappa\right]\left[1+(1-\alpha)\delta\kappa^2\right]}$$

$$\phi_1 = -\frac{(1-\alpha)(1+\delta\kappa)^2\left[3-\alpha-(1-\alpha)\delta\kappa+(1-\alpha^2)\delta\kappa^2+(1-\alpha)^2\delta^2\kappa^3\right]}{\left[1-(1-\alpha)\delta\kappa\right]\left[1+(1-\alpha)\delta\kappa^2\right]}$$

$$\phi_2 = \frac{(1-\alpha)^2(1+\delta\kappa)^2}{\left[1+(1-\alpha)\delta\kappa^2\right]}$$

and

$$\varphi_0 = \frac{(1+\delta\kappa)^2\left\{1+(1-\alpha)\delta\left[1-\alpha(1-\kappa)\right]\kappa\right\}}{\left[1+(1-\alpha)\delta\right]\left[1-(1-\alpha)\delta\kappa\right]\left[1+(1-\alpha)\delta\kappa^2\right]}$$

$$\varphi_1 = -(1-\alpha)(1+\delta\kappa)^2$$

$$\varphi_2 = 0.$$

This proves the second part of the lemma.

It is simple to verify that $V_t^I(\mu_t, \Delta_{t-1})$ is increasing in μ_t – a property the was used repeatedly in the proof of Lemma 1 and 2. ■

Proof of Proposition 1. From (16) and (20) we have

$$\frac{d\Delta_t^A}{d\mu_t} = -(1-\alpha)(1+\delta\kappa) < 0$$

and

$$\frac{d\Delta_t^A}{d\mu_t} = -(1+\delta\kappa) < 0. \quad \blacksquare$$

Proof of Proposition 2. From (16) and (20) we have

$$\Delta_t^A(\mu_t) - \Delta_t^{NA}(\mu_t) = \alpha(1+\delta\kappa)\mu_t > 0. \quad \blacksquare$$

Proof of Proposition 3. Using the steady state values (22) and (23) we get

$$\Delta_t^{NA}(\bar{\mu}^{NA}) - \Delta_t^A(\bar{\mu}^A) = (1+\delta\kappa)\left[1-(1-\alpha)\kappa - \frac{1}{(1+\kappa)}\right],$$

whence the result follows immediately. $\blacksquare$

Proof of Proposition 4. Simple algebra shows that

$$W_t^{NA} - W_t^A = (1+\delta\kappa)\,\frac{\left[\kappa\delta_S - \alpha(1+\kappa\delta_S)\right]\left[\delta_S(\kappa-\mu_t)+\mu_t\right]}{(1-\delta_S)^2(1+\kappa\delta_S)},$$

whence the result follows immediately. $\blacksquare$

Proof of Lemma 5. Proceeding as in the proof of Lemma 3, it is easy to verify that if acquisitions are permitted in period t, the derivative of the acquisition price $P_t = v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ with respect to Δ_t may depend only on whether acquisitions are prohibited or permitted in period $t+1$. We already know that if acquisitions are permitted in period $t+1$, we have

$$\begin{aligned} v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) \;&=\; (1-\mu_t)\Delta_{t-1} + \\ &\quad + (1+\delta\kappa)\left[(1-\mu_t)+\alpha\mu_t\right]\Delta_t + \text{ other terms} \end{aligned}$$

where the other terms depend on expected values $\Delta_{t+1}^{\exp}, \Delta_{t+2}^{\exp}, \Delta_{t+3}^{\exp}$ etc., and thus depend only on μ_t. If, on the other hand, acquisitions are prohibited in period $t+1$, the period-t acquisition price is

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) = (1-\alpha)\left[\pi_t^E + \delta\pi_{t+1}^I(\mu_{t+1}^{NA})\right] + \alpha\left[\pi_t^M + \delta\pi_t^I(\mu_{t+1}^A) - \pi_t^I\right].$$

That is, the continuation value reduces to the profits that the merged firm will reap in its capacity as the next-period incumbent. Simple calculations yield

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) = (1-\mu_t)\Delta_{t-1} + (1+\delta\kappa)\left[(1-\mu_t)+\alpha\mu_t\right]\Delta_t.$$

This immediately implies that if acquisitions are permitted in period t, $\Delta_t^A(\mu_t)$ does not depend on whether acquisitions are permitted or prohibited in the subsequent periods.

