

Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Mergers

Andrew Sweeting* Xuezhen Tao†

August 2016

Work in Progress - Do Not Cite

Abstract

Existing theoretical and structural empirical analyses of mergers assume that firms have complete information about their rivals' demand and marginal costs. On the other hand, if marginal costs are private information and serially correlated, firms may wish to use their price or quantity choices to signal information, in order to affect how their rivals will expect them to set prices in the future. We show that even quite small asymmetries of information can have very large effects on equilibrium pricing in concentrated markets, which can make merger simulations based on the complete information assumption misleading, and which are large enough to explain post-merger price increases that might otherwise be attributed to tacit collusion or 'coordinated effects'.

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

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

*Department of Economics, University of Maryland and NBER. Contact: sweeting@econ.umd.edu.

†University of Maryland. We thank Dan Vincent, Nate Miller and seminar participants at Bates White, the FCC, the Federal Reserve Board and the University of Arizona 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

In both theoretical and empirical analyses of mergers, it is standard to assume that firms operate in an environment of complete information, so that they choose prices or quantities with full information on the demand and marginal costs of their rivals, as well as their own demand and costs. In practice, the complete information assumption should probably be seen as a convenient modeling simplification in most real-world merger settings, as firm-level demand and cost information is usually treated as being commercially sensitive during the merger review process, and is usually assumed to be information that should not be shared with competitors, rather than being information that competitors already have. However, one might assume that as long as the degree of uncertainty about rivals' costs or demands is small, then prices would be close enough to the complete information case that it would be reasonable to assume complete information in order to benefit from the tractability that the complete information framework provides.

In this paper we show that this assumption may not be correct: even very small deviations from the complete information assumption, combined with the assumption that whatever is unobserved is somewhat persistent over time, can have very large effects on equilibrium prices. While the point that incomplete information about persistent demand or costs will affect equilibrium prices has been made in the theoretical literature (Mailath (1989), Mester (1992), Caminal (1990)) we believe that we are the first to show that these effects can be large enough to mean that a merger analysis that assumes complete information could arrive at conclusions that would be seriously misleading. We do this both by constructing and analyzing an example, and by examining the real-world setting of the 2008 Miller-Coors joint venture, showing that a small amount of asymmetric information could generate post-merger price increases almost as large as those found by Miller and Weinberg (2015) which those authors attribute to collusion.¹

We consider a multi-period discrete time oligopoly model where each firm knows its own marginal cost, but, without any additional information, is uncertain about exactly what the marginal costs of its rivals are. Demand is static. We will focus on the case where the degree of

¹To be clear, we are not, at least at this point, trying to run a horse-race between a tacit collusion story and an asymmetric information story for why prices rose following the completion of the joint venture. Instead, we are merely showing that an asymmetric information story can generate a very similar pattern where the merged firm raises its prices, despite benefiting from a large cost-reducing synergy, and a non-merging firm raises its price by a very similar amount.

uncertainty is small, in the sense that it will only concern the last one or two percent of costs, so it is quite consistent with the notion, with which we agree, that firms in the industry would be well-informed about their rivals' production processes, even if they do not know exactly the cost of labor or inputs, or exactly how much labor it takes for the firm to turn inputs into product. The surprising result will be how much this small uncertainty can affect equilibrium prices.

When the idiosyncratic component of each firm's marginal cost is positively serially correlated, the logic behind why equilibrium prices can rise when firms set prices is simple.² Suppose that in the next period all else equal, firm j will set a higher price if it believes that firm i has a higher marginal cost. This provides an incentive for firm i to signal that it has a higher cost in the current period, which it may be able to do by setting a higher price in the current period. Even ignoring its own signaling incentives, this will tend to make firm j want to set a higher price in the current period (based on the logic of prices as strategic complements), and, if j is also signaling, this may tend to further increase the price that firm i wants to set today. These effects can reinforce each other to generate large price effects. The specific framework that we consider has a known finite number of periods, and there is no linkage across periods apart from the correlation in marginal costs (for example, there are no menu costs). In this set-up, tacit collusion is unable to raise prices, because of standard backwards induction arguments, but we find that asymmetric information and signaling incentives can raise prices substantially.³ To illustrate, in a duopoly example we show that uncertainty about less than 1% of each firm's marginal cost can increase equilibrium prices in the early periods of a game, when strategies are approximately stationary, by a little under 20%.

Note that the previous paragraph is worded quite carefully, saying that this mechanism 'can' raise equilibrium prices. As we will also discuss, the characterization and computation of a well-behaved fully-separating equilibrium, depends on all firms' payoffs in each period of the game satisfying a number of conditions that are similar to those in a single-agent signaling model (see Mailath (1988)). It will turn out that once signaling effects become too large, with this threshold being a function of demand and cost parameters, these conditions will no longer be satisfied, and our computations will fail. These problems will arise partly because we will focus

²Note that significant effects arise from the combination of asymmetric information and incentives to signal information that arise in a dynamic model. Given the parameters that we consider, asymmetric information in a static model has almost no effects on prices.

³As we note below, we can also find equilibria with significant price increases in the infinite horizon version of our model. However, the finite horizon assumption also reflects how we solve for strategies.

on standard logit/nested logit demand structures where, once we are considering prices above static equilibrium levels, prices may no longer be strategic complements so that a firm that believed that one of its rivals would set a higher price would have more incentive to set a lower one, which will radically change firms' signaling incentives. With linear demand (and linear marginal costs, which we are also assuming) this problem would disappear, but we prefer to use demand systems which are typically used to model differentiated product markets and instead be up-front that this imposes some limitations on our conclusions. Analysis of what might happen in pooling or partial pooling equilibria, and whether asymmetric information could cause prices to rise in these equilibria as well, would be a fascinating extension and is left to future work. We plan to explore some alternative models where firms set prices but where there may be other linkages across periods, such as a model with stochastic learning-by-doing or capacity that can only be adjusted by incurring adjustment costs⁴, in future iterations of the current paper, as well as extending our current merger simulation example to allow for random coefficients in demand (which would allow for a more realistic demand model at the expense of increasing the computational burden).

