Should We Expect Merger Synergies To Be Passed Through to Consumers?

Mario Leccese
University of Maryland

Andrew Sweeting
University of Maryland, CEPR and NBER*

Xuezhen Tao
Shanghai University of Finance and Economics

Monday 11th March, 2024

Abstract

When reviewing horizontal mergers, antitrust agencies balance anticompetitive incentives, resulting from market power, with procompetitive incentives, created by efficiencies, assuming complete information and static, simultaneous move Nash equilibrium play. These models miss how a merged firm may prefer not to pass through efficiencies when rivals would respond by lowering their prices. We use an asymmetric information model, where rivals do not observe the size of the realized cost efficiency, to investigate how this incentive could affect post-merger prices. We highlight how the strength of this incentive will depend on the market structure of non-merging rivals and discuss alternative settings where similar issues arise.

JEL CODES: L1, L13, L4.

Keywords: oligopoly, horizontal mergers, vertical mergers, asymmetric information, pooling equilibria, firm conduct, pass-through, strategic incentives.

*Sweeting (corresponding author): 3114 Tydings Hall, University of Maryland, College Park MD 20742, USA. Email: atsweet@umd.edu. All estimates and analysis using IRI data are the work of the authors and not IRI. The counterfactual results reported in Section 3 are from theoretical simulations, and do not constitute real market statistics or measures of performance for the firms considered in this paper. This research was supported by US National Science Foundation Grant SES-1658670 (Sweeting) and National Natural Science Foundation of China Grant 72003120 (Tao). Xinlu Yao collaborated on this project when she was a graduate student at the University of Maryland, and we are very grateful for Xinlu's excellent work and insights. We thank Joe Harrington, Volker Nocke, Philipp Schmidt-Denger, Dan Vincent, Jonathan Wallen and seminar/conference participants at Mannheim (MaCCI), the Federal Trade Commission, West Virginia, Auburn, the CEPR VIO Seminar, IIOC 2022, the US Naval Academy and the US Department of Justice for useful comments. Two referees and Ryan McDevitt, the coeditor, provided helpful comments. All errors are our own.

1 Introduction

Merger review usually involves finding the net balance of anti-competitive effects, due to enhanced market power or foreclosure, and procompetitive effects, arising from marginal cost reductions, better products or, in the context of vertical mergers, supply chain coordination or the elimination of double marginalization. The pricing pressure and merger simulation calculations currently used for balancing assume that post-merger pricing will be determined as a complete information, static, simultaneous move Nash equilibrium (CISSNE). CISSNE calculations may miss how, in some circumstances, the merged firm may be in a position to make a choice about whether to pass-through, or realize, efficiencies and that passing-through may not be profitable once the fact that rivals may respond by lowering their prices is accounted for.

In this article, focusing on cost efficiencies after horizontal mergers, we show that the incentive not to pass-through efficiencies may be more significant than has been appreciated. By "significant," we mean that, for example, a merged firm that benefits from an efficiency that would be large enough to make the merger lower prices in a traditional analysis may prefer to price as if it has benefited from no efficiency at all. We find that pooling incentives are significant under a wide range of market structures for three alternative demand systems, and when we consider hypothetical mergers in the beer industry using the demand and marginal cost estimates from Miller and Weinberg (2017), a setting that has previously been used to illustrate merger efficiency calculations (Caradonna, Miller and Sheu (2021) and Nocke and Whinston (2022)).

To perform our analysis in a horizontal merger setting, we adapt the framework of Harrington (2021) who considers a repeated game setting where the merged firm has private information about the size of the realized cost synergy or the date when the post-merger synergy is realized.[1] The information asymmetry leads to the possibility of a pooling equilibrium where firms price as if no efficiency, or a small efficiency, has been realized even when the realized efficiency is large. The logic is straightforward. Suppose that there is common knowledge of demand and the costs of non-merging firms, but that, after a merger, rival firms attach positive probabilities to the merged firm's marginal costs being either high or low. Firms play a repeated pricing game after the merger and non-merging firms always best respond to the prices that they expect the merged firm to set. If the non-merging firms will interpret the merged firm setting a low price as implying that the efficiency must be large, so that they will lower their prices in future periods, then a patient merged firm with a large efficiency may prefer to set high prices.[2]

To quantify the potential for pooling we calculate the critical discount factors that can support pooling on no efficiency prices when efficiencies that would eliminate competitive concerns in a CISSNE analysis are realized, and, alternatively, the ranges of efficiencies that can support pooling when we assume a typical discount factor. We show that there tends to be more scope for pooling, and for pooling to raise prices, in mergers that do not involve the largest firms in the industry and when the market structure of rivals is more concentrated. The latter result reflects how rivals' prices will tend to drop by more in response to a price decrease by the merged firm if rivals' have more market power. This

[1] Harrington's paper was published at the time that our paper was being developed. As explained below, we will focus on issues of magnitude and the role of rivals' market structure, that are not emphasized in Harrington's paper.

[2] While the model implies a critical discount factor to sustain pooling, an important difference to a model of tacit collusion is that all of the non-merging firms use static best response pricing.

finding is policy relevant given the influential paper of Nocke and Whinston (2022) who show that the efficiencies required to offset anticompetitive effects in a CISSNE framework will depend only on the market shares of the merging firms. On this basis, they suggest that the agencies should place less weight on the measures of market-wide concentration, such as the level of the predicted post-merger Herfindahl-Hirschman Index (HHI), when screening for anticompetitive mergers. In our framework, the concentration of the rest of the market can play a very important role in determining post-merger price effects.

Efficiencies in horizontal mergers have been the subject of considerable recent debate, in light of findings from the merger retrospective literature that prices have risen after many horizontal mergers that have not been challenged (Ashenfelter, Hosken and Weinberg (2014), Kwoka (2014) and Asker and Nocke (2021)). Rose and Sallet (2019) argue that agencies tend to overestimate the efficiencies that are realized. However, we read the admittedly small retrospective literature on post-merger productivity changes as suggesting that cost efficiencies are realized (Haynes and Thompson (1999), Bitzan and Wilson (2007), Groff, Lien and Su (2007), Braguinsky et al. (2015), Kulick (2017), Grieco, Pinkse and Slade (2018), Walia and Boudreaux (2019) and Yan et al. (2019)), suggesting that the issue may be that merging firms do not pass through efficiencies to customer prices in the way that agencies' current models assume.[3] This idea is consistent with Ashenfelter et al. (1998) and Muehlegger and Sweeney (2019) who estimate, in non-merger settings, that pass-through of firm-specific cost reductions is very limited. We believe that the types of organizational changes that happen around mergers make it especially plausible that asymmetries of information could

[3]However, Blonigen and Pierce (2016), using a cross-industry panel, suggest productivity does not increase on average.

exist.[4] We discuss how we would incorporate our analysis into horizontal merger review, including when it may be appropriate to require merging parties to provide evidence of their pass-through incentives.[5]

Our analysis in the text focuses on cost efficiencies after horizontal mergers. However, the profitability of realizing other types of efficiencies, such as the introduction of a product improvement or the elimination of double marginalization after a vertical merger, will also depend on how much rivals would lower their prices in response. We illustrate these points in Section 5. We believe that our discussion of the pass-through of the elimination of double marginalization (EDM) is especially timely given that the question of whether agencies should assume that EDM will be realized and passed through has been at the heart of debates about vertical merger enforcement.[6]

The structure of the rest of the paper is as follows. After a brief description of related

[4]Section 10 of the 2010 Horizontal Merger Guidelines state that "[e]fficiencies are difficult to verify and quantify, in part because much of the information relating to efficiencies is uniquely in the possession of the merging firms. Moreover, efficiencies projected reasonably and in good faith by the merging firms may not be realized. Therefore, it is incumbent upon the merging firms to substantiate efficiency claims so that the Agencies can verify by reasonable means the likelihood and magnitude of each asserted efficiency, how and when each would be achieved (and any costs of doing so)". The 2023 Merger Guidelines contain a similar, but slightly shorter, discuission. Bernile and Bauguess (2011) estimate that pre-merger claims of synergies by U.S. public companies are not correlated with realized changes in firm performance once firm characteristics are controlled for. This suggests that even the merged firm may not have a clear view of how large synergies will be before consummation.

[5]Yde and Vita (1995) describe requiring parties to provide additional evidence of pass-through as "entirely at odds with the basic economic theory upon which modern antitrust law is based." Our proposal would move the burden of evidence onto firms in a limited set of circumstances where our economic model would suggest that incentives not to pass-through efficiencies might raise post-merger prices.

[6]When the merging parties are in a supply relationship and there are positive margins at both levels, it is common for economists to argue that there will be an inherent incentive to lower downstream prices as a result of EDM, and to use merger simulation models where the realization and pass-through of EDM are assumed. Non-economist practitioners are often unconvinced by EDM claims. For example, when withdrawing the 2020 Vertical Merger Guidelines, the FTC's Democratic majority argued that their "reliance on EDM ... [was] theoretically flawed because the economic model predicting EDM is limited to very specific factual scenarios" and that empirical evidence suggested "that we should be highly skeptical that EDM will even be realized—let alone passed on to end-users." See `https://www.ftc.gov/system/files/documents/public_statements/1596396/statement_of_chair_lina_m_khan_commissioner_rohit_chopra_and_commissioner_rebecca_kelly_slaughter_on.pdf` (accessed Feb 14, 2022).

models in the literature, Section 2 presents the model of a horizontal merger. Section 3 presents a multinomial logit example that fixes ideas, and then describes the results from considering a large number of simulated market structures using three demand systems. Section 4 presents the results using beer demand estimates. Section 5 presents extensions. Section 6 concludes. When we mention our working paper, we are referring to Sweeting, Leccese and Tao (2022).

Literature. As noted by Weyl and Fabinger (2013), the theoretical literature on oligopoly pass-through, which has been used to study many questions, including tax incidence, has traditionally relied on CISSNE assumptions. However, a recent literature, to which our paper contributes, has started to investigate the effects of relaxing these assumptions. Like ours, many of these papers are motivated by mergers. For example, MacKay and Remer (2022) and Brown and MacKay (2023) use the Markov Perfect Nash equilibrium concept of Maskin and Tirole (1988), where there is no persistent asymmetric information, to investigate how consumer inertia or the use of automated pricing algorithms may change the price effects of horizontal mergers. Farrell and Baker (2021) argue that merger analysis should consider models other than CISSNE competition and tacit collusion.

Shapiro (1986), Vives (2011) and Amir, Diamantoudi and Xue (2009), who consider an uncertain merger synergy, consider one-shot games with asymmetric information models where the strategic incentive to affect rivals' future prices that drives our results does not play a role. Aside from Harrington (2021), a handful of papers combine some form of asymmetric information and repeated pricing. For example, Bonatti, Cisternas and Toikka (2017) consider a model where quantity-setting firms learn about their rivals' costs, which are fixed, from a noisy market price. This creates signal jamming incentives, with quantities converging to their static complete information equivalents over time. In our repeated pricing game, pooling on higher prices can potentially last forever.

Sweeting, Tao and Yao (2023) consider a repeated price-setting game where several oligopolists simultaneously use separating pricing strategies to signal private information about their marginal costs, which can change from period-to-period. Prices in a fully separating Markov Perfect Bayesian equilibrium can be significantly higher than in either a CISSNE model or a static Bayesian Nash equilibrium model. Sweeting, Tao and Yao (2023) find that their model predicts larger price increases after horizontal mergers (that do not result in monopoly) than standard models. However, the model is not well-suited to consider uncertainty about merger efficiencies, because the equilibrium can only be characterized when each firm's marginal cost has a narrow support whereas agencies often identify a wide range for possible efficiencies. Solving the STY model also involves more computation than could plausibly be used in a merger investigation, whereas the framework proposed in this article can be applied quite easily with objects that agency economists already calculate.

2 Model of a Horizontal Merger

Before a horizontal merger, $F \geq 3$ firms sell $N \geq F$ differentiated, substitute products and set prices. $q_i(\mathbf{p})$ and c_i are the demand and marginal cost of product i, owned by $f(i)$, where, using bold to indicate a vector, $\mathbf{p}$ is the vector of all prices. Prices are assumed to be observed perfectly. "Equilibrium best response" prices for a subset of firms are the prices that maximize the static profits of each firm in the subset, given other subset firm prices, and specified prices for firms outside the subset. In order to quantify post-merger price changes, we assume CISSNE pricing before any merger. With multinomial logit demand, CISSNE and equilibrium best response prices for any subset of firms will be unique (Nocke and Schutz (2018)).

Post-Merger Game. A firm M is created by an exogenous one-off merger of firms 1 and 2, and M continues to sell all of their products. Demand and the marginal costs of the non-merging firms are unaffected, and they remain commonly known.

We assume that firms play a repeated post-merger pricing game, with no further mergers. M's discount factor is δ_M. The price of an individual product is p_i. $\mathbf{p}_j$ ($\mathbf{p}_{-j}$) denotes the prices of multi-product firm (collection of firms other than) j. $\mathbf{p}^{\dagger}(\mathbf{c_M})$ denotes post-merger CISSNE prices when M has post-merger costs $\mathbf{c}_M$. $\mathbf{c}_M^{\text{PRE}}$ are the pre-merger costs of products 1 and 2. $\mathbf{c}_M$ are M's post-merger costs, where $\mathbf{c}_M^{\text{CMCR}}$ are the costs that would keep post-merger CISSNE prices at exactly their pre-merger levels. The associated efficiencies, of magnitude $\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}}$, are known as the "compensating marginal cost reductions" (Werden (1996), Froeb, Tschantz and Werden (2005)), which we will call CI-CMCR efficiencies. $\pi_M(\mathbf{c}_M, \mathbf{p}_M, \mathbf{p}_{-M})$ are M's post-merger profits, as a function of its costs, prices and rival prices. $\mathbf{p}_M^{\text{BR}}(\mathbf{c}_M, \mathbf{p}_{-M})$ maximize M's one period profits given costs and rival prices.

We will consider incentives to pass through efficiencies using a model where M's post-merger marginal costs are uncertain to rivals. Specifically, suppose that $\mathbf{c}_M = \underline{\mathbf{c}}$ (large efficiency), with known probability q, $1 > q > 0$, or $\mathbf{c}_M = \overline{\mathbf{c}}$ (small efficiency), where every element of $\underline{\mathbf{c}}$ is less than $\overline{\mathbf{c}}$. After the merger, M knows the realized $\mathbf{c}_M$, but rivals do not observe the realization, but may make inferences from the prices that M sets.

We consider the existence of the following pooling Markov Perfect Bayesian Equilibrium (MPBE) (Toxvaerd (2008), Roddie (2012*a*)). An MPBE specifies expected payoff-maximizing prices for each firm as a function of its type and, where relevant, its beliefs, which should be consistent with Bayesian updating on the equilibrium path. The Markovian restriction is that past actions only matter through beliefs. An implication is that if M's cost was to be revealed to all players, CISSNE prices would be chosen in every period.

