

Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Mergers

Andrew Sweeting

Xuezhen Tao

University of Maryland and NBER*

Shanghai University of Finance and Economics

Xinlu Yao

University of Maryland

April 2018

Abstract

Existing theoretical and structural analyses of mergers in differentiated product markets assume that firms have complete information about both their own and their rivals' demand and marginal costs. While the complete information assumption simplifies the analysis, we show that equilibrium price levels may change significantly when there are small amounts of asymmetric information about serially-correlated marginal costs, as firms may have incentives to raise their prices in order to signal their costs, and therefore the likely level of their future prices, to competitors. We show that signaling effects could lead merger simulations that assume complete information to underpredict post-merger price rises quite significantly, and that they are also capable of explaining the magnitude of price increases often observed after actual mergers.

JEL CODES: D43, D82, L13, L41, L93.

Keywords: signaling, strategic investment, asymmetric information, oligopoly pricing, dynamic pricing.

*Contact: sweeting@econ.umd.edu. We thank Dan Vincent, Nate Miller and several seminar participants for useful comments. Carl Mela and Mike Krueger helped to provide access to the IRI academic data, and we benefited greatly from John Singleton's work with the beer data on another project. All errors are our own.

1 Introduction

Both theoretical and empirical analyses of mergers, as well as many other phenomena in Industrial Organization, typically assume complete information, i.e., that all firms know both their own and their rivals' demands and marginal costs exactly. While it is sensible that firms operating in the same industry will have good information about their rivals' competitive situation, it seems implausible that firms really know as much about these variables as their rivals do themselves. Indeed, the way that firm-level demand and cost information is treated as being highly confidential during the merger review process, and the merging parties themselves are often reluctant to share this information with each other before the merger is approved, suggests that asymmetries of information may have some significance.¹ A natural question is, therefore, whether oligopoly outcomes, and, in particular, the predictions of merger simulations, would change significantly if small amounts of asymmetric information were introduced. In this paper we answer this question in the context of a standard differentiated products oligopoly model where firms set prices, and we show that even very small amounts of asymmetric information about marginal costs can change equilibrium price levels by amounts that are large, policy-relevant and quite consistent with the magnitude of post-merger price increases that have been documented in the merger retrospectives literature.

We consider a fairly simple, finite-horizon dynamic model. In the simplest case, we consider a set of single-product oligopolists who simultaneously prices each period. Demand is determined by a static, nested logit demand model. It is common knowledge that each firm i 's marginal cost, c_i , will lie in a relatively narrow range, $[\underline{c}_i, \bar{c}_i]$, and that whatever value it currently takes, it evolves over time according to a positively serially correlated first-order Markov process. In this setting a firm's current price may signal its current marginal cost and therefore the level of prices it is likely to set in future periods. If there is a strategic incentive to signal that one's own price is likely to be higher in the future, this can provide a firm with an incentive to raise its price today, and, of course, if rivals are raising their current prices, this will tend to provide a firm with additional non-strategic incentives to raise its current prices when prices are strategic complements. We show that this cumulative effect can raise prices some periods from the end of the game quite substantially. For instance, in a duopoly example, we find that uncertainty about less than 1% of marginal costs can raise equilibrium prices by as much as 17%, even with fairly limited serial correlation. The large size of the effect reflects, in

¹It is standard for the antitrust authorities not to publicly release proprietary data that they gather during the review process publicly or to share it with any parties to the investigation. A Wall Street Journal article from May 23, 2016, describes how Anthem and Cigna, two health insurers whose merger is currently under review would not share information on cost-of-care or salary expenses because the firms would remain competitors if the merger was not approved, even though sharing this kind of data could be critical to building a compelling case for the merger.

part, the cumulative effects just suggested but also the fact that around a firm's static best response the firm's current period expected profit function is, of course, flat in its own price. This can make quite significant price increases valuable if they raise rivals future prices by even quite small amounts. However, it is important to note that what is raising prices is the combination of strategic signaling incentives, created by the serial correlation, and uncertainty. In our specifications, uncertainty about costs alone has almost no effects on price levels.

As usual with strategic incentives, they become stronger when there are fewer firms, as incentives to free-ride on the strategic behavior of other firms disappears. As a consequence, mergers, unless they lead to monopoly, will tend to strengthen strategic incentives and, therefore, raise prices by more than a static, complete information model with similar parameters would predict, and because we assume a finite horizon structure, prices identical to static Nash prices are the only prices that could be supported as equilibria in a complete information version of our dynamic model. Equivalently, a merger simulation based on complete information might underestimate the merger-specific synergy required to prevent prices from rising after the merger substantially. We illustrate the differences between complete information and asymmetric information predictions for mergers by considering both abstract examples, and examples based on trying to understand what happened after the 2008 joint venture between SAB Miller and Coors brewing companies, which effectively merged their US brewing and marketing assets. In recent work, Miller and Weinberg (2016) show that this JV was followed by significant, and nearly equal, price increases in the leading brands marketed by both MillerCoors and Anheuser-Busch, which was not a party to the transaction, even though MillerCoors was able to realize significant synergies, especially in transportation costs (Ashenfelter, Hosken, and Weinberg (2015)). A pattern where non-merging parties increase prices as much as merging ones is difficult to generate in a static Nash model with complete information, suggesting that there may have been some qualitative change in behavior after the JV. In this paper, we show that the increase in signaling incentives could generate the magnitude of price increases that are observed for a small amount (less than 2%) of uncertainty in marginal costs. This provides an alternative to explanations based on collusion in a complete information framework.

While we believe that we are the first to explore how large the effects of signaling could be in oligopoly models, and the quantitative implications of signaling for mergers, our approach builds on theoretical work by Mailath (1988), Mailath (1989) and Mester (1992).² Mailath (1988) considers conditions under which a separating equilibrium will exist in an abstract two-period game with simultaneous signaling by multiple players, and shows how the conditions required for uniqueness of each player's separating

²Caminal (1990) considers a two-period linear demand duopoly model where firms have private information about the demand for their own product, and also raise prices to signal that they will set higher prices in the final period.

best response function are similar to those required in classic single-agent signaling models. Mailath (1989) applies these results to consider a two-period pricing game with static linear demands where firms marginal costs are fixed but unknown, and where firms may raise their prices in the first period in order to try to raise their rivals' prices in the second period. He shows that, with linear demand, there will exist a unique, fully separating equilibrium and that, qualitatively, signaling will raise equilibrium prices. Mester (1992) considers a three-period quantity setting model where marginal costs can vary over time. Demand is linear. In this model, the fact that costs can change, mean that there is an incentive to signal in both the first and second periods, and Mester shows that qualitatively each firm produces more than the output that maximizes its current profits in the first period, reflecting the fact that quantities are typically strategic substitutes so that firms benefit when their rivals believe that their future marginal costs are likely to be small.

We build on this literature by extending the price-setting model to more periods (indeed, we will focus on the implications for prices in the early periods of a long finite horizon game), considering more standard non-linear demands for differentiated products, and focusing on the quantitative and policy implications of the results. However, because we use non-linear demands we are no longer able to prove, in general, existence or uniqueness of the equilibrium we look at. Indeed, we find that our equilibrium may no longer exist once signaling effects become too large. Instead we take a computational approach where we verify that the conditions required for the existence and separation of best responses are satisfied at each step of the solution process. Bonatti, Cisternas, and Toikka (2015) provide some recent work on oligopoly signaling using an elegant continuous-time model of a Cournot oligopoly where firms have private information about their own marginal cost, which is fixed over time, and only observe market prices, which are affected by unobserved demand shocks as well as each firm's output. They characterize both strategies and signaling and learning incentives when firms use symmetric linear Markov strategies. It would be interesting to extend our model to also incorporate noisy information transmission as, in practice when firms are able to offer discounts, it will often be difficult to rivals to observe prices exactly.

The paper is also connected to three other literatures. Kaya (2009), Toxvaerd (2014), Roddie (2012) and Sweeting, Roberts, and Gedge (2016) consider dynamic single-agent signaling models. Sweeting, Roberts, and Gedge (2016) use their model to show that limit pricing may provide a potential explanation for why incumbent legacy carriers lowered prices when faced by the threat of entry by Southwest.³ The advantage of single-agent signaling models is that it is more straightforward to relate the conditions

³In an on-going revision they incorporate two-sided asymmetric information, but not two sided signaling, into their model.

required for existence and uniqueness back to primitives of the model. On the other hand, limit pricing stories are only likely to be really plausible in a limited number of settings where there is a dominant incumbent but also a real threat of entry, whereas oligopoly pricing interactions are common to most differentiated product markets considered in IO.

A second literature has considered dynamic games of asymmetric information in applied settings. In particular, Fershtman and Pakes (2010) propose a tractable notion of Experience Based Equilibrium for games with discrete states, discrete choices and potentially limited recall of past states. Pooling is an almost necessary feature with discrete states and discrete actions. We consider a game with continuous states and actions, and focus only on separating equilibria. Some exploratory work, however, indicates that pooling might raise prices even higher, and it would be necessary to consider some degree of pooling if we introduced menu costs into the model.

The third literature is the recent literature on merger retrospectives (*inter alia* Borenstein (1990), Kim and Singal (1993), Peters (2009), Ashenfelter, Hosken, and Weinberg (2015) and Miller and Weinberg (2016)). Ashenfelter, Hosken, and Weinberg (2014) identify 36 of 49 studies as finding significant post-merger price increases when a merger increases concentration significantly, across a wide range of industries with the possible exception of banking. Given that the antitrust authorities try to prohibit mergers that will raise prices, this is a striking result, and it suggests that the conventional ex-ante analysis of mergers, which assumes complete information, either in the first-order conditions that underlie UPP calculations, or in more formal merger simulations (following Nevo (2000)), may be missing some ability or incentive to raise prices that becomes stronger after a merger. One explanation is that the remaining firms are more able to engage in some degree of tacit collusion after a merger. We view signaling as an alternative explanation, although it might also be considered under a broad interpretation of the types of coordinated effects described in the current *Horizontal Merger Guidelines*.

While we do not attempt to run a horse race between collusion and signaling theories here, it is worth noting that there are both similarities and differences between these theories.⁴ In both theories, firms raise prices above static best response levels in order to raise their rivals' future prices and both theories require that prices can be observed quite accurately in order for a firm not to want to return to a best response. On the other hand, signaling effects require some degree of asymmetric information, so that there clearly different implications for information sharing arrangements; it is unlikely that it would be possible to sustain joint-profit maximizing prices even when firms are patient; and, the model will

⁴Tacit collusion models may provide a better explanation for what is observed in some settings. For example, Sweeting (2007) found that leading generators in the England and Wales wholesale electricity Pool expanded their output aggressively when it was announced that the Pool would be replaced by a new trading system in several months time. This is consistent with a finite-time horizon substantially limiting collusive incentives. In our signaling model, signaling incentives can have large effects even in quite a short game.

imply qualitatively different effects when firms set quantities rather than prices. Indeed, Mester (1987) argued that a perpetual signaling quantity-setting model could explain why facts such as multi-market contact that are often associated with higher prices, seemed to be associated with more competitive outcomes in banking. Interestingly, banking is also an industry where evidence for post-merger price increases is weak (Ashenfelter, Hosken, and Weinberg (2014)).⁵

The remainder of this draft is structured as follows. Section 2 lays out the model and the equilibrium that is studied. Section 3 uses an example with a variable number of symmetric firms to show that price effects can be really large, and that a stylized merger analysis that assumes complete information could give misleading conclusions. Section 4 provides our analysis of the Miller-Coors joint venture, building off the analysis in Miller and Weinberg (2016). Future revisions will contain additional examples. Section 5 concludes.

