Horizontal Mergers and Innovation in Concentrated Industries

Brett Hollenbeck*

May 1, 2017

Abstract

An open question in antitrust economics is whether allowing rival firms to merge increases or decreases incentives to invest and innovate? I examine this in a dynamic oligopoly model with endogenous investment, entry, exit and horizontal mergers. Firms produce differentiated goods and may merge with rival firms to gain market power and increase the quality of their product. I extend previous work on dynamic mergers by allowing for products differentiated on quality with competition in prices and an endogenous long run rate of innovation. In equilibrium, horizontal mergers are mostly harmful to consumers in the short run, but the prospect of a buyout creates a powerful incentive for firms to enter the industry and invest to make themselves an attractive merger partner. The result is significantly higher total innovation with mergers than without and significantly higher long-run consumer welfare as well. This result also helps shed light on the larger question of the relationship between the competition and innovation. Further results show that long run welfare is increased most by mergers in industries with low consumer switching costs or brand loyalty, as well as in industries where large and rapid innovation is possible.

*UCLA Anderson School of Management, contact brett.hollenbeck@gmail.com.

1 Introduction

In a concentrated industry, does allowing rival firms to merge increase or decrease investment and innovation? Antitrust authorities increasingly deal with industries characterized by high levels of investment and where innovations cause rapid changes in firm market share and product quality.[1] For these industries, the effects of a merger on dynamic considerations such as investment, entry and exit are large relative to the standard considerations of market power and price increases when determining the merger's likely effect on consumer welfare. The relationship between industry concentration and innovation is itself complex and non-monotonic.[2] Furthermore, few things matter more for consumer welfare than the long run rate of innovation, and its determinant is a central question in economics. Nevertheless, the relationship between horizontal mergers and innovation remains poorly understood.

Mergers between rival firms may affect investment incentives in several ways. Investment typically imposes a negative externality on the industry, as some portion of the gains from a successful innovation come from stealing business from rival firms. By merging, firms will internalize this effect and reduce their investment accordingly.[3] Firms may also buy out a smaller rival to acquire its new innovation, and so use the merger as a substitute for investing in the new technology itself.[4] On the other hand, the prospect of being bought out may also encourage entry into the market by new firms, encouraging development of new products and technologies. A merger may also increase the new firm's ability to innovate via economies of scale or complementarities between the two firms' R&D capabilities. Because the relationship between mergers and innovation depends in a complex way on both pre and post-merger market structure, to determine the interplay between and relative importance of these effects requires modeling industry dynamics such as entry and investment along with endogenous mergers.

Empirical work on this question is limited and faces several challenges.[5] Instead, this

[1] Katz and Shelanski (2006) and Gilbert (2006) discuss the increasing importance of innovation in merger analysis. The 2010 Horizontal Merger Guidelines introduced a section dealing with innovation, and in the years 2013-2015, the Department of Justice challenged mergers in part due to concerns about innovation incentives in online platforms, online display advertising, chemicals, computer circuits, aircraft components and beer. For more see the Annual Report on Competition Policy Developments in the United States for those years, jointly produced by the DOJ and FTC.

[2] A long literature in economics considers this topic. Notably, Aghion et al. (2005) has shown a inverted-U shaped relationship between industry concentration and innovation.

[3] This is related to the notion of innovation markets developed by Gilbert and Sunshine (1995).

[4] This has been an issue in recent pharmaceutical mergers, such as the Merck purchase of Idenix in 2014.

[5] Mergers are frequently a response to a larger shock to technology, preferences or regulations that would cause firms or the entire industry to expand or contract in the absence of a merger Harford (2005), for in-

paper contributes to our understanding of the impact of mergers on R&D investment by modeling the dynamics of a concentrated industry with fully endogenous entry, exit, quality investment and horizontal mergers. While a model combining these elements presents many challenges, it is necessary in order to consider questions regarding innovation, which is inherently dynamic.[6] Future investment, exit, and entry, along with the potential for other future mergers, can have a dramatic effect on the welfare implications and profitability of today's potential merger, and despite the complexity it entails, a dynamic model containing each of these features is necessary to consider this question.

Despite the importance of these effects, the dynamic effects of horizontal mergers have rarely been studied in settings where mergers occur endogenously.[7] Endogeneity of mergers is crucial to understand their dynamic effects because current decisions regarding entry, investment and mergers should be effected by the set of likely future mergers. That mergers occur in waves within industries is well documented. Despite this, very few studies have been done where mergers arise endogenously in a dynamic context. Pesendorfer (2005), Gowrisankaran (1999), and **?** have explored this. Pesendorfer (2005) derives theoretical predictions from a Cournot model with entry, exit and mergers and finds that the standard Cournot result is overturned if firms expect the possibility of mergers in the future. Gowrisankaran (1999) presents a framework for studying mergers in a dynamic oligopoly with capacity-constrained, homogenous goods producers. In each period, firms make sequential bids to merge with smaller firms. Because firms only merge to increase market

stance, shows that industry level merger waves are primarily driven by "economic, regulatory or technological shocks." In addition, mergers strongly cluster over time and industries, and both the decision to merge and the decision to invest have strong strategic components that depend on rivals' actions. These factors make finding causal evidence from pre and post-merger R&D levels very difficult. Even if a plausible instrument could be found, it is unlikely the effects of mergers induced by this instrument would be generalizable to other settings. See Nevo and Winston (2010) for more on this point. Recent attempts have been made to estimate structural models of merger dynamics, including Jeziorski (2014) in the radio industry, Igami and Uetake (2015) the hard disk drive industry, Nishida and Yang (2014) in retail, and Stahl (2010) in broadcast television.

[6] Gowrisankaran (1999) discusses the challenges of solving a dynamic model with endogenous mergers and presents a lengthy discussion of the flaws inherent in static models and models of exogenously imposed mergers.

[7] Cheong and Judd (2006) and Chen (2009) present numerical results showing that the welfare conclusions of static models can be overturned in the long run. Cheong and Judd (2006) show that even in Cournot type industries mergers may be profitable in present value terms. In two recent companion papers Marshall and Parra (2016a) and Marshall and Parra (2016b) modify patent race models to study the relationship between market structure, mergers, and innovation. In these papers a set of firms invest in R&D to become the market leaded and they explore the implications of reducing the number of firms and combing their innovation capacities, finding in Marshall and Parra (2016a) that mergers that increase R&D capabilities are often welfare increasing and showing in Marshall and Parra (2016b) that a merger can increase or decrease consumer welfare depending on fairly straightforward and often observable features of product market competition. Each of these consider only exogenous mergers.

share, investment always declines as firms internalize the investment externality and the impact on total welfare is ambiguous.

Mermelstein et al. (2014) develop a investment technology that is merger neutral and use it to study the relationship between capital investment and mergers in a two firm, homogenous good model where investment lowers production costs. They find mergers decrease long term consumer surplus as well as incumbent profits, but that antitrust policy can increase aggregate value. Along with the different modeling choices described above and the different focus on innovation, I find very different results on the relationship between mergers and consumer welfare. Despite both involving investment, building physical capital and investing in innovations to improve product quality are quite distinct and interact with mergers distinctly as well.

Following this work by Mermelstein et al. (2014), I adopt a similar investment technology, described in more detail below, but study a model with a number of features that are better suited to the study of industries where innovation is of first order concern such as software, pharmaceuticals, semiconductors and other tech hardware, telecoms, etc. First, instead of firms producing homogenous goods and competing by setting quantity, firms produce differentiated goods and compete in prices. Second, I allow for more than 2 firms and hence mergers that do not result in monopoly. Third, the role of mergers is not to reduce production costs but to improve overall product quality. Finally, I alter the investment technology to allow for an endogenous and variable long run rate of innovation.

I proceed by embedding an endogenous merger stage game into an Ericson-Pakes style dynamic oligopoly model where firms produce differentiated goods and compete in prices. They engage in entry, exit, and invest in future product quality. In each period firms may enter merger negotiations with one another. At this time they observe a private signal of "synergy" value, reflecting the complementarities between their products. If the firms merge, in the following period they will produce a new, higher-quality product. Embedded in the model is an antitrust authority, who evaluates mergers and solves for the optimal antitrust policy. This model represents a substantial increase in generality over previous attempts to model the dynamic effects of mergers, which have been limited to exogenous mergers or mergers to monopoly between two firms, a practice rarely allowed in reality, while also extending the setting to differentiated goods and an endogenous rate of innovation.

I take advantage of two recent methodological advances necessary to approach this topic. First, I follow Goettler and Gordon (2014) in modifying the Pakes and McGuire (1994) framework to allow for a long run rate of innovation that is endogenous. In addition, Mermelstein et al. (2014) show that in the Pakes and McGuire (1994) framework, the industry-wide investment opportunity set is reduced by mergers, which necessarily reduce the number of firms. I adapt their investment framework that allows for rich and flexible investment patterns that are merger neutral and allow entrants to endogenously choose the quality of their product at the time of entry. I then solve numerically for Markov Perfect Equilibrium for several types of counterfactuals. I use the stochastic algorithm of Pakes and McGuire (2001) to which I contribute a modification that substantially improves stability and performance.