Definition 1 ***Pooling MPBE.*** *M sets prices $\mathbf{p}_M^{\dagger}(\overline{\mathbf{c}})$ for its products in the first period and when it has done so in each previous period, and otherwise it sets prices $\mathbf{p}_M^{\mathrm{BR}}(\mathbf{c_M}, \mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}}))$. Rival firms believe that $\mathbf{c}_M = \underline{\mathbf{c}}$ with probability q and set prices $\mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})$ in the first period and when M has set prices $\mathbf{p}_M^{\dagger}(\overline{\mathbf{c}})$ in every period. Otherwise, rival firms believe that $\mathbf{c}_M = \underline{\mathbf{c}}$ with probability 1 and set prices $\mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}})$.*

In this equilibrium, the non-merging firms, and a $\overline{\mathbf{c}}$-cost M, always set prices that are equilibrium best responses to the prices that they expect other firms to set. The existence of a pooling equilibrium therefore depends on whether a $\underline{\mathbf{c}}$-cost M is willing to set prices $\mathbf{p}_M^{\dagger}(\overline{\mathbf{c}})$ (earning $\pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\dagger}(\overline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))$ per period) that are above its best response prices, $\mathbf{p}_M^{\mathrm{BR}}(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))$, in order to prevent its rivals' prices dropping from $\mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})$ to $\mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}})$ (causing M to earn $\pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\dagger}(\underline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}}))$ in all subsequent periods. If the post-merger game is infinitely repeated, the condition for the pooling equilibrium to exist is

$$\delta_M \geq \frac{\pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\mathrm{BR}}(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})) - \pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\dagger}(\overline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))}{\pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\mathrm{BR}}(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})) - \pi_M(\underline{\mathbf{c}}, \mathbf{p}_M^{\dagger}(\underline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}}))}, \tag{1}$$

where we will call the right-hand side value the "critical discount factor", δ_M^*.

Harrington (2021) derives the same condition in a closely related model.[7] One could change the model by assuming that there is some probability, ϕ, that nature publicly reveals M's marginal cost, transforming the game to one of complete information, in any period. In this case, one would replace δ_M in (1) with $\delta_M \times (1-\phi)$. As we will see in our examples, the δ_M^*s implied by (1), when we assume $\phi = 0$, are often much smaller than the discount factors that would be assumed in empirical applications. Therefore, our results imply that pooling could be supported in this changed model for plausible discount factors (say, for quarterly pricing) even if ϕ was moderately large (e.g., 0.1 or 0.2).

Our working paper shows that an equilibrium with pooling on $\mathbf{p}^\dagger(\overline{\mathbf{c}})$ will exist under the same condition on δ_M if the merged firm's realized marginal costs can be any convex combination of $\underline{\mathbf{c}}$ and $\overline{\mathbf{c}}$, i.e., it is not necessary to assume that costs can only have two possible values, as long as the supports are known. The working paper also shows that, with multinomial logit, CES or linear demand, the condition will always be satisfied for some δ_M close to 1 when the differences between $\underline{\mathbf{c}}$ and $\overline{\mathbf{c}}$ are small enough. However, small differences between $\underline{\mathbf{c}}$ and $\overline{\mathbf{c}}$ would lead to small effects on prices. In this article, we are interested in whether the price effects of pooling could be large.

We need to specify values of $\underline{\mathbf{c}}$ and $\overline{\mathbf{c}}$ to test whether a pooling equilibrium could exist. We will focus on two alternative, but complementary, specifications. The first specification has $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$. We calculate the critical discount factor to support pooling if $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$ (no efficiency) or some alternative $\overline{\mathbf{c}}$ involving either a smaller efficiency or a marginal cost increase. While merger analysis usually assumes that, at worst, a horizontal merger will

[7] Harrington assumes that the merged firm realizes the large efficiency in some (random) time period after the merger, at which time, the discount factor condition determines whether the merged firm would want to lower its prices, indicating to its rivals that the efficiency has been realized.

leave marginal costs unchanged, some corporate integrations may carry at least some risk that marginal costs will rise even if fixed cost savings are certain to be realized.

The second specification assumes that possible post-merger marginal costs are $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} - \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} + \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$ where $0 \leq \kappa \leq 1$. We identify the highest value of κ, denoted κ^*, that can support a pooling MPBE for a given discount factor by evaluating inequality (1) for a grid of κs between 0 and 1.[8] When $\kappa^* = 1$, a pooling equilibrium would lead to firms pricing as if the merged firm benefits from no cost efficiency even though the efficiency is twice as large as the CI-CMCR.

In both specifications, pooling play results in certain harm to consumers. As long as $q < 1$, the first specification would result in expected consumer surplus falling with CISSNE play, whereas, in the second specification, expected consumer surplus would increase if q is large enough.[9] However, merger case law can also be interpreted as implying that mergers should be blocked when harm to competition is more likely than not. In this case, a CISSNE analysis would imply allowing the merger in either specification if $q \geq \frac{1}{2}$.

3 Illustrative Example and Simulations.

We begin with an example which assumes multinomial logit demand and particular demand and marginal cost parameters. The example illustrates the logic behind the pooling equilibrium and our calculations, and provides intuition for how the slope of rivals' reaction functions determines whether a pooling equilibrium exists, and how this slope is related to

[8] For all of our examples, if pooling is supported for κ' it is also supported for all smaller κs.

[9] The exact value of q required to make a merger with CISSNE pricing consumer surplus neutral will depend on the exact shape of the demand curve as well as the exact value of κ. However, in all of the examples that we have looked at, the merger is almost neutral when $q = 0.5$.

rival market structure. We describe the results of simulation exercises that explore how far these intuitions generalize.

3.1 Example

Suppose that, pre-merger, there are $N = F = 3$ single-product firms with marginal costs of 4. There is multinomial logit demand with consumer c's indirect utility for good i, with price p_i, $u_{ic} = a_i - 0.25p_i + \varepsilon_{ic} = \delta_i + \varepsilon_{ic}$. The ε_{ic}s are i.i.d. logit draws. For our example, $a_1 = a_2 = 4$ and $a_3 = 6$. $u_{0j} = \varepsilon_{0j}$ for the outside good. Symmetric firms 1 and 2 merge to become firm M.

The black solid line in Figure 1 depicts the pre-merger static equilibrium best response prices of firms 1 and 2 to a value of $\log(1+\exp(\delta_3))$ (the "inclusive value" of good 3 and the outside good). The solid magenta line shows the $\log(1 + \exp(\delta_3))$ implied by best response pricing of firm 3 given $p_1 = p_2$. We describe firm 3's strategy in this way so that we can easily extend the analysis to allow for additional rivals. The pre-merger equilibrium is at A with $p_1^\dagger = p_2^\dagger = 9.03$ and $p_3^\dagger = 13.01$.

The red solid line shows M's static best response if $\mathbf{c}_M = \mathbf{c}_M^{\mathrm{PRE}}$. The new CISSNE would be at B ($p_1^\dagger = p_2^\dagger = 10.33$ and $p_3^\dagger = 13.69$). M's best response function would be the red dashed line if $\mathbf{c}_M = \mathbf{c}_M^{\mathrm{CMCR}} = 2.27$. This passes through A by construction. It has a steeper slope than the pre-merger equilibrium best response function, reflecting how M would respond to a firm 3 price increase by raising its prices by more than pre-merger firms 1 and 2, which would be competing with higher costs, would have done.

If $\underline{\mathbf{c}} = \mathbf{c}_M^{\mathrm{CMCR}}$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\mathrm{PRE}}$, then pooling would involve the firms pricing at B, even

Figure 1: Horizontal Merger Example. "BRF" is a "best response function", and "Eqm. BRF" is an "equilibrium best response" function (see text). The values on the isoprofit curves show differences in the per-period profits of a merged firm that benefits from a CI-CMCR synergy from its profits at A.

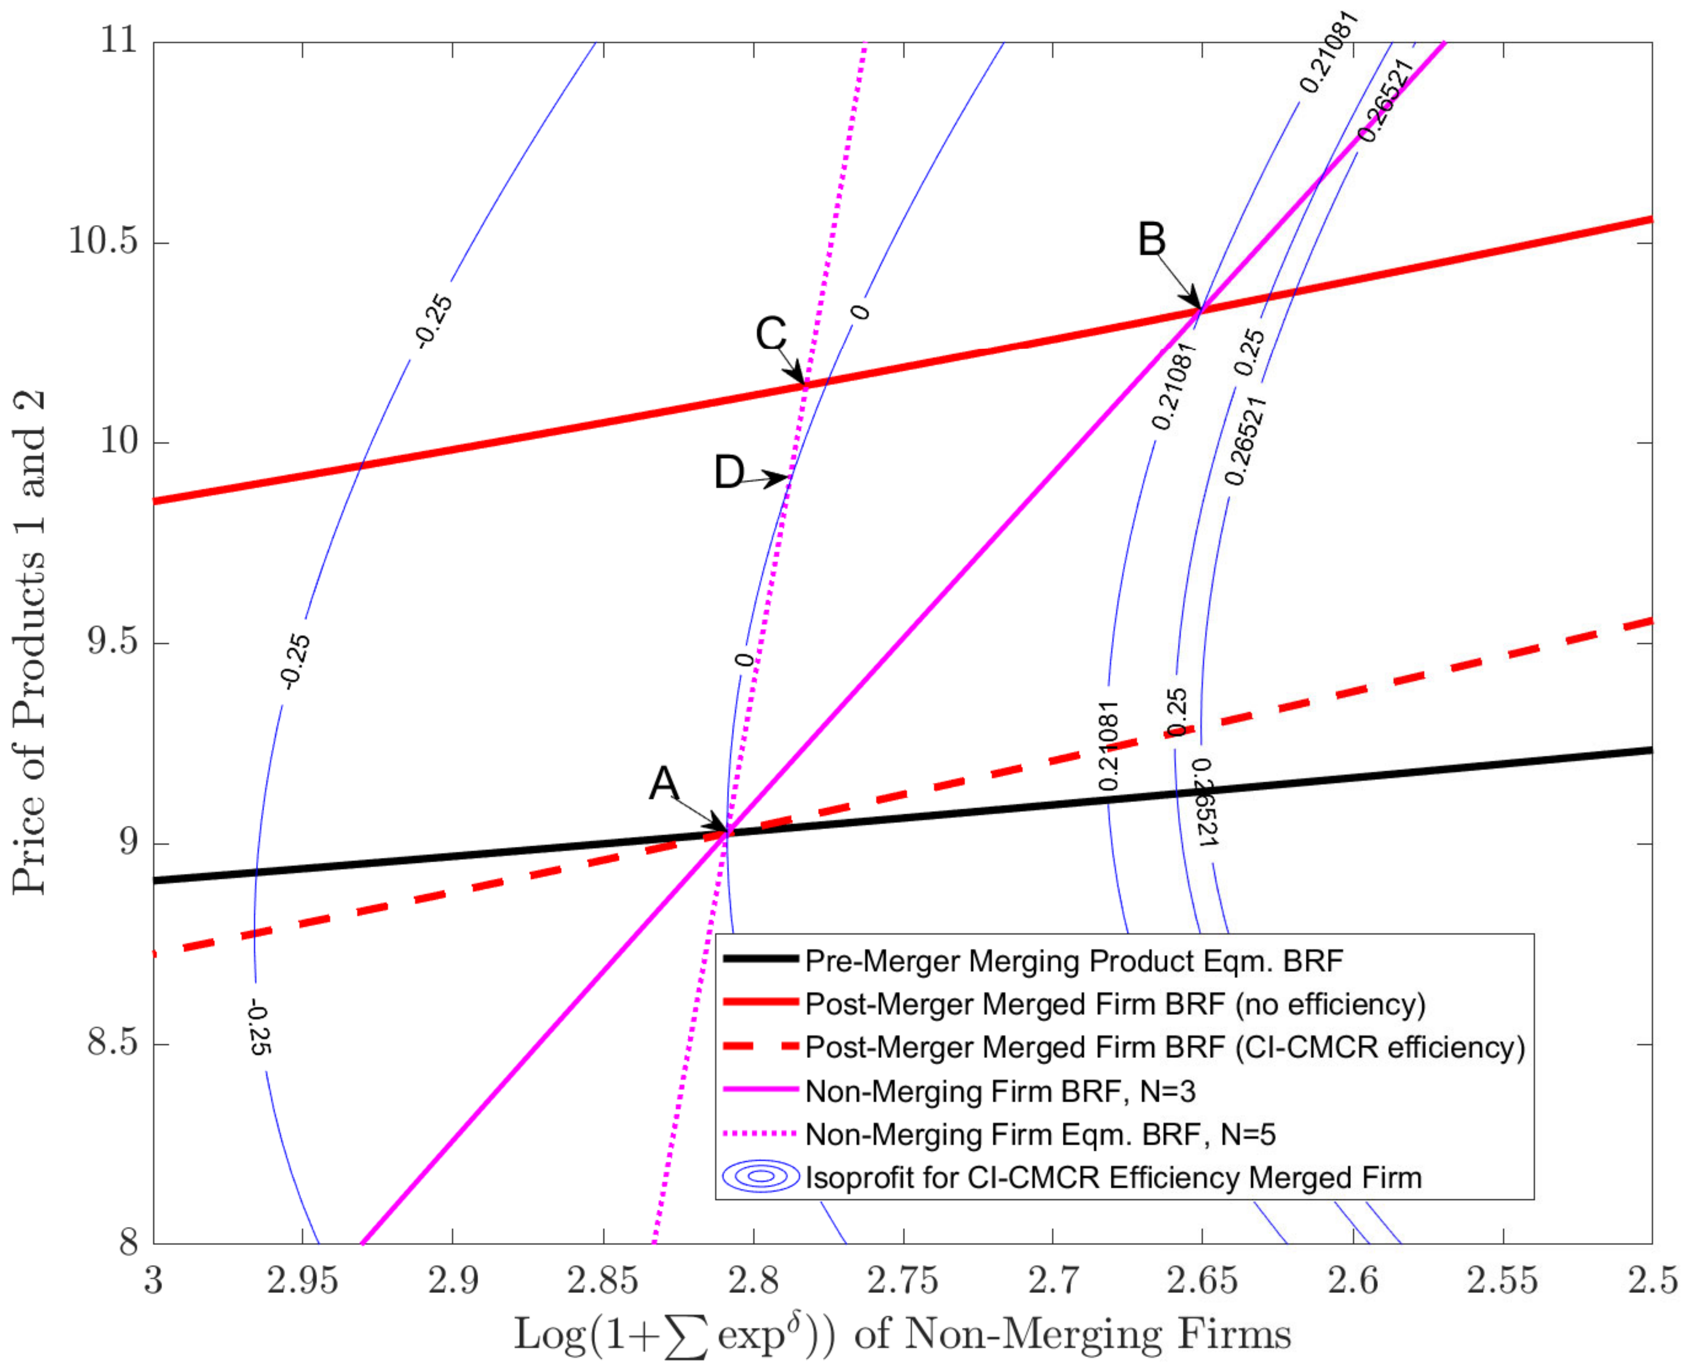

when the merged firm realizes the large efficiency. Compared with A, the outcome at B lowers consumer surplus by 0.89 per potential consumer. The blue isoprofit curves show the differences in a $\mathbf{c}_M^{\mathrm{CMCR}}$-$M$'s profits from its profits at A. Play at B raises M's per-period profits by 0.21081 (by 7.6% relative to A, and 10.2% relative to pre-merger profits). If M deviates from play at B, it could raise its profits by 0.26521-0.21081 for one period, but subequent play would be at A. Pooling can therefore be sustained if and only if $\delta_M > \frac{0.26521-0.21081}{0.26521-0} = 0.21$, so only a very low degree of patience is required.