Next suppose that acquisitions are prohibited in period t. If acquisitions are prohibited also in period $t+1$, the entrant's profit is

$$\begin{aligned}\Pi_t^{E,NA} &= \pi_t^E(\mu_t) + \delta\pi_{t+1}^I(\mu_{t+1}^{NA}) \\ &= (1-\mu_t)(\Delta_{t-1}+\Delta_t) + \delta\kappa(1-\mu_t)\Delta_t,\end{aligned}$$

whence (20) follows. If, on the other hand, acquisitions are permitted in period $t+1$, the entrant's profit is

$$\begin{aligned}\Pi_t^{E,NA} &= \pi_t^E(\mu_t) + \delta\left[v_{t+1}^M(\mu_{t+1}^{NA},\Delta_t,\Delta_{t+1}^{\exp}) - v_{t+1}^E(\mu_{t+1}^A,\Delta_t,\Delta_{t+1}^{\exp})\right] \\ &= \pi_t^E(\mu_t) + \delta\left\{\pi_{t+1}^I(\mu_{t+1}^{NA}) + (1-\alpha)\left[\pi_{t+1}^M - \pi_{t+1}^I(\mu_{t+1}^{NA}) - \pi_{t+1}^E(\mu_{t+1}^{NA})\right]\right\} + \\ &\quad + \text{other terms that depend only on } \mu_t.\end{aligned}$$

Plugging the equilibrium profits into this formula, one obtains

$$\begin{aligned}\Pi_t^{E,NA} &= (1-\mu_t)(\Delta_{t-1}+\Delta_t) + \delta\left[\alpha\mu_{t+1}^{NA}\Delta_t + (1-\alpha)\mu_{t+1}^{NA}\left(\Delta_t + \Delta_{t+1}^{\exp}\right)\right] \\ &= (1-\mu_t)\Delta_{t-1} + \left[(1-\mu_t) + \delta\mu_{t+1}^{NA}\right]\Delta_t + \text{other terms} \\ &= (1-\mu_t)\Delta_{t-1} + \\ &\quad +(1+\delta\kappa)(1-\mu_t)\Delta_t + \\ &\quad +\text{other terms that depend only on } \mu_t\end{aligned}$$

so the optimal level of innovation is still given by (20). ■

Proof of Proposition 5. We verify that the non state-contingent policy is optimal by applying the single-deviation principle. To begin with, suppose that condition (36) holds, so that always prohibiting acquisitions is the optimal non-contingent policy. Consider an alternative policy where acquisitions are permitted in period t and then are always prohibited from period $t+1$ onwards. With this policy, social welfare is:

$$W_t^{1,0}(\mu_t) = \Delta_t^A(\mu_t) + \delta_S W_t^{NA}(\mu_{t+1}^A).$$

Using (16) and (35), this rewrites as:

$$W_t^{1,0}(\mu_t) = (1+\delta\kappa)\left[(1-\mu_t)+\alpha\mu_t\right] + \frac{\delta_S(1+\delta\kappa)\left[1-(1-\delta_S)\kappa\right]}{(1-\delta_S)(1+\delta_S\kappa)}.$$

Simple algebra shows that $W_t^{1,0}(\mu_t) \leq W_t^{NA}(\mu_t)$ when (36) holds.

Next suppose that condition (36) fails, so that always permitting acquisitions is the optimal non-contingent policy. Consider an alternative policy where acquisitions are prohibited in period t and then are always permitted from period $t+1$ onwards. With this policy, social welfare is:

$$W_t^{0,1}(\mu_t) = \Delta_t^{NA}(\mu_t) + \delta_S W_t^A(\mu_{t+1}^{NA}).$$

Using (20) and (34), this rewrites as:

$$W_t^{0,1}(\mu_t) = \frac{1-\delta_S\kappa\,(1-\alpha)+\delta\kappa\,[1-(1-\alpha)\kappa\delta_S]}{(1-\delta_S)\,(1+\delta_S\kappa)}\,[1-(1-\delta_S)\mu_t]\,.$$

Simple algebra shows that $W_t^{0,1}(\mu_t) \leq W_t^A(\mu_t)$ when (36) fails.