2 Model

2.1 Set-Up

We consider the following model. There a finite number of discrete time periods, $t = 1, \dots, T$, and a common discount factor $0 < \beta < 1$. The emphasis will be on games where T is large, and we will focus on strategies that appear stationary, or very close to stationary, some periods from the end of the game. There are N firms, and no entry and exit. Firms are assumed to be risk-neutral and to maximize the current discounted value of current and future profits. The marginal costs of firm i can lie on a compact interval $[\underline{c}_i, \bar{c}_i]$, and evolve, exogenously, from period-to-period according to a first-order Markov process, $\psi_i : c_{i,t-1} \rightarrow c_{i,t}$ with full support (i.e., $c_{i,t-1}$ can evolve to any point on the support in the next period).⁶ We will think of the range $\bar{c}_i - \underline{c}_i$ as being a measure of how much uncertainty there is about costs. The conditional pdf is denoted $\psi_i(c_{i,t}|c_{i,t-1})$.

Assumption 1 *Marginal Cost Transitions*

1. $\psi_i(c_{i,t}|c_{i,t-1})$ is continuous and differentiable (with appropriate one-sided derivatives at the boundaries).

⁵There are potentially other differences between the models as well. For example, if there is some asymmetry of information then a tacitly collusive arrangement might lead to prices being quite rigid as firms would be reluctant to ever cut prices if doing so would lead to a price war. On the other hand, the signaling model can potentially explain prices (including for example, observable discounts and promotions) that move quite frequently without starting price wars.

⁶We have also considered models where it is the intercept of an increasing marginal cost curve that is uncertain, and we also find large price-increasing effects in this case. Increasing marginal costs can also help to relax some of the problems that arise in satisfying the conditions required for fully-separating best responses as the incentive to undercut when a rival has a very high price are softened when a firm's marginal cost is increasing in its own output.

2. $\psi_i(c_{i,t}|c_{i,t-1})$ is strictly increasing i.e., a higher cost type in one period implies a higher cost type in the following period will be more likely. Specifically, we will require that for all $c_{i,t-1}$ there is some c' such that $\frac{\partial \psi_i(c_{i,t}|c_{i,t-1})}{\partial c_{i,t-1}}|_{c_{i,t}=c'} = 0$ and $\frac{\partial \psi_i(c_{i,t}|c_{i,t-1})}{\partial c_{i,t-1}} < 0$ for all $c_{i,t} < c'$ and $\frac{\partial \psi_i(c_{i,t}|c_{i,t-1})}{\partial c_{i,t-1}} > 0$ for all $c_{i,t} > c'$. Obviously it will also be the case that $\int_{\underline{c}_i}^{\bar{c}_i} \frac{\partial \psi_i(c_{i,t}|c_{i,t-1})}{\partial c_{i,t-1}} dc_{i,t} = 0$.

The increasing nature of the transition may provide a firm with an incentive to signal that it has a high cost in the current period if this will cause other firms to raise their prices in future periods. If there was negative serial correlation then the strategic incentives would change significantly, but positive serial correlation seems far more plausible in practice. The transition is assumed to be independent across firms, although one could also allow for a common, observed and time-varying component of marginal costs, and it would be interesting to consider, for example, how the introduction of asymmetric information would affect the cost-pass through of common shocks in oligopoly, given the large effects that asymmetric information has on mark-ups.⁷

In each period t , timing is as follows.

1. Firms enter the period with their marginal costs from the previous period, $t - 1$. These marginal costs then evolve exogenously according to the processes ψ_i .
2. Firms simultaneously set prices, and there are no menu costs preventing price changes.⁸ A firm's profits are given by

$$\pi_{i,t} = (p_{i,t} - c_{i,t}) Q_{i,t}(\underline{p}_t)$$

where $Q_{i,t}(\underline{p}_t)$ is a static demand function and \underline{p}_t is the vector of all firms prices. In our examples and application we will use nested logit demand. When making its price choice, a firm observes its own marginal cost, but not the current or previous marginal cost of other firms.⁹ It is, however, able to observe the complete history of prices in previous periods. Formally we will assume that prices are chosen from some compact support, $[\underline{p}, \bar{p}]$ where the bounds are wide enough to satisfy support conditions.¹⁰ Note that in the absence of menu costs, a firm's price only has a direct effect on its current period profits. This separability will be helpful in interpreting the conditions required for the existence and uniqueness of separating best responses.

⁷Nakamura and Zerom (2010) consider the effect of menu costs on pass-through in an oligopoly setting. They find that menu costs have little effect on the level of prices but do affect how quickly costs are passed-through.

⁸With menu costs, one would expect some pooling where firms with different cost realizations choose the same price. These may be difficult to analyze.

⁹Given that we consider a fully-separating equilibrium in each period and a first-order Markov process with full support for costs, everything would still work as presented if a firm was able to observe its rival costs with a delay of two periods.

¹⁰In particular, the lower support needs to be below prices that the firms might ever want to charge if they were pricing statically, and the upper bound needs to be so high that no firm would ever want to charge it whatever effects it could have on the beliefs of rivals.

2.2 Equilibrium

Under complete information, there would be a unique subgame perfect Nash equilibrium where each firm sets its static Nash equilibrium price, given the realization of costs, in every period as long as the equilibrium in the static game is unique, as will be the case with single-product firms and constant, with respect to quantity, marginal costs under most commonly-used demand systems such as logit or nested logit. In a one-period asymmetric information game, firms would play a static Bayesian Nash equilibrium (BNE) where each firm maximizes its profits given its prior beliefs about the distribution of other firms' marginal costs, and the strategies that those firms are using. As we will illustrate below, when $\bar{c}_i - \underline{c}_i$ is small, average BNE prices will tend to be very close to average complete information Nash prices.¹¹

In the dynamic game with asymmetric information, we assume that firms play a Markov Perfect Bayesian Equilibrium (MPBE) (Roddie (2012), Toxvaerd (2008)). This requires, for each period:

- a time-specific pricing strategy for each firm as a function of its contemporaneous marginal cost, its beliefs about the marginal cost of the other firms, and what it believes to be those firms' beliefs about its own marginal costs; and,
- a specification of each firm's beliefs about rivals' marginal costs given all possible histories of the game, which here means the prices that other firms have set.

Note that in this equilibrium history can matter, even though it is only the current costs of firms that are directly payoff-relevant, because observed history can affect beliefs about rivals' current costs, and these beliefs are directly relevant for expected current profits. We will assume that, following any history of prices, all rivals will have similar beliefs about a firm's marginal cost.

2.2.1 Final Period

In the final period, each firm price will use static Bayesian Nash equilibrium strategies given their beliefs, so that they maximize their expected final period profits, as there are no future periods to be concerned about. Considering the duopoly case ($N = 2$) for simplicity, if firm i believes that firm j 's $T-1$ marginal cost is distributed with a density $g_{j,T-1}^i(c_{j,T-1})$, then it will set a price $p_{i,T}^*(c_{i,T}, g_{j,T-1}^i(c_{j,T-1}))$ as

$$p_{i,T}^*(c_{i,T}, g_{j,T-1}^i(c_{j,T-1})) = \arg \max_{p_i} (p_i - c_{i,T}) \int_{\underline{c}_j}^{\bar{c}_j} \int_{\underline{c}_j}^{\bar{c}_j} Q_{i,t} \left(\frac{p_i}{p_{j,T}^*(c_{j,T})} \right) \psi_j(c_{j,T} | c_{j,T-1}) g_{j,T-1}^i(c_{j,T-1}) dc_{j,T-1} dc_{j,T}$$

¹¹Shapiro (1986) qualitatively compares a complete information oligopoly outcome, modeled as being played when firms share cost information via a trade association, with an incomplete information outcome, showing that complete information tends to lower expected consumer surplus, while raising firm profits and total efficiency.

where $p_{j,T}^*(c_j)$ is the pricing function for firm j given its marginal cost, implicitly conditioning on its beliefs about i . Note that, unlike in Mailath (1989) where costs are fixed over time, final period prices will not be exactly identical to complete information prices even if equilibrium play in the previous period has fully revealed all firms' marginal costs, so in that case $g_j^i(c_{j,T-1})$ would have all of its mass at a single point, as the stochastic innovations in j 's marginal cost, represented by the $\psi_j(c_{j,T}|c_{j,T-1})$ function, are j 's private information.

Given equilibrium strategies, conditioned on beliefs, we can define the value of each firm at the beginning of the final period, before marginal costs have evolved to their current values. For example, again in the duopoly case, $V_{i,T}(c_{i,T-1}, g_{j,T-1}^i, g_{i,T-1}^j)$ where the second term reflects i 's beliefs about j 's costs (which may depend on historical pricing) and the final term reflects j 's beliefs about i 's costs, which should affect j 's equilibrium pricing. We will assume that for any set of beliefs, there is a unique final period BNE pricing equilibrium.¹²

2.2.2 Penultimate Period, $T - 1$

In the penultimate period, firms may want to not only maximize their current period profits, but also signal information to their rivals about what their costs are likely to be in the final period. We write the so-called ‘‘separating signaling payoff function’’ of firm i , at the time when it is making its pricing choice (so it knows its $T - 1$ marginal cost), in the duopoly case, as $\Pi^{i,T-1}(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i})$ where the second term $\left(\widehat{c_{i,T-1}^j}\right)$ represents a point belief that j will have about i 's $T-1$ cost at the beginning of the **next** (final) period. The fourth term reflects the pricing strategy that i expects j to use in period $T - 1$ which will reflect i 's beliefs about j 's prior marginal cost (contained in $g_{j,T-2}^i(c_{j,T-2})$), as well as j 's pricing strategy. Writing i 's expected payoffs in this way is convenient when expressing conditions for i 's best response function, incorporating any signaling incentives, to be well-behaved.

To be more explicit about the form of $\Pi^{i,T-1}$, assume that j 's $T-1$ pricing strategy is fully separating so that i will be able to infer j 's $T - 1$ cost exactly when entering period T . Then,

$$\begin{aligned} \Pi^{i,T-1}(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i}) &= (p_{i,T-1} - c_{i,T-1}) \times \dots \\ &\int_{\underline{c_j}}^{\overline{c_j}} \int_{\underline{c_j}}^{\overline{c_j}} \left\{ Q_{i,t} \left(\begin{array}{c} p_{i,T-1} \\ p_{j,T-1}^*(c_{j,T-1}) \end{array} \right) + \beta V_{i,T}(c_{i,T-1}, c_{j,T-1}, \widehat{c_{i,T-1}^j}) \right\} \psi_j(c_{j,T-1}|c_{j,T-2}) g_{j,T-2}^i(c_{j,T-2}) dc_{j,T-2} dc_{j,T-1} \end{aligned} \quad (1)$$

Note that because $p_{i,T-1}$ has no effect on period T payoffs holding j 's belief at the end of period $T-1$ fixed, the signaling payoff function is separable in price across periods, as in Mailath (1988) and Mailath (1989).

¹²Existing uniqueness results are proved for complete information. However, as the introduction of incomplete information tends to smooth reaction functions, it is reasonable to believe that uniqueness would carry over to models with a small degree of asymmetric information about marginal costs.