While the issue of asymmetric information has been ignored in the merger literature, our paper is related to an older theoretical work on oligopoly models and a very recent strand of literature on dynamic models with persistent asymmetric information. Mailath (1988) considers an abstract two-period game, and shows that, under the assumption that firms' flow payoffs are separable across periods (also assumed here), existence of a separating equilibrium follows under almost the same conditions on firms' payoffs that are required in a game where there is only one firm with private information. As Mailath comments, it is not straightforward to relate these requirements back to the primitives of the model, and as we will show it is quite possible that some of the conditions that are typically satisfied almost trivially in one-shot or single agent signaling models, such as belief monotonicity, can fail in a multi-period oligopoly game even with standard forms of demand.⁴ Mailath (1989) considers a more specific two-period model where each firm's cost is drawn from a commonly known distribution, but is then fixed across periods. Mailath considers a fully separating equilibrium where firms' first period prices reveal their costs, so that the equilibrium outcome in the second period is the same as if the costs were

⁴In the current setting, belief monotonicity would mean that a firm always benefits when its rivals believe that it has higher costs. But this can fail in the current model if, for example, rivals will signal more aggressively in the future when they face a firm with a low marginal cost.

public information. Assuming linear demand, the first period outcome can be shown to be the unique fully separating equilibrium. In the current paper, we extend this model to allow for multiple periods and serially correlated (but not perfectly correlated) costs. In a theoretical paper, Mester (1992) does a similar exercise in a three-period model showing uniqueness of the equilibrium given linear demand when firms set quantities, showing that equilibrium output in the first and second periods are above static, complete information levels, consistent with how strategic incentives change when firms compete in quantities rather than prices. Caminal (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.⁵ Considering this type of demand-signaling would be a natural extension. Much more recently, Bonatti, Cisternas, and Toikka (2015) consider 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.⁶

In the empirically-oriented IO literature, Fershtman and Pakes (2010) propose a framework for analyzing dynamic oligopoly games with asymmetric information where firms have discrete types and take discrete actions, which can naturally lead to pooling equilibria, where players of the same type choose the same action. Rather than trying to use more standard Perfect Bayesian concepts adapted to a dynamic setting, they propose an alternative concept, called Experienced Based Equilibrium, which potentially makes analysis computationally tractable by specifying each player’s beliefs in terms of their expectations about their payoffs from choosing different actions rather than in terms of their beliefs about other players’ types. This comes at the cost of possibly increasing the number of equilibria, and makes it less easy to identify signaling incentives. In the current paper, we consider continuous choices (prices), and attempt to stick more closely to standard concepts. This is also the approach taken in Sweeting, Roberts, and Gedge (2016), where a dynamic version of the Milgrom and Roberts (1982) limit pricing model is developed and argued to be a plausible explanation for why incumbent airlines, which dominate their routes, lowered prices significantly when threatened with entry by Southwest. In

⁵Caminal considers a model with two discrete demand types (High or Low) for each firm.

⁶There are also connections to the literature on signaling in auctions, which has focused on settings with resale or aftermarkets. For example, Haile (2003), Goeree (2003) and Molnár and Virág (2008).

that paper, only the incumbent has private information and may want to signal that a potential entrant’s post-entry profits will be low, where the post-entry game is assumed to be one of complete information. In a simple model where the incumbent has a serially correlated, linear marginal cost, one can show that a fully separating Perfect Markov Bayesian Equilibrium will exist and be unique, under refinement, under several easy-to-check conditions on the primitives. In a more complicated model with endogenous capacity investment and asymmetric information about the incumbent carrier’s connecting demand, conditions for existence and uniqueness have to be verified computationally, but limit pricing effects can remain large, or actually be larger, than in the exogenous cost case.⁷ In the current paper, entry plays no role and the focus is on signaling between oligopolists. With logit-based demand, theoretical conditions only allow one to show existence and uniqueness of firms’ best responses, given strategies of other firms, rather than directly providing results about the nature of the equilibrium. In this sense, the characterization here is much less complete than in Sweeting, Roberts, and Gedge (2016). However, we focus on a setting of broad and practical interest, mergers, and show that even small degrees of uncertainty about costs can generate large price effects.⁸

Our results are closely related to the antitrust literatures studying the effects of mergers and discussing coordinated effects (Whinston (2008)). Weinberg (2008) finds that, even for the set of selected mergers that regulators have allowed to be completed, prices of both merging and the leading non-merging firms have tended to rise after completion (see also Peters (2009), Kim and Singal (1993) and Borenstein (1990) for evidence from the airline industry; and also see Ashenfelter, Hosken, and Weinberg (2015) and Miller and Weinberg (2015) for discussion of brewer mergers which will be the focus of the empirical example below). One explanation for this pattern is that models that assume only unilateral effects underpredict price increases because, once the industry is more concentrated, coordinated effects, usually interpreted to mean tacit collusion, tend to give an additional boost to equilibrium prices (Jayaratne and Ordovery (2015)). Our model provides an alternative theory for why prices increase, because increasing concentration tends to make signaling effects much larger. Of course, both theories

⁷One intuition for why effects can become larger in a richer model is that the firm has more margins that it can use to reduce the cost of signaling. In equilibrium this requires large price reductions for the signals to be credible.