In a baseline with no mergers allowed, the industry exists primarily in a state of duopoly with one firm near the technological frontier and another offering an inferior product and investing little. Entry and net innovation are rare. When mergers are allowed, they frequently arise, and there is substantially more entry. This includes firms who enter in states where their static profits are negative, because the prospect of a buyout is so lucrative. New entrants occasionally generate rapid and large innovations that make them the leading firm. As a result the rate of innovation is dramatically higher than in a setting without mergers. So in a setting where the vast majority of mergers increase prices and harm consumer welfare in the short term, the ultimate long term effect of allowing these anti-competitive mergers is much higher average consumer welfare as consumers are made better off by the increased entry and innovation.

Next I explore the parameter space in several dimensions, and do so to consider what industry characteristics make this result more or less likely. Specifically, I examine the role of "contestability" as described by Shapiro (2010), meaning the degree to which firms who successfully innovate can capture higher market share as a result.[8] This could be thought of as representing consumer switching costs, brand loyalty, or the ease of distribution. I show that when contestability is high mergers (including anti-competitive mergers) increase long-run consumer surplus by increasing innovation. When contestability is low, however, this is no longer true. While this paper's main result effectively argues for leniency in horizontal merger review, this secondary result suggests greater attention should be paid by antitrust authorities to actions taken by firms that decrease contestability. These include long-term

[8]Specifically, Shapiro (2010) uses the definition "The prospect of gaining or protecting profitable sales by providing greater value to customers."

contracts and bundling requirements that increase switching costs or other practices that make it difficult for innovative products to be accessed through dominant platforms.

I also consider a related industry characteristic that I call "disruptability," reflecting how capable firms are of generating large innovations quickly, including entering firms. While contestability is a feature of demand, disruptability is a feature of the industry's investment and innovation technology. I show that in highly "disruptable" industries, mergers generate long-run increases in consumer surplus. In less disruptable industries the reverse is true and consumer welfare is lowered by mergers in the long run. In practice, this feature can be observed by policymakers by considering relative rates of patent filings, product life-cycles, and the underlying technology.

This paper also helps answer the larger question of what is the relationship between competition and innovation?[9] Theoretical work on this question dates back to Schumpeter (1942) and Arrow (1962) who gave dueling arguments for why innovation should be higher under monopolistic and competitive industries. Recently, Aghion et al. (2005), suggests an inverted-U shaped relationship with low rates of innovation in highly competitive and monopolistic settings, and high innovation in intermediate settings. Goettler and Gordon (2014) find a similar result. Segal and Whinston (2007) contribute to this literature by showing in a general model that antitrust policy that protects entrant profits leads to higher innovation. They demonstrate this result for competition policy related to exclusive contracts and network externalities, I show a result for horizontal mergers that is contradictory in the sense that stricter antitrust policy slows innovation, but via a complementary mechanism. In this case, while mergers are anti-competitive, they increase the value of entry by allowing for potentially lucrative buyouts of small firms.

The rest of this paper will be organized as follows, section 2 describes the model environment and lays out the nature of mergers and investment, section 3 describes the nature of equilibrium and method of computation, and section 4 presents results on the static and long-term effects of mergers for a range of parameter values.

[9] Shapiro (2010) calls this "arguably the most important question in the field of industrial organization."

2 Model

Industry dynamics are based on the Ericson and Pakes (1995) framework, in which a set of firms invest, enter, and exit endogenously in discrete time with an infinite horizon. This model and its properties and many applications are reviewed at length in Doraszelski and Pakes (2008), and will be given a shorter treatment here with more emphasis on the model's novel elements. In the model, a set of constant marginal cost firms produce differentiated goods and compete in prices. The goods differ with respect to their level of quality and firms can invest in future product quality using a stochastic R&D technology that combines features of Mermelstein et al. (2014) and Goettler and Gordon (2014). Importantly, the total set of possible investment is not necessarily reduced by a merger and the long run rate of innovation is endogenous. Here and throughout, innovation refers to an increase in technological frontier or an increase in product quality beyond what has been available before. Each period, firms are allowed to enter merger negotiations with any other firm following a random sequence. Firms will attempt to merge if the net gain to the acquiring firm is greater than the reservation value of the acquired firm. Mergers are quality-increasing, in that the merger results in a new, higher quality product.

2.1 Incumbent Firms

Product Market Competition At any given time there are $n \leq \overline{n}$ firms active in the market, each producing a good of quality $\omega_i \in \{\omega_1, ..., \overline{\omega}_{max}\}$. This "quality" can be thought of broadly, including as a function over a bundle of characteristics. For instance, the quality of a wireless company's product is a function of the size and quality of its coverage network, the quality and variety of handsets, the retail distribution network, etc. The set of firms' qualities will be referred to as $\Omega = \{\omega_1, ..., \omega_n\}$. This is public information and represents the state of the industry.

Consumer preferences are represented by $u(\cdot)$, where consumer k's utility from good i is given by $u_{k,i} = \omega_i + log(y - p_i) + \epsilon_{k,i}$, where $\epsilon_{i,k}$ represents consumers' differing tastes. Each consumer purchases one unit of the product that gives them the highest utility. They may also purchase an "outside option" whose utility is normalized to 0. Following the work of McFadden (1974), if ϵ is drawn from an extreme value distribution with dispersion parameter ϕ_ϵ^{-1}, this results in the logit demand system:

$$q_i(p_1,...p_n;\Omega) = M\frac{exp(\phi_\epsilon(\omega_i + log(y-p_i)))}{1+\sum_j exp(\phi_\epsilon(\omega_j + log(y-p_j))))} \quad (1)$$

where $q_i(\cdot)$ is firm i's demand and M is the size of the market, or the measure of consumers. In this setting, ϕ_ϵ can be thought of as the degree of horizontal differentiation in consumer preferences, such that an increase in ϕ_ϵ translates into a higher market share for the highest quality product in the market. Firms face symmetric marginal costs mc and choose prices conditional on the set of goods in the market to maximize profits, such that:

$$\pi(p_i, p_{-i}) = q_i(p_1, ..., p_n; \Omega)(p_i - mc) \quad (2)$$

Investment Firms invest to increase product quality using a stochastic R&D technology. This technology follows recent work by Mermelstein, Nocke, Satterthwaite and Whinston (2014) (hereafter MNS&W.) Their key insight was to recognize that, despite the value in studying the dynamic effects of mergers, in the widely used Pakes and McGuire (1994) dynamic oligopoly framework mergers directly reduce the industry-wide investment possibility set by reducing the number of firms who can invest. In the MNS&W framework, the set of possible investments and investment costs are purely a function of a firm's current state ω_i. Thus, when firms merge and combine products, this action does not necessarily reduce the total set of possible investments. If the new state is $\omega_i' = \omega_i + \omega_j$, the firm's investment problem is unchanged from the pre-merger problems of both firms, except that the business stealing externality has been internalized. This is crucial for examining the relationship between mergers and innovation, since the standard model mechanically generates a negative relationship.

In each period, firms may invest in an attempt to increase their product quality. Let this quality ω_i takes an integer value. Firm i then draws a set of investment costs $\{c_j\}_{j=1}^{\omega_i} \in [\underline{c}, \bar{c}]$ for each unit that makes up ω_i. This is the cost of upgrading that unit by 1. MNS&W refer to this technology as *capital augmentation* although in this context it might better be thought of as *quality augmentation.* In addition, firms draw another cost, which MNS&W refer to as a *greenfield cost*, from some distribution $[\bar{c}, c_g]$. This determines the cost of product improvement for investment levels above ω_i. Thus, any firm can reach any greater state in each period, and firms with higher quality products are more likely to get low cost draws for some number of innovations and increase their quality. In some counterfactuals that follow, I cap the amount of innovation a firm can achieve in a single period.

As in Pakes and McGuire (1994), the industry as a whole also faces an exogenously improving outside good. This happens with probability δ each period, and is equivalent to reducing all firms product qualities by 1. If the endogenous long-term rate of innovation exceeds δ, the set of potential good qualities Ω becomes unbounded. To avoid the problems this would imply, I follow Goettler and Gordon (2014) in setting some $\overline{\omega}_{max}$ as the industry frontier. If in any period a firm innovates to a quality level $\omega_i' > \overline{\omega}_{max}$, the result is that all firms experience a downward shock equal to $\omega_i' - \overline{\omega}_{max}$. This keeps the frontier firm at level $\overline{\omega}_{max}$ and preserves the relative differences between the product qualities of all active firms. Because only these relative differences matter for profits, this does not effect equilibrium outcomes.[10]

Firms also face a flat, fixed operating cost FC that must be paid each period. In the beginning of each period, after observing investment costs $\{c_1, ..., c_\omega\}$, firms choose whether to remain in business and pay FC or exit. They then choose investment level $x_i \in \{0, ..., \overline{\omega}_{max}\}$.

2.2 Merger Stage

The bulk of previous research studying the implications of horizontal mergers has examined the results of exogenously imposed mergers. Although studying the pattern of mergers that would arise endogenously has clear benefits, modeling endogenous mergers poses a challenge. In many industries there may exist a set of profitable but mutually exclusive merger arrangements. The mergers in this set represent multiple equilibria and there is no clear equilibrium selection mechanism. One solution to this problem is to model the merger stage as a non-cooperative game, where firms propose buyout offers according to a defined sequence that provides a unique equilibrium in each stage.