We will focus on a fully separating equilibrium, if one exists, so that each firm's pricing decision exactly reveals its marginal cost to the other firms. While there may be other equilibria that involve some degree of pooling, Mailath (1989) argues that, if it exists, the separating equilibrium is the natural one to look at. Following Mailath (1989), under a set of conditions on firms' separating signaling payoff functions, to be described in a moment, the equilibrium strategies can be characterized as follows.

Characterization of Strategies in a Period $T - 1$ Separating Equilibrium. Holding beliefs about j 's pricing fixed, each firm's best response pricing strategy will be given as the solutions to a set of differential equations where

$$\frac{\partial p_{i,T-1}^*(c_{i,T-1})}{\partial c_{i,T-1}} = - \frac{\Pi_2^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)}{\Pi_3^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)} > 0 \quad (2)$$

where the subscript n in $\Pi_n^{i,T-1}$ means the partial derivative of the separating signaling payoff function with respect to the n^{th} argument; and, an initial value condition, where $p_{i,T-1}^*(\underline{c}_i)$ is the solution to

$$\Pi_3^{i,T-1} \left(\underline{c}_i, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) = 0 \quad (3)$$

i.e., it should be used a static best response to j 's pricing policy. Note that given separability in how prices and beliefs enter the separating signaling payoff function, the numerator in the differential equation represents the future benefit from raising a rival's belief, and the denominator is the effect of an increase in price on current profit. Above the static best response price, the denominator will be negative.

Given these strategies, a firm that observes firm i setting a price $p_{i,T-1}$ will infer i 's $T - 1$ marginal cost by inverting the pricing function if the price is within the range of the solution given by the differential equation. If it is outside the range of the pricing function, we assume, for completeness, that the other firms infer that $c_i = \underline{c}_i$ (i.e., they infer the lowest possible cost), although other configurations of beliefs could be used.

Firms' best response functions will be unique and strictly increasing under the following conditions on their signaling payoffs (assuming that support conditions on prices are satisfied).

Condition 1 For any $(c_{i,T-1}, \widehat{c_{i,T-1}^j}, \widehat{\zeta_{j,T-1}^i})$, $\Pi_2^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)$ has a unique optimum in $p_{i,T-1}$, and, for all $c_{i,T-1}$, for any $p_{i,T-1} \in [\underline{p}, \bar{p}]$ where $\Pi_3^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) >$

0, there is some $k > 0$ such that $\left| \Pi_3^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) \right| > k$.

Remark Given separability, this is a condition that each firm's static profit function should be well-behaved, e.g., strictly quasi-concave, given the expected pricing of rivals.

Condition 2 *Type Monotonicity*: $\Pi_{13}^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) \neq 0$ for all $(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1})$.

Remark Given separability of the payoff function, this condition simply requires that it is always less expensive, in terms of forsaken current profits, for a higher marginal cost firm to raise its price, which is natural as a lost unit of output will be less costly when the firm's margin is smaller.

Condition 3 *Belief Monotonicity*: $\Pi_2^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) \neq 0$ for all $(c_{i,T-1}, \widehat{c_{i,T-1}^j}$.

Remark In our context, this condition requires that a firm should always benefit when its rivals believe that it has a higher marginal cost. In a two-period price setting game this condition is natural as a rival's final period best response price will tend to increase if it believes one of its rivals' marginal costs is higher. However, this condition is not necessarily satisfied with general demand forms when future prices are far above static best response levels, as a rival may have a stronger incentive to lower its price towards its static best response when i 's expected future price rises.

Condition 4 *Single Crossing*: $\frac{\Pi_3^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)}{\Pi_2^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)}$ is a monotone function of $c_{i,T-1}$ for all $\widehat{c_{i,T-1}^j}$ and for relevant $p_{i,T-1}$ above the static best response price.

Remark This condition implies that a firm with a higher marginal cost should always be willing to increase its price slightly more than a firm with a lower marginal cost in order to increase the belief of rivals about its cost by the same amount. Whether this will be satisfied will depend on the exact parameters of the model, including the degree of serial correlation about costs and the length of the support of costs, because it is quite possible that a firm with lower current marginal costs will actually benefit more from raising its rivals' prices in the future, even if it is giving more up in terms of current profits, because it expects its margins in future periods to be larger.

As discussed in Mailath (1988), these conditions parallel those in a single-agent signaling problem. In that paper, Mailath also solved a technical problem to prove that conditions on best responses could also imply that a fully-separating **equilibrium** exists.¹³ Unfortunately, two limitations are associated with these conditions. First, it is difficult to express them in terms of cost and demand primitives, so that, even to show uniqueness of best responses it is necessary to verify that they hold while computing the equilibrium recursively. Second, they do not guarantee uniqueness of an equilibrium because they only imply uniqueness of a best response conditional on other firms' strategies. In a two-period model with linear demand, Mailath (1989) was able to overcome these problems to show equilibrium uniqueness (within the class of fully separating equilibria), as was Mester (1992) in the case of a three period, quantity-setting duopoly model, also with linear demand. In order to examine more realistic demand settings it is necessary to forsake proving equilibrium uniqueness, and the results that follow will be conditional on the method used to solve for the equilibrium. This being said, the iterative algorithm described below appears to converge to the same fully-separating solution from several different sets of starting points for several sets of parameters that we have tried. We have not tried to solve for pooling or partial pooling equilibria: while in a dynamic single agent signaling model it is possible to eliminate pooling equilibria under similar conditions by applying a refinement (e.g., Sweeting, Roberts, and Gedge (2016)), this is not generally possible with several signaling firms, even with linear demand (Mailath (1989)).

2.2.3 $T - 2$ and Earlier Periods

Now consider period $T - 2$. If equilibrium play in $T - 1$ is known to have the fully separating form just described, then the beginning-of-period $T - 1$ values, $V_{i,T-1}(c_{i,T-2}, g_{j,T-2}^i, g_{j,T-2}^i)$ can be calculated. Given these continuation values, we can then apply the same logic as in $T - 1$ to derive the form of best-response pricing strategies in a separating $T - 2$ equilibrium. $V_{i,T-2}(\cdot)$ can then be calculated, and the same procedure applied to $T - 3$, etc..

The calculation of the signaling payoff functions in each period assumes that pricing behavior in the previous period was separating. This raises the question of what is believed in the first period of the game. Note that results from $t = 2$ onwards are not affected by our assumptions on this point as long as whatever we assume about beliefs in the beginning of the first period supports first period separation. One possibility is that we assume that firms know rivals' marginal costs in some fictitious $t = 0$ period. Alternatively we could assume that each player assumes a rival's marginal cost is being drawn from

¹³Mailath (1987) laid out the conditions for a unique separating signaling strategy in a single-agent model. Mailath and von Thadden (2013) present a more tractable version of the required conditions that are more straightforward to check in applications.

some suitable distribution. In any case, we do not focus on what happens in the first period of the game when interpreting the results. Instead, we simply solve the game backwards for a large number of periods, and focus on prices that appear to be (very close to) stationary across periods once we have gone far enough back.¹⁴

2.3 Computation

We use the following computational steps to solve the model. In the case where firms are symmetric it is possible to ignore the ‘repeat for each firm’ steps that are described below.

2.3.1 Preliminaries

We start by specifying discrete vectors of points for the actual and for the perceived marginal costs of each firm (we will use interpolation and numerical integration to deal with the fact that actual costs will likely lie between these isolated points). For instance, in the symmetric example below, each firm’s marginal cost will lie on $[8, 8.075]$ and we will use 10 equally spaced points $\{8, 8.0083, 8.0167, 8.0250, 8.0333, 8.0417, 8.0500, 8.0583, 8.0667, 8.0750\}$.¹⁵ As the number of players expands to four or more, one has to reduce the number of points considered for each firm in order to prevent the computation time growing too rapidly, especially when firms are asymmetric.

2.3.2 Period T

Assuming that play at $T - 1$ has been fully separating, we solve for BNE pricing strategies for each possible combination of beliefs about firms costs entering the final period. A strategy for each firm is an optimal price given each realized value of its own cost on the grid, given the pricing strategy of each firm.¹⁶ Trapezoidal integration is used to integrate expected profits over the gridpoints given the pdf of each firm’s cost transitions. We then use these strategies to calculate $V_{i,T}(c_{i,T-1}, c_{j,T-1}, \widehat{c_{i,T-1}^j})$ (assuming the duopoly case for simplicity of exposition) for each firm where we are allowing for the possibility that i ’s actual $T - 1$ marginal costs are different from those perceived by firm j .¹⁷

¹⁴Once strategies have converged, one can also examine whether these strategies would form a stationary MPBE in the infinite horizon game. For the examples we have studied, this has been the case. However, there are some cases with quite large price increases, where pricing strategies seem not to quite converge but instead to cycle slightly. More analysis of these cases is required.

¹⁵For the $N = 2$ example below, the strategies differ by less than one cent (compared to marginal costs of around \$8 if we use 20 or 40 gridpoints, although computation time increases exponentially).

¹⁶So, for example, in the duopoly case, for a given pair of beliefs about each firm’s marginal cost, we have to solve for 20 prices (1 for each realized cost gridpoint for each firm).

¹⁷Under duopoly with a 10 point actual and perceived cost grid, $V_{i,T}$ is stored as a 10 x 10 x 10 array.

2.3.3 Period $T - 1$ (and earlier steps)

In $T - 1$ we use the following procedure.

Step 1. (a) compute $\beta \frac{\partial V_{i,T}(\widehat{c_{i,T-1}}, \widehat{c_{i,T-1}^j}, \widehat{c_{i,T-1}^i})}{\partial \widehat{c_{i,T-1}^j}}$ by taking numerical derivatives at each of the grid-points. This array provides us with a set of values for the numerator in the differential equation (2) ($\Pi_2^{i,T-1}$) as, because of separability, it does not depend on period $T - 1$ prices.

(b) We verify belief monotonicity using these derivatives.

Step 2. For each set of beginning of period point beliefs about each firm's previous period marginal costs on the grid, ($\widehat{c_{i,T-2}^j}, \widehat{c_{i,T-2}^i}$), where we are implicitly assuming separating play in the previous period, we use the following iterative procedure to solve for equilibrium fully separating prices. For simplicity of exposition, we will assume duopoly.¹⁸

(a) Use BNE prices (i.e., those calculated in period T) as an initial guess. Set the iteration counter, $iter = 0$.

(b) Given the current guess of the strategy of firm j , calculate the derivative of expected current flow profits with respect to i 's price on a fine grid of prices, which extends above the maximum current guess of prices. In the example below we use a 0.01 steps for prices when the average price is around 20. This vector will be used to calculate the denominator in the differential equation ($\Pi_3^{i,T-1}$).¹⁹

(c) We verify single-crossing and type monotonicity properties of the payoff function at the cost and price gridpoints.

(d) Solve $\Pi_3^{i,T-1}(\underline{c_i}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{c_{j,T-1}^i}) = 0$ to find the lower boundary condition for i 's pricing function (using a cubic spline to interpolate the vector calculated in (b)).²⁰

(e) Using the boundary condition as the starting point of the pricing function when $c_{it} = \underline{c_i}$, solve the differential equation to recover i 's best response pricing function. This is done using `ode113` in MATLAB.²¹ We then use cubic spline interpolation to get values for the pricing function at the points on the cost grid.

¹⁸We do not claim that this iterative procedure is computationally optimal, although it works well in our examples which have relatively small scale. There are some parallels between our problem and the problem of solving for equilibrium bid functions in asymmetric first-price auctions where both the lower and upper bounds of bid functions are endogenous. Hubbard and Paarsch (2013) provides a discussion of the types of methods that are used for these problems.