Pooling at B could also be sustained if $\overline{\mathbf{c}} = \mathbf{c}_M^{\mathrm{PRE}}$ and $\underline{\mathbf{c}} = \mathbf{c}_M^{\mathrm{PRE}} - 2(\mathbf{c}_M^{\mathrm{PRE}} - \mathbf{c}_M^{\mathrm{CMCR}})$ as long as $\delta_M > 0.37$. Therefore, our calculated value of κ^* would be 1 using any plausible discount factor. If $\kappa = 1$, a CISSNE analysis would conclude the merger would increase expected consumer surplus if $q > 0.51$. If $\underline{\mathbf{c}} = \mathbf{c}_M^{\mathrm{CMCR}}$ we could sustain pooling with $\delta_M = 0.95$ for values of $\overline{\mathbf{c}}$ as high as 11.58 (i.e., a very large loss of efficiency relative to the pre-merger marginal cost of 4).

Market Structure of Rivals. The size of CI-CMCR efficiencies does not depend on the slope of the equilibrium reaction function of the non-merging rivals. Nocke and Whinston (2022) use this fact, together with the property that the magnitude of the CI-CMCR is a function of the merging firms' market shares for demand systems such as the multinomial logit and CES, to suggest that a market structure screen for whether a transaction is anticompetitive should focus primarily on the market shares of the merging parties. This contrasts with standard practice which looks just as closely at the post-merger level of the HHI index (based on pre-merger market shares), which depends on the market shares of all of the firms in the market.[10]

However, the slope of the equilibrium best response function of the non-merging firms is key to the existence of a pooling equilibrium. To see this, suppose that the reaction function of firm 3 in Figure 1 rotates around point A to become the magenta dotted line, with the isoprofit curves, CI-CMCRs and pre-merger equilibrium best response function and post-merger best response function of M held fixed. If $\mathbf{c}_M = \mathbf{c}_M^{\text{CMCR}}$, M's profits at C are lower than those at A, so that pooling on $\mathbf{c}_M = \mathbf{c}_M^{\text{PRE}}$ prices could not be sustained for any $\delta_M < 1$. Pooling, with an outcome at D, could be sustained for a δ_M close enough to 1, if the possible efficiencies were the CI-CMCRs and a small marginal cost reduction of 0.375.

A screen for the possible existence of a pooling equilibrium would, therefore, want to include factors that would make rivals' prices more sensitive to those of the merging firms. The prices of competitive rivals with constant marginal costs would not be sensitive, so, intuitively, one would expect rivals' prices to be more sensitive when, all else equal, the rivals have more market power (e.g., larger shares and larger margins). This intuition would be consistent with M's best response function becoming steeper than the pre-merger equilibrium best response function of firms 1 and 2. It is also reflected in how we actually constructed the dotted magenta line in Figure 1: it is rivals' equilibrium best response function when $F = N = 5$, so that there are three single-product rivals to M, each with $a_{-M} = 6$ but with marginal costs increased so that the pre-merger equilibrium remains at A.

[10] Nocke and Whinston (2022) has been influential, and may be one reason why the December 2023 Merger Guildelines (`https://www.justice.gov/d9/2023-12/2023%20Merger%20Guidelines.pdf`, downloaded January 2, 2024) include a criterion where a merger is presumed harmful if one party has a share of more than 30% and the change in HHI, which equals twice the product of the parties' market shares, is more than 100. The market-level HHI may, of course, be correlated with the ability of firms to sustain collusion (Loertscher and Marx (2021)).

3.2 Generalizing the Example via Simulations.

As pointed out by a referee, our intuition for what determines the slope of rivals' equilibrium best response functions, and how this affects pass-through incentives, is not discussed in the existing literature. We now use simulations to understand how far our intuition generalizes beyond the demand system and parameters assumed in our example. This will complement our analysis of hypothetical beer mergers in Section 4, which will use a random coefficients logit demand system.

We will discuss simulations that illustrate the role of rival market structure (measured by product ownership or HHI) and the role of the relative size of the merging firms. The text provides enough information on the simulations to interpret the results. Full details are provided in Appendix A. The simulation and calibration procedure is similar to the one used by Miller et al. (2016) to study pass-through and merger price effects in a CISSNE framework.

Design. All of our simulations assume that there are six products in the market. To investigate the role of rival market structure, we will consider a merger between products 1 and 2, which are symmetric and owned by single-product firms before the merger. We draw 120 random sets of shares for these products and the outside good, and the pre-merger margin of the merging firms. For each of these sets, we draw 20 alternative sets of shares for products 3 to 6, and consider all 15 possible ownership structures of these rival products. This gives us 36,000 alternative market structures in total, with the merging firms larger than rivals in some simulations but not in others.

For each market structure that we consider, we calibrate marginal costs and the parameters of three alternative demand systems, multinomial logit, linear and the Almost Ideal Demand System (AIDS) of Deaton and Muellbauer (1980), so that, with CISSNE pricing, the pre-merger outcome would produce the specified pre-merger shares and margins when all prices are normalized to one.[11] Given the calibrated parameters, we calculate

- the slope, at the pre-merger equilibrium, of an index that summarizes rivals' equilibrium best response function.[12]
- for a firm 1-firm 2 merger, the δ_M^* that will support a pooling equilibrium if $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$, using inequality (1). If $\pi_M(\mathbf{c}_M^{\text{CMCR}}, \mathbf{p}_M^\dagger(\mathbf{c}_M^{\text{PRE}}), \mathbf{p}_{-M}^\dagger(\mathbf{c}_M^{\text{PRE}})) < \pi_M(\mathbf{c}_M^{\text{CMCR}}, \mathbf{p}_M^\dagger(\mathbf{c}_M^{\text{CMCR}}), \mathbf{p}_{-M}^\dagger(\mathbf{c}_M^{\text{CMCR}}))$, the calculated δ_M^* will be greater than 1. While such a value indicates that pooling on no efficiency prices cannot be supported, we will still use the calculated δ_M^* in our comparisons below.
- for a firm 1-firm 2 merger and assuming $\delta_M = 0.95$, the value of κ^*, identified using a grid of κ values in steps of 0.01 from 0.01 to 1.

Results: Rival Market Structure. We examine the role of rival market structure (i.e., their ownership structure and shares) holding fixed the pre-merger market shares and margin of the merging firms.

Our first set of tests hold fixed the market shares of all products and the margin of the merging parties, so that variation comes from the ownership of rival products. Specifically, for each of the 2,400 combinations of product market shares and the margin of the merging

[11] If one of the calibrated marginal costs is negative, we take another set of draws.

[12] See Appendix A for details. In the case of the mulitnomial logit demand system, the index is the negative of the inclusive value of the rival firms and the outside good when they choose best response prices. For the other demand system, the index is a weighted average of rival prices, where the weights reflect parameters in the demand system.

firm, we compare the slopes of the rival best response function, the δ_M^*s and the κ^*s for 45 pairwise comparisons where the second market structure in the pair involves the divestment of one or more products from the first market structure (and no acquisitions).[13]

For all 108,000 (120×20×45) pairwise comparisons for all three demand systems, we find that, for the more concentrated ownership structure, (i) the rivals' equilibrium best response prices are more sensitive to the merged firm's prices, (ii) the δ_M^* required to support pooling on no synergy prices is lower, and (iii) the value of κ^* is higher. These findings are consistent with our interpretation of how rival market structure would affect the slope of rivals' equilibrium best response functions and pooling incentives in our example.

Our second set of tests hold fixed the outside good market share and merging firm market shares and margins, but allow both rival product market shares and rival product ownership to vary. We examine how our three outcomes vary with an HHI measure of rivals' market structure, calculated as $\sum_{j=3,..,6}\left(\frac{s_j}{\sum_{k=3,..,6}s_k}\right)^2$. As we are holding the merged firms' shares fixed, this HHI measure will vary monotonically with the market-level HHI used by the agencies. While there is no theoretical reason to think that this HHI measure is the ideal statistic in general, we use it because HHI measures are common in merger analysis and it depends only on information on non-merging products (shares and ownership) that are likely to be available to agencies during the initial phase of merger investigation.

The relationships between the outcomes and the rival HHI measure are not perfectly monotonic, but the relationships are very strong and consistent with less concentrated rivals having less sensitive prices and the conditions to support pooling being tighter. As an

[13] For example, if the first ownership structure of products 3 to 6 is {A,A,B,B}, where the letter denotes the identity of the owner, then {A,C,B,B},{A,A,B,C}, {A,C,B,D} would be the alternatives for comparison, but {A,C,C,B} would not as this would involve two separately owned products becoming commonly owned. If the first ownership structure is {A,A,A,A}, every other ownership structure is an alternative for comparison.

Figure 2: Critical Discount Factors and HHIs of Non-Merging Firms for 300 Alternative Market Structures for Three Alternative Calibrated Demand Systems. In all cases the pre-merger market share of the outside good is 17.7%, and the merging products each have 18.2% shares of the units sold and their margins are 0.51.

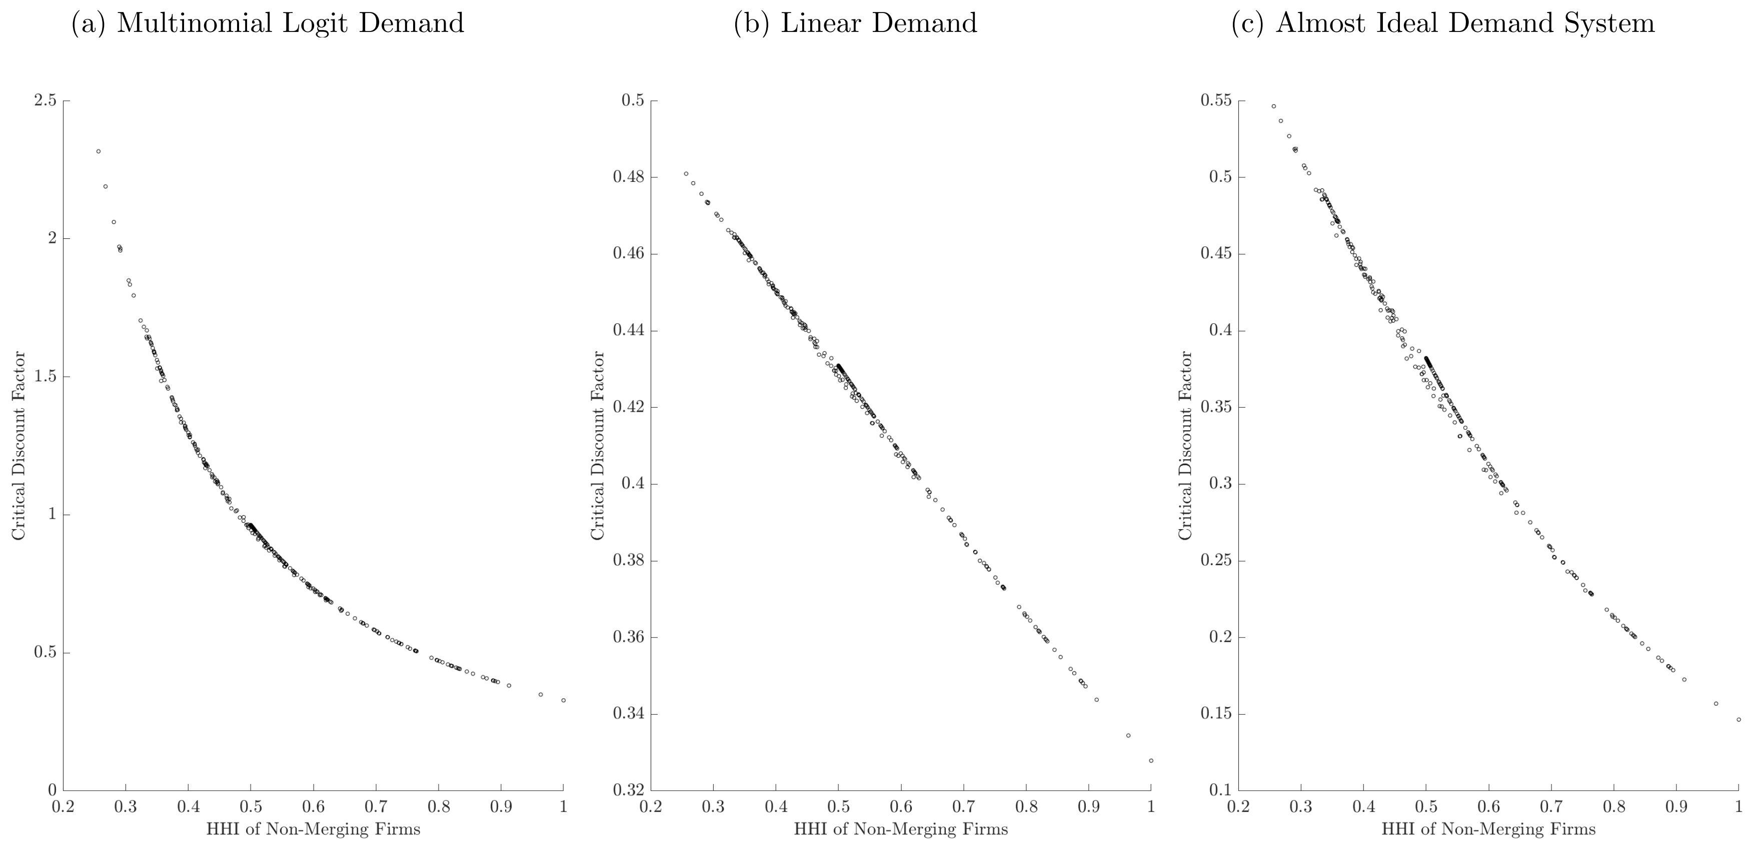

example, Figure 2(a)-(c) show scatter plots of the rival HHI and the δ_M^* to support pooling when $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$ are the possible post-merger costs, the outside good has share 17.7%, products 1 and 2 each sell 18.2% of the total units sold and have margins of 0.51. The CI-CMCR requires marginal costs to fall from 0.49 to 0.38. As the rival HHI increases the critical discount factor falls.[14] Appendix A shows the relationships for the slopes of the best response functions, and reports that the negative correlations are just as strong for different merging firm market shares and margins.

Results: Size of the Merging Firms. When we analyze hypothetical beer mergers in Section 4 we calculate larger δ_M^*s and smaller κ^*s for mergers involving the largest firm in the market. This is consistent with our rival market structure findings as the market structure of rivals will tend to be more concentrated when the largest firm in the market is one of those rivals. However, we perform additional simulations to examine more directly how δ_M^* and κ^* vary with the relative size of the firms that are merging.

In our first set of simulations to examine this relationship, we continue to assume that merging firms are symmetric. Specifically we assume that all products have independent owners before the merger, and that there are three pairs of symmetric products. We take 1,000 draws of the share of the outside good and the market shares for each pair, and the margin for the products with the largest shares. We then calibrate demand parameters and marginal costs as before, and calculate the δ_M^*s and κ^*s for mergers involving each pair of symmetric firms.

[14] One other feature of these figures is also worth noting. When a single rival owns all four non-merging products, the critical discount factors for the three demand systems have a similar order of magnitude but, when ownership is less concentrated, the δ_M^*s for the multinomial logit calibrations are larger.

Comparing across the mergers, the δ_M^*s and κ^*s show a clear pattern that is consistent with what we find in our beer examples. For the linear and AIDS demand models, for all 1,000 draws, the δ_M^* for the merger involving the largest firms is bigger than the δ_M^* for the "middle-sized" firm pair merger, and this δ_M^* is bigger than the δ_M^* for the smallest pair merger. The order for the κ^*s is reversed for every pair. In this sense, pooling equilibria are easier to sustain for mergers involving smaller firms. For the multinomial logit demand system, the δ_M^* always have their highest values, and the κ_M^*s their smallest values, for the largest pair merger, but there are 2% (3%) of simulations where the order of the δ_M^*s (κ^*s) for the middle-sized pair and smallest pair mergers are reversed. However, the general pattern remains quite clear.