Gowrisankaran (1999) follows this approach, embedding in an Ericson-Pakes model a stage game wherein the largest firm acts first. It has the ability to propose a merger to any other firm. If it chooses not to the second largest firm may propose, and so on. The stage game employed here is similar although the sequence by which firms may propose mergers

[10]When the market share of the outside option is large, absolute levels of ω, and not just relative values, do matter. Consequently, $\overline{\omega}_{max}$ is set so that the outside goods share is less than 1%. This can be thought of as a weak notion of industry technology spillovers. Goettler and Gordon (2014) describe the implications of this at length.

is random. While this adds to the difficulty of solving the model, it should result in a richer pattern of outcomes.

At the beginning of each period, a firm is randomly chosen and allowed to enter merger negotiations with any other firm. Before choosing its partner, the offering firm observes a random "synergy" value for each rival firm σ_{ij} that is uniform i.i.d. This represents the degree to which their products can be integrated into a new future product. The period following the merger, the new, combined firm will produce a product of quality $\omega' = \sigma_{sb}(\omega_B + \omega_S)$, where subscripts indicate the buyer and seller. The degree of synergy might reflect the amount of overlap between the two firms' products pre-merger, particularly as products represent a bundle of characteristics or services. In addition, within a period $\sigma_{ij} = \sigma_{ji}$, meaning synergy values are fixed for any pair of firms, and $1 \geq \sigma_{sb} \geq \frac{\omega_B}{\omega_B + \omega_S}$ such that a merger can never result in a lower quality product than the buying firm.

These mergers are unlikely to benefit consumers in a static sense except in cases where σ_{ij} is quite high and other firms in the market constrain price increases. This is in part because there are no cost efficiencies that accompany the merger. This modeling choice is intentional, since mergers with large efficiency gains that increase static welfare are uncontroversial regardless of their effects on innovation.

Having firms combine their products rather than continuing to produce both and simply adjusting prices is not how mergers are typically treated in the literature on competition in Bertrand settings. This combination of product qualities can be thought of as representing the total utility consumers receive from a firm's offerings, which reflect a bundle of characteristics, and where σ_{ij} represents the degree of overlap in these characteristics. As one example in the case of software services, a company may offer a mobile app product that has a number of map-based features including driving directions, links to online shopping, information on public transit options, information on specific business locations including consumer reviews, pictures, and links, as well as varying degrees of integration with other devices. If a new firm enters that offers one of these specific features that a large dominant firm currently lacks, they may acquire the new entrant and integrate the new feature into their existing product, increasing it's quality while also eliminating a rival in the online maps market.

A similar mechanism applies in some pharmaceutical mergers, a setting which has attracted much interest with regards to the effects of these mergers on long term innovation.

For instance, when Merck bought Idenix in 2014 they stated that they planned to "combine the drug with two of its own drugs that work by different mechanisms for a triple-drug regimen that could potentially cure most types of hepatitus C in less than two months (Loftus (2014)).

An added benefit of this approach to modeling mergers is that it fits better within the quality ladder setting, where firms and products are defined by their position in vertical quality space. This is a more natural setting to study innovation and particularly the relationship between competition and innovation, and has been used by Greenstein and Ramey (1998), Chen and Sappington (2011), Acemoglu and Akcigit (2012), Goettler and Gordon (2014), and Acemoglu and Cao (2015), among others, for this purpose. In this case vertical considerations are more important than static horizontal ones like the diversion ratio between firms and the pricing externality. While the product synergy approach used here sacrifices realism with respect to static horizontal features, it does so to focus more on what is relevant to the study of long-term innovation.

During the merger stage firms are fully forward-looking. Conditional on the set of synergies (or lack thereof), the proposing firm will either propose a merger with the firm offering the highest return in the merger stage or pass on the option. If the firm passes, a new firm is chosen at random and given the opportunity to offer a merger. The process continues until all firms have had an opportunity or a merger occurs. Because firms know that if they refuse a buyout offer they may be the next firm with the power to propose a merger, they may have the incentive to turn down a profitable merger foreseeing another, more profitable merger with some other firm. Similarly, they may accept or propose a less valuable merger to prevent two other firms in the market from merging and becoming too powerful. Because synergy values are random and private information, this merger stage game can be quite complex but results in a rich and fully endogenous pattern of mergers.

To evaluate a possible acquisition, firms consider the potential surplus that would result from a merger. The merger's surplus is the difference between the combined firm's presented discounted value and the sum of the separate firms' values. If there is a positive surplus from the firms' merger, it will be split between the two parties. This split results from Nash bargaining where reservation price of the firm being acquired is its value if the negotiation fails.[11] The value to the acquiring firm is the difference in values between the combined firm

[11] One reason previous papers have resisted considering dynamic models of mergers in industries with more than 2 firms is that there is no fully satisfactory solution to the bargaining problem merging firms

and its value if negotiations fail. Let $V^B(\cdot)$ and $V^S(\cdot)$ be the values of the buyer and the seller at the beginning of the following period at market structure Ω, which are described in greater detail in the following section, and let m_{ij} indicate whether or not a merger was agreed to by both parties, with 1 meaning it was. The size of the buyout offer τ_{ij}solves:

$$\max_{\tau_{ij}} \left(V^B(\Omega'|m_{ij}=1,\sigma_j) - \tau_{ij} - V^B(\Omega'|m_{ij}=0) \right)^{\rho_b} \left(\tau_{ij} - V^S(\Omega'|m_{ij}=0) \right)^{\rho_s} \tag{3}$$

where ρ_b and ρ_s represent buyer and seller bargaining power parameters. The result is a payment equal to

$$\tau_{ij} = \rho_b V^S(\Omega'|m_{ij}=0) + \rho_s[V^B(\Omega'|m_{ij}=1,\sigma_j) - V^B(\Omega'|m_{ij}=0)] \tag{4}$$

Because $\rho_s = 1 - \rho_b$ by definition, this can equivalently be written as

$$\tau_{ij} = V^S(\Omega'|m_{ij}=0) + \rho_s[V^B(\Omega'|m_{ij}=1,\sigma_j) - V^B(\Omega'|m_{ij}=0) - V^S(\Omega'|m_{ij}=0)] \tag{5}$$

The first term is the reservation value of the seller and the second term is the share of the surplus from the buyer that is paid out.

2.3 Antitrust Policy

In the results that follow I consider specifications both with and without antitrust policy. In specifications with antitrust policy, all mergers are reviewed by an antitrust authority who determines whether or not to allow the merger to proceed. The antitrust authority has full information, observing the state $s = \{\Omega, i, j, \sigma_{ij}\}$ where i and j are the identities of the merging firms.

The antitrust authority evaluates the post-merger state of the world against the alternative. Denote these Ω'_M and Ω'_{NM} for the state where the merger proceeds and where the merger is blocked, respectively. When active, the authority calculates consumer surplus under both alternatives and decides its policy according to the rule:

must solve due to the fact that there is a positive externality being conferred on the non-merging firm. I use the bilateral Nash bargaining outcome, effectively ignoring the effect of the merger on the non-merging firm. The externality still effects outcomes however, as firms might strategically wait or turn down a merger opportunity in cases where they would benefit more from their two rivals merging.

$$a_{it} = \begin{cases} reject & \text{if } CS(\Omega'_{NM}) > CS(\Omega'_M) \\ approve & \text{if } CS(\Omega'_{NM}) \leq CS(\Omega'_M) \end{cases} \tag{6}$$

where $CS(\cdot)$ represents static consumer surplus calculated by integrating over the demand system described in equation 1.[12]

2.4 Potential Entrants

In each period, a single firm may enter the market. The potential entrant lives for a single period and must pay an entry cost to join the industry, becoming an incumbent and competing in the product market in the following period. Potential entrants face the same investment cost function as incumbents, but where greenfield costs begin at ω_{max}, allowing entrants to innovate up to any possible ω level if they were to receive a favorable cost draw. The timing of the model is such that potential entrants make their entry and investment decision at the beginning of the period, simultaneous with incumbent firms making their exit and investment decisions.

2.5 Timing

1. Incumbent firms observe investment costs and potential entrants observe entry costs.
2. Incumbents choose whether or not to exit, their investment level if continuing, and entrants decide to enter or not and at what quality level. Their product qualities adjust as a result.
3. Firms enter the merger stage:
   (a) Some firm i is selected and observes merger synergies σ_{ij} for all other firms.
   (b) Firm i selects its most profitable potential partner, and if the total surplus from the merger is positive, the two firms agree to merge.
   (c) If no merger is agreed to in step (b), the merger stage repeats until all firms have had a chance to propose mergers or an agreement occurs.

[12]See Small and Rosen (1981) for more detail on calculating welfare in this model.

4. If a merger agreement was reached and the antitrust authority is active, it evaluates the proposed merger and approves or rejects. If the merger is approved, firm i's product updates to $\omega_B + \sigma_{ij}\omega_S$ and pays τ_{ij} to firm j.

5. Firms compete and earn profits $\pi(\Omega')$

3 Equilibrium and Computation

3.1 Firm Policies

In this section I formally describe firm policies over entry, exit, investment and mergers as well as the antitrust authority's merger review policy. I describe the conditions for a symmetric, Markov perfect equilibrium and the computational algorithm for finding it.