Thus, the optimal non-contingent policy always survives one-shot deviations. This confirms that such policy remains optimal even when state-contingent policies are feasible. ■

Proof of Lemma 6. As in the proof of Lemma 2, let us begin by assuming that the entrant prices myopically ($\delta = 0$). Given the behavior of the fringe, consider the incumbent's best response to p_t^E. Captive consumers choose to purchase from the incumbent if $U_t^I > U_t^F$, that is if $p_t^I < \Delta_{t-1}$. Free consumers choose to purchase from the incumbent if $U_t^I > \max\left\{U_t^E, U_t^F\right\}$, that is, if $p_t^I < \min\left\{p_t^E - \Delta_t, \Delta_{t-1}\right\}$. Therefore, the incumbent must always undercut the fringe by setting a price no higher than $p_t^I = \Delta_{t-1}$, for otherwise it would make no sales. (Likewise, the entrant must always undercut the fringe to make positive sales, so we can restrict attention to the case $p_t^E \leq \Delta_t + \Delta_{t-1}$ with no loss of generality.) If the incumbent prices exactly at Δ_{t-1} (with a tiny discount to break the captive consumers' indifference, if necessary), it gets a profit of $\pi_t^I = \mu_t\Delta_{t-1}$. If it further reduces the price to undercut not only the fringe but also the entrant, pricing at $p_t^I = p_t^E - \Delta_t$, it earns $\pi_t^I = p_t^E - \Delta_t$. Therefore, the incumbent's best response is

$$p_t^I(p_t^E) = \begin{cases} \Delta_{t-1} & \text{if } p_t^E \leq \Delta_t + \mu_t\Delta_{t-1}, \\ \\ p_t^E - \Delta_t & \text{if } p_t^E > \Delta_t + \mu_t\Delta_{t-1}. \end{cases}$$

Next, consider the entrant's strategy as a price leader. As noted, the entrant makes no sales if $p_t^E > \Delta_t + \Delta_{t-1}$. It also makes no sales if $p_t^E > \Delta_t + \mu_t\Delta_{t-1}$, as in this case the incumbent would undercut it. Therefore, the entrant's equilibrium price is $p_t^E(\mu_t) = \Delta_t + \mu_t\Delta_{t-1}$. At this price, the entrant serves all free consumers: $x_t^E = 1$.

Finally, suppose that $\delta > 0$ so that the entrant is forward looking. Since the continuation value V_{t+1}^I is increasing in μ_{t+1}, the entrant would have a further incentive to reduce the price to increase x_t^E if that were possible. But since x_t^E is already equal to 1, the myopic price remains optimal also for a forward-looking firm. ■

Derivation of equation (40). Proceeding as in the proof of Lemma 3, but using the current-period profit functions

$$\begin{aligned} \pi_t^M &= \Delta_t + \Delta_{t-1} \\ \pi_t^I &= \mu_t\Delta_{t-1} \\ \pi_t^E &= (1-\mu_t)(\Delta_t + \mu_t\Delta_{t-1}) \end{aligned}$$

one obtains

$$
\begin{aligned}
v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) \;\; = \;\; & (1-\mu_t)\left[\alpha + (1-\alpha)\mu_t\right]\Delta_{t-1} + \\
& + \{\ \alpha + \delta + (1-\alpha)(1-\mu_t) - (1-\alpha)\delta\left[1-\kappa(1-\mu_t)\right]\left[\alpha+\kappa(1-\alpha)(1-\mu_t)\right] + \\
& -\alpha\delta\,(1-\kappa)\left[\alpha+\kappa(1-\alpha)(1-2\mu_t)\right]\ \}\,\Delta_t \\
& + \text{terms that depend only on } \mu_t
\end{aligned}
$$

From this expression, (40) follows immediately.