¹⁹A fine grid is required because it is important to evaluate it accurately around the static best response, where the derivative will be equal to zero.

²⁰In practice, the exact value of the derivative will be zero at the static best response, so that the differential equation will not be well-defined if this derivative is plugged in. We therefore solve for the price where $\Pi_3^{i,T-1} + 1e - 4 = 0$, and use this as the starting point. Pricing functions are essentially identical if we use $1e - 5$ or $1e - 6$ instead.

²¹In our example, we use an initial step size of $1e-4$ and a maximum step size of 0.005, when prices are in the range of 18 to 28.

(f) update the current guess of i 's pricing strategy using the updating formula:

$$p_{i,t}^{iter=1} = p_{i,t}^{iter=0} + \frac{1}{1 + iter^{\frac{1}{6}}} p'_{i,t}$$

where $p'_{i,t}$ is the best response price that has just been found.

(g) Repeat for each firm as required by asymmetry.

(h) Update the iteration counter to $iter = iter + 1$.

(i) Repeat steps (b)-(h) until the price functions change by less than $1e - 6$ at every point on the price grid.

Step 3. Compute beginning of period values,

$$V_{i,T-1} \left(c_{i,T-2}, \widehat{c_{i,T-2}^j}, \widehat{c_{j,T-2}^i} \right) = \dots$$

$$\int_{\underline{c}}^{\bar{c}} \int_{\underline{c}}^{\bar{c}} \left\{ \begin{array}{l} \pi_i(\zeta_{i,t}^*(c_{i,T-1}), \zeta_{j,t}^{i*}(c_{j,T-1})) + \dots \\ \beta V_{i,T}(c_{i,T-1}, c_{i,T-1}, c_{j,T-1}) \end{array} \right\} \psi_i(c_{i,T-1} | c_{i,T-2}) \psi_i(c_{j,T-1} | \widehat{c_{j,T-2}^i}) dc_{j,T-2} dc_{j,T-1}$$

This process is then repeated for earlier periods. The results in this version of the paper are computed using games where $T = 25$, or $T = 30$ in cases where strategies had not converged so that prices changed by less than one cent at the beginning of the $T = 25$ game. We have also solved several examples with $T = 50, 100$ and 200 periods to verify that strategies do not change when we extend the game.

3 Example

We now consider an example, which we present with several objectives in mind: (i) to illustrate the solutions to the model; (ii) to show that the pricing effects of asymmetric information can be very large in a dynamic model, and to develop intuition for why this is the case; (iii) to provide some simple examples of how merger counterfactuals that ignore the effects of asymmetric information may go astray; and (iv) to give some intuition about why the conditions required to characterize best response functions that were laid out above can fail for the types of demand that we consider.

3.1 Parameterization

We assume the following parameterization with $N = 2, \dots, 4$ single-product symmetric firms. We restrict ourselves to single product firms because multi-dimensional signaling problems are generally intractable

even when there is a single informed agent. In the application we will introduce multi-product firms but in a restrictive way which allows us to sidestep these problems. The discount factor β is 0.99, which is consistent with firms setting prices every 1-2 months. The marginal costs of each firm lies on the interval $[8, 8.075]$ (so the range of costs is less than 1% of the mean level of costs), and they evolve according to independent AR(1) processes where

$$c_{i,t} = \rho c_{i,t-1} + (1 - \rho) \frac{\underline{c} + \bar{c}}{2} + \eta_{i,t}, \text{ where } \rho = 0.8.$$

The distribution of η is truncated so that marginal costs remain on their support, and the underlying non-truncated distribution is assumed to be normal with mean zero and standard deviation 0.025. Given the moderate value of serial correlation, and the reasonably large standard deviation of the innovations relative to the possible range of costs, firm costs can change quite quickly from high to low values, or vice-versa. We discuss what happens with more serial correlation below.

Demand has a single one-level nested logit structure, where the single-products of the N firms are all included in a single nest (the other nest contains only the outside good). Indirect utility for a *person* choosing good i is

$$u_{person,i} = 5 - 0.1p_i + \sigma \nu_{person, nest} + (1 - \sigma) \varepsilon_{person,i} \tag{4}$$

with the nesting parameter, $\sigma = 0.25$. As usual the indirect utility of the outside good is normalized to

$$u_{person,0} = \varepsilon_{person,0}.$$

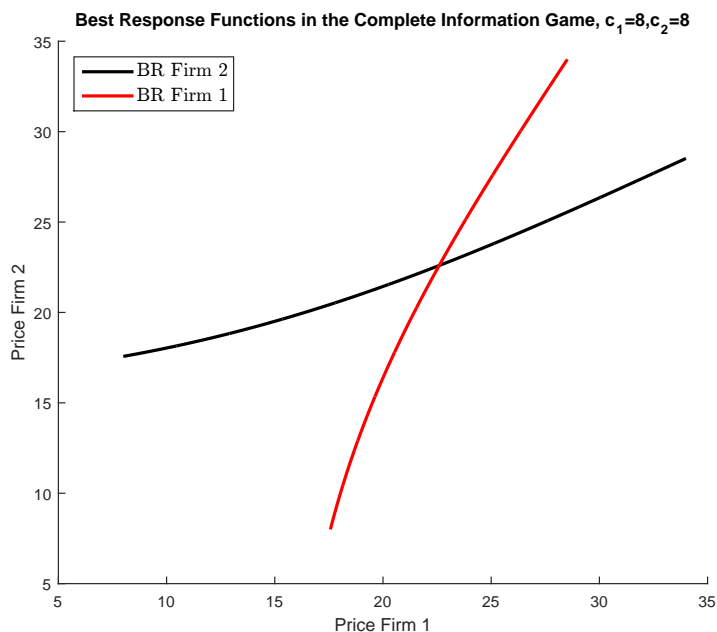
We use nested logit here because it provides a transparent way to separately control the elasticity of firms' residual demand curves and the degree of substitution to the outside good. However, for the results presented in this draft we could equally well use a multinomial demand with suitable chosen parameters.

3.2 Analysis with $N = 2$

An important feature of demand in this example is that the included goods effectively "cover the market" so that their combined market shares are close to 1, even when there are only two firms. In this sense the example resembles an auction with no reserve price, or the type of demand system focused on by Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) who also examine dynamic strategic interactions in a pricing game (in their case a complete information model with learning-by-doing

and forgetting, although this would also be an interesting setting in which to introduce asymmetric information). As a result, demand gained by one firm is largely being taken from its rival. On the other hand, the price parameter is quite small so that mark-ups are quite high: the average complete information Nash equilibrium price is 23.63. Figure 1 illustrates this by showing the firms' reaction functions when they both have, and are known to have, the lowest marginal cost of 8. Each firm's optimal price is quite sensitive to the price charged by its rival.

Figure 1: Duopoly Reaction Functions in the Static Complete Information Game, with $c_1 = 8$ and $c_2 = 8$



Consider strategies in the final period T , assuming that strategies in $T-1$ are fully revealing. Figure 2 shows the BNE pricing strategies for firm 2, when $\widehat{c_{2,T-1}} = 8$, for different values of $\widehat{c_{1,T-1}}$. As one would expect, the firm 2's optimal price is increasing in $\widehat{c_{1,T-1}}$ for any realization of $c_{2,T}$. However, the small scale on the y-axis reflects the fact that with limited cost uncertainty, the range of prices that can be observed with BNE pricing in the final period is small, and, to two decimal places, the average BNE price is equal to its complete information Nash counterpart. In the final period, then, asymmetric information has little effect.

Things get more interesting in the penultimate period. Assume again, that strategies in period $T-2$ are fully revealing, and that entering the period $\widehat{c_{1,T-2}} = \widehat{c_{2,T-2}} = 8$. First, we show how firm 1 would respond if firm 2 used its static BNE pricing strategy. In this case, firm 1 would still have an incentive to signal in order to try to increase firm 2's final period price. Given firm 2's assumed pricing strategy, firm 1 will want to set the static BNE price when its own realized cost, $c_{1,T-1} = 8$, however for higher costs its optimal pricing schedule will be given as the solution to the differential equation (2).

Figure 2: Firm 2's Final Period Bayesian Nash Equilibrium Pricing Functions for Different Beliefs About $c_{1,T-1}$ When $c_{2,T-1} = 8$




Figure 3: Firm 1's $T - 1$ Best Response Signaling Strategy When $\widehat{c_{1,T-2}} = \widehat{c_{2,T-2}} = 8$ And Firm 2 Uses Its Period T/Static Bayesian Nash Equilibrium Pricing Strategy

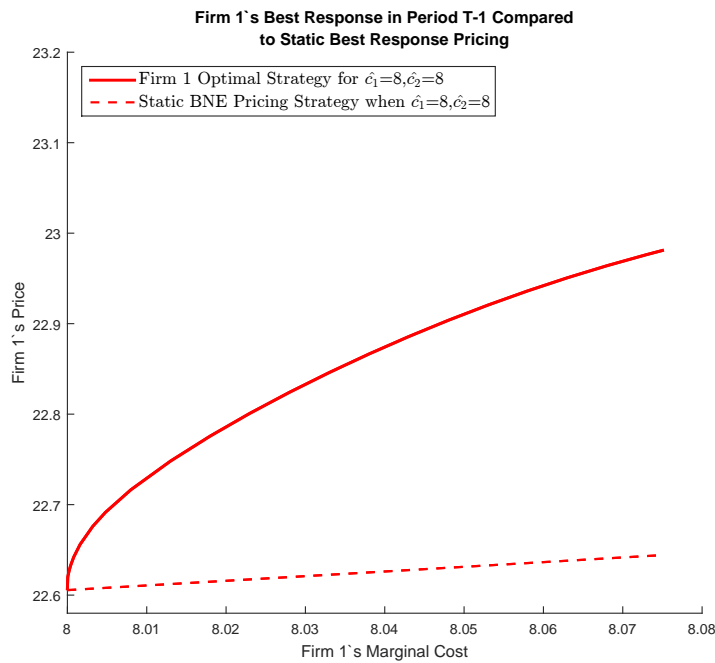


Figure 3 shows the solution to the differential equation, as well as the static BNE pricing strategy as a comparison. Signaling incentives lead firm 1 to increase its price substantially for almost all levels of cost. Although the conditions laid out above guarantee that the IC constraints will be satisfied, one can also manually verify that the prices implied by the differential equation are indeed best responses. For example, suppose that $c_{1,T-1} = 8.05$. If it chooses the price implied by the differential equation then its $T - 1$ cost type will be correctly inferred by firm 2, and it will have an expected profit of 7.0200 in the final period (the effect of discounting included), while it will have an expected $T - 1$ profit of 7.0115. On the other hand, if it deviates to the lower static best-response price, its expected $T - 1$ profit increases by 0.0025, but in period T it will be expected to have a lower marginal cost and its expected profits will fall by 0.0028. Therefore, deviation is not optimal.