We use an additional simulation analysis to consider whether this pattern holds when we consider mergers between firms with different market shares. Some practitioners (e.g., Rose and Shapiro (2022)) have argued that mergers involving the largest firm in a market may be anticompetitive even when the change in the HHI is small.[15] For this analysis, we take 1,000 sets of draws for the shares of three firms and the share of each of three symmetric smaller firms.

Having calibrated the demand and marginal cost parameters, we calculate the δ_M^*s and κ^*s for mergers involving each of the largest firms and one of the three small symmetric firms. For all three demand models, and for all 1,000 draws, the δ_M^* is largest, and the κ^* smallest, for the merger involving the largest pre-merger firm, and these measures have their

[15] Rose and Shapiro (2022) tentatively suggest a presumption against a firm with a 50% market share buying any firm with a market share larger than 1%, while noting that the 1982 Merger Guidelines contained a presumption that could have caught leading firms with market shares of more than 35%. The 2023 Merger Guidelines have a presumption against mergers where the firm has a post-merger share of more than 30% if the increase in HHI is more than 100.

second largest/smallest values for the merger involving the second largest firm. Therefore, once again, the simulation points towards the existence of pooling equilibria being more of a potential concern for mergers that involve smaller firms.

4 Application: Horizontal Mergers in the Beer Industry

This section investigates the existence and effects of pooling MPBEs for simulated mergers using the random coefficient nested logit demand and marginal cost estimates from Miller and Weinberg (2017). Caradonna, Miller and Sheu (2021) and Nocke and Whinston (2022) use these estimates to calculate the efficiencies needed to prevent post-merger price increases, assuming that firms use CISSNE strategies.

The Miller and Weinberg (2017) sample, taken from the IRI Academic Dataset (Bronnenberg, Kruger and Mela (2008)), covers 5 beer manufacturers (brewers) with 13 brands (39 brand-size combinations) in 39 local markets. We will assume, as Miller and Weinberg (2017) do, that retailers passively pass through costs, so that retail prices can modeled as being chosen by brewers. Before the 2008 Miller-Coors joint venture, the firms are Anheuser-Busch (AB), Miller and Molson-Coors, together with two importers, Grupo Modelo and Heineken.

Our simulations, with step-by-step details in Appendix B, use observed prices and market shares from Q3 2007, immediately before the joint venture was announced, and the quarterly random coefficients nested logit demand (RCNL-2) estimates from Miller and Weinberg (2017) Table IV, column (iii). The estimates imply some income heterogeneity in preferences over prices, calories and the included products. Imported beers sell at higher price points than the domestic brands so that the heterogeneity in price tastes will also imply that imported brands are systematically more attractive to some consumers. Consistent with our previous assumptions, we use standard CISSNE first-order conditions to infer pre-merger marginal costs.

Mergers in a Representative Market. The first part of our analysis considers the local ("representative") market where firm market shares are closest to their national averages. The pre-merger HHI (out of 10,000), based on pre-merger volume market shares for Miller and Weinberg (2017) sample products, is 2,850 (i.e., we are looking at mergers in an already fairly concentrated market). For the ten possible firm-pair mergers, the first columns of Table 1 report the change in HHI for the merger (using the same market shares), the pre-merger HHI of rivals and the share-weighted average CI-CMCRs in dollars and percentages. As there are only three rivals for each merger, the lowest possible value of the HHI of rivals would be 3,333.

The next column reports the δ^*_M to support a pooling equilibrium when $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$. The following columns are based on our κ^* calculations using a grid of 0.01 values for κ and assuming $\delta_M = 0.9^{\frac{1}{4}}$.[16] The columns report the value of κ^* and show how much pooling would raise per-period prices and the profits of rival firms, and how much it would lower consumer surplus compared to a CISSNE equilibrium when $\kappa = \kappa^*$ and the realized efficiency is large. The definition of κ^* implies that pooling would only increase the merged firm's per-period profits by a small amount.

Consider the Coors/Miller merger. It would raise the market HHI by 711 and, under a

[16] Lower δ_Ms imply smaller κ^*s but the changes are not dramatic. For example, for the Coors/Miller merger in the representative market, the κ^*s are 0.88 and 0.83 for $\delta_M = 0.95$ and $\delta_M = 0.9$, compared to 0.9 in the baseline. This calculation also implies that pooling could have significant effects on prices even if the asymmetry of information is only likely to last a few years. Merger analysis would usually be concerned about such medium-run price effects.

Table 1: Analysis of Pooling Equilibria in the Representative Market.

Merger	ΔHHI	HHI Rivals	$\overline{\text{CI-CMCR}}$ (\$)	$\overline{\text{CI-CMCR}}$ (%)	Critical Discount Factor (δ_M^*)	κ^*	$\Delta\overline{p_M}(\kappa^*)$ (%)	$\Delta\overline{p_{-M}}(\kappa^*)$ (%)	ΔCS(κ^*) (\$000)	$\Delta\pi_{-M}(\kappa^*)$ (\$000)
AB/Coors	1,273	4,264	1.09	16.5%	2.02	0.19	1.91%	0.16%	-17.3	7.8
AB/GM	925	4,034	0.70	9.3%	1.49	0.28	1.91%	0.24%	-16.8	8.5
AB/Heineken	613	3,729	0.45	6.3%	1.39	0.30	1.34%	0.15%	-11.4	6.3
AB/Miller	2,103	3,609	1.97	31.1%	4.56	0.08	1.58%	0.06%	-12.4	4.2
Coors/Miller	711	5,482	0.58	8.8%	0.56	0.90	9.48%	1.54%	-47.6	33.2
GM/Coors	312	4,520	0.27	2.9%	0.45	1.00	3.99%	0.65%	-19.2	16.7
GM/Heineken	150	3,967	0.29	2.4%	0.75	0.68	3.08%	0.40%	-9.9	9.1
GM/Miller	516	5,067	0.39	4.9%	0.53	0.98	6.00%	1.06%	-32.5	24.4
Heineken/Coors	207	4,218	0.19	2.2%	0.44	1.00	3.05%	0.39%	-12.0	11.0
Heineken/Miller	342	4,686	0.26	3.5%	0.52	0.97	4.47%	0.64%	-20.2	16.1

Notes: the representative market is the regional market in the Miller and Weinberg (2017) sample with market shares closest to the national averages in Q3 2007. The pre-merger HHI is 2,850. HHIs (out of 10,000) based on market shares in Q3 2007. The HHI Rivals is calculated using the market shares of the three sample firms not involved in the merger. The reported $\overline{\text{CI-CMCR}}$s are share-weighted averages across the merging firms' products. The % $\overline{\text{CI-CMCR}}$s are relative to pre-merger costs. The critical discount factor is the lowest discount factor that would support pooling on prices that would be the small efficiency CISSNE if $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$ and $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$. The remaining columns consider outcomes in a pooling equilibrium where a merging firm with marginal costs equal to $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} - \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$ sets prices that would be a CISSNE equilibrium if its costs were $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} + \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$. Prices, profits and surplus are compared with the CISSNE outcome where M's marginal costs are $\underline{\mathbf{c}}$. AB=Anheuser-Busch and GM=Grupo Modelo.

CISSNE analysis, would require an average of 9% marginal cost efficiencies to offset anti-competitive effects. If the merged firm benefits from CI-CMCR efficiencies, an equilibrium where it pools on no efficiency prices could be sustained if $\delta_M > 0.56$. The κ^* implies that, if the two possible efficiencies were marginal cost reductions of 0.9% and 16.7%, pooling could lead to M increasing prices by 4.6% and 9.5%, relative to pre-merger and CISSNE prices respectively, even if realized synergies are 90% *larger* than the CI-CMCRs. Compared to CISSNE pricing with such a 16.7% efficiency, pooling would lower quarterly consumer surplus by $47,600 and raises quarterly rivals' profits by $33,200. These dollar effects sound small, but they are non-trivial given that these 13 brands had only $1.3 million of retail sales at IRI sample stores in this market-quarter.

Considering the rest of the table, hypothetical mergers involving the largest pre-merger firm, Anheuser-Busch, have the largest δ^*_Ms and small κ^*s. This is especially true for mergers with Miller or Coors, the second and third largest firms. This is consistent with the pattern in our simulations with simpler demand systems.

One might think that this pattern simply reflects how mergers involving the largest firms have the largest CI-CMCRs (recall that κ multiplies the CI-CMCR efficiencies). However, the table shows that the κ^*s are so much smaller for the Anheuser-Busch/Coors and Anheuser-Busch/Miller mergers that, if $\kappa = \kappa^*$, the consumer surplus reductions due to pooling after these mergers ($17,300 or $12,400) would be smaller than for (say) a Heineken/Miller merger ($20,200).

The smallest firms in this analysis are the importers Heineken and Grupo Modelo. Based on our simulation results, one might have therefore expected that this merger would have had the lowest δ^*_M and the largest κ^*. This is not the case, and the reason appears to lie in the heterogeneity in tastes for imported brands implied by the random coefficients in the demand system. This taste heterogeneity implies higher demand diversion between Heineken and Grupo Modelo products, which will make the CI-CMCRs larger than for a merger between a domestic firm and an importer with similar market shares. However, pooling incentives can exist even when demand diversion between the merging products is small, so that pooling by a CI-CMCR domestic/importer merged firm on no efficiency prices may be easy to sustain.[17]

Mergers Across All Markets. We repeat the analysis across all 390 possible firm-pair-geographical market mergers, treating each market-merger as independent.[18] The box-and-whiskers plot in Figure 3 shows the ranges of calculated κ^*s for each possible merger.

For a Coors/Miller merger, the median κ^* is around 0.9, but there is one market where $\kappa^* = 0.03$. Coors and Miller's combined pre-merger market share is over 75% in this market, so the small κ^* is consistent with our earlier results.

Adding consumer surplus effects across the 39 markets, pooling after a Coors/Miller merger, with market-specific $\kappa = \kappa^*$s, would lower quarterly consumer surplus by $1.2 million, compared with total quarterly sample sales of $58.3 million. Given national off-premise beer sales in 2007 of $51.1 billion (i.e., over 850 times sales in one quarter of the Miller and Weinberg (2017) sample), scaling our results suggest that post-merger pooling would

[17]To understand why pooling incentives can exist even when there is limited demand diversion, consider a Salop circle model example where pre-merger all firms are independent and two non-neighbors propose to merge. There will be no demand diversion between the merging products and CI-CMCRs would be zero. However, the merged firm's profits would still rise if all firms price as if the merging firms have slightly higher marginal costs than they actually do.

[18]We have also calculated the κ^* that could support pooling when the range of synergies (in percentage terms relative to the market-specific CI-CMCRs) is assumed to be the same in every market, and deviation in any market would lead to CISSNE pricing in all markets. The common κ^* is typically somewhat lower than the median market-specific $\kappa^* s$ shown in Figure 3.

Figure 3: Distribution (Across Markets) of the Largest κs that Support Pooling MPBEs for Each Merger.

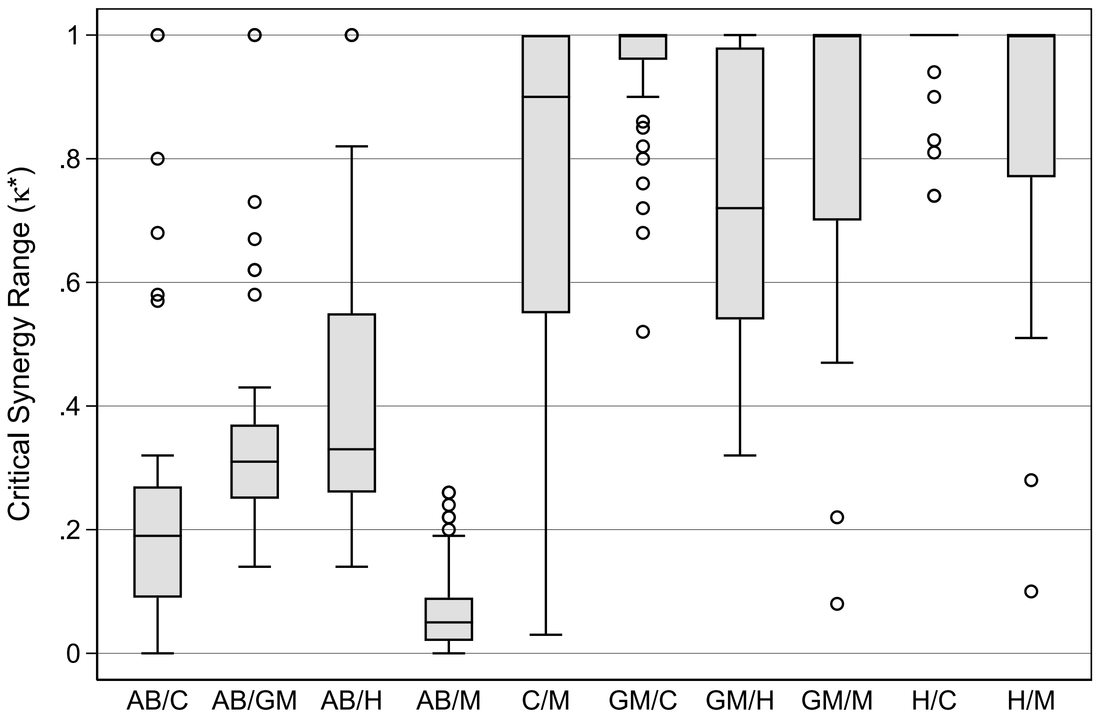

Notes: mergers identified on the x-axis using the following abbreviations: AB=Anheuser-Busch, C=Coors, GM=Grupo Modelo, M=Miller, H=Heineken. For each merger, the black center line indicates the median value of κ^*, the limits of the grey box indicate the 25th to 75th percentiles, and the black whiskers indicate the adjacent values. The circles indicate values lying outside the range of the adjacent values. The value of κ^* indicates that pooling on low synergy prices can be supported when the possible synergies are $(1 + \kappa^*)\times$ and $(1 - \kappa^*)\times$ the CI-CMCRs.

have the potential to lower annual consumer surplus by hundreds of millions of dollars after several mergers that might plausibly not be challenged, including the Coors/Miller merger which was allowed to proceed.[19] The pattern that $\kappa^* s$ are higher for mergers involving Coors or Miller and an importer than they are for mergers between the two importers also holds on average across markets in the sample.

We perform some additional analysis to test how our finding that δ^*_Ms are smaller and κ^*s are larger when the rest of the market is more concentrated generalizes to this richer demand system. Specifically, for every market-merger, we re-calculate δ^*_M and κ^* holding fixed pre-merger product market shares and prices, but changing the assumed ownership structure of the non-merging products.

As an extreme example, suppose that we assume that all non-merging products are sold by independent firms. This assumption implies smaller pre-merger markups for the non-merging products but does not change the CI-CMCRs for any merger.[20] For all but eight (out of 390 market-mergers), the κ^* is less than 0.001, implying that pooling could only be sustained for a tiny range of marginal costs. Suppose that we, alternatively, assume that all non-merging products are sold by a single firm. Once again, the CI-CMCRs are unchanged, but now $\kappa^* = 1$ (the largest value considered) for 363 out of 390 market-mergers, including 17 Anheuser-Busch/Miller market-mergers, for which the median κ^* given the actual ownership of other products was less than 0.05. Appendix B supplements these results by presenting the relationship between the critical discount factor and the rival HHI, based on varying the

[19] Value of total sales taken from Beverage Information Group (2021), p. 34 (as reported by `https://www.statista.com/statistics/483537/us-off-premise-beer-retail-dollar-sales/`), deflated to January 2010 dollars using the CPI-U to be consistent with the monetary values used by Miller and Weinberg (2017).