⁸In Sweeting, Roberts, and Gedge (2016) large effects require what is private information to the incumbent to potentially have a significant effect on the entry decision of the potential entrant. This is unlikely to be the case where the degree of uncertainty is as small as in the examples that we consider here.

involve a role for dynamics, but there are several important differences between the collusive and signaling theories that are worth stressing. First, tacit collusion stories require firms' strategies to involve some type of retaliation in response to opponent deviations, which is not true in the current model. Second, significant signaling effects can arise in finite-horizon model whereas no degree of tacit collusion can be supported in a finite, complete information game.⁹ Third, due to folk theorems, tacit collusion models can potentially explain a very wide range of outcomes, including coordination on joint-profit maximizing prices if firms are patient enough, whereas, even when prices rise, it is rarely claimed that prices are close joint profit-maximizing levels after mergers. In contrast, an asymmetric information model is only likely to be able to support smaller price increases like those observed in the data. Fourth, while it is often argued that complete information about other firms' demand and costs will tend to make collusion easier to achieve, in our model it would return prices to lower, static levels. Finally, in our model, asymmetric information would tend to create pro-competitive effects if firms competed in quantities, rather than prices. Indeed, Mester (1992) was partly motivated by her empirical observation that in some industries, such as banking, multi-market contact actually seemed to lead to output expansion, rather than output reduction, as would be expected in a collusive model.

The remainder of the paper 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 (2015). Future revisions will contain additional examples. Section 5 concludes.

⁹Of course, this may be a good feature for tacit collusion models to have. For example, Sweeting (2007) found that extent to which leading generators withheld output decreased substantially when it was announced that the England and Wales wholesale electricity 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.

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 $0 < \beta < 1$. 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).¹⁰ 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).
2. $\psi_i(c_{i,t}|c_{i,t-1})$ is strictly increasing i.e., a higher type in one period implies a higher 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 if this will imply that other firms will raise their prices in response in future periods. 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 cost-pass through in oligopoly, given the large effects that asymmetric information has on mark-ups.

¹⁰We 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.

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.¹¹
- 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.¹² 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.¹³

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 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 complete information Nash prices.¹⁴

¹¹With menu costs, one would expect some pooling where firms with different cost realizations choose the same price. These may be difficult to analyze.

¹²Given that I 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.

¹³In 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.

¹⁴Shapiro (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.

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 $p_{i,T}^*(c_{i,T})$ as

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

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 innovations in 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.¹⁵

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 ‘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 the beliefs that j will have about i 's $T-1$ cost at the beginning of the **next** (final) period, and 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, 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,

$$\Pi^{i,T-1}(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i}) = (p_{i,T-1} - c_{i,T-1})x \dots \quad (1)$$

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

Note that, in this form, the signaling payoff is separable between periods, as in Mailath (1988) and Mailath (1989), because price and output conditions in period $T - 1$ only affect the flow payoff from that period, holding j 's inference about i 's cost fixed.

We will focus on a fully separating equilibrium, so that each firm's pricing decision exactly reveals its marginal cost to the other firms. While there may be other equilibria that involve

¹⁵Existing 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.

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' 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. Each firm's best response pricing strategy, will be given, holding beliefs about j 's pricing fixed, 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 with respect to the n^{th} argument, and an initial value condition, where $p_{i,T-1}^*(c_i)$ 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 is the static best response, given that i has the lowest possible marginal cost, to the other firms' expected pricing strategies). 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 that the other firms infer that $c_i = \underline{c}_i$ (i.e., they infer the lowest possible cost).

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 $(c_{i,T-1}, \widehat{c_{i,T-1}^j}, \widehat{\zeta_{j,T-1}^i})$, $\Pi^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right)$ 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 $\Pi_{33}^{i,T-1} \left(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i} \right) > 0$, 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 The signaling payoff function is additively separable so that, holding the future beliefs of the rival fixed, the current price only affects a firm's payoffs in the current period. In the current setting, where a firm may want to signal that its marginal cost is high by raising its price, this condition implies 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 $(c_{i,T-1}, \widehat{c_{i,T-1}^j}, p_{i,T-1})$.

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 when future prices are above static best response levels, as it could be the case that a higher a rival's incentive to drop its price towards the static best response becomes stronger when it expects its rival to set a higher price.

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 all $p_{i,T-1}$ 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.

These conditions parallel those in a single-agent signaling problem, as discussed in Mailath (1988), who solved a technical problem to prove that conditions on best responses can be used to show that a fully-separating equilibrium exists.¹⁶ Unfortunately, two limitations are associated with these conditions. First, it is difficult to express them in terms of model primitives (such as demand and costs), 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 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 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. $V_{i,T-2}$ can then be calculated, and the same procedure applied to $T - 3$, etc.. In the first period of the game, one can assume that firms enter the game knowing some prior, fictitious marginal costs (in which case equilibrium prices could be like those in the second period). As long as $t = 1$ strategies are fully separating what is assumed about these initial beliefs does not affect the rest

¹⁶Mailath (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.

of the game. In practice, once one has gone some way from the end of the game (say, 15 to 25 periods), pricing strategies tend to converge to being almost perfectly stationary (i.e., the same in period t and $t+1$) as long as the conditions stated above are always satisfied. It will be pricing strategies in these periods that will be the subject of our analysis below.¹⁷

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 $\{8, 8.0083, 8.0167, 8.0250, 8.0333, 8.0417, 8.0500, 8.0583, 8.0667, 8.0750\}$.¹⁸ 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.¹⁹ 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

¹⁷Once 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.

¹⁸For the $N = 2$ example below, the strategies differ by less than one cent if we use 20 gridpoints.

¹⁹So, 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).

each firm where are allowing for the possibility that i 's actual $T - 1$ marginal are different from those perceived by firm j .²⁰

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

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

Step 1. 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. We verify belief monotonicity at this point.

Step 2. For each set of beginning of period point beliefs about each firm's prior 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.²¹

(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 significantly 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 $(\Pi_3^{i,T-1})$.²²

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

(d) Solve $\Pi_3^{i,T-1}(\widehat{c_i}, \widehat{c_{i,T-1}^j}, p_{i,T-1}, \widehat{\zeta_{j,T-1}^i}) = 0$ to find the lower boundary condition for i 's pricing function (using a cubic spline to interpolate the vector calculated in (b)).²³

(e) Using the boundary condition as the starting point of the pricing function when $c_{it} = \underline{c_i}$,

²⁰Under duopoly with a 10 point actual and perceived cost grid, $V_{i,T}$ is stored as a 10 x 10 x 10 array.