Incumbent exit and investment policies: At the beginning of a period, each incumbent draws a set of investment costs equal in number to their product quality ω, which takes an integer value. For simplicity, I will describe the policies of one representative firm. Firm i with product quality ω_i takes ω_i draws uniformly from the distribution $[\underline{c}, \overline{c}]$. In addition they draw a greenfield cost from the distribution $[\overline{c}, c_g]$.

Let $\overline{V}(\omega_i, \omega_{-i})$ represent the interim value of being in state ω_i while your rivals have states ω_{-i} after entry, exit, and investment have taken place but before the merger stage. After observing its set of cost draws $\tilde{c_i}$, firm i chooses its exit policy $\chi^{EX} \in \{0,1\}$ and, if not exiting, the amount of investment to undertake $x_i \in \{0, ..., \overline{\omega}_{max}\}$. Simultaneously, the industry-wide depreciation shock $\eta \in \{0,1\}$ is realized, taking value 1 with probability δ. After investing at level x_i, a firm's state updates to $\omega_i' = \omega_i + x_i - \eta$. The firm therefore solves:

$$\max_{x_i}\{-FC - C(\tilde{c_i}, x_i) + \beta \sum_{\eta} \sum_{\omega_{-i}} \overline{V}(\omega_i + x_i - \eta, \omega'_{-i}) p(\eta) h(\omega'_{-i}|\omega_{-i})\} \tag{7}$$

Where $h(\cdot|\cdot)$ represents beliefs over rival firms investment outcomes, including potential entry and exit. Let x_i^* represent the solution to this problem. The firm exits and $\chi^{EX} = 1$ if

$\{-C(\tilde{c_i}, x_i^*) + \beta \sum_{\eta} \sum_{\omega_{-i}} \overline{V}(\omega_i + x_i^* - \eta, \omega'_{-i}) p(\eta) h(\omega'_{-i}|\omega_{-i})\} < FC.$

The investment level x_i^* is determined by equating the marginal cost of an additional unit of investment to the increase in the expected value upon reaching the merger stage.

The potential entrant's problem is very similar to that of an incumbent. It draws $\overline{\omega}_{max}$ investment cost draws from the distribution $[\underline{c}, \overline{c}]$. It then decides whether or not to enter based on the expected value of pursuing the optimal level of investment. Consequently, the product quality of the entrant is endogenous and can take any value in $\{1, ..., \overline{\omega}_{max}\}$.

Mergers: When deciding whether or not to propose a merger with another firm, firm i must evaluate a set of potential future outcomes. Denote as $V(\omega_i, \omega_{-i})$ the value of being in state ω_i with rivals in states ω_{-i} at the beginning of a period, before cost shocks have been observed. To take expectations over the future state if firm i does not propose a merger, they must consider the probability that there is a merger between other firms, which occurs with probability $\sum_k \sum_j \int_{\sigma_{jk}} Q(m_{jk}|\Omega)$ where m_{jk} represents a merger between firms k and j. This probability represents the joint probability that firms k or j are next to propose a merger as well as the distribution of potential synergy values over each pair that is reached in equilibrium. If the antitrust authority is active, they must also consider the its decision rule χ^{AA} when considering proposing a merger. Because the antitrust authority faces a deterministic decision, in equilibrium firms will never propose mergers that are subsequently rejected.

If firm i is the proposing firm, they observe synergy values σ_{ij} for all firms in the market, including a new entrant if entry occurred earlier in the period. They then choose

$$\max \left\{ \max_j \{\chi^{AA}(\Omega, i, j, \sigma_{ij}) \Big[-\tau_{ij}(\omega_i, \omega_{-i}, \sigma_{ij}) + \pi(\omega_i + \sigma_{ij}\omega_j, \omega_{-i}|m_{ij}) + \beta V(\omega_i + \sigma_{ij}\omega_j, \omega_{-i}|m_{ij}) \Big] \}, \right.$$
$$\left. \sum_k \sum_j \int_{\sigma jk} Q(m_{jk}|\Omega) \chi^{AA}(\Omega, j, k, \sigma jk) \Big[\pi(\omega_i, \omega'_{-i}|m_{jk}, \sigma_{jk}) + \beta V(\omega_i, \omega'_{-i}|m_{jk}, \sigma_{jk}) \Big] d\sigma \right\} \tag{8}$$

The first term inside the max operator is the firm's choice of merger partner. For each potential partner, conditional on the observed σ_{ij}, they evaluate the size of the buyout payment and post-merger profits and continuation value. The second term is the expected value of not proposing a merger and potentially seeing rival firms merge in the same period.

3.2 Equilibrium

I will consider Markov Perfect Equilibria (MPE) for this model. If $s \in \mathcal{S}$ represents some element of the state space, a MPE consists of:

- A subset $\mathcal{R} \subset \mathcal{S}$;
- Strategies χ^* for every $s \in \mathcal{R}$, where $\chi^* = (\chi^E, \chi^{EX}, m_{ij}, \tau_{ij}, x_i, xe_i)$ respectively governing entry, exit, mergers, buyout offers, and investment.
- Expected discounted values conditional on these strategies, $V^E(\Omega, ce_i), V(\omega_i, \omega_{-i})$, $V^M(\Omega, i, j, \sigma_j)\forall j$, and $V^I(\omega_i, \omega_{-i})$.

Such that:

1. The Markov process defined by any initial condition s_0 and the strategies χ^* has $\mathcal{R}$ as a recurrent class.
2. For every $s \in \mathcal{R}$, strategies are optimal given $V^E(\cdot)$, $V(\cdot)$, $V^M(\cdot)$, and $V^I(\cdot)$. That is, $\chi^*(\Omega)$ solves:

$$\max_{\chi^E} V^E(\Omega, ce_i),\ \max_{\chi^{EX}, x_i} V^I(\omega_i, \omega_{-i}),\ \max_{\chi^M, m_{ij}, \tau_{ij}} V^M(\Omega, i, j, \sigma_j)$$

3. Values are consistent on $\mathcal{R}$. For every Ω and Ω' that are components of $s \in \mathcal{R}$:

$$V(\omega_i, \omega_{-i}) = \pi(\omega_i, \omega_{-i}) - FC + \sum_i \sum_j \int_{[0,1]} Q(m_{jk}|\Omega) V^I(\omega_i', \omega_{-i}'|m_{ij}, \sigma_j) \mathrm{d}\sigma$$

$$V^I(\omega_i, \omega_{-i}) = \max\{0, \max_{x_i} -C(\omega_i, x_i) + \beta \sum_\eta \sum_{\omega_{-i}} V(\omega_i + x_i - \eta, \omega_{-i}') p(\eta) h(\omega_{-i}'|\omega_{-i})\}$$

$$V^E(\chi^{E*}, xe_i|\Omega, ce_i) = \max\{0, max_{xe_i} - C(\omega_{max}, xe_i) + \beta \sum_\eta \sum_{\omega_{-i}} V(xe_i - \eta, \omega_{-i}') p(\eta) h(\omega_{-i}'|\omega_{-i})\}$$

An MPE for this model can be shown to exist following Doraszelski and Satterthwaite (2010). For a discussion of potential multiplicity, see Doraszelski and Pakes (2008). Generally, there is no way to rule out the possibility of multiple equilibria, which poses a challenge for counterfactual policy analysis. Given that multiple equilibria have been found to exist in similar models without a merger stage, a more complex model also plausibly suffers from this problem. Borkovski et al. (2012) show multiplicity in a quality ladder model, although they conclude that "the differences between equilibria tend to be small and may matter little in practice." Here, I offer the standard disclaimer that in the course of solving the model large numbers of starting states and policies were tested and never produced meaningfully different outcomes. In addition, when computing the equilibrium with mergers, I tested initiating the policies and values using the results from the equilibrium without mergers, and vice versa. In both cases, the algorithm converged to the same equilibrium. Because, as described in more detail below, the computational algorithm used to compute equilibria uses reinforcement learning, it generates stable equilibria only, which are likely to be more relevant for practical purposes, and the comparisons between counterfactuals can actually be thought of as finding a single equilibrium in which antitrust policy suddenly switches regimes.

To compute the model, I map the measure of product quality ω onto the integers $\{0, ..., 10\}$. There is no limit on the number of firms allowed in the market although under the parameters chosen there are never more than 4 firms active in equilibrium. Most prior work in this literature caps the number of active firms at 2 for computational reasons. This limitation is potentially costly, as it necessarily restricts attention to mergers to monopoly, which are rarely allowed in practice and which always reduce consumer welfare in this model, unlike mergers from 3 to 2 firms. A binding cap could be thought of as imposing an infinite entry cost at states with 2 firms in the industry, even if a third firm could profitably operate.

The model is too complex to allow an analytic solution, instead, it is solved computationally using a modified version of the stochastic algorithm of Pakes and McGuire (2001). The potential computational burden of the model described is enormous. The size of the state space grows exponentially in the number of firms and potential good qualities, and for each state, the integral over potential future states required to calculate the expected discounted value of different actions involves probability distributions over the random sequence of merger proposers, synergy values, exit and entry behavior, and the outcomes of investment. The computational burden of this high-dimensional integral and state space is the reason there has been little work done on this type of analysis to date. The stochastic algorithm method substantially reduces this burden.