[20] For example, for the representative market, average non-merging product pre-merger margins are \$0.37 and \$0.90 when we assume they are owned by independent firms or a single firm respectively.

ownership of rival products, for a Coors/Miller merger in the representative market.

Comparison to Collusion (Coordinated Effects). In our model, pooling raises the prices of all firms relative to a CISSNE equilibrium. Of course, tacit collusion in a complete information model, supported, for example, by Nash reversion trigger strategies, would also raise prices, with the difference that, with collusion, at least some of the non-merging rivals would also charge prices that are higher than best responses. This behavior can potentially lead to prices that are much higher than could be supported in any pooling equilibrium. For example, suppose that the merging firms realize CI-CMCR efficiencies. In the representative market, tacit collusion between all five firms on joint-profit maximizing prices, which would increase prices by as much as 131% from pre-merger prices, could be sustained after nine out of the ten hypothetical mergers with a discount factor of $0.9^{\frac{1}{4}}$.[21] In contrast, even if we allow the merged firm to act as a Stackelberg leader, choosing the outcome that is most profitable for it given that other firms will best respond, a possibility that we will discuss in Section 5, the merged firms prices would rise by no more than 4.4%.[22]

Observed Price Changes after the MillerCoors Joint Venture. A definitive test of whether firms play the type of equilibrium that we have described requires data on the size of realized post-merger marginal cost efficiencies as well as post-merger pricing. Constructing these types of datasets is beyond this current project, but Appendix C presents evidence that, after its creation in 2008, MillerCoors, a joint venture that was an effective merger

[21] The exception would be the Anheuser-Busch/Miller merger, where the joint-profit maximizing prices with CI-CMCR synergies would allocate so much profit to the merged firm that Coors would deviate. However, collusive prices equal to 0.88 × the joint profit-maximizing prices + 0.12 × CISSNE prices could be sustained given a $0.9^{\frac{1}{4}}$ discount factor.

[22] This would be the increase for the merged firm after a Coors/Miller merger. The price increases of other firms would be smaller.

of Miller's and Coors' brewing and marketing assets, changed its pricing in a way that is consistent with either a significant decrease in the efficiency of its transportation network, an outcome that might be viewed as possible but not likely in merger analysis, or the type of pooling behavior that we have described.[23]

5 Application and Extensions

In this section, we discuss some extensions or variants of our model, including an application to vertical mergers, and we discuss how one might apply our model in the context of a horizontal merger investigation.

Finite Horizon Game. One might wonder whether the assumption that pooling would cause the asymmetry of information about the merged firm's type to potentially last forever is important to our results, especially given that, for comparison purposes, we assume that pre-merger prices are those from a complete information game. In fact, one can generate significant effects even in games where players know that M's type will be revealed after a fixed and small number of periods, although the exact form of the equilibrium will change.

To illustrate, consider our Section 3 multinomial logit $N = 3$ example, and suppose that whether M's marginal costs are $\underline{\mathbf{c}} = 2.27$ (i.e., CI-CMCR efficiency) or $\overline{\mathbf{c}} = 4$ will be revealed after two periods. If $\delta_M = 0.95$ and the probability that $\mathbf{c}_M = \underline{\mathbf{c}}$ is less than 0.8, then a

[23]The analysis in Appendix C shows that after the JV, MillerCoors priced as if its per unit per shipping mile cost of shipping beer had increased. The interpretation that the efficiency of its distribution network actually decreased is at odds with public claims that the JV exceeded its overall target for cost savings (e.g., Molson-Coors annual report, `https://ir.molsoncoors.com/news/news-details/2009/SABMiller-and-Molson-Coors-Report-MillerCoors-First-Quarter-Earnings/default.aspx`, accessed February 12, 2024) and more specific evidence that innovations reduced the volume of lost products in its distribution network (e.g., `https://www.logisticsmgmt.com/article/millercoors_taps_dunnage_innovation`, accessed February 12, 2024). However, these discussions are not precise about how synergies affected fixed costs and how they affected marginal or variable costs.

pooling MPBE where firms set $\mathbf{p}_M^{\dagger}(\overline{\mathbf{c}})$ prices in the first period, and static Bayesian Nash (BNE) prices in the second period, will exist if firm 3 interprets *any* first period deviation by M from its no efficiency CISSNE price as reflecting a low marginal cost.[24]

Product Quality Synergies. While the academic literature has focused on marginal cost reductions, parties often argue that their merger will generate other types of synergies.[25] We can adjust our model to think about a product-improving synergy.

Consider the Section 3 multinomial logit $N = 3$ example. Assume the marginal cost of each product is 4, before and after the merger. Recall that consumer c's indirect utility for good i, with price p_i, is $u_{ic} = a_i - 0.25p_i + \varepsilon_{ic} = \delta_i + \varepsilon_{ic}$ with $a_1 = a_2 = 4$ and $a_3 = 6$. As before, firms 1 and 2 merge, but now assume that, with probability $0 < q < 1$, M will have the ability to raise the indirect utility intercepts of both of its products from 4 to 5 by paying a per-period fixed cost of 1.4. With probability $1 - q$ the fixed cost associated with raising quality will be prohibitively high. In each post-merger period, M chooses whether to increase quality and its prices simultaneously with the other firm choosing its prices. After the period the rival will know whether M increased its quality because its demand will be lower.

Given this specification, we can construct a pooling equilibrium where M does not invest if it has a low investment cost, because if it did the rival would infer that M's products

[24] As the probability that the efficiency is large increases, second period BNE prices after first period pooling will approach large efficiency CISSNE prices, reducing the incentive of a large efficiency M to pool in the first period. We assume that **any** deviation in the first period is interpreted as reflecting a low cost (large efficiency) so that a high cost M does not have an incentive to set a price above the pooling price to signal its type.

[25] For example, when the MillerCoors joint venture was announced the parties stated that "[g]reater scale and resources will allow additional investment in brands, product innovation and sales execution", so that consumers and retailers would "benefit from greater choice and access to brands", as well as the reduction in operating costs (`https://www.sec.gov/Archives/edgar/data/24545/000110465907073971/a07-26272_1ex99d2.htm`, accessed January 23, 2024).

will have higher quality, which would result in the rival setting lower prices, in every future period.[26] On the other hand, if rivals prices were less sensitive to M's quality, as is the case if $N = 5$, implementing the improvement would be profitable.

Choice of Whether to Implement an Efficiency. We have assumed that the realization of the post-merger marginal cost is exogenous. In practice, realizing efficiencies may require investments by M. The model can be adapted to allow for this possibility. The simplest way, which leaves the analysis essentially unchanged, assumes that low marginal costs will be realized if an investment is made, but to assume that rivals cannot observe whether a sunk investment is made and that the investment cost is either low (for example, zero) or prohibitively high. A firm's prices are observed. In this case, one could construct an equilibrium where a firm with a low investment cost will invest but price as if its marginal costs are high, pooling with the type that has the high investment cost.

One might argue that such a model is unattractive because rivals might well find out if the investment has been made (for example, they might receive job applications from laid-off workers). If this is the case, M's investment decision involves a choice between a complete information game where it has a high marginal cost, and a complete information game with low marginal cost. This is different to the game considered so far, but the incentive to invest to lower costs will still be stronger when rivals' prices are less responsive to those set by the merging firm.

[26] If the rival correctly expects M's products to have higher quality, it will set a price of 12.21. M will make a per-period profit (net of the fixed cost) of 2.01, and consumer surplus will be higher than before the merger (14.46 vs. 13.33). On the other hand, in the pooling equilibrium where M does not implement the improvement, even when it is able to do so, the rival's price will be 13.69, M will make a per-period profit of 2.33, and consumer surplus will be lower than before the merger (12.44 vs. 13.33). The pooling equilibrium can be sustained if $\delta_M \geq 0.53$.

To illustrate, consider the Section 3 multinomial logit $N = 3$ example and suppose that achieving the $\mathbf{c}_M = \mathbf{c}_M^{\text{CMCR}}$ requires a per-period fixed cost investment of 0.5. Otherwise M's $\mathbf{c}_M = \mathbf{c}_M^{\text{PRE}}$. If $\mathbf{c}_M = \mathbf{c}_M^{\text{CMCR}}$ and the equilibrium is at A, M's variable profits are 2.76, and its total profit 2.26. If, on the other hand, $\mathbf{c} = \mathbf{c}^{\text{PRE}}$ and the equilibrium is at B, M's profits will be 2.33 (note that the isoprofit curves in the figure no longer indicate differences in profits for different outcomes because M's marginal costs are now higher). Therefore, in a subgame perfect equilibrium, M would not choose to invest. On the other hand, if $N = 5$ and the rival equilibrium best response function is given by the magenta dotted line, M is choosing between the pricing outcome at A and the outcome at C (variable profits 2.14), so that the cost-reducing investment is profitable.

Stackelberg Price Leadership. A related alternative model would involve M pricing as a Stackelberg leader every period. For example, this might happen if the proposed pooling model was preceeded by a stage where M could announce the value of $\overline{\mathbf{c}}$ or if, as in Fudenberg and Levine (1989), rivals attach some probability to M having a behavioral type where it would always play the Stackelberg strategy.[27] Stackelberg-like solutions also arise in recent work, such as Brown and MacKay (2023), where a firm with a pricing algorithm that updates less frequently will end up pricing like a Stackelberg leader to raise the price that its rival sets.[28]

[27] Fudenberg and Levine (1989) show that a patient rational firm, with a known cost-type, should be able to achieve a payoff at least as large as a Stackelberg leader in equilibrium. Roddie (2012*b*) develops an alternative approach, without behavioral types but where the possiblity that a firm's cost evolves over time leads to perpetual signaling, and outcomes that, in some circumstances and some types, are similar to those of Stackelberg game.

[28] If a merger gives the parties access to an algorithm that updates more quickly than other firms, the logic of Brown and MacKay (2023) would be that it may be rivals that respond by setting prices that are above best response levels, whereas in our model it is only the merged firm that engages in non-best response pricing.

Consider the Section 3 multinomial logit example. A $\mathbf{c}_M^{\mathrm{CMCR}}$-M's Stackelberg leader would choose a price where its rivals' equilibrium best response function is tangent to its isoprofit curves. With one rival firm and the solid magenta rival best response function, the outcome selected would be off the top right of the figure with a merged firm price of 12.95 (a post-merger increase of 43%) and a rival price of 15.06 (increase of 16%). These prices are the same as in a CISSNE outcome where M has marginal costs of 7.3 (even though its actual costs are 2.3).[29] With three rivals and the dotted best response function, the merged firm's price would increase to 9.45 (a smaller 4.7% increase).

Stackelberg price leadership would also lead to significant price increases in our beer examples. For example, suppose that a merged Coors and Miller merger in the representative market benefits from the 8.8% CI-CMCR cost efficiency. Acting as a Stackelberg price leader it would choose to raise its prices by an average of 4.4%, which would be the CISSNE outcome if it benefited from only a small 0.7% efficiency. Further calculations show that if a merged representative market Coors/Miller would always set Stackelberg leader prices, its marginal costs would have to fall by at least 24.5%, a huge amount, for consumer surplus not to decrease after the merger. Once again rival market structure matters for this prediction: if we calibrate the model assuming all non-Coors/Miller products are owned by independent firms, the required cost reduction is much smaller (9.1%).[30]

[29]If the merged firm believed that setting a lower price would be followed by reversion to CI-CMCR CISSNE prices, pooling on these much Stackelberg prices could be sustained if $\delta_M \geq 0.55$.

[30]These calculations assume that actual cost efficiencies must be proportional to the CI-CMCRs. Stackelberg prices are found by assuming that M prices as if it has a cost type which equals $\mathbf{c}^{\mathrm{PRE}}\pm$ efficiencies proportional to the CI-CMCRs. We use 0.001 grids to find both proportionality factors. If all rival products were owned by a single firm, the required cost reduction to prevent prices from rising would be 33.4%.

Application to Horizontal Merger Review. We believe that it would be relatively straightforward to incorporate key components of our model in horizontal merger investigations.

An investigation usually proceeds in two stages. At the first screening stage, there is a preliminary attempt at market definition, and market structure, together with the concerns of large customers, is then used to assess whether it is plausible that the merger could have substantial anticompetitive effects. As already noted, our results suggest that the possibility of pooling provides a justification for why agencies should continue to think that the potential for anticompetitive effects will be positively correlated with higher concentration in the market as a whole as well as larger market shares of the merging firms, at least in markets where firms set prices (see below for comments about quantity-setting).

Second stage investigations involve much more extensive document review, economic modeling, often including calibrated merger simulations, by staff economists and the assessment of likely and cognizable efficiencies by agency financial analysts. The economic and efficiency analyses currently tend to proceed independently of each other, whereas our model suggests there should be more interaction.

A simple example illustrates this logic. Suppose that agency economists calculate, using pre-merger market share and margin data, that marginal cost reductions of 8% (CI-CMCRs) would keep CISSNE prices from rising. Agency financial analysts estimate a 5-9% range for "likely" efficiencies.[31] Our experience is that an agency would be unlikely to challenge this merger, reasoning that efficiencies in the middle or upper half of the 5-9% range would

[31] Section 4D of Sweeting et al. (2020) discusses a merger where a range of likely efficiencies was produced, and, in our experience, efficiency estimates usually have a range of several percentage points.

predict price decreases or price increases too small to meet the standard of a substantial loss of competition.

Our model would suggest that strategic concerns could limit the pass-through of larger efficiencies. To analyze the merger through the lens of our model, the agency would, as a first step, try to establish whether their own 5% estimate is a reasonable estimate of the lower bound efficiencies that rivals expect (recognizing that rivals will not have access to the types of internal data that agency analysts are able to use) and which would likely be achieved without a significant expenditure of resources by the merged firm. Assuming that the updated lower bound estimate is still below the CI-CMCRs, agency economists would then use rivals' best response functions, which are already calculated to perform calibrated merger simulations, to compute whether passing through realized efficiencies close to or above the CI-CMCRs would be more profitable, accounting for the reaction of rivals, than only passing through the lower bound. An agency could account for fixed costs of implementing efficiencies in its profitability analysis.

If the pass-through of CI-CMCR efficiencies would be profitable then it would appear reasonable not to challenge the merger (absent evidence that collusion is possible). However, if this is not the case, parties could be asked to present evidence that pass-through would actually be profitable or should be assumed likely given the history of the industry. This would amount to introducing a limited version of the controversial "passing-on requirement" (Yde and Vita (1995)), but it seems plausible to us that many companies would be able to demonstrate historical pass-through rates using their own cost and pricing data.[32]

[32] Froeb, Tschantz and Werden (2005) criticize the FTC's use of empirical evidence suggesting low pass-through of cost savings (Ashenfelter et al. (1998)) in the original *Staples* merger litigation. This is the type of evidence that our model would suggest that it is appropriate to consider. If the firms have been involved in earlier mergers, it may be possible for parties to demonstrate more specifically that merger efficiencies

Asymmetric Information in Oligopoly and Quantity Competition. Our model assumes that, after a merger, only M's marginal cost is private information, but if some elements of marginal costs really are opaque to rivals it is likely that all firms may have some private information both before and after mergers.