²¹We do not claim that this iterative procedure is optimal, although it works well in our examples. There are close parallels between our problem and variants of asymmetric first-price auction problems where both the lower and upper bounds of bid functions are endogenous. See, for example, Hubbard and Paarsch (2013) for discussion.

²²A 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.

²³In 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.

solve the differential equation to recover i 's best response pricing function. This is done using `ode113` in MATLAB.²⁴ We then use cubic spline interpolation to get values for the pricing function at the points on the cost grid.

(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,

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

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

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 $T = 50$ and $T = 100$ 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 solution to the model, to show that the pricing effects of asymmetric information can be very large in a dynamic model; (ii) to provide some simple examples of how merger counterfactuals that ignore the effects of asymmetric information may go astray; and (iii) to give some intuition

²⁴In 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 26.

about why the conditions laid out above can fail for the types of demand that we consider.

3.1 Parameterization

We assume the following parameterization with $N = 2, \dots, 4$ symmetric firms. 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

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

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 standard deviation of the innovations and the limited range of costs, firm costs can change quite quickly from high to low values, or vice-versa.

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

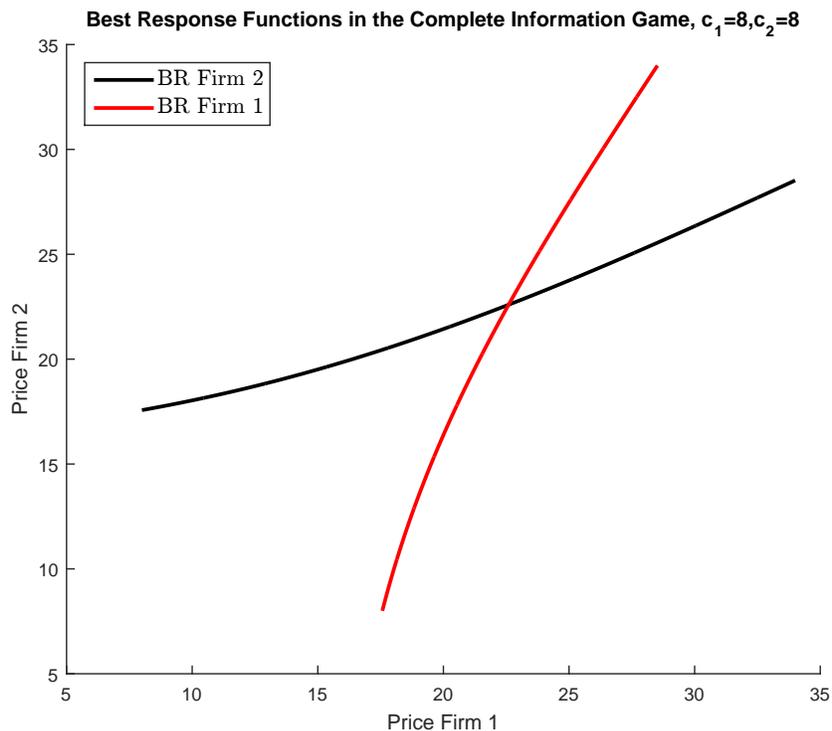
$$u_{person,0} = \varepsilon_{person,0}.$$

3.2 Analysis with $N = 2$

A feature of 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). 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. For example, the average complete information Nash equilibrium price is 23.63.

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

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

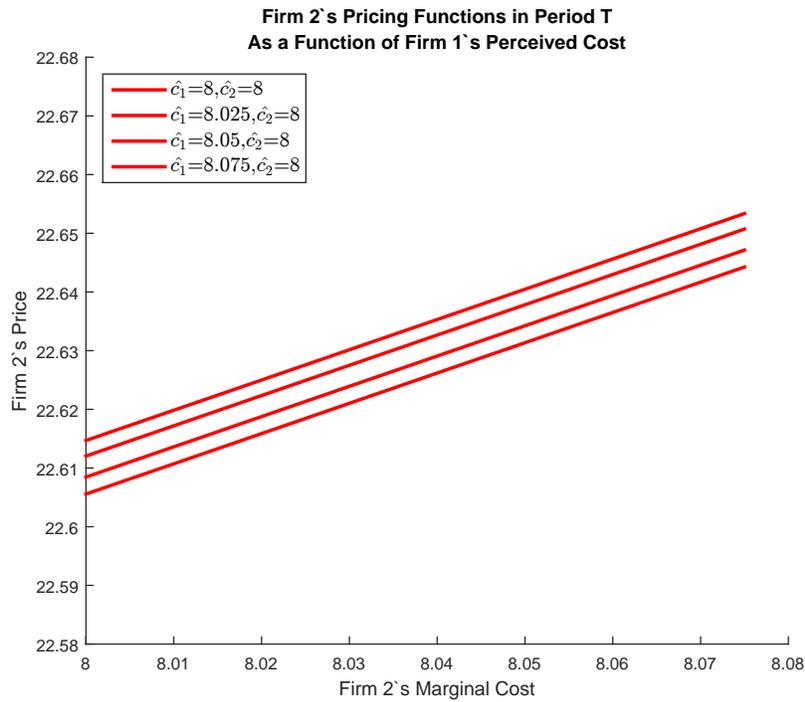
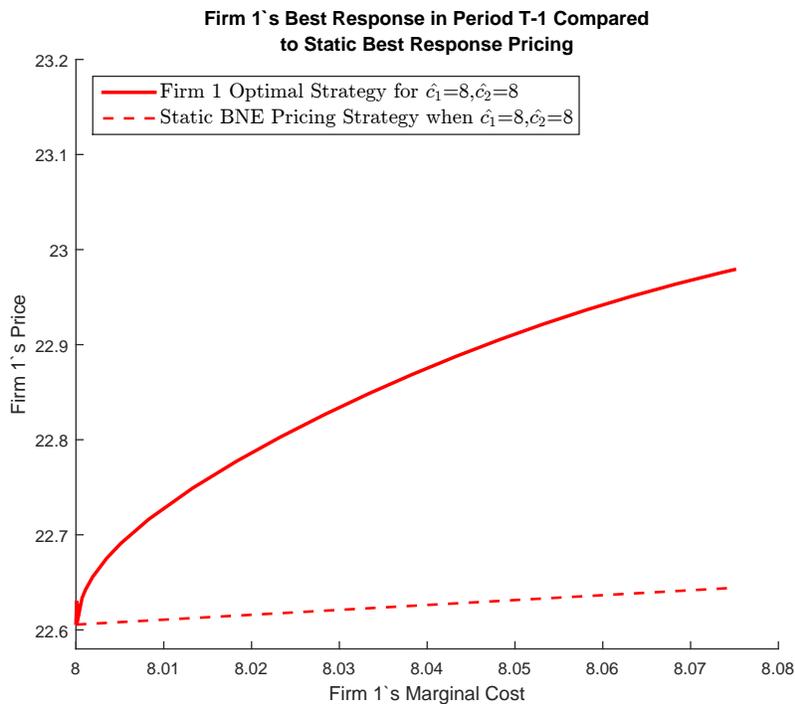