Figure 4: Firm 2's $T - 1$ Best Response Signaling Strategy When $\widehat{c}_{1,T-2} = \widehat{c}_{2,T-2} = 8$ And Firm 1 Uses Its Signaling Strategy From Figure 3

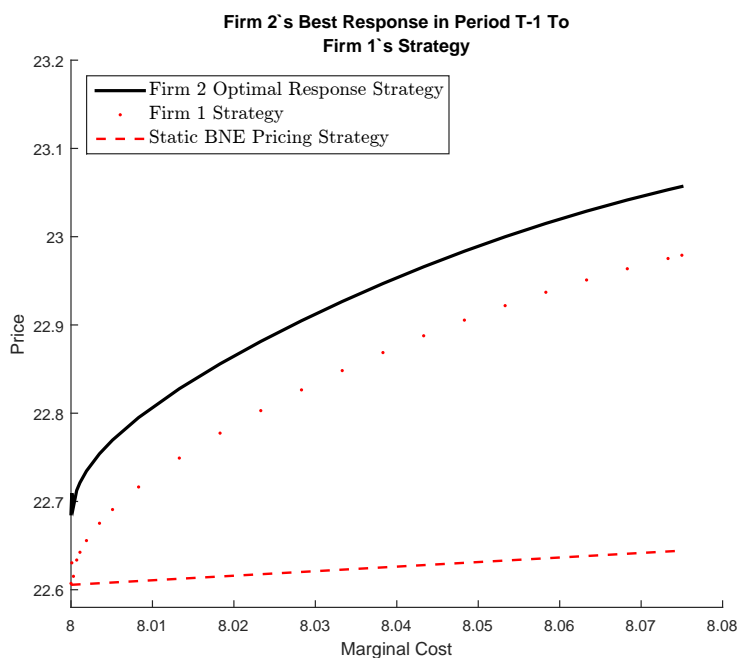



Figure 4 shows firm 2's best response signaling strategy when firm 1 uses the strategy shown in Figure 3. Now, because firm 1's prices have risen for almost all costs, firm 2's static best response will be higher, and the lower boundary condition of firm 2's pricing function is translated upwards, and higher prices are the optimal signaling response at all cost realizations. Of course, this iterative process can be continued. Figure 5 shows Firm 2's equilibrium signaling strategies in period $T - 1$ for different beliefs about firm 1's marginal cost entering the period. For comparison, the BNE pricing functions for period T are also shown (these are the narrow group of dashed lines near the bottom of

Figure 5: Firm 2's Equilibrium $T - 1$ Signaling Strategies (Red) For Different $\widehat{c}_{1,T-2}$ Compared With T Strategies (Black), Conditional on $\widehat{c}_{2,T-2} = 8$. Firm 1's Strategies Given $\widehat{c}_{1,T-2} = 8$ Are Symmetric.



the figure). As can be seen in the figure, the average $T - 1$ price is significantly higher than in period T (23.04 vs. 22.63). Equilibrium prices are also more heterogeneous, which potentially makes it even more attractive for a firm to signal that its cost is high in $T - 2$ than it was in $T - 1$, as a rival may increase its price by more in $T - 1$, than T , when it has a higher belief.

Figure 6 shows the same set of equilibrium pricing functions for $T - 2$, and the average price is higher (23.93), and once again, prices will be more dispersed. This process continues in this example until one reaches $T - 19$ at which point the equilibrium strategies have essentially converged and do not change when one moves to earlier periods. Figure 7 shows the equilibrium strategies in $T - 19$ and $T - 20$, when the average price is 26.42, or 17% above its static complete information or BNE level.

Figure 6: Firm 2's Equilibrium $T - 2$ Signaling Strategies (Red) For Different $\widehat{c_{1,T-2}}$ Compared With T (Black) and $T - 1$ (Red) Strategies, Conditional on $\widehat{c_{2,T-3}} = 8$. Firm 1's Strategies Given $\widehat{c_{1,T-3}} = 8$ Are Symmetric.





Figure 7: Firm 2's Equilibrium $T - 19$ (Vertical Pink Dashes) And $T - 20$ (Green) Signaling Strategies For Different $\widehat{c}_{1,T-x}$, Conditional on $\widehat{c}_{2,T-20} = 8$ or $\widehat{c}_{2,T-21} = 8$, Compared to Equilibrium Strategies for $T - 2$ (Blue), $T - 1$ (Red) and T (Black).



3.3 $N = 3$ or 4, and an Illustrative Merger Counterfactual

Table 1 reports average MPBE prices for $T - 1$ and $T - 25$ for $N = 2, 3$ and 4, and compares them to complete information and static BNE average prices, which are almost identical. While the price increases that signaling creates are smaller with more firms, which reflects the fact that any individual firm’s signal will have less effect on the pricing decision of other firms in the future, the effects are still quite significant in percentage terms even for four firms.

Table 1: Nested Logit Example with 2 to 4 Symmetric Firms: Average Prices

	$N = 2$	$N = 3$	$N = 4$
Complete Information Nash Eqm.	22.63	19.19	18.00
Static Bayesian Nash	22.63	19.20	18.00
T-1 Signaling Equilibrium	23.04	19.41	18.13
T-25 Signaling Equilibrium	26.42	20.32	18.71
Complete Info \rightarrow T-25 Signaling			
Δ Price	3.79 (17%)	1.13 (5.9%)	0.71 (3.9%)
Δ P-MC Markup	3.79 (26%)	1.13 (10.1%)	0.71 (7.1%)

To illustrate this point further, we consider the following very stylized merger simulation counterfactual (which we will make more realistic in future versions by relaxing the maintained post-merger symmetry assumption). Suppose that firms use strategies from $T - 25$ (or earlier) in the signaling game, and that a merger is proposed that may reduce the number of firms from either 4 to 3, or 3 to 2. We assume that the analyst is able to accurately and precisely estimate the demand parameters, and that he then proceeds to estimate pre-merger average marginal costs by inverting the standard complete information first-order conditions and plugging in average observed prices (e.g., 18.71 if $N = 4$). The implied average marginal costs are shown in the second row of Table 2. With four firms, estimated marginal cost would equal 8.83, compared with true average marginal costs of 8.0375.

Based on these estimates, and maintaining the incorrect assumption that firms play a symmetric, static complete information Nash equilibrium, the researcher could then calculate the marginal cost synergy required to keep average prices from increasing after a merger that reduces the number of firms by one. Here we make the restrictive assumptions (relaxed in the empirical application) that the firms remain symmetric single-product firms after the merger so that the merger eliminates a product and that it is known that *all* firms will benefit from any synergies. Note, however, that this assumption will tend to reduce the size of the synergy required to prevent prices from rising.

Because of the small number of firms in these examples, the required synergies are substantial (see the third row of Table 2 which reports the required level of post-merger marginal costs), amounting to

Table 2: Hypothetical Merger Counterfactuals

	3 to 2 Firm Merger	4 to 3 Firm Merger
Pre-Merger Price (Signaling Eqm)	20.32	18.71
Marginal Cost Under Complete Information	9.19	8.83
Required Marginal Cost to Prevent Price Increases (Compl. Info.)	5.65	7.53
Post-Merger Price with Synergy (Signaling Eqm)	24.54 (+15.9%)	19.87 (+6.2%)
Actual Marginal Cost Req'd to Prevent Price Increases (Sgng. Eqm)	1.70 (-79%)	6.30 (-22%)

more than one-third of estimated marginal costs in the ‘3 to 2’ case.

The fourth row reports average realized prices if these required post-merger marginal costs are realized but, after the merger, the firms play the symmetric equilibrium with a small number of firms. Recall that the analyst would be expecting prices to stay the same in this case, but they actually increase substantially, by 6.2% in the ‘4 to 3’ case and by 15.9% in the ‘3 to 2’ case. This reflects the fact that prices increase more rapidly when the number of firms falls in the signaling equilibrium than in the complete information Nash equilibrium assumed by the researcher. Of course, this also implies that even if the researcher knew the true pre-merger average marginal costs, for example from engineering studies of the industry, he would still tend to underpredict how the merger will increase prices, or, putting it another way, underestimate the marginal cost synergies required to prevent price increases. This is also illustrated in the final row of the table that shows how low marginal costs would have to be in the post-merger game (and the % decrease in costs from their true average pre-merger levels). In the ‘3 to 2’ case, the marginal costs would have to fall by almost 80% to prevent average prices from rising, which is surely larger than almost any synergies that merger applicants have suggested are plausible.

3.4 Cases When the Conditions do not Hold

While we have illustrated that signaling equilibria can produce prices that are substantially above those supported in a complete information Nash equilibrium, it is important to be clear that the conditions laid out in Section 2 can fail, especially once one goes to longer games. While we have not thoroughly explored all possible causes of failure, the most common problem we have seen so far is that the belief monotonicity fails when prices get significantly above static BNE levels i.e., for some costs firms may

begin to prefer to signal that their marginal costs are lower.

With logit-based demand, one reason why this can happen is that prices are not necessarily strategic complements (even when restricting oneself to look at static profits. Recall that, with two firms, the definition of strategic complementarity is that (Bulow, Geanakoplos, and Klemperer (1985), Tirole (1988))

$$\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} > 0 \text{ for all } p_i, p_j$$

i.e., the marginal profitability of firm i increasing its price increases in j 's price. While strategic complementarity may hold for prices close to static best response levels, so that static best response functions are upward sloping, it may not hold when considering prices substantially above static best response levels.²² This is illustrated in Figure 8 which shows the value of the second derivative firm 1's static flow profits in duopoly with respect to both prices, given the nested logit demand parameters assumed in the example, together with the static best response functions of the two firms. When prices are strategic complements the arrows point upwards. When p_1 is significantly above a static response and p_2 is less than p_1 strategic complementarity can fail, and firm 1's incentive to drop back towards its static best response becomes larger as p_2 increases slightly, as a price cut would allow it to attract more consumers from the outside good as competition from firm 2 is reduced. As suggested by the relatively steep slope of the reaction functions in the signaling game (see 7), these types of price combinations can become relevant in an equilibrium where there is lot of signaling.

Indeed, we can illustrate failure holding the demand function fixed, but increasing signaling incentives by either increasing the range of uncertainty of marginal costs or increasing the degree of serial correlation. This is done in Table 3 for duopoly, where we include our baseline parameters as well for comparison. These changes generate more signaling and significantly higher prices in the final periods of the game, compared to the baseline. However, once average prices rise above 28 dollars, failure of belief monotonicity (and also single crossing) can occur for some cost realizations. What appears to happen is that, for example, at $T - 7$ the equilibrium strategies will imply that, in that period, a firm may actually set a lower price when it believes its rival had a higher marginal cost at $T - 8$. But in this case, signaling incentives at $T - 8$ are quite different and our characterization does not work. Note that signaling incentives would also increase if we made firms more patient. With a lower discount factor,

²²For example, consider a simple aggregate data logit demand model where there are two products/firms, 1 and 2, and the mean utility of product i is $\beta_i - \alpha p_i$, and its marginal cost is c_i . Then one can show that

$$\frac{\partial^2 \pi_1}{\partial p_1 \partial p_2} = \alpha s_1 s_2 + (p_1 - c_1) \alpha^2 s_1 s_2 (2s_1 - 1)$$

where s_i represents i 's market share. The first term will be positive when demand slopes downwards. But the second term will be negative when $s_1 < 0.5$ and $p_1 > c_1$, and when this term is large enough, which will be the case when p_1 is large and s_1 is fairly small, prices will no longer be strategic complements.

Figure 8: Value of The Second Derivative of Static Flow Profits with Respect to Prices Given the Demand Parameters in the Example




Table 3: Average Price Paths For Different Parameters

	Baseline Cost Range [8,8.075] $\rho = 0.8, \sigma_\eta = 0.025$	Increased Range Cost Range [8,8.150] $\rho = 0.8, \sigma_\eta = 0.025$	Increased Serial Correlation Cost Range [8,8.075] $\rho = 0.95, \sigma_\eta = 0.015$
T	22.63	22.66	22.63
$T - 1$	23.04	23.40	23.16
$T - 3$	24.88	26.76	26.47
$T - 5$	26.06	28.78	28.90
$T - 6$	26.25	29.02	29.27
$T - 7$	26.34	28.82	29.13
$T - 8$	26.37	conditions fail	conditions fail

signaling incentives would be reduced, but, by the same token, we could preserve significant signaling by increasing the degree of serial correlation or increasing the cost range appropriately.