Mailath (1989) considers a two period duopoly game with linear demand where each firm has private information about its marginal cost. He shows that signaling incentives will support a separating price equilibrium with higher first period prices than a complete information subgame equilibrium. Sweeting, Tao and Yao (2023) (STY) find price increases in asymmetric information games with more firms, more periods, nested logit demand and marginal costs that are positively, but not perfectly, serially correlated, and that this framework predicts larger price increases after horizontal mergers unless they are mergers to monopoly.[33]

We see the results in this article as more specifically relevant to merger review than those in STY for two reasons. First, merger review often identifies quite wide ranges for possible cost efficiencies, but the equilibria in STY can only be characterized when the supports for marginal costs are narrow. For example, STY compute examples where an equilibrium exists when marginal costs lie have a support of \$8-\$8.05, and CISSNE prices are on the order of \$24, but not when the support is \$8-\$8.15. Second, solving the STY model involves considerable computation and could not easily be used by agency economists within the short periods allowed for merger review, whereas the results in this model come from a framework that is much simpler and more transparent.

have also been passed through.

[33]STY's assumption that the costs of all firms can change from period-to-period is important because it implies that firms always have new private information to signal. In a game where cost types are fixed, separating play with prices that are perfectly observed would reveal cost types after a single period.

If firms set quantities, rather than prices, incentives will clearly be different.[34] In particular, a merged firm with private information would want to reduce its rivals' output by overproducing, whereas in our model the merged firm has the incentive to raise its price and produce less. On-going work is examining how much these incentives could lead to higher outputs in equilibrium, and how these effects would vary with the number of firms.

Vertical Mergers and the Elimination of Double Marginalization. So far we have only considered horizontal mergers. The reaction of rivals could also make it unprofitable for the parties to a vertical merger to pass through efficiencies. A model with a privately known marginal cost efficiency could follow the model in Section 2 quite closely.

However, the profitability of realizing pricing efficiencies such as the "elimination of double marginalization" (EDM) (Spengler (1950)), where an integrated firm internalizes how a lower downstream price increases upstream sales, can also depend on how much rivals' prices would respond. Appendix D provides a stylized example. The example is motivated by how agencies such as the FTC analyze transactions between upstream healthcare providers and downstream insurance companies, such as UnitedHealth Group's (UHG)'s 2019 purchase of DaVita's Medical Group in several states, including Colorado.[35] The example allows the merger to have a "raising rivals' costs" effect (Salop and Scheffman (1983)), where the

[34]As discussed in Sweeting, Tao and Yao (2023), even though reaction functions are upward sloping, prices are not necessarily "strategic complements" with logit demand as $\frac{\partial^2 \pi_i}{\partial p_i \partial p_j}$ can be negative when considering prices that are significantly above best response levels.

[35]UnitedHealth Group (UHG) has an insurance business (UnitedHealth) and a provider business through its Optum division. The FTC chose not to challenge the acquisition in the state of Colorado, where Kaiser Permanente is a significant competitor, because the Republican majority believed that the benefits of EDM would outweigh anticompetitive effects, or that the balance would be so close that the agency would not prevail in Court. `https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/joint-statement-commissioners-phillips-wilson-concerning-unitedhealth-group-davita` (accessed May 20, 2023).

integrated firm will raise the provider prices that it charges to rival insurance companies, and it also includes competition from a fully integrated insurance/provider that does not contract with rivals (e.g., Kaiser Permanente). The Appendix example shows that the way in which the integrated rival's prices would change in response to EDM means that the merged firm's profits can be higher when EDM is not realized.[36]

The example does not develop a mechanism by which the merged firm could avoid realizing EDM in equilibrium. However, a few possible mechanisms suggest themselves. An integrated rival like Kaiser, which does not deal with other providers or insurers, may in practice have limited information about the size of pre-margins and therefore the scope of EDM. This might allow an asymmetric information mechanism to be constructed. Alternatively, the merged firm may make organizational choices that have the effect of limiting internalization, or which create opaqueness about how far or when incentives have been aligned. For example, UHG executives have emphasized that Optum is operated at "arm's length, ... separately reported and separately managed" from its insurance business.[37] Kwoka and Slade (2019) discusses possible organizational or principal-agent barriers to EDM being realized, and our results suggest that merging firms may have strategic reasons to keep these barriers in place. Empirical evidence also suggests that, in practice, EDM incentives are only partially realized, and, in our example, it is partial realization which is most profitable. For

[36]The example is interesting partly because Kaiser's presence is often a feature of the market that merging parties claim will limit their market power, and which may require firms to consolidate in order to compete (Ross et al. (2019)).

[37]`https://thehill.com/policy/healthcare/133507-conflict-of-interest-concerns-raised-as-obama-races-to-implement-health-reform/` (accessed January 30, 2024), and it is noticeable that UHG and Optum have not attempted the type of complete integration that Kaiser Permanente has adopted even though competition with Kaiser has been used as a rationalization for vertical integration. Transparent organizational separation may be required to convince third party insurers and providers that their data will not be misused for competitive purposes.

example, Crawford et al. (2018) estimate that, on average, business units within vertically integrated cable TV providers account only for approximately 79 cents of every dollar of the effect that their decisions have on the profits of other units in the same firm.

6 Conclusion

Textbook presentations of strategic incentives in price-setting games (e.g., chapters 8 or 9 of Tirole (1988)) suggest that asymmetric information, or some ability by firms to make commitments, may support equilibria with higher prices than models where firms simultaneously set prices under complete information. Agency merger review and the academic modeling of mergers have ignored these effects, even though mergers are the type of event that could create significant uncertainties about merged firm costs or incentives from the perspective of rivals. We believe that practitioners likely presume that the effects of these strategic incentives will be smaller than the types of competitive effects that merger analysis usually focuses on. Our examples suggest that this presumption will often be incorrect, and that strategic effects could provide an explanation for why non-trivial price increases have been observed after many mergers that the agencies expected, after thorough reviews, to create significant efficiencies. Our model is simple enough that an agency could use it to identify when merging parties should be asked to provide additional evidence of their incentives to pass-through efficiencies.

Section 4 and Appendix C describe price changes around the MillerCoors joint venture that are consistent with the mechanism in our model or a reduction in the efficiency of the merged firm's distribution network. However, we believe that detailed empirical work, combining cost and price data, is required to provide a better of understanding of cost pass-through both in this example and in general, building on the work of Ashenfelter et al. (1998) and Muehlegger and Sweeney (2019) who estimate low pass-through rates in retail and industrial settings respectively, and, more specifically, the pass-through, as well as the realization, of merger efficiencies.

References

Amir, Rabah, Effrosyni Diamantoudi, and Licun Xue. 2009. "Merger Performance under Uncertain Efficiency Gains." *International Journal of Industrial Organization*, 27(2): 264–273.

Ashenfelter, Orley C, Daniel S Hosken, and Matthew C 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, Orley C, David Ashmore, Jonathan B. Baker, and Signe-Mary McKernan. 1998. "Identifying the Firm-Specific Cost Pass-Through Rate." Bureau of Economics, Federal Trade Commission Working Paper No. 217.

Asker, John, and Volker Nocke. 2021. "Collusion, Mergers, and Related Antitrust Issues." In *Handbook of Industrial Organization.* Vol. 5, 177–279. Elsevier.

Bernile, Gennaro, and Scott W Bauguess. 2011. "Do Merger-Related Operating Synergies Exist?" University of Miami.

Beverage Information Group. 2021. *2021 Liquor Handbook.* Norwalk, CT:Beverage Information Group.

Bitzan, John D., and Wesley W. Wilson. 2007. "Industry Costs and Consolidation: Efficiency Gains and Mergers in the U.S. Railroad Industry." *Review of Industrial Organization*, 30(2): 81–105.

Blonigen, Bruce A, and Justin R Pierce. 2016. "Evidence for the Effects of Mergers on Market Power and Efficiency." National Bureau of Economic Research WP No. 22750.

Bonatti, Alessandro, Gonzalo Cisternas, and Juuso Toikka. 2017. "Dynamic Oligopoly with Incomplete Information." *Review of Economic Studies*, 84(2): 503–546.

Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson. 2015. "Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry." *American Economic Review*, 105(7): 2086–2119.

Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela. 2008. "The IRI Marketing Dataset." *Marketing Science*, 27(4): 745–748.

Brown, Zach Y, and Alexander MacKay. 2023. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics*, 15(2): 109–156.

Caradonna, Peter, Nathan Miller, and Gloria Sheu. 2021. "Mergers, Entry, and Consumer Welfare." Georgetown University.

Crawford, Gregory S, Robin S Lee, Michael D Whinston, and Ali Yurukoglu. 2018. "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica*, 86(3): 891–954.

Deaton, Angus, and John Muellbauer. 1980. "An Almost Ideal Demand System." *American Economic Review*, 70(3): 312–326.

Farrell, Joseph, and Jonathan B Baker. 2021. "Natural Oligopoly Responses, Repeated Games, and Coordinated Effects in Merger Analysis: A Perspective and Research Agenda." *Review of Industrial Organization*, 58(2): 1–39.

Froeb, Luke, Steven Tschantz, and Gregory J Werden. 2005. "Pass-Through Rates and the Price Effects of Mergers." *International Journal of Industrial Organization*, 23(9-10): 703–715.

Fudenberg, Drew, and David K. Levine. 1989. "Reputation and Equilibrium Selection in Games with a Patient Player." *Econometrica*, 57(4): 759–778.

Grieco, Paul, Joris Pinkse, and Margaret Slade. 2018. "Brewed in North America: Mergers, Marginal Costs, and Efficiency." *International Journal of Industrial Organization*, 59: 24–65.

Groff, James E, Donald Lien, and Jiwei Su. 2007. "Measuring Efficiency Gains from Hospital Mergers." *Research in Healthcare Financial Management*, 11(1): 77–90.

Harrington, Joseph E. 2021. "There May Be No Pass Through of a Merger-Related Cost Efficiency." *Economics Letters*, 208: 110050.

Haynes, Michelle, and Steve Thompson. 1999. "The Productivity Effects of Bank Mergers: Evidence from the UK Building Societies." *Journal of Banking & Finance*, 23(5): 825–846.

Kulick, Robert B. 2017. "Ready-to-Mix: Horizontal Mergers, Prices, and Productivity." US Census Bureau Center for Economic Studies Paper No. CES-WP-17-38.

Kwoka, John. 2014. *Mergers, Merger Control, and Remedies: A Retrospective Analysis of US Policy.* MIT Press.

Kwoka, John, and Margaret Slade. 2019. "Second Thoughts on Double Marginalization." *Antitrust*, 34: 51.

Loertscher, Simon, and Leslie M Marx. 2021. "Coordinated Effects in Merger Review." *Journal of Law and Economics*, 64: 705–744.

MacKay, Alexander, and Marc Remer. 2022. "Consumer Inertia and Market Power." Harvard Business School.

Mailath, George J. 1989. "Simultaneous Signaling in an Oligopoly Model." *Quarterly Journal of Economics*, 104(2): 417–427.

Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, 56(3): 571–599.

Miller, Nathan H, and Matthew C Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica*, 85(6): 1763–1791.

Miller, Nathan H, Marc Remer, Conor Ryan, and Gloria Sheu. 2016. "Pass-Through and the Prediction of Merger Price Effects." *Journal of Industrial Economics*, 64(4): 683–709.

Muehlegger, Erich, and Richard L Sweeney. 2019. "Pass-Through of Own and Rival Cost Shocks: Evidence from the US Fracking Boom." *Review of Economics and Statistics*, 1–32.

Nocke, Volker, and Michael D. Whinston. 2022. "Concentration Thresholds for Horizontal Mergers." *American Economic Review*, 112(6): 1915–48.

Nocke, Volker, and Nicolas Schutz. 2018. "Multiproduct-Firm Oligopoly: An Aggregative Games Approach." *Econometrica*, 86(2): 523–557.

Roddie, C. 2012*a*. "Signaling and Reputation in Repeated Games, I: Finite Games." University of Cambridge.

Roddie, C. 2012*b*. "Signaling and Reputation in Repeated Games, II: Stackelberg Limit Properties." University of Cambridge.

Rose, Nancy L, and Carl Shapiro. 2022. "What Next for the Horizontal Merger Guidelines?" *Antitrust Magazine*, 36(2): 4–13.

Rose, Nancy L, and Jonathan Sallet. 2019. "The Dichotomous Treatment of Efficiencies in Horizontal Mergers: Too Much? Too Little? Getting it Right." *University of Pennsylvania Law Review*, 168: 1941–1984.

Ross, Douglas, James Harlan Corning, David Maas, and Douglas Litvack. 2019. "When Providers Merge, Is Kaiser a Competitor?" *Competition Policy International*.

Salop, Steven C, and David T Scheffman. 1983. "Raising Rivals' Costs." *American Economic Review*, 73(2): 267–271.

Shapiro, Carl. 1986. "Exchange of Cost Information in Oligopoly." *Review of Economic Studies*, 53(3): 433–446.

Spengler, Joseph J. 1950. "Vertical Integration and Antitrust Policy." *Journal of Political Economy*, 58(4): 347–352.

Sweeting, Andrew, David J Balan, Nicholas Kreisle, Matthew T Panhans, and Devesh Raval. 2020. "Economics at the FTC: Fertilizer, Consumer Complaints, and Private Label Cereal." *Review of Industrial Organization*, 57(4): 751–781.

Sweeting, Andrew, Mario Leccese, and Xuezhen Tao. 2022. "Should We Expect Merger Synergies to Be Passed Through to Consumers?" Center for Economic Policy Research Discussion Paper No. 17059-2.

Sweeting, Andrew, Xuezhen Tao, and Xinlu Yao. 2023. “Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers.” University of Maryland.

Tirole, Jean. 1988. *The Theory of Industrial Organization.* MIT Press.

Toxvaerd, Flavio. 2008. “Strategic Merger Waves: A Theory of Musical Chairs.” *Journal of Economic Theory*, 140(1): 1–26.

Vives, Xavier. 2011. “Strategic Supply Function Competition with Private Information.” *Econometrica*, 79(6): 1919–1966.

Walia, Bhavneet, and Christopher J Boudreaux. 2019. “Hospital Mergers, Acquisitions and Regulatory Policy Implications: Price, Cost, Access and Market Power Effects.” *Managerial Finance*, 45(10/11): 1354–1362.

Werden, Gregory J. 1996. “A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products.” *Journal of Industrial Economics*, 409–413.

Weyl, E Glen, and Michael Fabinger. 2013. “Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition.” *Journal of Political Economy*, 121(3): 528–583.

Yan, Jia, Xiaowen Fu, Tae Hoon Oum, and Kun Wang. 2019. “Airline Horizontal Mergers and Productivity: Empirical Evidence from a Quasi-Natural Experiment in China.” *International Journal of Industrial Organization*, 62: 358–376.

Yde, Paul L, and Michael G Vita. 1995. “Merger Efficiencies: Reconsidering the Passing-on Requirement.” *Antitrust Law Journal*, 64: 735–747.

Online Appendices for "Should We Expect Merger Synergies To Be Passed Through to Consumers?"