Figure 3: Firm 1's $T - 1$ Best Response Signaling Strategy When $\widehat{c}_{1,T-2} = \widehat{c}_{2,T-2} = 8$ And Firm 2 Uses Its Period T/Static Bayesian Nash Equilibrium Pricing Strategy



as the solution to the differential equation (2). 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 BNE 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

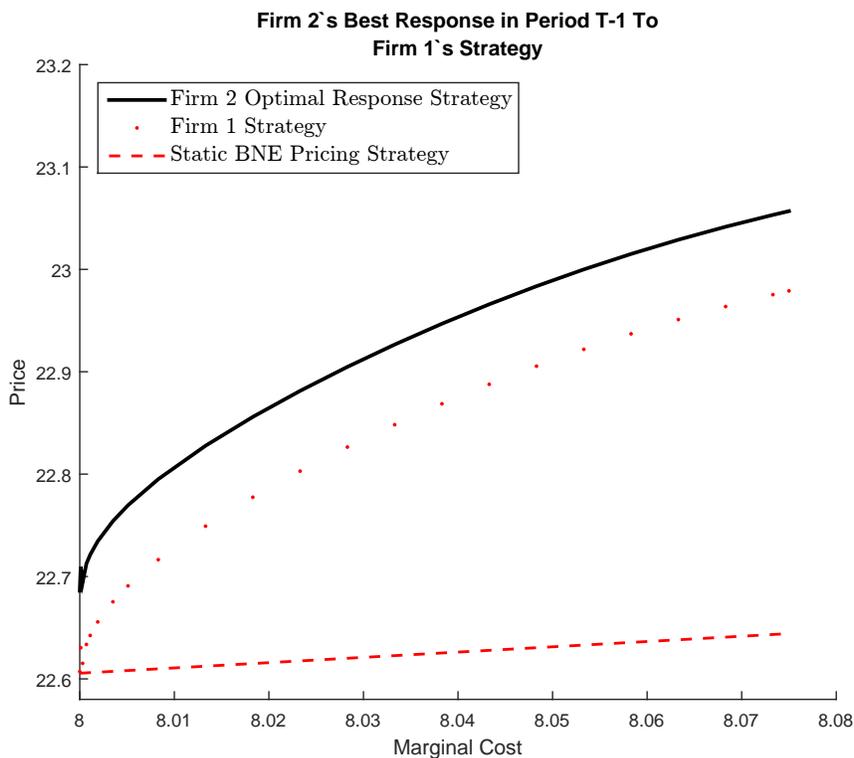
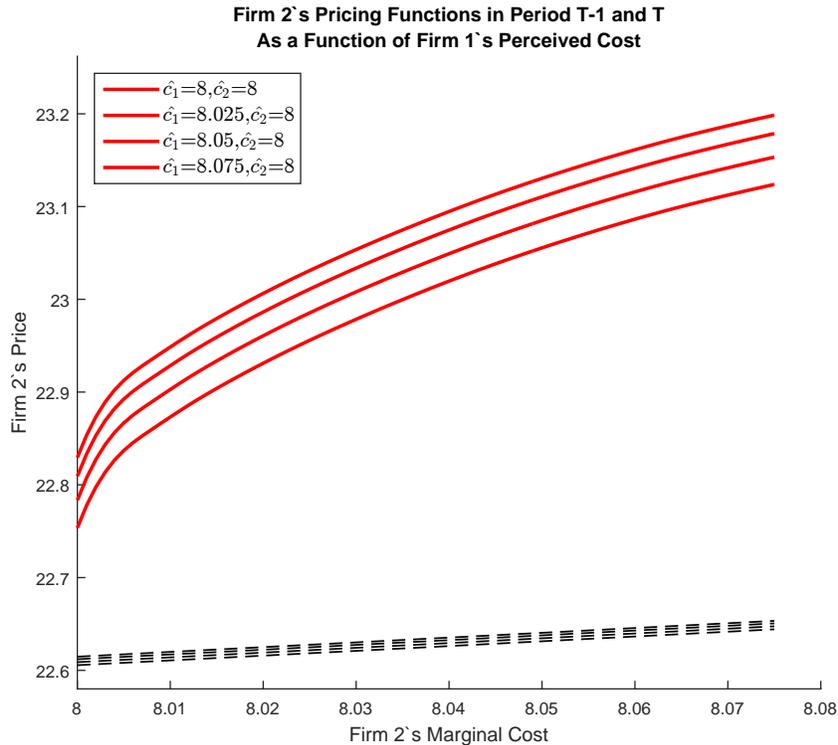


Figure 4 shows firm 2's best response signaling strategy when firm 1 uses the strategy shown in Figure 3. Now, because firm 1's prices have risen for almost all costs, firm 2's static best response will be higher, and the lower boundary condition of firm 2's pricing function is translated upwards, and higher prices are the optimal signaling response at all cost realizations. Of course, this

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



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 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 $T - 1$. 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 For Different $\widehat{c_{1,T-2}}$ Compared With T Strategies, Conditional on $\widehat{c_{2,T-3}} = 8$. Firm 1's Strategies Given $\widehat{c_{1,T-3}} = 8$ Are Symmetric.