A detailed description of how the model is solved can be found in Appendix A. This method solves the model asynchronously using the technique of reinforcement learning. The model is simulated for a very high number of periods, with firms' value functions being updated with the observed results of their actions. Over time, the average of a firm's experiences becomes equal to its true expected value. The method offers several advantages. The first is that equilibrium policy and value functions are only computed over a subset of the state space. This subset, $\mathcal{R} \subset \mathcal{S}$, is the recurrent class of the Markov process formed by equilibrium strategies. While the state space grows exponentially in the number of potential firms, its possible for $\mathcal{R}$ to grow linearly or even not grow at all.

The other advantage is that by simulating the model rather than solving it directly, it is not necessary to solve any high-dimensional integrals except once, in the limit. To briefly describe the algorithm; for each visit to a state, firms solve the optimal policy based on their estimate of the value function. Once they choose, pseudo-random numbers are drawn for any stochastic component and the state is updated. The value function estimate at the original state is then updated to include the profit realized and value at the new state. The process then repeats at the new state. To improve performance, policy functions are randomly perturbed in a small share of periods that slowly declines to zero.[13] Periodically, a test of the equilibrium conditions is conducted, this test is described in detail in Fershtman and Pakes (2012). The algorithm performs well, converging to the same equilibrium outcome from very different initializations of value of policy functions.

4 Results

4.1 Mergers and Innovation Incentives

The model is solved for MPE numerically with parameter values initially taken from Pakes and McGuire (1994) with the addition of merger fixed costs set at .5 and merging firms having equal bargaining power.[14] The full set of parameters can be seen in Table 1. The results that follow simulate the industry for 500,000 periods from the equilibrium's recurrent class of states. This is a sufficient number of simulation periods that the results displayed do not vary over simulations and so no standard deviations are presented alongside them.

Table 2 summarizes the key equilibrium outcomes in three different settings: one in which no mergers are allowed, one with no restrictions on mergers, and one in which mergers are allowed except that a simple antitrust policy blocks those that would create a monopoly.

[13] This prevents the algorithm from getting "stuck" in non-equilibrium values. It is referred to in the machine learning literature as an epsilon-greedy or epsilon-decreasing strategy. More detail on how this is implemented can be seen in Appendix A.

[14] The empirical finance literature finds inconclusive results on the shares of a merger's surplus going to either party, but most work finds the shares roughly equal. See, for instance, Ahern (2012).

Table 1: Base Parameterization

β	.925
$\overline{\omega}$	10
FC	.6
c_M	.5
ρ_b, ρ_s	.5
δ	.6
$\underline{c}$	.1
$\overline{c}$	15
y	15
M	10

In the benchmark case without firm acquisitions shown in column 1, the industry forms a relatively stable duopoly with little entry or exit. The duopoly takes the form of a leader and a laggard on quality where the leading firm invests enough to typically maintain its position. The rate of innovation is measured as the share of periods in which investment advances the industry past ω_{max}, effectively defining innovation as the rate at which the industry's technological frontier advances. In the no-merger duopoly outcome, there is little incentive for the industry leader to advance ω_{max}, and so innovation only occurs in .3% of periods.

When mergers are allowed with no restrictions, they occur frequently, in over one third of periods. Figures 1 and 2 show the distribution of mergers in equilibrium when the total number of firms active pre-merger is 2 and 3, respectively. The vertical axis shows the state of the acquiring firm and the horizontal axis shows the state of the selling firm. Only states below the diagonal are represented because the larger firm is the acquirer by definition. When $N = 2$, the vast majority of mergers are between the leading firm and either a firm of state 1 or of state 8. When $N = 3$, these are again the most common merger states but there is a greater range of outcomes.

The most notable fact about the mergers that occur is that they are almost entirely harmful to consumers in the short run. Static consumer surplus is always decreased by a merger at $N = 2$ to create a monopoly. When $N = 3$, the short run welfare effects of a merger depend on the size of the firms involved, the synergy value σ, and the size of the additional firm in the market. Table 3 shows the share of mergers that increase welfare for every buyer-seller pair. Of all the mergers that occur in equilibrium, only 18.8% are welfare increasing, the vast majority harm consumers in the short run. This is not surprising, the

Figure 1: Distribution of Mergers when N=2

Buyer state \ Seller state	1	2	3	4	5	6	7	8	9
1	0								
2	0	0							
3	0	0	0						
4	0	0	0	0					
5	0	0	0	0	0				
6	0	0	0	0	0	0			
7	0.001	0	0.000	0	0	0	0		
8	0.002	0.000	0.000	0.000	0	0	0	0	
9	0.503	0.009	0.006	0.002	0.002	0.011	0.002	0.461	0

Seller state

Note: This figure shows the relative frequency of mergers occurring in states in which 2 firms are active entering the merger stage, where the horizontal axis is the state of the seller and the vertical axis is the state of the buyer.

Figure 2: Distribution of Mergers when N=3

Buyer state \ Seller state	1	2	3	4	5	6	7	8	9
1	0								
2	0.007	0							
3	0.038	0.007	0						
4	0.008	0.003	0.002	0					
5	0.013	0.004	0.000	0.000	0				
6	0.003	0.001	0.000	0.000	0.000	0			
7	0.025	0.003	0.000	0.000	0.000	0.000	0		
8	0.007	0.002	0.001	0.000	0.000	0.000	0.006	0	
9	0.236	0.029	0.026	0.017	0.031	0.013	0.488	0.029	0

Seller state

Note: This figure shows the relative frequency of mergers occurring in states in which 3 firms are active entering the merger stage, where the horizontal axis is the state of the seller and the vertical axis is the state of the buyer.

model essentially stacks the deck against mergers that benefit consumers in the short run in order to highlight the difference between the static outcome and the long-term outcome. Mergers harm consumers because they reduce the number of products available, reduce price competition, and have no offsetting cost efficiencies.[15]

The most important result is that when mergers are allowed the rate of entry increases substantially from only 6.5% of periods to roughly 49% of periods. One of the results of this entry is that allowing mergers results in only slightly fewer firms in the market, on average, decreasing from 2.06 to 1.84 and while the share of periods spent in monopoly grows substantially, mergers actually increase the share of periods with 3 firms active. The result is an industry spending roughly 11% of the time with three firms and 27% as a monopoly.

The first main result is that, despite the fact that the vast majority of mergers reduce

[15]As noted by Ackerberg and Rysman (2005) and others, the way taste heterogeneity is represented in logit style models means that reducing the number of products available reduces total consumer surplus, even holding all else equal.

Figure 3: Share of Mergers increasing welfare when N=3

	1	2	3	4	5	6	7	8	9
1	0								
2	0.33	0.11							
3	0.38	0.33	0.125						
4	0.5	0.4	0.25	0.12					
5	0.67	0.5	0.375	0.23	0.11				
6	0.71	0.62	0.46	0.34	0.21	0.1			
7	0.75	0.67	0.563	0.425	0.31	0.2	0.09		
8	0.83	0.7	0.61	0.51	0.39	0.29	0.19	0.1	
9	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.19	0.09

Seller state

Note: This table shows the share of mergers that reduce static consumer surplus compared to no merger when 3 firms are present, as a function of the state of the non-merging 3rd firm. The horizontal axis is the state of the seller and the vertical axis is the state of the buyer.

welfare in the short run, allowing mergers increases the average consumer surplus in the industry and the rate of innovation substantially. The share of periods with innovations increases from .3% to 85%. Consumers benefit in perpetuity from an increase in the technological frontier and so the increase in the endogenous rate of innovation dramatically increases long-term consumer welfare. Interestingly, while there is more entry, average firm profits also increase when mergers are allowed, probably because of the higher share periods where one firm holds a monopoly. The difference between the outcomes with unlimited mergers and mergers up to duopoly are qualitatively similar, so for the remainder of the results section, for simplicity I consider only the unlimited mergers and no mergers equilibria.

Beyond average market outcomes, we can observe firm policy functions directly to explore the effects of mergers on innovation. Figure 4 shows firm R&D policy functions plotted over quality level. These are calculated from simulations and weighted by states visited. For each equilibrium, we see an inverse-U shaped relationship between firm size/quality and investment. Firms invest the most at intermediate states and reduce investment when they reach the highest quality states. Two notable results emerge when comparing investment functions with and without mergers. First, at all quality states, investment is higher when mergers are allowed. Second, with fewer restrictions on mergers, total investment is higher and it peaks at a lower state.

Note: This figure shows the average amount of investment by firms in each state for the equilibria with and without mergers. This is computed as the average across all market states that appear in equilibrium across many simulations.

Figure 5 shows that the rates of entry occurring for different numbers of incumbent firms. When there is only one incumbent in the market, entry occurs 86% of the time in

Table 2: Comparison of equilibrium with and without mergers

	No Mergers	Mergers	Mergers up to Duopoly
Firm Characteristics			
Mean number of firms	2.06	1.84	2.09
Mean firm quality	4.80	6.08	4.99
Share of periods with entry	6.50%	48.98%	25.74%
Share of periods with mergers		37.68%	16.59%
Investment			
Total investment	.80	1.29	1.15
Mean investment	.39	.70	.55
Rate of innovation	.003	.85	.55
Mean entrant investment	2.10	4.01	4.93
Mean investment by market leader	.79	1.09	1.07
Surplus			
Mean consumer surplus	5.89	11.26	9.67
Mean total profit	12.89	13.23	12.94
Firm Distribution			
Share of periods in Monopoly	.05	27.24%	.18%
Share of periods in Duopoly	93.49%	61.75%	90.66%
Share of periods with 3 firms	6.45%	11.01%	9.15%
Share of periods with 4 firms	5e-4%	.001%	.004%

the mergers equilibria. Entry occurs more rarely in the no-mergers equilibrium, regardless of number of incumbent firms.