The equilibrium value functions v_t^i and V_t^i can be calculated by the method of undetermined coefficients, as in the proof of Lemma 3. They are now polynomials of degree 4 in μ_t; the solution can be found in a mathematical appendix available from the authors upon request. ■

Proof of Proposition 3′. The critical value $\tilde{\kappa}$ is implicitly defined by the condition

$$
\mathcal{H}(\alpha, \delta, \kappa) \equiv 1 - (1-\alpha)\kappa + \delta\left[1 - \kappa\,(1-\kappa) - (1-\alpha)\kappa^2(1-\kappa)^2\right] - \frac{1+\delta\kappa}{1+\kappa} = 0
$$

Simple algebra shows that

$$
\begin{aligned}
\left.\frac{\partial \mathcal{H}}{\partial \kappa}\right|_{\kappa=\tilde{\kappa}} &< 0 \\
\left.\frac{\partial \mathcal{H}}{\partial \alpha}\right|_{\kappa=\tilde{\kappa}} &> 0 \\
\left.\frac{\partial \mathcal{H}}{\partial \alpha}\right|_{\kappa=\tilde{\kappa}} &> 0.
\end{aligned}
$$

The result then follows by implicit differentiations. ■

Proof of Proposition 6. Along a period-2 cycle, μ_t oscillates between $\mu_t = \kappa$, where acquisitions are prohibited, and $\mu_t = \kappa\,(1-\kappa)$, where acquisitions are permitted. By Lemma 5, the level of innovation then oscillates between $\Delta_t^{NA}(\kappa)$ and $\Delta_t^{A}(\kappa\,(1-\kappa))$. On average, the level of innovation is

$$
\begin{aligned}
\Delta^{SC} \;\; = \;\; & \frac{2+\delta\,(1-\alpha)}{2} - \frac{2-\alpha(1+2\delta)}{2}\kappa + \frac{1-\alpha+\delta\left[1-2\alpha\,(2-\alpha)\right]}{2}\kappa^2 + \\
& -\frac{3-(7-4\alpha)\,\alpha\delta}{2}\kappa^3 + \delta\,(1-\alpha)^2\,\frac{3-\kappa+\kappa^2}{2}\kappa^4
\end{aligned}
$$

Next, note that for $\delta = 0$ and $\kappa = \tilde{\kappa} = \frac{\alpha}{1-\alpha}$, we have $\Delta^{NA}\left(\bar{\mu}^{NA}\right) = \Delta^{A}\left(\bar{\mu}^{A}\right) = \Delta^{SC}$. It may then be verified that if one increases δ and at the same time increases κ so that $\kappa = \tilde{\kappa}$, Δ^{SC} increases more rapidly than either $\Delta^{NA}\left(\bar{\mu}^{NA}\right)$ or $\Delta^{A}\left(\bar{\mu}^{A}\right)$ (which in fact, by construction, increase at the same rate). This implies that there exists a neighborhood of $\delta = 0$ where the state contingent policy is optimal. This however requires that $\tilde{\kappa} < 1$, and hence that α and δ are not too large.

Suppose now that the merged firm rations demand so as to keep its degree of market dominance below the threshold $\hat{\mu}$. To this end, it must set an output level of $x^M = \frac{\hat{\mu}}{\kappa}$, earning a profit of $\pi^M = \frac{\hat{\mu}}{\kappa}(\Delta_t + \Delta_{t-1})$. Proceeding as in the proof of Lemma 3, one finds that the corresponding level of innovation is

$$\begin{aligned}\Delta^R &= (1-\alpha)\{1-\delta\kappa[1-\alpha(2-\kappa)]\} + \\ &- \left[1-\alpha+\alpha\delta(1-2\alpha) - \frac{\alpha+\delta(1-\alpha)}{\kappa} - (1-\alpha)(1-2\alpha)\delta\kappa + 2(1-\alpha)^2\delta\kappa^2\right]\hat{\mu} + \\ &+ (1-\alpha)\left[\alpha+(1-\alpha)\kappa^2\right]\delta\hat{\mu}^2.\end{aligned}$$

Tedious algebra shows that Δ^R is always lower than $\max\left[\Delta^{NA}\left(\bar{\mu}^{NA}\right), \Delta^A\left(\bar{\mu}^A\right)\right]$. This implies that a state contingent policy that induces the merged firm to ration its demand is never optimal. ■

Omitted derivations for section 7.1. Proceeding as in the proof of Lemma 3, one finds

$$\begin{aligned}v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\mu_t)\Delta_{t-1} + \\ &+ [(1+\delta\kappa)(1-\mu_t) + \alpha\mu_t(1+\delta\xi)]\Delta_t \\ &+ \text{terms that depend only on } \mu_t\end{aligned}$$

From this expression, one gets

$$\Delta_t^A(\mu_t) = (1+\delta\kappa)(1-\mu_t) + \alpha(1+\delta\xi)\mu_t.$$

In a steady state where $\bar{\mu}^A = \frac{\kappa}{1+\kappa-\xi}$, the level of innovation is

$$\Delta^A(\bar{\mu}^A) = \frac{(1+\delta\kappa)(1-\xi) + \alpha\kappa(1+\delta\xi)}{1+\kappa-\xi}.$$

Contrasting it with (25), condition (46) in the main text follows.