Of course, because we are reliant on a particular method of solving for equilibrium strategies, it is not necessarily the case that no stationary separating equilibrium exists in the earlier periods of the game. We will analyze this question using alternative techniques moving forward.

4 Empirical Example: MillerCoors Joint Venture

The example illustrates that small amounts of asymmetric information can have large effects, but the symmetry assumptions, and the assumption that the market is close to covered, are restrictive. In this section we look at a stylized model of the US Light Beer market, and investigate whether our model provides a potential explanation for why the prices of Bud Light, Miller Lite and Coors Light increased significantly after the Miller-Coors joint venture (JV) as documented by Miller and Weinberg (2016) (MW hereafter).

MW's preferred explanation for these price increases is that there was an increase in tacit collusion after the JV, based on a framework where there is complete information about demand and costs, although formally what their framework allows them to do is to reject static, complete information Nash pricing after the JV, assuming that this type of pricing prevailed before it. Critical to the rejection of static, complete information Nash is that they observe Anheuser-Busch increasing its prices as much as MillerCoors raised the prices of its brands. This result is very hard to generate in the static, complete information framework with linear marginal costs and no cost effects on parties that do not participate in the merger. It is, however, explicable by either tacit collusion, or a signaling story, as the incentives of a non-merging party to increase its prices rises directly when the merging firms are more likely to respond to them. In the current draft we only try to make the point that signaling is also a potential

explanation for what is observed, not that it is necessarily a better explanation than tacit collusion.²³ However, we plan to dig deeper into cross-market heterogeneity in market shares to understand whether there may be places where the signaling model may predict that price increases should be smaller or larger than average that we might be able to compare to the data.

There will also be some differences between our analysis and that of MW. In particular, we will make modeling choices that allow us to side-step the problem that modeling multi-dimensional signaling problems is exceedingly difficult without using particularly restrictive functional forms (Armstrong (1996)). Instead, we will make assumptions so that, after the JV, MillerCoors' products will have identical demands and marginal costs, and we will require that they set the same price for both products. In this context, this is not too unreasonable as the JV led to both flagship brands being brewed in the same facilities, so they are likely to have very similar costs. Their national market shares are also very similar. Of course, applying our approach to a wider range of differentiated product mergers would either require making less reasonable assumptions in those contexts, or tackling the hard multi-dimensional signaling problem directly.

4.1 Application, Data and Summary Statistics

Our data is drawn from the same IRI Academic panel (Bronnenberg, Kruger, and Mela (2008)), necessarily restricted to regions where beer is sold in grocery stores, as Miller and Weinberg (2016), although we will restrict attention to the years 2003 to 2011. Prices are converted to January 2010 dollars. The JV was announced in October 2007, and was completed, following US Department of Justice approval, in June 2008.²⁴ One reason for approval was that it was believed that Coors-branded products would enjoy significant JV-related efficiencies from being brewed in Miller breweries located around the country rather than just being produced in Golden, CO and Elkton, VA, and that both brands might benefit from economies in distribution. This was expected to allow Miller and Coors brands to compete more effectively with the many beers produced by Anheuser-Busch, and, it was hoped, constrain Anheuser-Busch's pricing more effectively.

In our first major departure from MW, we restrict attention to the market for light beers. This is partly driven by a desire for simplicity (a smaller choice set and number of prices in the signaling game) but it is appropriate to the extent if, when buying from grocery stores, consumers view full-calorie beers as being fairly poor substitutes to lower-calorie light beers, and light beers are known to be particularly

²³Note however that, at least based on retail prices, there is a fair amount of churn in the prices of the prices of the leading brewers both before and after the JV, which, in a collusion model, one might believe would require a fair amount of active coordination if it was not to trigger a price war. In practice, active coordination has not been alleged.

²⁴In 2009, Anheuser-Busch was bought by InBev to form ABI. Given our sample selection this does not directly affect our analysis.

popular with female drinkers.²⁵ Light beers have come to account for more retail sales than their full-calorie brand partners, especially for domestic brewers.

Figure 9: Light Beer Market Shares in 2007

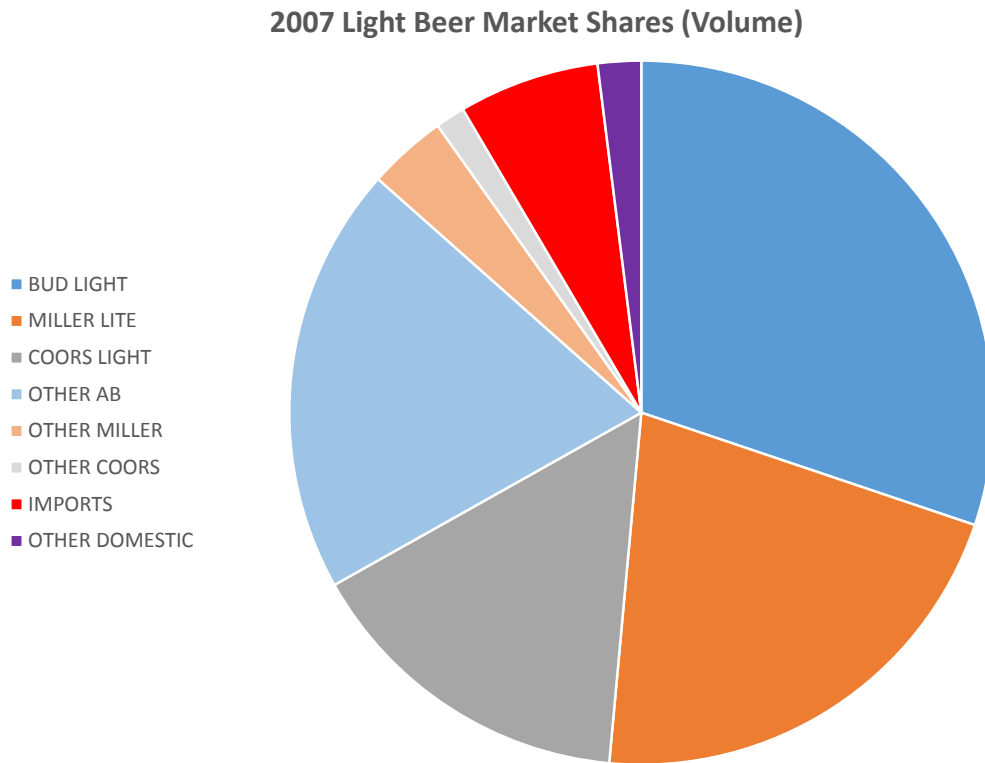


Figure 9 shows market shares in the year before the JV was completed. Bud Light, Miller Lite and Coors Light were the top three retail brands among all beers in 2007, with traditional Budweiser placing fourth. On the other hand, for imported brands, full-calorie sales tend to be larger (for example, Corona Extra places fifth in all beer sales, whereas Corona Light places seventeenth). This also suggests some differences between imported and domestic brands, which may be viewed quite differently by many consumers. In this analysis, we focus on Bud Light, Miller Lite, Coors Light, Corona Light, Amstel Light and Heineken Premium Light, with the last three counting as imported brands. In doing so, we recognize that we are ignoring a number of light beer products produced by the leading domestic brewers that have higher sales than those of the imported brands, including Natural Light and Busch Light (Anheuser Busch) and Keystone Light (Coors). The market shares of these products are represented by the lighter blue, orange and grey segments in the pie chart. These beers have significantly lower price

²⁵Substitution patterns and market shares for beers sold in restaurants, bars and specialist liquor stores may look quite different.

Table 4: Summary Statistics for Included Light Beer Brands

Product	Share of Light	Price Per 24-Pack Equivalent	
	Beer, 2007	Mean	Std. Dev.
Amstel Light	1%	29.55	2.93
Bud Light	31.7%	19.69	2.38
Coors Light	15.9%	19.79	2.47
Corona Light	2.1%	29.13	2.95
Heineken Premium	1.1%	28.77	2.78
Miller Lite	18.7%	19.66	2.52

points than the flagship brands, so incorporating them in the model (at least if they have uncertain marginal costs) would mean that we would have to confront the multi-dimensional signaling problem, although, assuming that there is some substitution between these lower-priced brands and the flagship brands, excluding them (which for us means including them in the outside good) will likely have the effect that we will underestimate the mark-ups that the leading domestic brewers would like to charge under any informational assumptions.

MW focus on 12 and 24 packs for each brand, and treat them as separate products. Given our need to avoid a multi-dimensional signaling problem, we aggregate pack sizes to the brand/product level, treating a 12-pack as being half of the volume of a 24-pack, and then calculating prices as total dollars spent on the product divided by the number of equivalent 24-packs. Table 4 shows summary statistics for the brands included in the analysis. The dominance of the leading domestic brands in the light beer market is clear: Bud Light, Coors Light and Miller Lite account for over 66% of all retail light beer sales. It is also noticeable that these beers sell at very similar average prices (this is true both before and after the JV), whereas the imported brands sell at a much higher, although also similar, price points.

Using this selection of brands, we can also repeat MW's analysis of the price effect of the JV. The baseline specification is

$$\text{Log}(\text{Real Retail Price}_{i,t,s}) = \mathbb{I}\{\text{AB or MC Product}\} \times \mathbb{I}\{\text{Post-Merger}\} \theta + \dots \\ \text{Time Effects}_t + \text{Product X Store FE}_{i,s} + \varepsilon_{i,t,s}$$

where prices are measured at the store-brand-month-level and we include product x store fixed effects to allow for the fact that some stores may tend to price domestic and imported products at different relative prices throughout the time that the store is in the sample. For this regression we aggregate weekly IRI observations to the monthly level (following MW) by dividing total monthly revenues by total volume-adjusted units sold.