It is clear from Figure 5 that the additional entry generated in the mergers equilibrium is not merely replacement entry after a merger reduces the number of firms below the stable duopoly number. This is because with mergers allowed, the average number of firms in the market actually increases, the increase in entry takes place even when mergers to monopoly are not permitted, and entry rates are higher even when 2 or 3 incumbents are present, as shown in Figure 5. Entry rates are higher at every market state. While antitrust economists have long known that entry can mitigate the anticompetitive effects of mergers, and entry is discussed in the Horizontal Merger Guidelines for this reason, the argument for entry here is distinct. The value created by entry is not about reducing market power of large, post-merger incumbents, but instead the prospect of a future buyout is generating new, additional entry that is increasing competition and innovation while also increasing consumer surplus.

Figure 4: Investment levels with and without mergers

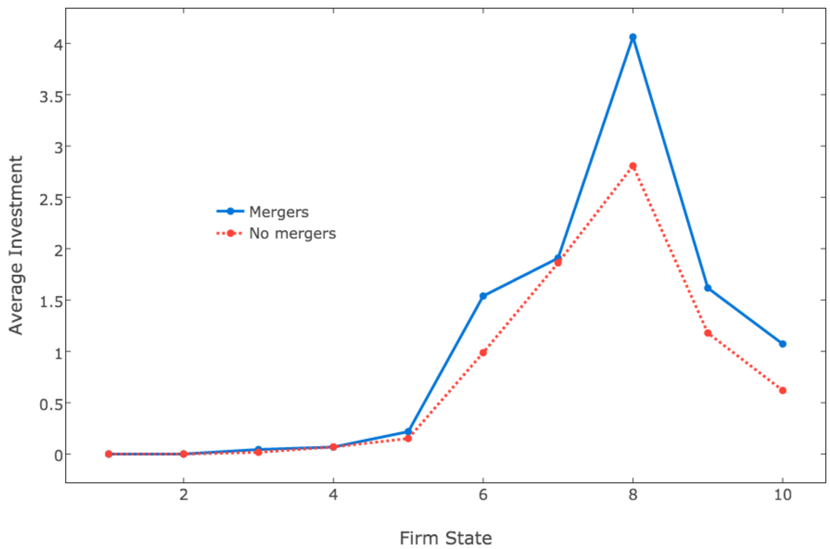

Figure 5: Entry Rates

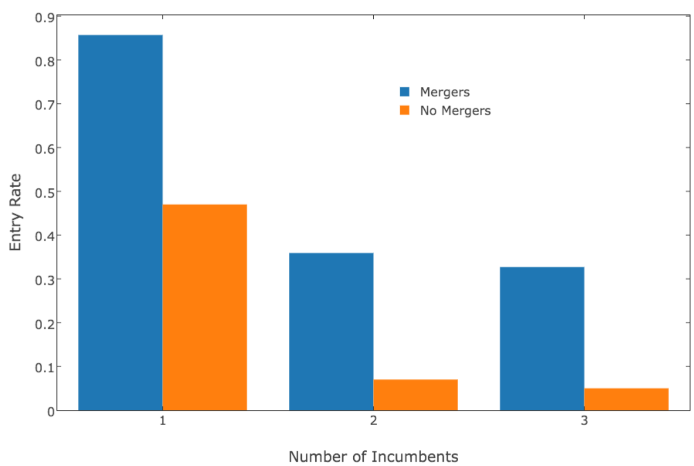

Note: This figure shows the average rate of entry as a function of number of incumbent firms for the equilibria with and without mergers. This is computed as the average across all market states that appear in equilibrium across many simulations.

4.2 Antitrust Policy

This section describes the equilibrium outcome when the antitrust authority is active, blocking any merger that would decrease static consumer surplus. I compare it to the benchmark cases where all mergers are allowed and where none are. Results are shown in Table 3. The primary result is that very few mergers are proposed and approved, occurring in less than 1% of periods. This shows that the set of mergers that are simultaneously profitable, static welfare increasing, and reachable from the equilibrium path is very small. Almost all mergers have high synergies meaning $\frac{\sigma_{ij}}{\omega_j} > .5$. Because there is no uncertainty in merger review, in equilibrium no mergers are proposed that would be rejected, but of all hypothetical mergers possible along the equilibrium path, 19% would be approved and the rest rejected.

Average consumer welfare is only slightly higher than in the equilibrium with no mergers at all and innovation is essentially unchanged. The industry is essentially in a very stable duopoly with one leader and one laggard. The leading firm invests more in this equilibrium than in one without mergers, but less than when mergers are unrestricted. Innovation is therefore quite rare.

4.3 Contestability and Disruptability

In this section I explore what industry characteristics are driving the above results. This also serves to explore the parameter space and consider the robustness of the equilibrium outcome. I consider two main characteristics. First, I examine the role of "contestability" in the sense of Shapiro (2010), meaning the degree to which firms who successfully innovate can capture higher market share as a result.[16] This is represented by the parameter ϕ_ϵ in the consumers utility function, the dispersion of the random component of utility. When this parameter is high, preferences exhibit less heterogeneity, consumers agree more on which product offers the highest utility and firms with higher ω capture a higher market share. When ϕ_ϵ is low, preferences are more horizontally differentiated.

When contestability is high, a successful innovation translates into a large increase in market share and profits. Thus this could be thought of as representing low consumer switching costs, brand loyalty, or the ease of distribution.[17] In industries with high contestability,

[16] This is distinct from how the term contestable is used by Baumol et al. (1982) who used it to referred to a market with frictionless reversible entry.

[17] The scale of the dispersion parameter represents taste heterogeneity or unobserved product features, which are not dynamic in in the way switching costs are, but they do have similar implications for changes

Table 3: Comparison of equilibrium with and without mergers

	No Mergers	Mergers	Optimal Antitrust Policy
Firm Characteristics			
Mean number of firms	2.06	1.84	2.00
Mean firm quality	4.80	6.08	4.99
Share of periods with entry	6.50%	48.98%	.08%
Share of periods with mergers		37.68%	.04%
Investment			
Total investment	.80	1.29	1.00
Mean investment	.39	.70	.50
Rate of innovation	.003	.85	.003
Mean entrant investment	2.10	4.01	5.04
Mean investment by market leader	.79	1.09	1.00
Surplus			
Mean consumer surplus	5.89	11.26	5.93
Mean total profit	12.89	13.23	13.09
Firm Distribution			
Share of periods in Monopoly	.05	27.24%	1e-5%
Share of periods in Duopoly	93.49%	61.75%	99.91%
Share of periods with 3 firms	6.45%	11.01%	.001%
Share of periods with 4 firms	5e-4%	.001%	3e-6%

the benchmark rate of innovation will naturally be higher, but the relative effects of horizontal mergers is potentially ambiguous. To resolve this, I compare equilibria with and without horizontal mergers in industries with varying levels of contestability. The results are shown in Table 4.

As contestability rises the effects of horizontal mergers change dramatically. When $\phi_\epsilon = .8$, there is no difference in rate of innovation or consumer welfare between the mergers equilibrium and the no-mergers equilibrium. When contestability increases to $\phi_\epsilon = 1.33$, the rate of innovation and average consumer welfare are both substantially higher in the mergers equilibrium and the no-mergers equilibrium. This result is somewhat counterintuitive, because when switching costs or loyalty are lower, a higher share of the gains from innovation come from business stealing. Merging firms should internalize this incentive and invest less. Indeed, they do, investment by leading firms falls when contestability increases, relative to the no-mergers case. But this effect is outweighed by a large increase in entry

in market share resulting from innovation, although not for dynamic pricing policies that are not allowed here.

Table 4: Higher Contestability Increases Innovation Disparity

	$\phi_\epsilon = 1.33$		$\phi_\epsilon = 1$		$\phi_\epsilon = .87$		$\phi_\epsilon = .8$	
	NM	M	NM	M	NM	M	NM	M
Rate of Innovation	.07	1.83	.003	.85	.01	.09	0	0
Consumer Surplus	4.26	16	5.89	11.26	5.94	6.86	6.35	6.35

Each column shows the rates of innovation and long run average consumer surplus with mergers and with no mergers.

and innovation by new entrants. The reason is that when contestability is high, monopoly is relatively more valuable. The result is that the equilibrium buyout offer τ at the modal merger state is 2.4 times higher when $\phi = 1.33$ than when $\tau = .8$.

While contestability is a feature of consumer demand, I consider a related industry characteristic that is a feature of firm capacity to innovate. In the baseline model, firms are able to innovate to achieve any level of ω in any period. Indeed, in rare circumstances when a firm receives a set of very low investment costs, we do observe large changes in ω in a single period. This investment technology follows Mermelstein et al. (2014) and is important because it is "merger neutral" in the sense that it avoids mergers that mechanically reduce the industry's technological possibility set and hence investment. While important, it may not be realistic for all industries. In some industries, it is impossible regardless of investment costs to dramatically increase the quality of a product quickly.