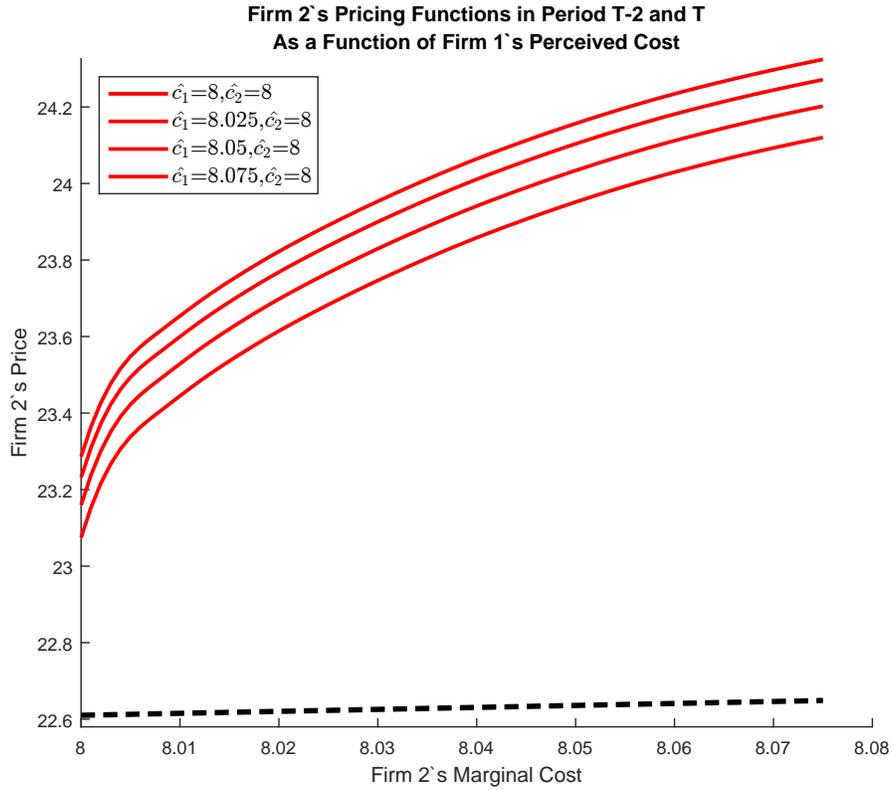


Figure 7: Firm 2's Equilibrium $T - 19$ And $T - 20$ Signaling Strategies For Different $\widehat{c_{1,T-x}}$, Conditional on $\widehat{c_{2,T-20}} = 8$ or $\widehat{c_{2,T-21}} = 8$.

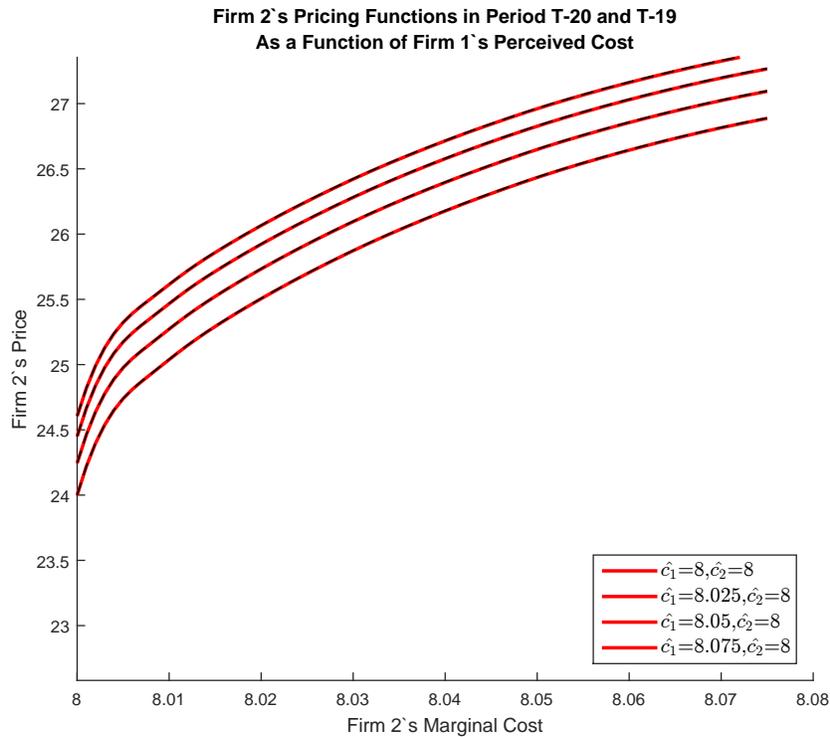


Table 1: Nested Logit Example with 2 to 4 Symmetric Firms: Average Prices

	$N = 2$	$N = 3$	$N = 4$
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%)	0.71 (3.9%)
Δ P-MC Markup	3.79 (26%)	1.13 (10%)	0.71 (7.1%)

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 prices 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.