Figure 10: Price Paths for Leading Domestic Light Beer Brands, 24 Pack Equivalents

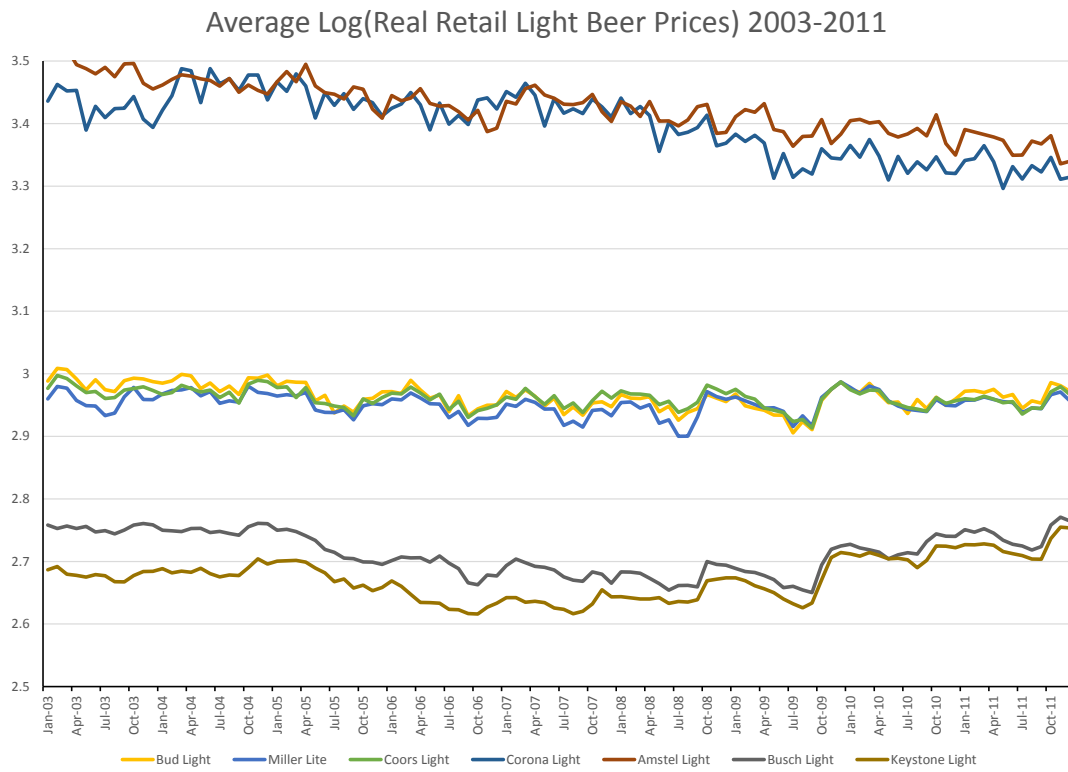


Table 5: Light Beer Estimated Merger Effects

	(1)	(2)	(3)	(4)
Controls	Product x Store FEs, Month FEs		Product x Store FEs Month FEs, Brand Time Trends	
x Post-Merger				
Bud Light	0.062 (0.012)	0.061 (0.012)	0.034 (0.010)	0.034 (0.009)
ML/CL	0.073 ($\approx 140\text{¢}$) (0.008)	0.073 (0.008)	0.039 ($\approx 75\text{¢}$) (0.009)	0.039 (0.009)
Busch Light		0.072 (0.010)		0.047 (0.012)
Natural Light		0.120 (0.011)		0.055 (0.015)
Michelob Ultra		0.027 (0.014)		0.021 (0.018)
Observations	523,096	813,605	523,096	813,605

Visual inspection of price-paths in Figure 10 suggests that there are some differences in price trends across products, so we experiment with different time controls. In the baseline we simply control for month fixed effects that are common across products. The results are in Table 5, with standard errors clustered on the IRI region, where we present results when we include the brands only included in the structural model and including the main other light beer brands sold by Anheuser-Busch.

The table indicates that when we control for only month effects, the effect of the JV on the prices of the leading brands is around 7.3%, or \$1.40 per 24 pack equivalent, for Miller Lite and Coors Light, and just a little less for Bud Light. When we include brand time trends, all of these effects are smaller (just under 4% or 75 cents), but they are still statistically and economically significant. One interesting feature of the extended results is that there are slightly larger effects, although differences are not statistically significant, for Bush Light and Natural Light, which are more insulated from import competition, than for the leading brands, or Anheuser Busch’s Michelob Ultra, which retails at a price point between the leading brands and the imported brands.

4.2 Demand

While we could use the type of random coefficient demand model implemented by MW, and we plan to do so, here we choose to use a nested logit demand model as it is convenient, when solving our model, to use as many analytic derivatives as possible. We will also use observations at the store-product-month level whereas MW aggregate observations to the IRI region level, which, potentially, may create some issues as the set of stores in the IRI sample varies over time and different stores in the same region may be heterogeneous in the consumers that they serve and the products that they stock.

We assume a one-level nested logit structure where the nests are ‘domestic’, ‘imported’ and the outside good, which contains the options of not purchasing or purchasing one of the brands not included in the dataset. The estimation baseline specification is standard, following Berry (1994),

$$\begin{aligned} \log(s_{i,t,s}) - \ln(s_{0,t,s}) = & \sigma \ln(s_{i,t,s}^{nest}) + FE_{product} - \dots \\ & \alpha p_{i,t,s} + FE_{month} + FE_{region} + \varepsilon_{i,t,s} \end{aligned} \quad (5)$$

where i represents the product, t the month and s the store. $s_{i,t,s}^{nest}$ is i ’s share of its nest. $p_{i,t,s}$ is the 24-pack equivalent price in dollars. The definition of market size is, as is typical when estimating discrete-choice demand models, somewhat arbitrary and we define it as 110% of all light beer sales observed at the store. The months between the announcement and the completion of the JV are excluded, although the results are very similar if these are included.

Our choice of instruments for $\ln(s_{i,t,s}^{nest})$ and $p_{i,t,s}$ follows MW, with adaptation to reflect the nesting structure. Specifically we use: (i) product-specific interactions between the distance from the relevant brewery to store’s region interacted with the contemporaneous price of diesel²⁶; (ii) the average value of the distance*diesel price variable for other products in the same nest; (iii) a dummy for time periods after the MillerCoors JV interacted with dummies for MC products, Anheuser-Busch products and imports (the validity of these instruments implicitly assumes that pricing changes after the JV reflect supply-side changes rather than demand changes); and (iv) number of other products available in the nest.²⁷

Column (1) of Table 6 reports the estimates from the baseline specification. The nesting coefficient is highly significant, and implies that domestic products and imports are poor substitutes for many consumers. The magnitude of the nesting and price coefficients also implies that each product’s demand is very price elastic: the average implied domestic elasticity for Bud Light is about -6. This elasticity is higher than is typically estimated for beer, although it is consistent with the fact that the domestic brands have very similar prices, both before and after the JV. Obviously one explanation for the high elasticity is that we have ignored ‘consumer stockpiling’ (Hendel and Nevo (2006)), although one would have expected these effects to be reduced by our aggregation of sales to the monthly level and our use of instruments that reflect longer-run effects to deal with the endogeneity of prices. In

²⁶For domestic beers we assume that the beer is shipped from the firm’s closest brewery (which changes for Coors before and after the JV), which is an appropriate assumption for the flagship domestic light beers (see for example, <http://www.millercoors.com/breweries/brewing-locations>, which lists which beers are produced in which plant under the JV). For imports, we follow MW in calculating distances from the Corona’s closest Mexican plant, and using the closest major ports where Heineken imports into the US.

²⁷The entry of Heineken Premium Light provides some nationwide variation in the number of products in the imported nest, plus there is some limited variation in the products available across different stores.

Table 6: Light Beer Nested Logit IV Demand Estimates (SEs Clustered on DMA)

	(1)	(2)	(3)
	As Above	Add Store FEs	Separate Nests
Price Coefficient (α)	-0.203 (0.044)	-0.204 (0.040)	-0.173 (0.055)
Nesting Coefficient (σ)	0.698 (0.143)	0.680 (0.132)	-
σ Domestic	-	-	0.878 (0.160)
σ Imports	-	-	0.272 (0.259)
Number of Observations	496,889	496,889	496,889

the counterfactual below we will consider the effects of signaling on prices using both these demand estimates, and an alternative set of parameters where we halve the price coefficient so that the price elasticity is more consistent with earlier estimates.

The remaining columns indicate that the price and nesting coefficients are qualitatively robust to introducing store fixed effects, and to allowing for separate nesting coefficients on the domestic and imported nests. In the later case, we estimate that the included domestic beers are particularly close substitutes for each other. For the rest of the analysis, we will assume, as a baseline, a nesting coefficient of 0.7 and price coefficient of -0.2, reflecting the coefficients in column (1). We will also, however, present some results when we halve the magnitude of the price coefficient (adjusting the value of the brand dummy coefficients appropriately, of course), so that we get a mean elasticity for Bud Light of around -3, which implies larger mark-ups in the complete information model.

4.3 Can an Asymmetric Information Model Generate the Observed Price Increases?

We now turn to our main question of whether a dynamic asymmetric information/signaling model can generate price increases similar to the price increases that followed the MillerCoors JV. As already indicated, we do the exercise in a very stylized way to side-step the problem that multi-dimensional signaling models are generally intractable, although we plan to try to add some additional flexibility in later iterations of the paper. In modeling a counterfactual we treat the ‘market’ as being defined by a single representative store that sells all of our products, with demand as defined above, completely ignoring the fact that there are differences in both demand and costs across geographic areas. We also ignore the presence of a retailer who also takes active pricing decisions, consistent with a model where a retailer just passes through brewer prices. MW allow for an active retailer but estimate that the brewer’s prices are passed through almost perfectly with the addition of a small retail margin. Like

MW, we also ignore the existence of beer distributors who exist between brewers and retailers.

4.3.1 Pre-JV Scenario

Before the JV we view the market as consisting of three firms that might be using signaling strategies (AB (for Bud Light), Miller (for Miller Lite) and Coors (for Coors Light)) and two import brewers (Corona and Heineken) that we allow to respond to other firms' expected prices (i.e., they interpret signals and have correct equilibrium expectations about the prices that other firms will set in the next period) but which we assume are not engaged in active signaling themselves.²⁸

We assume that the pre-JV marginal cost of AB (for a 24 pack volume-equivalent) may range from \$15.45 to \$15.65; for Miller (for a 24 pack volume-equivalent) from \$16.25 to \$16.45; and for Coors (for a 24 pack volume-equivalent) from \$16.25 to \$16.45, so the uncertainty is around approximately the last $1\frac{1}{2}\%$ of marginal costs for each brewer. Heineken and Corona have fixed marginal costs of \$25, which applies to both Heineken Premium and Amstel Light, and \$25.50 respectively. The marginal costs mean that we would approximately match the nominal average 2007 prices of each of the brands under static, complete information pricing, and the pre-JV market structure.

We will assume that the marginal costs of each firm evolve independently according to AR(1) processes:

$$c_{i,t} = 0.9c_{i,t-1} + (1 - \rho)\frac{c + \bar{c}}{2} + \eta,$$

where η is truncated so that the marginal cost remains on its support, with an untruncated normal distribution that has mean zero and standard deviation 0.05.²⁹ The mean utilities of the different brands (less price effects) are set at 5.85 (for AB), 5.65 for Miller and Coors, 6 for the two Heineken products and 6.2 for Corona. As with the level of marginal costs, these qualities were chosen so that, at average equilibrium complete information prices, we approximately match the average market shares of the products around the time of the JV, and can explain why prices are approximately identical for the beers in each nest.

4.3.2 Post-JV Scenario

We assume that, after the JV, product qualities are unchanged, so any price changes must reflect some supply-side adjustment. We then choose to lower the costs of Miller Lite and Coors Light so that

²⁸For example, their small market shares would tend to imply that their ability to affect the future prices of the domestic brands would be very limited.

²⁹While brand prices display significant positive serial correlation, the choice of these parameters for the unobserved component of costs is arbitrary. Indeed estimation of the supports would remain difficult even with a strategy for estimating the serial correlation parameter.

Table 7: Light Beer Application: Predicted Average Prices under Complete Information

Product	Pre-JV Complete Information	Post-JV Complete Information
Bud Light	18.31	18.30
Miller Lite	18.31	18.30
Coors Light	18.31	18.30
Corona Extra	27.78	27.77
Amstel Light	27.54	27.53
Heineken Premium	27.54	27.53

under complete information, average prices of all products would be almost exactly the same as they were before the JV. This is done by dropping the range of their marginal cost from $[\$16.25, \$16.45]$ to $[\$15.40, \$15.60]$, although, to be clear, we are now assuming that both brands have an identical marginal cost whereas prior to the JV we allowing their realized marginal costs to be different even if the supports were the same. Serial correlation and innovations remain unchanged. Table 7 shows the predictions based on complete information.