I refer to this industry feature as "disruptability" and examine the results after varying the amount of product innovation firms are capable of from the baseline of no limit down to a limit of one incremental unit per period. Let $\tilde{I}$ represent this cap, such that $\omega' - \omega \leq \tilde{I}$, or equivalently $c_i = \infty$ for $\omega' > \omega_i + \tilde{I}$.

Results are presented in Table 5. We see that, as disruptability falls and firms can only advance by 1 in each period, the previous results are reversed. There is essentially no innovation and long run consumer welfare is lower with horizontal mergers than without. For middle levels of disruptability, mergers produce slightly higher levels of innovation and welfare. Finally, when there is no cap on innovation, the previous result is seen again, mergers create the incentive for much higher innovation and long run consumer welfare. The discontinuous jump in welfare and innovation when comparing the rightmost two columns in Table 5 suggest much of the effect is driven by infrequent but very large innovations by new entrants and small firms. It may be that firms who enter the industry at a small

Table 5: Higher Disruptability Increases Innovation Disparity

	$\tilde{I} = 1$		$\tilde{I} = 3$		$\tilde{I} = 5$		$\tilde{I} = 7$		No Cap	
	NM	M	NM	M	NM	M	NM	M	NM	M
Rate of Innovation	0	0	.002	.01	.003	.05	.003	.05	.003	.45
Consumer Surplus	4.8	3.8	5.8	6.1	5.8	6.1	5.8	6.1	5.9	8.9

Each column shows the rates of innovation and long run average consumer surplus with mergers and with no mergers.

level, prompted to do so by the prospect of a buyout when it otherwise would not have been profitable to enter, occasionally see very low cost draws giving them an opportunity to generate a very large innovation.

Combined, the results on contestability and disruptability give some guidance to antitrust policymakers as to when there is potential long run benefit of allowing mergers that are harmful to consumers in the short run. This is most likely to occur in dynamic industries where consumers have low switching costs or brand loyalty and when innovation is occasionally rapid and disruptive. When consumers face switching costs or most innovation is of the incremental variety, the short run harm to consumers from anticompetitive mergers is likely to be the dominant force.

While the nature of innovation in an industry is a feature of underlying technology and thus out of the scope of policy, it may be observable and therefore a useful factor for a merger authority to consider. Contestability, however, is potentially under the influence of policymakers, who can seek a more contestable market by restricting the use of long-term contracts, bundling requirements, and other practices that raise switching costs. In innovative industries this policy may be more effective in promoting consumer welfare and innovation than strict merger review.

5 Conclusion

The relationship between horizontal mergers and innovation is increasingly important but poorly understood. To examine it requires simultaneously modeling endogenous entry and investment behavior and endogenous mergers. In addition, firms must have a broad ability to innovate, the long-run rate of innovation must be made endogenous, and the mergers technology must allow for a flexible and rich pattern of mergers.

This paper shows that in a model with these features, traditional horizontal mergers policy based on static welfare analysis may be counterproductive in the long run. While the vast majority of equilibrium mergers fail such a test, the long term result when they are allowed is substantially higher innovation and consumer welfare. The prospect of a windfall gain from a buyout offer by the leading firm generates additional entry that otherwise would not occur. This is distinct from the replacement entry post-merger discussed in the Horizontal Merger Guidelines. In addition, firms find it profitable to invest to improve their prospects as a merger partner and some of the new entrants generate substantial innovations and become the leading firm in the market.

This pattern is most likely when the product market is more "contestable" in the sense that consumer switching costs or brand loyalties are low. It is also more likely when the industry is more "disruptable" meaning large and rapid innovations are possible. In industries with these features, antitrust policymakers should place more weight on the long run gains from the incentives provided by merger prospects and less weight on the immediate harm done to consumers by these mergers.

References

D. Acemoglu and U. Akcigit. Intellectual Property Rights Policy, Competition and Innovation. *Journal of the European Economics Association*, 2012.

D. Acemoglu and D. Cao. Innovation by Entrants and Incumbents. *Journal of Economic Theory*, 157, 2015.

D. Ackerberg and M. Rysman. Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects. *RAND Journal of Economics*, 36: 771–788, 2005.

P. Aghion, N. Bloom, R. Blundell, R. Griffith, and P. Howitt. Competition and Innovation: An Inverted-U Relationship. *Quarterly Journal of Economics*, 120:701–728, 2005.

K. Ahern. Bargaining Power and Industry Dependence in Mergers. *Journal of Financial Economics*, 103:530–550, 2012.

K. Arrow. Economic Welfare and the Allocation of Resources for Invention. In R. Nelson, editor, *The Rate and Direction of Inventive Activity*. Princeton University Press, 1962.

J. Asker, C. Fershtman, J. Jeon, and A. Pakes. The Competitive Effects of Information Sharing. *Working Paper*, 2016.

W. Baumol, J. Panzar, and R. Willig. *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, New York, 1982.

R. Borkovski, U. Doraszelski, and Y. Kryukov. A Dynamic Quality Ladder Model with Entry and Exit: Exploring the Equilibrium Correspondence Using the Homotopy Method. *Quantitative Marketing and Economics*, 10:197–229, 2012.

J. Chen. The Effects of Mergers with Dynamic Capacity Accumulation. *International Journal of Industrial Organization*, 27:92–109, 2009.

Y. Chen and D. Sappington. Exclusive Contracts, Innovation and Welfare. *American Economic Journal: Microeconomics*, 3, 2011.

K. Cheong and K. Judd. Mergers and Dynamic Oligopoly. *Journal Economic Dynamics and Control*, 2006.

U. Doraszelski and A. Pakes. A Framework for Applied Dynamic Analysis in I.O. *The Handbook of Industrial Organization*, 2008.

U. Doraszelski and M. Satterthwaite. Computable Markov Perfect Industry Dynamics. *RAND Journal of Economics*, 42:215–243, 2010.

R. Ericson and A. Pakes. Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *Review of Economic Studies*, 62:53–82, 1995.

C. Fershtman and A. Pakes. Dynamic Games with Asymmetric Information: A Framework for Empirical Work. *Quarterly Journal of Economics*, 2012.

R. Gilbert. Looking for Mr. Schumpeter: Where Are We in the Competition-Innovation Debate? In A. Jaffe, J. Lerner, and S. Stern, editors, *Innovation Policy and the Economy*. MIT Press, Cambridge, 2006.

R. Gilbert and S. Sunshine. Incorporating Dynamic Efficiency Concerns in Merger Analysis: The Use of Innovation Markets. *Antitrust Law Journal*, 1995.

R. Goettler and B. Gordon. Competition and Product Innovation in Dynamic Oligopoly. *Quantitative Marketing and Economics*, 12:1–42, 2014.

G. Gowrisankaran. A Dynamic Model of Endogenous Horizontal Mergers. *RAND Journal of Economics*, 30:56–83, 1999.

S. Greenstein and G. Ramey. Market Structure, Innovation, and Vertical Product Differentiation. *International Journal of Industrial Organization*, 16, 1998.

J. Harford. What Drives Merger Waves? *Journal of Financial Economics*, 2005.

M. Igami and K. Uetake. Mergers, Innovation and Entry-Exit Dynamics: The Consolidation of the Hard DIsk Drive Industry (1976-2014). *Working Paper*, 2015.

P. Jeziorski. Estimation of Cost Synergies from Mergers: Application to U.S. Radio. *RAND Journal of Economics*, 45:37–53, 2014.

M. Katz and H. Shelanski. Mergers and Innovation . *Antitrust Law Journal*, 2006.

P. Loftus. Merck agrees to buy idenix for $3.85 billion. *The Wall Street Journal*, June 9 2014.

G. Marshall and A. Parra. Mergers in Innovative Industries: A Dynamic Framework. *Working Paper*, 2016a.

G. Marshall and A. Parra. Mergers in Innovative Industries: The Role of Product Market Competition. *Working Paper*, 2016b.

B. Mermelstein, V. Nocke, M. Satterthwaite, and M. Whinston. Internal versus External Growth in Industries with Scale Economies: A Computational Model of Optimal Merger Policy. *Working paper*, 2014.

A. Nevo and M. Winston. Taking the Dogma Out of Econometrics: Structural Modeling and Credible Inference. *Journal of Economic Perspectives*, 24:69–82, 2010.

M. Nishida and N. Yang. Better Together? Retail Chain Performance Dynamics in Store Expansion Before and After Mergers. *Working Paper*, 2014.

A. Pakes and P. McGuire. Computing Markov Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model. *RAND Journal of Economics*, 25: 555–589, 1994.

A. Pakes and P. McGuire. Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the "Curse" of Dimensionality. *Econometrica*, 69:1261–1281, 2001.

M. Pesendorfer. Mergers Under Entry. *RAND Journal of Economics*, 36:661–679, 2005.

J. Schumpeter. *Capitalism, Socialism and Democracy.* Harper, New York, 1942.

I. Segal and M. Whinston. Antitrust in innovative industries. *American Economic Review*, 97:1703–1730, 2007.

C. Shapiro. Competition and Innovation: Did Arrow Hit the Bull's Eye? In J. Lerner and S. Stern, editors, *The Rate and Direction of Inventive Activity Revisited*. University of Chicago Press, 2010.