Leccese, Sweeting and Tao

March 2024

A Section 3 Simulations

Setion 3 presents a multinomial logit demand example of pooling incentives after one particular merger. We use simulations to understand how the results generalize, focusing on (i) how pooling incentives vary with the slope of rivals' equilibrium best response and how this slope varies with market structure, and (ii) how pooling incentives vary with the relative size of the merging firms. This Appendix details these simulations and presents some additional results. Our simulation design is similar to those used by Miller et al. (2016), Miller et al. (2017) and Panhans and Taragin (2023), who also use calibrated demand systems to analyze the performance of measures used to predict the effects of horizontal mergers.

A.1 Simulations to Test How Rival Market Structure Affects The Slope of Rivals' Equilibrium Best Response Function and Incentives to Pool.

We will detail our simulations that examine the effects of rival market structure. The other simulations will make similar, but slightly different, assumptions.

A.1.1 Market Structures.

We consider industries with six products before and after a merger. The merging parties are the single-product owners of products 1 and 2, which will be symmetric.

Step 1. We take 120 alternative sets of random draws for the share of the outside good, the share of the merging products and the pre-merger margin of the merging products. These draws are taken as follows:

- draw the pre-merger market share of the outside good, $s_0 \sim U[0.05, 0.50]$.
- draw the pre-merger market share of each of the merging products as $s_{1,2} \sim U[0.05, 0.45]$ multiplied by $(1 - s_0)$.
- draw the pre-merger margin of each of the merging products as $m_{1,2} \sim U[0.1, 0.7]$.

Step 2. For each of these sets of draws, we take 20 alternative sets of draws of the market shares of the non-merging products (labeled 3 to 6). Specifically, we take four draws from a Dirichlet distribution, Dir(2.5), which determines how the remaining share is split. The market shares are then the Dirichlet draws multiplied by $(1 - s_0) \times (1 - s_1 - s_2)$.

Step 3. We calibrate costs and the multinomial logit demand model assuming all pre-merger prices equal 1 and that products 3 to 6 are owned a single firm, as described below, to verify that all pre-merger marginal costs and the post-merger marginal costs of M when it benefits from CI-CMCR efficiencies, are greater than zero. If this condition fails, we take another set of Dirichlet draws.[1]

In our simulations, we consider all 15 possible ownership structures of the rival products. We therefore have 36,000 $(120 \times 20 \times 15)$ alternative sets of pre-merger shares and margins.

[1] Our code would take another set of draws for the shares and margin of the products 1 and 2 if we rejected 81 out of the first 100 sets of Dirichlet draws, but this never happens.

A.1.2 Calibration.

We consider three demand systems. For any given alternative market structure (shares and ownership), the implied margins of each product will be the same across the three demand systems.

Multinomial Logit Demand Calibration. In a multinomial logit model, the mean utility of product i is $\theta_i - \alpha p_i$. We find the θs, α and marginal costs for each product, including $\mathbf{c}_M^{\mathrm{PRE}}$, that match the drawn market shares and product 1 margin as a CISSNE with the prices of all products equal to one, given the assumed pre-merger ownership structure.

Linear Demand Calibration. In a linear demand model, the demand for product i is $q_i = a_i - \sum_{j=1}^{N} b_{ij} p_j$. Assuming $q_i = s_i$ (i.e., that market size equals 1), we set b_{ij} to equal the $\frac{\partial q_i}{\partial p_j}$s implied by the logit demand model (at prices of 1). The a_is are then calculated from $a_i = q_i + \sum_{j=1}^{N} b_{ij} p_j$. Prices of 1 will then be a CISSNE using linear demand when we use marginal costs from the logit model.

Almost Ideal Demand System Calibration. We follow Miller et al. (2016) in how we calibrate the AIDS model (Deaton and Muellbauer (1980)). The expenditure share of product i is

$$w_i = \alpha_i + \sum_{j=0}^{N} \gamma_{ij} \log p_j + \beta_i \log \frac{x}{P} \tag{2}$$

where x is total expenditure, including on the outside good, and P is a price index. The price of the outside good is normalized to 1, so that γ_{i0} is not identified, and we set $\beta_i = 0$.

$$\log(x) = \text{constant} + \sum_{k=1}^{N} \alpha_k \log(p_k) + \frac{1}{2} \sum_{k=1}^{N} \sum_{j=1}^{N} \gamma_{kj} \log(p_k) \log(p_j) \tag{3}$$

and

$$\begin{aligned} \frac{\partial q_i}{\partial p_j} &= \frac{x}{p_i^2}(\gamma_{ii} - w_i + w_i^2) \text{ for } i = j \\ &= \frac{x}{p_i p_j}(\gamma_{ii} - w_i w_j) \text{ for } i \neq j \end{aligned} \tag{4}$$

We calculate the implied expenditure shares (w) and x from the drawn market shares and prices of 1. The γs are set to match the $\frac{\partial q_i}{\partial p_j}$ from the logit calibration at these prices and shares. The outside good γs are set so that $\sum_{i=0}^{N} \gamma_{ij} = 0$ for all j. The αs are then calculated from (2), and the constant from (3).

A.1.3 Calculation of Pooling Incentives and the Slope of Rivals' Equilibrium Best Response Function.

For each market structure and calibrated demand system we do the following:

- calculate the HHI and ΔHHI measures (based on quantities) and the concentration of the non-merging firms (i.e., the HHI for only these firms). Consistent with standard practice, the outside good is not included in these calculations (unless we specifically indicate that it is included).
- calculate the CI-CMCR efficiency, and $\mathbf{c}_M^{\text{CMCR}}$, for a merger of firms 1 and 2.
- calculate CISSNE prices with a CI-CMCR efficiency and no efficiency, and use inequal-

ity (1) to calculate the critical discount factor that would support a pooling equilibrium if $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}}$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{PRE}}$. We allow for critical discount factors greater than 1.

- calculate the slopes of the equilibrium best response function prices of products 1 and 2 before the merger, and the best response of M after the merger, and calculate the slope of the equilibrium best response function of the non-merging firms. This evaluation happens at pre-merger prices. More details below.

- calculate κ^* using a grid of values of κ from 0.01 to 1 in 0.01 steps. We verify that if pooling can be supported for κ' it can also be supported for any smaller κ (this is always true). κ^* is the highest κ for which pooling can be supported when $\underline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} - \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$ and $\overline{\mathbf{c}} = \mathbf{c}_M^{\text{CMCR}} + \kappa(\mathbf{c}_M^{\text{PRE}} - \mathbf{c}_M^{\text{CMCR}})$.

The calculation of the slopes has two steps. The first step is the calculation of a single index from the best response prices of the non-merging products to a given price of the symmetric merging products. For multinomial logit demand, we follow our example by calculating the index as $-\log(1 + \sum_{j=3,..,6} \exp^{\delta_j(p_j^*(p_M))})$, so that the best response function is upwards sloping. For linear demand, we use $\frac{\sum_{j=3,..,6} b_{1j} p_j^*(p_M)}{\sum_{j=3,..,6} b_{1j}}$ where b_{1j} is the calibrated parameter that measures how the demand for product 1 changes with the price of product j. For AIDS demand, we similarly use $\frac{\sum_{j=3,..,6} \gamma_{1j} p_j^*(p_M)}{\sum_{j=3,..,6} \gamma_{1j}}$ where γ_{1j} is the coefficient on the price of $\log(p_j)$ in the calibrated equation for the expenditure share of good 1. We can then calculate the slope of the index given a small change in the prices of the merging products.

Given the vectors of best response prices of the non-merging products, we can also calculate the best response prices of the merging products before and after the merger, and calculate the slope of these best response functions with respect to a small change in the value of the index that summarizes the rivals' prices. For all mergers and all demand systems we observe that the merger increases this slope.

A.1.4 Comparisons

We investigate the role of rival market structure using several comparisons. In this subsection we detail these comparisons as well as providing additional results.

Pairwise Comparisons. Our first set of tests holds fixed the market shares of the non-merging products, and looks at variation in their ownership structure. Specifically, for each of the 2,400 combinations of product market shares and the margin of the merging firm, we compare the slope of the rival best response function, the critical discount factor and κ^* for 45 pairwise comparisons where the second market structure in the pair involves one or more divestments from the first market structure. For example, if the first ownership structure of the non-merging products 3 to 6 is {A,A,B,B}, where the letter denotes the identity of the owner, then {A,C,B,B}, {A,A,B,C}, {A,C,B,D} would be the alternatives for comparison. If the first ownership structure is {A,A,A,A}, every other ownership structure is an alternative for comparison.

Comparisons Allowing Ownership Structure and Shares to Vary. Our second set of tests hold fixed the merging firm market shares and margins, and examines how the three outcomes (critical discount factor, κ^* and slope) vary with an HHI of rivals measure calculated as $\sum_{j=3,..,6}\left(\frac{s_j}{\sum_{k=3,..,6} s_k}\right)^2$ across the 300 (200 $\times$ 15) simulations where ownership structure and the shares of the merging products may vary. We use the HHI as the type of measure that can be calculated with the types of data available in the early stages of a

Table A.1: Average Correlations Between Rival HHI Measure and Outcome Measure Across 120 Draws of Merging Firm Market Shares and Share of Outside Good . Eqm. BRF = equilibrium best response function. For κ^*, markets where the statistic is the same for all rival HHI measures are not included in the calculation.

Demand Model	Critical Discount	Rival Eqm. BRF Slope	κ^*
Multinomial Logit	-0.919	0.997	0.993
Linear	-0.999	0.998	0.886
AIDS	-0.997	0.995	0.958

merger investigation.

Because there is no theoretical reason why this HHI measure would exactly determine the values of the outcome variables, our analysis involves observing the strengths of the correlations. Text Figure 2 shows, for each demand system, the relationships between the HHI index and the critical discount factor for a given assumptions on the merging firms (pre-merger market share of the outside good is 17.7%, and the merging products each have 18.2% shares of the units sold and their margins are 0.51). Figure A.1 shows similar scatterplots for the slope of rivals' equilibrium best response functions.

For the linear and AIDS models, κ^* equals one (the highest value considered) for almost all of the draws of rival shares and ownership structures. Figure A.2 shows the scatterplot for logit demand, where more of the κ^*s are smaller. The relationship is again clear, with $\kappa^* = 1$ once the rival HHIs are sufficiently high.

Relationships are also strong for the remaining 119 draws of the merging firm shares and margins. For each draw, we calculate the correlation of the outcome and the rival HHI. Table A.1 presents the average correlations, which have signs consistent with our earlier results.[2]

[2] The correlation is smaller for the κ^*s for linear demand because more κ^* have the maximum value of 1.

Figure A.1: Slopes of Rival Equilibrium Best Response Function and the HHIs of Non-Merging Firms for 300 Alternative Market Structures for Three Alternative Calibrated Demand Systems. The pre-merger market share of the outside good is 17.7%, and the merging products each have 18.2% shares of the units sold and their margins are 0.51. Variation in HHIs comes from variation in the distribution of shares across products and variation in non-merging product ownership structure. See the text for a description of how the slopes are measured.

(a) Multinomial Logit Demand (b) Linear Demand (c) Almost Ideal Demand System

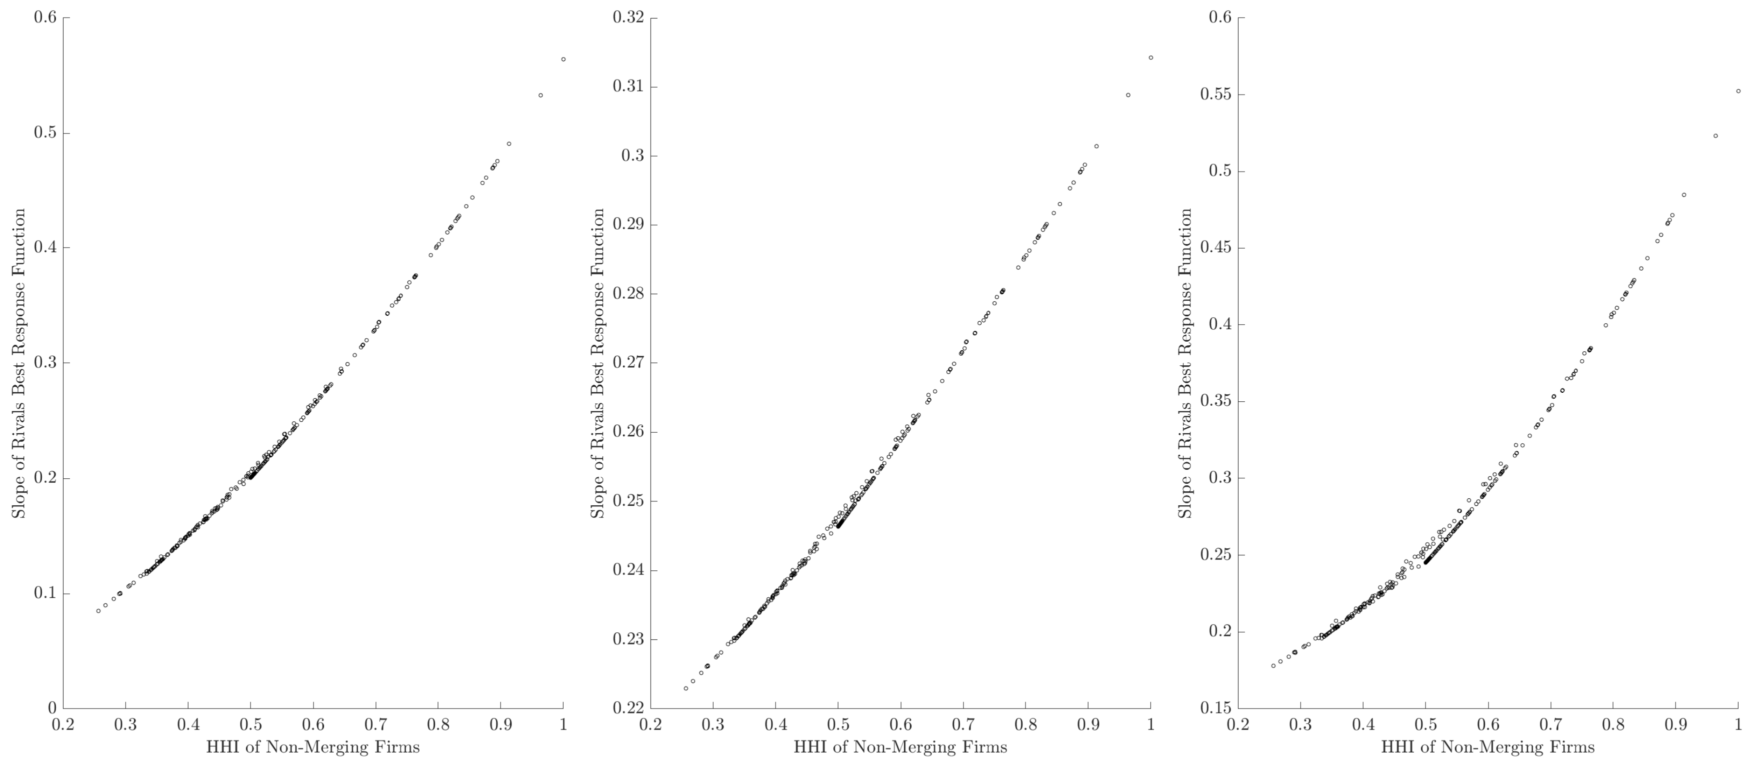

Figure A.2: κ^* and the HHIs of Non-Merging Firms for 300 Alternative Market Structures for the Multinomial Logit Demand Model. The pre-merger market share of the outside good is 17.7%, and the merging products each have 18.2% shares of the units sold and their margins are 0.51.