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 the researcher assumes that firms are symmetric and observes the true form of demand and the average price prior to a merger that will take the industry either from 4 to 3 firms, or from 3 to 2 firms. Based on the standard procedure for inverting the pricing first-order conditions to find marginal costs, he would calculate that (average) marginal costs are 8.83 with four firms or 9.19 with three firms, compared with the true average marginal costs of 8.0375. The researcher’s calculations are shown in the second row of Table 2.

Based on this estimate, and maintaining the assumption that firms play a symmetric complete information Nash equilibrium, the researcher can calculate the marginal cost reduction that would be required to keep average prices from increasing after a merger that reduces the number of firms by one. Here we make the strong and unrealistic assumption that the firms remain symmetric, single product firms after the merger so that the synergy is assumed to be realized for all firms, not just those involved in the merger. Because of the small number of firms in these examples, the required synergies are substantial, amounting to more than one-third of estimated marginal

Table 2: Hypothetical Merger Counterfactual with Symmetric, Single Product Firms

	3 to 2	4 to 3
Pre-Merger Price (Signaling Eqm)	20.32	18.71
Marginal Cost Under Complete Information	9.19	8.83
Required Synergy to Prevent Price Increases (Compl. Info.)	3.55	1.30
Post-Merger Price with Synergy (Signaling Eqm)	24.54 (+15.9%)	19.87 (+6.2%)

costs in the ‘3 to 2’ case.

The final row reports average realized prices if this synergy is realized (so that, for example, in the ‘3 to 2’ case, average marginal costs fall from 8.0375 to $8.0375 - 3.55 = 4.4875$), but firms play the signaling equilibrium ex-post. Recall that the researcher would expect 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.

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., firms 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))

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

i.e., the marginal profitability of firm i increasing its price increases in j 's price. While this may hold for prices close to static best response prices, it may not hold when considering prices substantially above static best response levels. Suppose that p_i is substantially above the static best response level, because, for example, i is signaling. If p_j is low, the cost to firm i of increasing its high price even further may be small, because its quantity is small, and it might even be small if it responded by setting its static best response price. On the other hand, if p_j is very high, i might be able to get a much larger increase in both quantity and profit by lowering its price towards the static best response. If so, the strategic complementarity condition will not hold.

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 JV as documented by Miller and Weinberg (2015) (MW hereafter).

MW attribute these price increases to an increase in tacit collusion after the JV, based on a framework where there is complete information about demand and marginal costs. Tacit collusion is a potential explanation for the observed pattern, and there is nothing in the present draft that will argue that asymmetric information provides a superior explanation. However, we will suggest that it at least comes close to providing an alternative explanation, which, if correct, would have implications for further antitrust actions that might be pursued in this industry (for example, concerning facilitating practices that might increase transparency). 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 areas where beer is sold in grocery stores, as Miller and Weinberg (2015), 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 2008.²⁵ 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 but it is appropriate to the extent that, 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 popular with female drinkers.²⁶ Light beers have come to account for more retail sales than their full-calorie brand partners, especially for domestic brewers. Bud Light, Miller Lite and Coors Light were the top three retail brands in 2008, with traditional Budweiser placing fourth. On the other hand, for

²⁵In 2009, Anheuser-Busch was bought by InBev to form ABI. Given our sample selection this does not directly affect our analysis.

²⁶Substitution patterns and market shares for beers sold in restaurants, bars and specialist liquor stores may look quite different.

Table 3: Summary Statistics for Included Light Beer Brands

Product	Share of Light	Price	
	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

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 important 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). These beers have significantly lower price points than the flagship brands, so incorporating them in the model (at least if they have uncertain marginal costs) would also introduce the multi-dimensional signaling problem, although, assuming that there is some substitution between these lower-priced brands and the flagship brands, excluding them 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 3 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.

Table 4: Light Beer: Estimated Price Effects of the JV

Dep. Var.	(1)	(2)
	$\mathbb{I}\{\text{ABI or MC}\} \times \mathbb{I}\{\text{Post-JV}\}$ Log(Price)	coefficient Price (\$)
Product-Store FEs &		
1. Month FEs	0.076 (0.008)	2.160 (0.179)
2. Product x Month Trends	0.023 (0.007)	0.467 (0.143)
3. Month FEs, Import Time Trend	0.047 (0.008)	1.104 (0.192)
4. Month FEs, Product Time Trend	0.048 (0.007)	1.156 (0.173)
Observations	578,967	

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

$$\text{Log}(\text{Price}_{i,t,s}) \text{ or } \text{Price}_{i,t,s} = \mathbb{I}\{\text{brand of ABI or MC}\} \times \mathbb{I}\{\text{Post-JV}\} \theta + \dots \\ \text{Month FE}_t + \text{Product X Store FE}_{i,s} + \varepsilon_{i,t,s}$$

where prices are measured at the store-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. Visual inspection of price-paths 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 4, with standard errors clustered on the IRI region. The estimates of θ , which measures the increase in prices for the domestic brands after the completion of the JV relative to imported brands, indicate post-JV prices increases of 7.6% or \$2.16 per 24-pack. Instead of month fixed effects, the second row includes product-specific time trends. The estimate of the JV effect falls to 2.3 or 47 cents/24 pack, although it remains statistically significant at the 1% level. The third and fourth rows include month FEs, together with an import-specific time trend, or product-specific time trends (the most general specification in the table). These results indicate

price effects of a little under 4% or just over \$1 per 24 pack. We now turn to the question of whether asymmetric information about serially-correlated marginal costs can generate this type of effect.

4.2 Demand

While we could use the type of random coefficient demand model preferred by MW, 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 quite heterogeneous.