K. Small and H. Rosen. Applied Welfare Economics with Discrete Choice Models. *Econometrica*, 49:105–130, 1981.

J. Stahl. A Dynamic Analysis of Consolidation in the Broadcast Television Industry. *Working Paper*, 2010.

A Appendix A: Algorithm for solving the model

This appendix describes the algorithm used to solve for an equilibrium to the model described in section 3. This builds on the work of Pakes and McGuire (1994) and Fershtman and Pakes (2012) but adds a novel element from the reinforcement learning literature to improve convergence properties. It has been noted that the basic stochastic algorithm often performs badly on convergence. Indeed, without the changes described below the basic algorithm almost always fails to reach an equilibrium. Here I review some of the reasons for this and how they can be fixed. For greater detail on the basic algorithm see Pakes and McGuire (1994) and Fershtman and Pakes (2012).[18]

The algorithm proceeds iteratively, simulating the dynamic game using a stored value function that firms use to generate policies regarding entry, exit, investment, and mergers. At each step of this simulation, the value function is updated with the payoffs realized for each action taken or not taken. The key components that are stored in memory are the current state of the industry at each iteration k, called Ω_k, the stored value functions defining payoffs for each action at that state: $W_k^I(\Omega_k, x_i)$, $W_k^E(\Omega_k, xe_i)$, and $W_k^M(\Omega_k, i, j, \sigma_{ij})$, and a counter that stores the number of prior visits to state Ω_k to that point: $h_k^I(\Omega_k, x_i)$, $h_k^E(\Omega_k, xe_i)$, and $h_k^M(\Omega_k, i, j, \sigma_{ij})$. If $h_k^I(\Omega_k, x_i) = 0$, $W_k^I(\Omega_k, x_i)$ is empty. When the state (Ω, x_i) is reached for the first time, $W_k^I(\Omega, x_i)$ is set to an initial value and updated from there.

Profits $\pi(\Omega)$ are computed offline for all states. Each value function is initialized at some level that I discuss in more detail below. The timing is as follows, at each iteration k:

1. At state Ω_k draw from memory:$W_k^I(\Omega_k, x_i)$, $W_k^E(\Omega_k, xe_i)$, $W_k^M(\Omega_k, i, j, \sigma_{ij})$, $h_k^I(\Omega_k, x_i)$, $h_k^E(\Omega_k, xe_i)$, and $h_k^M(\Omega_k, i, j, \sigma_{ij})$.

2. For all incumbent firms i and for the potential entrant, draw investment costs $\tilde{c}_i$

3. Incumbents solve:

$$\max_{x_i}\{-C(\tilde{c}_i, x_i) + W_k^I(\Omega_k, x_i)\} \tag{9}$$

   and exits if the max of this term is less than zero.

   Entrants solve:

$$\max_{xe_i}\{-C(\tilde{c}_i, xe_i) + W_k^E(\Omega_k, xe_i)\}$$

[18] For another application, see also Asker et al. (2016).

and enter if the max of this term is greater than zero.

4. Randomly draw the industry-wide depreciation shock η.

5. Using η and the investment, entry and exit decisions of incumbents and the potential entrant, update the market state from Ω_k to Ω'_k.

6. Begin the merger stage by drawing a random ordering of firms to act as merger proposers. Draw synergy values σ_{ij} for all pairs i, j.

7. Loop over all firms, for each solve for the best merger partner as $max_j W_k^M((\Omega'_k, i, j, \sigma_{ij})$. If the value of merging for both firms is higher than the option value of letting the next firm in the merger order proceed, they agree to merge, τ_{ij} is calculated, Ω'_k is updated to Ω_{k+1}, and the merger stage ends. Once this has happened or all firms have had a chance to propose, the stage ends.

8. Profits for all firms are calculated as $\pi(\Omega_{k+1})$.

9. Stored value functions are updated as:

$$W_{k+1}^I(\Omega_k, x_i^*) = \alpha^I(\Omega_k)\beta[\pi(\omega'_i) - FC + \mathbb{E}W_k^I(\Omega_{k+1})] + (1 - \alpha^I(\Omega_k))W_k^I(\Omega_k, x_i^*) \quad (10)$$

$$W_{k+1}^E(\Omega_k, xe_i^*) = \alpha^E(\Omega_k)\beta[\pi(\omega'_i) - FC + \mathbb{E}W_k^I(\Omega_{k+1})] + (1 - \alpha^E(\Omega_k))W_k^E(\Omega_k, xe_i^*)) \quad (11)$$

If firm i is acquired by firm j:

$$W_{k+1}^M(\Omega'_k, i, j, \sigma_{ij}) = \alpha^M(\Omega'_k, i, j, \sigma_{ij})\tau_{ij} + (1 - \alpha^M(\Omega'_k, i, j, \sigma_{ij}))W_k^M(\Omega'_k, i, j, \sigma_{ij}) \quad (12)$$

If firm i acquires firm j:

$$W_{k+1}^M(\Omega'_k, i, j, \sigma_{ij}) = \alpha^M(\Omega'_k, i, j, \sigma_{ij})[\mathbb{E}W_k^I(\Omega_{k+1}) - \tau_{ij}] + (1 - \alpha^M(\Omega'_k, i, j, \sigma_{ij}))W_k^M(\Omega'_k, i, j, \sigma_{ij}) \quad (13)$$

where $\alpha^I(\cdot)$, $\alpha^E(\cdot)$, and $\alpha^M(\cdot)$ are weighting functions to be described in detail below. In addition, counters $h_k^I\cdot)$, $h_k^E(\cdot)$, and $h_k^M(\cdot)$ are incremented by 1.

10. Return to step 1 at state Ω_{k+1}.

The algorithm is periodically paused to test for whether an equilibrium has been reached. This test follows Fershtman and Pakes (2012) and checks whether the value functions are consistent with equilibrium notions described in section 3.2. This simulates a sample path of the model and keeps a separate memory of the distribution of outcomes reached at each state on this sample path. The mean squared difference between these outcomes and the value function stored in memory is used to calculate bias. This is done separately for $W_k^I(\cdot)$, $W_k^E(\cdot)$, and $W_k^M(\cdot)$ and the highest bias value of the three is compared to .001 to determine if an equilibrium has been reached.

A key consideration in the algorithm is what function to use to weight realized payoffs in iteration k and how much to weight the current estimate of $W_k(\cdot)$. One alternative would be to simply use $\alpha^I(\cdot) = \frac{1}{h_k^I(\Omega_k)}$, ie the number of previous visits to that state. This ultimately would give value functions equal to the arithmetic mean of realized payoffs across all visits to that state. One problem with this approach is that if the initialized value functions are far from their true values, it could take a very long time for the algorithm to converge. This is because a large amount of weight would initially be placed on the incorrect initialized value function and increasingly smaller weight placed on the realized payoffs.

A second and more serious problem with the algorithm as described above is that for discrete choices such as entry and exit, it can get "stuck" at a suboptimal choice. For example, if the value functions for firm investment are set high in order to encourage exploration, firms will initially invest large amounts and rarely exit. At the same time as they make these choices, potential entrants are exploring entry strategies and updating their entry value function with the realized outcomes. In the case where incumbents invest highly entry is rarely profitable and so in certain states potential entrants may update all entry options as having negative value. Once they have done so, entry will cease at those states and as incumbent firms investment policies converge towards equilibrium the potential entrant will have stopped testing entry even though it may be profitable to do so.

I solve this second problem by implementing a strategy from the reinforcement learning literature known as ϵ-greedy exploration. In this case, firms will take what they perceive as being the optimal action with probability $1-\epsilon$ and with probability ϵ they will choose a policy at random from their set of possible actions. The researcher sets the initial value of ϵ to encourage exploration and as the algorithm proceeds it declines slowly to 0. In states in which firms take suboptimal policies as a result of this process, they update values for those policies but the other firms in the market do not update at those states.

This approach is simple but has the advantages that it ensures each action will be taken a large number of times and that policies ultimately converge to the optimal ones. In practice, without implementing this strategy the algorithm almost always converged to an outcome where a suboptimal discrete choice was taken and which was strongly rejected as an equilibrium outcome by the testing procedure. This would remain true regardless of how long it ran. Even if this extreme case of non-convergence were not possible, it is highly likely the ϵ-greedy exploration improves the speed to convergence of even a simple model.

A.1 Computational details

Here I provide specific details on the implementation of the algorithm. For incumbent and entrant value functions, I initialize the values above the level of discounted profits if the state they entered were permanent. A high initialization is useful for ensuring firms explore their strategy space early on. That is:

$$W_0^I(\Omega) = \frac{1.1\pi(\Omega)}{1-\beta}. \tag{14}$$

Merger value functions are set at a flat constant value of 5, which is high enough to encourage exploration. For all policies, ϵ is set initially at .1 and declines such that $\epsilon' = .9\epsilon$ every $200,000$ iterations.

The weighting functions used are:

$$\alpha^I(\Omega_k) = \frac{1}{min(h_k^I(\Omega_k, x_i), \bar{h}^I})} \tag{15}$$

where $\bar{h}$ is a cap on high $h_k(\cdot)$ for weighting purposes. This effectively places more weight on the more recent $\bar{h}$ observations. In practice, $\bar{h}$ begins at 100 and doubles every $1,000,000$ iterations until it ceases to bind. In all cases, the test concluded that an algorithm had been reached before at most 350 million iterations and in some cases much sooner.