Next, consider the social welfare (33). When acquisitions are permitted, the share of captive consumers evolves over time as follows

$$\mu_t^A = \frac{\kappa}{1+\kappa-\xi} + \left(\mu_0 - \frac{\kappa}{1+\kappa-\xi}\right)(\xi-\kappa)^t,$$

and thus the level of innovation is

$$\begin{aligned}\Delta_t^A &= \frac{1+\kappa[\alpha(1+\delta\xi)+\delta(1-\xi)]-\xi}{1+\kappa-\xi} + \\ &- [1+\delta\kappa-\alpha(1+\delta\xi)]\left(\mu_0 - \frac{\kappa}{1+\kappa-\xi}\right)(\xi-\kappa)^t.\end{aligned}$$

It follows that social welfare becomes:

$$W_t^A(\mu_t) = \frac{(1+\delta\kappa)(1-\xi\delta_S) - \alpha\kappa\delta_S(1+\delta\xi)}{(1-\delta_S)[1-\delta_S(\xi-\kappa)]} - \frac{1-\alpha+\delta\kappa-\alpha\delta\xi}{1-\delta_S(\xi-\kappa)}\mu_t.$$

Comparing $W_t^{NA}(\mu_t)$ and $W_t^A(\mu_t)$, condition (47) in the main text follows immediately.

We conclude the analysis of the extension in Section 7.1 by looking at the case of state-dependent policy. First of all, consider a generic cycle of length ℓ; that is, suppose that there are $\ell - 1$ consecutive periods in which acquisitions are permitted, followed by one period in which they are prohibited. Focusing only on the beginning and the end of such a cycle, the dynamics of μ_t can be described by the following difference equation

$$\mu_{t+\ell} = \kappa - \kappa^2 \sum_{n=0}^{\ell-2} (\xi-\kappa)^n - \kappa(\xi-\kappa)^{\ell-1}\mu_t.$$

The characteristic root of this difference equation is, in absolute value,

$$\left| -\kappa^{\frac{1}{\ell}}(\xi-\kappa) \right| < 1.$$

This shows that all cycles of length ℓ converge to the limit cycle of length ℓ, irrespective of the initial condition.

Next consider a generic cut-off $\hat{\mu} \in \left[\bar{\mu}^{NA}, \bar{\mu}^A\right]$. We show that starting from an arbitrary μ_0, after at most one cycle, all subsequent cycles have length $\ell(\hat{\mu})$. With no loss of generality, assume that $\mu_0 < \hat{\mu}$. (If $\mu_0 \geq \hat{\mu}$, it suffices to apply the argument below starting from μ_1, which will then be lower than $\hat{\mu}$.) The value of μ_t for which acquisitions are first prohibited may then range from $\hat{\mu}$ to $\kappa(1-\hat{\mu}) + \xi\hat{\mu}$. Therefore, the value of μ_t at the beginning of the next acquisition cycle may range from $\kappa(1-\kappa) - \kappa(\xi-\kappa)\hat{\mu}$ to $\kappa(1-\hat{\mu})$. Tedious but simple algebra shows that for all starting points in this interval, μ_t will cross the threshold $\hat{\mu}$ in exactly $\ell(\hat{\mu})$ periods.

Combining the fact that under a cut-off policy rule with cut-off $\hat{\mu}$ all cycles will eventually have length $\ell(\hat{\mu})$, and that all cycles of that length converge to the limit cycle of length $\ell(\hat{\mu})$, it follows that the system converges to the limit cycle of length $\ell(\hat{\mu})$. The function $\ell(\hat{\mu})$ is stepwise increasing, as shown in Figure 5.