When interpreting our results, note that MW predict that if the joint venture had realized synergies and only (static) unilateral effects, which matches the complete information assumption here because the finite nature of our game rules out collusive coordinated effects, the prices of Miller and Coors would have increased significantly: by about 30 cents for a 12-pack, or a little under one-third of the price increase they predict under the combination of unilateral effects, coordinated effects and synergies (see, for example, their Figure 5 for Coors). They predict that the price of Bud Light would have increased slightly with only unilateral effects (their Figure 6). Therefore by choosing a benchmark where, without signaling effects, prices remain the same after the JV we may be biasing ourselves towards not being able to explain the price increases observed in the data.

4.3.3 Predicted Prices Under Asymmetric Information

We now solve for equilibrium prices both before and after the JV (with the synergy) allowing for asymmetric information. To do so, we adapt the solution procedure described previously to allow for asymmetries between the signaling firms and for two non-signaling firms, one of which has two products. We solve the game for $T = 25$ periods, and focus on the prices at the beginning of this game by which point they had converged (average prices change by less than $\frac{1}{2}$ cent over the last few periods). The predictions of average prices are contained in Table 8.

Prior to the JV, signaling has quite small effects on prices given the estimated demand and assumed

Table 8: Light Beer Application: Average Prices Under Complete Information and Asymmetric Information with Baseline Demand

Product	Pre-JV		Post-JV	
	Complete Information	Signaling Equilibrium	Complete Information	Signaling Equilibrium
Bud Light	18.31	18.45	18.30	19.04
Miller Lite	18.31	18.43	18.30	19.03
Coors Light	18.31	18.43	18.30	19.03
Corona	27.78	27.78	27.77	27.79
Amstel	27.54	27.54	27.53	27.55
Heineken	27.54	27.54	27.53	27.55

cost parameters, raising the prices of Miller and Coors, for instance, by 12 cents or just under 1%, while the price of Bud Light rises only slightly more. The importers are sufficiently differentiated that their prices are hardly affected at all by the small increase in the price of domestic beers.

After the JV, the prices of all of domestic beers increase by 73 cents relative to complete information, or by around 60 cents relative to the pre-JV signaling equilibrium. The price increase is almost exactly the same for MillerCoors and Anheuser-Busch, reflecting the fact that these firms are now very close to being symmetric in terms of combined market share. We regard the fact that the model predicts a large price increase for the non-merging domestic firm as significant, as the fact that a non-merging firm that (by assumption) did not experience a synergy raised its price roughly as much as the merging firms, lies at the heart of MW's rejection of a static, complete information post-JV equilibrium. While tacit collusion might provide one explanation, here we are arguing that it could also be predicted by a non-collusive model with a small, and plausible, amount of asymmetric information concerning serially correlated marginal costs and dynamic signaling.

It is natural to ask whether we could generate larger predicted price increases in this example, by increasing the degree of uncertainty about costs or increasing the degree of serial correlation/reducing the standard deviations of the cost innovations. We have not, so far, done an exhaustive search on this question, but if we change parameters in these directions significantly, without changing demand, we find that although the price increases in the final ($T - 1, T - 2, \dots$) periods are larger (more than \$1.20 in some cases), the conditions required to characterize best responses fail once we move several periods from the end of the game. This is consistent with what was observed in the example.

Table 9 reports the results from the same counterfactual exercise but with less elastic demand, where the halve the baseline price coefficient. The brand-specific marginal cost, synergy and mean utility are adjusted so that under complete information outcomes are the same as in the baseline case (specifically,

Table 9: Light Beer Application: Average Prices Under Complete Information and Asymmetric Information with Less Elastic Demand

	Pre-Merger		Post-Merger	
	Complete	Signaling	Complete	Signaling
	Information	Equilibrium	Information	Equilibrium
		T-25		T-25
Bud Light	18.31	18.47	18.31	19.34
Miller Lite	18.31	18.46	18.31	19.32
Coors Light	18.31	18.46	18.31	19.32
Corona	27.78	27.78	27.77	27.80
Amstel	27.54	27.54	27.53	27.56
Heineken	27.54	27.54	27.53	27.56

mark-ups are now larger). This is obviously an ad-hoc adjustment but it is intended to illustrate that the previous results are not being driven by an unrealistically high estimate of the demand elasticity. We observe that before the JV, we get increases in equilibrium prices that are similar to those in Table 8. However, after the JV increases are larger, so that, assuming signaling both before and after the JV, the change in market structure should cause prices to rise by about 86 cents per 24-pack equivalent, an increase which is not inconsistent with the descriptive results in Table 5.

5 Conclusion

In an environment where firms compete in prices, it is natural that firms might want to increase their prices today if doing so will make their rivals respond by setting higher prices in future periods. We develop this idea in a finite horizon dynamic model where firms sell differentiated products, with small uncertainty about marginal costs. This extends the simple, linear demand and two or three period settings considered by Mailath (1989) and Mester (1992). Our examples show that small uncertainty about persistent marginal costs can have very large effects on equilibrium prices, especially when a market comes to be dominated by two, relatively symmetric firms. We show that a merger analysis that assumes firms play a complete information, static Nash equilibrium may substantially under-predict the synergies required to prevent prices from rising, and we also suggest that post-merger price increases that have been attributed to tacitly collusive coordinated effects might equally be explained by dynamic, but non-collusive, signaling behavior.

The examples in this draft are preliminary, and we hope and plan to extend the framework in several directions in both revisions and future work. Several questions remain outstanding. First, the model may suggest a number of comparative statics (for example, how price increases may change with

differences in firm market shares or costs across geographies), which could be used to provide additional support for the model. Second, examples where firms compete in quantities should be investigated. In this case, signaling incentives should tend give rise to more competitive outcomes. Similarly, we should try to understand what happens when there is asymmetric information about product demand. Third, one could extend the model to allow for uncertainty in some shared variable such as market demand growth, where each firm may try to signal its own beliefs about demand by using publicly observed capacity choices. Some initial examples in this direction suggest that with two or three firms signaling can generate something which looks like ‘capacity discipline’ where firms invest less in capacity than they would in a model with complete information.

References

- ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica*, 64(1), 51–75.
- ASHENFELTER, O., D. HOSKEN, AND M. WEINBERG (2014): “Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers,” *Journal of Law and Economics*, 57(S3), S67–S100.
- ASHENFELTER, O. C., D. S. HOSKEN, AND M. C. WEINBERG (2015): “Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry,” *RAND Journal of Economics*, 46(2), 328–361.
- BERRY, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25(2), 242–262.
- BESANKO, D., U. DORASZELSKI, Y. KRYUKOV, AND M. SATTERTHWAITTE (2010): “Learning-by-Doing, Organizational Forgetting, and Industry Dynamics,” *Econometrica*, 78(2), 453–508.
- BONATTI, A., G. CISTERNAS, AND J. TOIKKA (2015): “Dynamic Oligopoly with Incomplete Information,” Discussion paper, MIT, Available at <http://economics.mit.edu/files/10575>.
- BORENSTEIN, S. (1990): “Airline Mergers, Airport Dominance, and Market Power,” *American Economic Review*, 80(2), 400–404.
- BRONNENBERG, B. J., M. W. KRUGER, AND C. F. MELA (2008): “The IRI Marketing Dataset,” *Marketing Science*, 27(4), 745–748.
- BULOW, J. I., J. D. GEANAKOPOLOS, AND P. D. KLEMPERER (1985): “Multimarket Oligopoly: Strategic Substitutes and Complements,” *Journal of Political economy*, 93(3), 488–511.
- CAMINAL, R. (1990): “A Dynamic Duopoly Model with Asymmetric Information,” *The Journal of Industrial Economics*, pp. 315–333.
- FERSHTMAN, C., AND A. PAKES (2010): “Oligopolistic Dynamics with Asymmetric Information: A Framework for Empirical Work,” Discussion paper, Working Paper, Harvard University.
- HENDEL, I., AND A. NEVO (2006): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74(6), 1637–1673.
- HUBBARD, T. P., AND H. J. PAARSCH (2013): “On the Numerical Solution of Equilibria in Auction Models with Asymmetries within the Private-Values Paradigm,” in *Handbook of Computational Economics*, ed. by K. Schmedders, and K. L. Judd, vol. 3, pp. 37–115. Elsevier: North-Holland.

- KAYA, A. (2009): “Repeated Signaling Games,” *Games and Economic Behavior*, 66(2), 841–854.
- KIM, E. H., AND V. SINGAL (1993): “Mergers and Market Power: Evidence from the Airline Industry,” *American Economic Review*, 83(3), 549–69.
- MAILATH, G. (1987): “Incentive Compatibility in Signaling Games with a Continuum of Types,” *Econometrica*, 55(6), 1349–1365.
- MAILATH, G. J. (1988): “An Abstract Two-Period Game with Simultaneous Signaling: Existence of Separating Equilibria,” *Journal of Economic Theory*, 46(2), 373–394.
- (1989): “Simultaneous Signaling in an Oligopoly Model,” *Quarterly Journal of Economics*, 104(2), 417–427.
- MAILATH, G. J., AND E.-L. VON THADDEN (2013): “Incentive Compatibility and Differentiability: New Results and Classic Applications,” *Journal of Economic Theory*, 148(5), 1841–1861.
- MESTER, L. J. (1987): “Multiple Market Contact Between Savings and Loans: Note,” *Journal of Money, Credit and Banking*, 19(4), 538–549.
- (1992): “Perpetual Signalling with Imperfectly Correlated Costs,” *RAND Journal of Economics*, 23(4), 548–563.
- MILLER, N. H., AND M. C. WEINBERG (2016): “The Market Power Effects of a Merger: Evidence from the US Brewing Industry,” Discussion paper, Georgetown University, Available at <http://www.nathanhmilller.org/miller%20weinberg%202016%2007%2025%20revised.pdf>.
- NAKAMURA, E., AND D. ZEROM (2010): “Accounting for Incomplete Pass-Through,” *Review of Economic Studies*, 77(3), 1192–1230.
- NEVO, A. (2000): “Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry,” *RAND Journal of Economics*, 31(3), 395–421.
- PETERS, C. (2009): “Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry,” *Journal of Law and Economics*, 49(2), 627–649.
- RODDIE, C. (2012): “Signaling and Reputation in Repeated Games, I: Finite Games,” Discussion paper, University of Cambridge, Available at <http://ssrn.com/abstract=1994378>.
- SHAPIRO, C. (1986): “Exchange of Cost Information in Oligopoly,” *The Review of Economic Studies*, 53(3), 433–446.

- SWEETING, A. (2007): “Market Power in the England and Wales Wholesale Electricity Market 1995–2000,” *The Economic Journal*, 117(520), 654–685.
- SWEETING, A., J. W. ROBERTS, AND C. GEDGE (2016): “A Model of Dynamic Limit Pricing with an Application to the Airline Industry,” Discussion paper, University of Maryland, Available at http://econweb.umd.edu/~sweeting/SWEETING_DLP_JAN2016.pdf.
- TIOLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.
- TOXVAERD, F. (2008): “Strategic Merger Waves: A Theory of Musical Chairs,” *Journal of Economic Theory*, 140(1), 1–26.
- TOXVAERD, F. (2014): “Dynamic Limit Pricing,” Discussion paper, University of Cambridge.