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),

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

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 typically when estimating discrete-choice demand models, somewhat arbitrary and we define it as 110% of the maximum light beer sales observed at the store in the calendar year, which assumes that the month fixed effects are sufficient to pick up the fact that demand varies across months (for example, it will be higher around the Superbowl and national holidays). 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 diesel²⁷; (ii)

²⁷For domestic beers we assume that the beer is shipped from the firm’s closest brewery (which changes for

Table 5: 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 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.²⁸

Column (1) of Table 5 reports the estimates from the baseline specification. The nesting coefficient is highly significant, and implies that domestic products and imports are poor substitutes for most consumers. The magnitude of the nesting coefficient also implies that each product's demand is very price elastic: the average domestic elasticity is about -7.²⁹ 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 a nesting coefficient of 0.7 and price coefficient of -0.2, reflecting the coefficients in column (1).

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.

²⁸The 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.

²⁹This elasticity is higher than is typically estimated for these products, reflecting the introduction of a nesting structure that identifies the leading domestic brands as being close substitutes. It is, however, consistent with the fact that the domestic brands have very similar prices, both before and after the JV, in different cities (even if the level of those prices varies) where different brands should have different distribution costs. 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.

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 hard, although we plan to try to add some greater richness 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.

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) but which we assume are not engaged in active signaling themselves.³⁰

We assume that the pre-JV marginal cost of AB (for a 24 pack volume-equivalent) may range from \$15.45 to \$15.60; for Miller (for a 24 pack volume-equivalent) from \$16.25 to \$16.40; and for Coors (for a 24 pack volume-equivalent) from \$16.25 to \$16.40, so the uncertainty is around approximately the last 1 $\frac{1}{4}$ % 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. 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

³⁰For 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.

Table 6: 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

normal distribution that has mean zero and standard deviation 0.05.³¹ 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. The cost ranges and qualities were chosen so that, at average 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. Average complete information prices are slightly lower than is observed immediately before the JV, although this was partly done in order to see if asymmetric information prices were closer to those that are observed.

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 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,\$16.40] to [\$15.40,\$15.55], 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 6 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

³¹While brand prices display 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: Average Prices Under Complete Information and Asymmetric Information

Product	Pre-JV		Post-JV	
	Complete Information	Signaling Equilibrium	Complete Information	Signaling Equilibrium
Bud Light	18.31	18.43	18.30	19.03
Miller Lite	18.31	18.46	18.30	19.03
Coors Light	18.31	18.46	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

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 setting 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 increase 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 7.

Prior to the JV, signaling has quite small effects on prices given the estimated demand and assumed cost parameters, raising the prices of Miller and Coors, for instance, by 15 cents or just under 1%. Interestingly in this example, the prices of the firms with smaller market shares are slightly more affected by signaling than those of Anheuser-Busch which has roughly the same market share as these firms combined. 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 identification of collusive behavior, as ordinarily it would not be predicted by a static differentiated products model with only unilateral effects. 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, the conditions required to characterize best responses once we move several periods from the end of the game.

5 Conclusion

In an environment where firms compete in strategic complements, 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 differentiated products framework where there is small uncertainty about firms marginal costs, which extends the simple, linear demand 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 underpredict 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. Three examples can illustrate. First, we would like to investigate how large the effects are in a setting where firms compete in quantities, as they might be expected to do in commodity or electricity markets, for example, as they should work in the opposite direction to our price-setting example. Second, we would like to see the interaction between price signaling and choices of capacity. In a single-agent signaling framework, Sweeting, Roberts, and Gedge (2016) show that when a dominant incumbent is able to respond to the threat of entry by both limit pricing and investing in observable excess capacity to deter entry, the firm may engage, in equilibrium, in Milgrom and Roberts (1982)-style limit pricing without investing in additional capacity, partly because extra capacity may make signaling more costly. It would be interesting to see, maybe motivated by recent claims of capacity coordination between major airlines, how capacity choices in a concentrated oligopoly situation when costs or demand are uncertain. Finally, we plan to look at what happens in models such as the learning-and-forgetting duopoly model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) where the possibility of forgetting, which is assumed to be observed, introduces some equilibria with very aggressive pricing in certain states. It may be more realistic to assume that while cumulative sales, which drive learning, are observed by other firms, forgetting is not, which may provide firms with incentives to use prices in order to signal information about their costs, which could potentially either exacerbate or soften incentives to price aggressively.

References

- ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica*, pp. 51–75.
- ASHENFELTER, O. C., D. S. HOSKEN, AND M. C. WEINBERG (2015): “Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry,” *The 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.
- 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): “Database paper-The IRI marketing data set,” *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.
- GOEREE, J. K. (2003): “Bidding for the Future: Signaling in Auctions with an Aftermarket,” *Journal of Economic Theory*, 108(2), 345–364.
- HAILE, P. A. (2003): “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 108(1), 72–110.

- 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, volume 3*, vol. 3, pp. 37–115. Newnes.
- JAYARATNE, J., AND J. ORDOVER (2015): “Coordinated Effects: Evolution of Practice and Theory,” in *The Oxford Handbook of International Antitrust Economics, Volume 1*, ed. by R. D. Blair, and D. D. Sokol, pp. 509–528. Oxford University Press.
- 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,” *The 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. (1992): “Perpetual Signalling with Imperfectly Correlated Costs,” *The Rand Journal of Economics*, pp. 548–563.
- MILGROM, P., AND J. ROBERTS (1982): “Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis,” *Econometrica*, 50(2), 443–459.
- MILLER, N. H., AND M. C. WEINBERG (2015): “Mergers Facilitate Tacit Collusion: Empirical Evidence from the US Brewing Industry,” Discussion paper, Georgetown University.
- MOLNÁR, J., AND G. VIRÁG (2008): “Revenue Maximizing Auctions with Market Interaction and Signaling,” *Economics Letters*, 99(2), 360–363.

- 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.
- 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.
- 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.
- WEINBERG, M. (2008): “The Price Effects of Horizontal Mergers,” *Journal of Competition Law and Economics*, 4(2), 433–447.
- WHINSTON, M. D. (2008): *Lectures on Antitrust Economics*. MIT Press.