Next, we have analyzed the optimal length of the limit cycle, ℓ^*, by means of numerical calculations. It turns out that ℓ^* is (i) a stepwise decreasing function of the entrenchment-of-monopoly parameter ξ, (ii) a stepwise increasing function of the entrant's bargaining power α, (iii) a stepwise decreasing function of the private discount factor δ, and (iv) a stepwise decreasing function of the fidelization rate κ. Since ℓ^* may be interpreted as an index of the leniency of acquisition policy, these results confirm the comparative statics of the case of state-dependent policy. Figure 6 illustrates the effect of the entrenchment of monopoly parameter ξ.

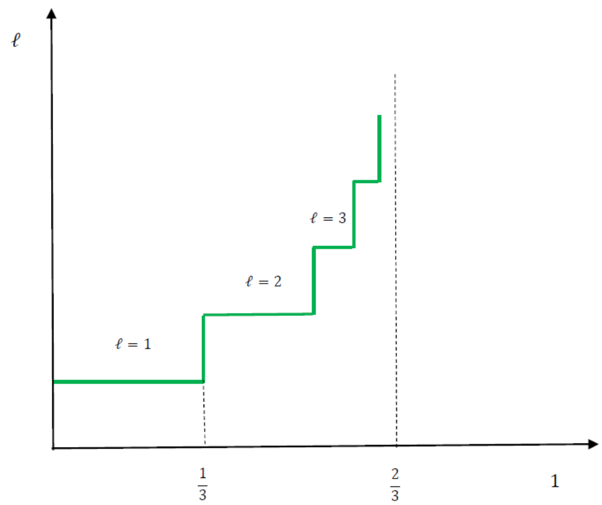

Figure 5: The relationship between the length of the limit cycles ℓ and the cut-off $\hat{\mu}$. The figure has been drawn for $\kappa = \frac{1}{2}$ and $\xi = \frac{3}{4}$. With these values, $\bar{\mu}^{NA} = \frac{1}{3}$ and $\bar{\mu}^{A} = \frac{2}{3}$, and the length of each step after the first one is two thirds of the remaining distance to $\bar{\mu}^{A}$. For example, the threshold where ℓ jumps from 1 to 2 is $\hat{\mu} = \frac{5}{9}$, the threshold where ℓ jumps from 2 to 3 is $\hat{\mu} = \frac{17}{27}$, and so on.

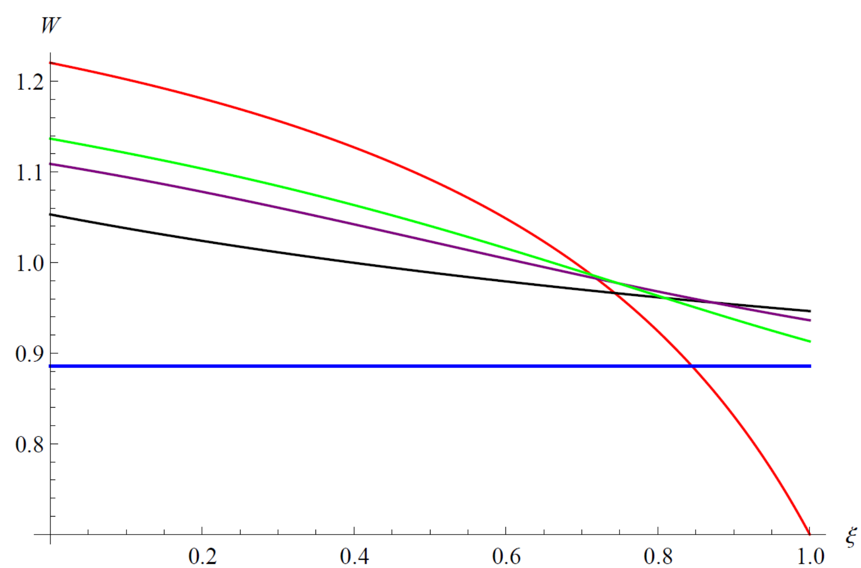

Figure 6: Social welfare under different lengths of the limit cycles. The degenerate cases of non-contingent policies are depicted in blue ($\ell = \infty$) and red ($\ell = 0$). The other curves represent the case of non-degenerate limit cycles for $\ell = 1$ (black), $\ell = 2$ (purple) and $\ell = 3$ (green). The figure has been drawn for $\alpha = 0.1$, $\delta = 0.6$ and $\kappa = 0.4$.