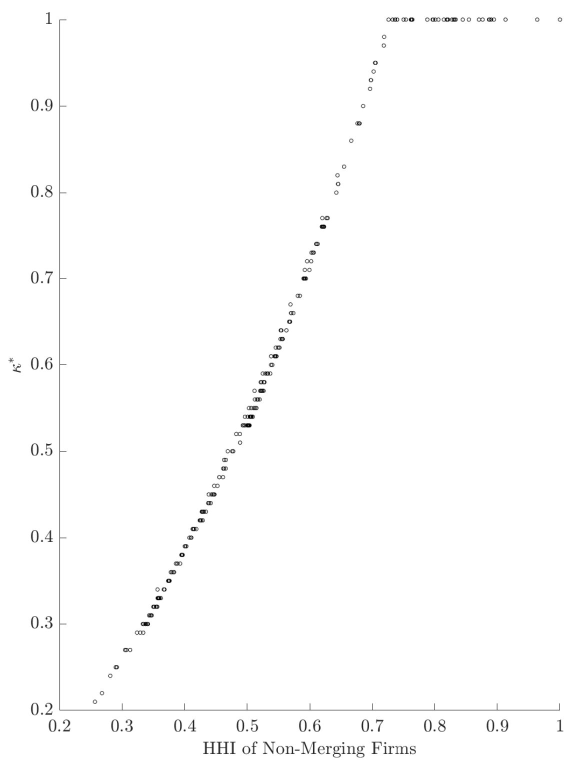

Allowing the Share of the Outside Good to Vary. We can also examine what the relationships look like when we keep the shares of the merging firms the same, but allow the share of the outside good, as well as the distribution of sales among the non-merging firms, to vary.

To do so, we assume that the merging firms have the same pre-merger shares (of the potential market) and the same margins as in the previous example (i.e., the margin is 0.51). We then take 100 uniform random draws of the outside good share allowing it to be between 0% and 95% of the remaining potential demand, and, for each of these draws, draw the shares of the non-merging products as before (i.e., using the Dirichlet distribution). We compute the statistics for each possible ownership structure of the non-merging products, so we have 1,500 alternatives in total.

Figure A.3 shows the scatterplot of the relationship between the critical discount factor and the rival HHI for each demand system. In these figures, the rival HHI is calculated using shares that include the outside good in the denominators, so that the HHI is low when the share of the outside good is high. The correlations are clearly negative, but the relationships, especially for linear and AIDS demand, are more heterogeneous than when the share of the outside good was held fixed, reflecting how the outside good, which has a fixed price, has a different price response to rival sellers. The heterogeneity in the multinomial logit case is obscured by the scale of the differences in the critical discount factor.

Figure A.3: Critical Discount Factors and HHIs of Non-Merging Firms for 1,500 Alternative Market Structures for Three Alternative Calibrated Demand Systems. In all cases, the pre-merger market shares of the merging products (with the outside good in the denominator) are 15.0% ($(1 - 0.177) \times 0.182$) and their margins are 0.51. In these examples the HHI of the non-merging firms is calculated as $\sum_{j=3,..,6} \left(\frac{s_j}{1-s_1-s_2}\right)^2$.

(a) Multinomial Logit Demand

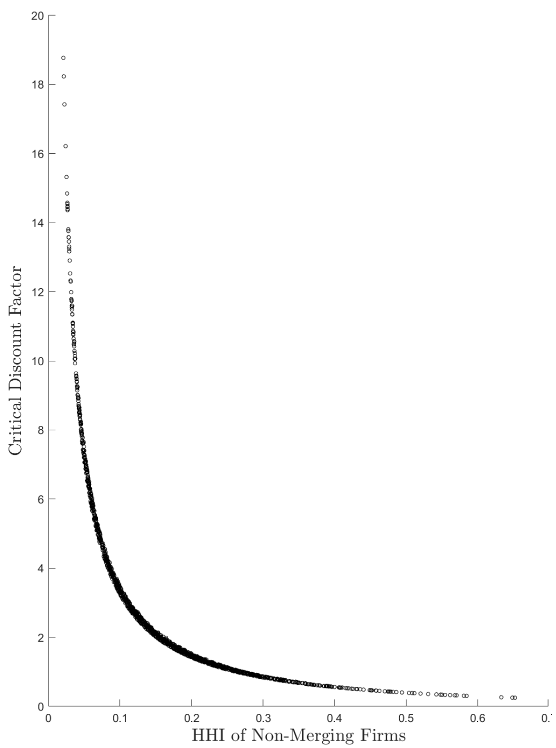

(b) Linear Demand

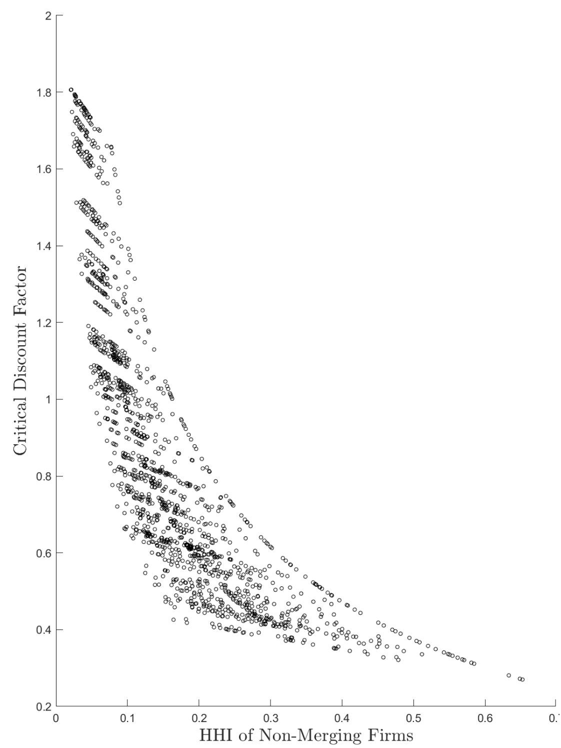

(c) Almost Ideal Demand System

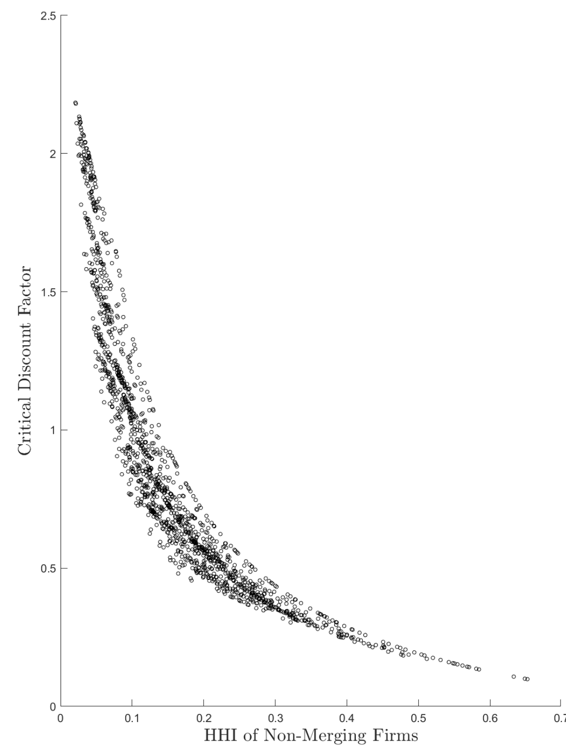

A.2 Simulations to Test How the Size of the Merging Firms Affects Pooling Incentives.

Our analysis of hypothetical beer mergers in text Section 4 finds higher critical discount factors and lower $\kappa^* s$ for mergers involving the largest firm in the market. This is consistent with the rival market structure results, in the sense that the market structure of rivals will be more concentrated when the largest firm is one of the rivals. We run two sets of additional simulations to look more directly at how the outcome measures vary with the size of the merging firms.

Symmetric Mergers. These simulations assume that pre-merger, all six products have independent owners but that there are three pairs of symmetric products. We take 1,000 sets of draws where

- the share of the outside good is $s_0 \sim U[0.05, 0.50]$.
- the shares of the remaining market for the three firm pairs are drawn from a Dirichlet distribution, Dir(2.5), so that the pair's combined share is $(1 - s_0) \times s_{pair}$, and we assume that each firm in the pair has an equal share.
- the pre-merger margin of each firm in the largest pair is $\sim U[0.1, 0.7]$.

We calibrate the demand and cost parameters for the three demand systems as before. We identify the largest pair, the smallest pair and the "middle" pair based on market shares. We then consider mergers between firms in each of these pairs and compute the critical discount factor and κ^* statistics. We then look at the pairwise comparisons between the statistics for the largest, smallest and middle pair mergers. The share of the outside good and the pre-merger large firm margin is therefore fixed across the comparisons.

For the linear and AIDS demand models, for every draw, the critical discount factors for the merger between the largest firms are bigger than the critical discount factors for the "middle-sized" pair merger, and these are bigger than the critical discount factors for the smallest pair merger. The same pattern holds in reverse for the κ^*s (although some values are equal as we identify κ^* using a discrete grid).

For the multinomial logit demand system, the critical discount factors are always bigger ($\kappa^* s$ smaller) for the largest pair merger than the other mergers, but there are 2-3% of simulations where the ordering is different for the medium and small pair merger. However, the general pattern remains clear, with stronger pooling incentives for mergers involving smaller firms.

Asymmetric Mergers. We use an adapted simulation to consider how this logic extends to mergers involving firms of different sizes and a smaller rival. We take 1,000 sets of four draws from the Dirichlet distribution, Dir(2.5), and order them. The three largest draws correspond to the shares of three asymmetric "larger" firms, while we replicate the fourth draw as the share of three symmetric "small" firms.[3] We then consider mergers between each of the three "larger firms" and one of the small firms, and calculate the critical discount factor and κ^*.

For each of the 1,000 simulations for each of the demand systems, the δ^*_M is bigger and the κ^* smaller for the merger involving the largest firm than for the merger involving the second largest firm, and similarly for the merger involving the second largest firm than the

[3] We rescale the shares so that, when combined with the share of the outside good, drawn, as before from a $U[0.05, 0.5]$ distribution, they add up to 1.

merger involving the third largest firm. Therefore, the pattern that there are stronger pooling incentives for mergers involving smaller firms also holds for asymmetric mergers.

B Numerical Simulations of Brewer Mergers

This Appendix details our numerical simulations of hypothetical horizontal mergers in the U.S. beer industry. They are based on the data, sample selection and the demand estimation of Miller and Weinberg (2017). We briefly describe this data, before detailing our analysis.

Data. The Miller and Weinberg (2017) data is taken from the IRI Academic Database (Bronnenberg, Kruger and Mela (2008)). This database contains revenues and unit sales at the UPC-week-store level for a sample of grocery stores between 2001 and 2011. We select and aggregate data in exactly the same way as Miller and Weinberg (2017) to get observations at the brand-size-region-quarter level for 39 brand-size combinations (referred to as "products", e.g., "Bud Light 12 pack"). Considered pack sizes are 6 packs, 12 packs and an aggregation of 24 and 30 packs. There are 13 included brands produced by the following five brewers: Anheuser-Busch, SABMiller, Coors, Grupo Modelo and Heineken. Miller and Weinberg (2017) use a sample of stores from 39 markets, excluding IRI-sample markets where grocery stores sell only limited quantities or ranges of beer. Prices, defined as product revenues divided by units sold, are deflated, using the CPI-U series, to be in January 2010 dollars. Miller and Weinberg (2017) define market size by inflating the highest volume sales that are observed in the geographic market, with purchases of non-sample brands included in the outside good. As noted in the text, the percentage of off-premise national sales included in the sample is small, partly because convenience stores and alcohol stores are not included.

Representative Market. For some of our analysis, we focus on a single "representative" geographic market. We identify this market as the market where, in Q3 2007, brewer (volume) market shares are closest to their national averages. For the 39 products in the data, we aggregate volume market shares to the firm-market level. We then calculate the difference between firm shares and their national averages using the sum of squared differences in shares or the sum of absolute differences in shares. For Q3 2007 we identify the same market as representative using both of these measures, whether or not we include the outside good. However, in some other quarters these measures would identify different markets.

Demand. Our analysis uses the same demand specifications as Miller and Weinberg (2017). Our hypothetical merger simulations uses the quarterly "RCNL-2" specification (column (iii) of Miller and Weinberg (2017) Table IV). This model allows for a nesting parameter over the included brands, and preferences over the included brands, price and calorie content that vary with household income, as well as brand-pack size and quarter fixed effects.

Hypothetical Merger Simulations. For each hypothetical market-merger, we proceed as follows:

1. Given the demand system estimates, and observed Q3 2007 prices and market shares, we calculate implied unobserved product qualities and the marginal costs of each product using the first-order conditions implied by CISSNE pre-merger pricing. This is done for each market separately. For example,

$$\mathbf{c}^{\hat{\mathrm{PRE}}} = \mathbf{p} + (\Omega(\mathbf{p}))^{-1}\,\mathbf{s}$$

   where $\mathbf{c}$, $\mathbf{p}$ and $\mathbf{s}$ are the marginal cost and observed price and market share vectors, and Ω is the dot product of an indicator ownership matrix (element $\{i, j\}$=1 if and only if products i and j have the same owner) and the matrix of demand derivatives (element $\{i, j\} = \frac{\partial s_i}{\partial p_j}$).

2. For a given merger, we calculate the CI-CMCRs for each product,

$$\textbf{CI-CMCR} = \left((\Omega(\mathbf{p}))^{-1} - (\Omega'(\mathbf{p}))^{-1} \right) \mathbf{s}.$$

These are the marginal cost reductions that would keep CISSNE prices at exactly their observed levels after a merger that transforms Ω to Ω', reflecting the change in the ownership matrix. The CI-CMCR elements for all of the non-merging products will be zero and the elements for the merging products will be positive. Denote the subvector of the pre-merger marginal cost vector for the merging firms' products as $\mathbf{c}_M^{\hat{\text{PRE}}}$ and the subvector of CI-CMCRs as $\mathbf{CMCR}_M$.

3. Assuming $\overline{\mathbf{c}} = \mathbf{c}_M^{\hat{\text{PRE}}}$ and $\underline{\mathbf{c}} = \mathbf{c}_M^{\hat{\text{PRE}}} - \mathbf{CMCR}_M$, we calculate the critical discount factor that would support a pooling equilibrium. This follows the steps laid out below for finding the critical discount for a given κ.

4. We find the largest scalar κ, which we will denote κ^*, on a 0.01 step grid between 0.01 and 1 that will support a pooling MPBE, when we define the merged firm's possible marginal costs as $\overline{\mathbf{c}} = \mathbf{c}_M^{\hat{\text{PRE}}} - (1-\kappa)\mathbf{CMCR}_M$ and $\underline{\mathbf{c}} = \mathbf{c}_M^{\hat{\text{PRE}}} - (1+\kappa)\mathbf{CMCR}_M$. If $\kappa = 1$ then marginal costs are either equal to pre-merger marginal costs, or they are equal to pre-merger marginal costs less *twice* the CI-CMCRs. The test for whether a pooling MPBE can be supported for given κ proceeds in the following four steps.

   (a) compute CISSNE prices (by solving the standard first-order conditions) when the merged firm's costs are either $\underline{\mathbf{c}}$ or $\overline{\mathbf{c}}$, and other firms' products have the marginal costs calculated in step 1. Denote these prices $\mathbf{p}^{\dagger}(\underline{\mathbf{c}})$ and $\mathbf{p}^{\dagger}(\overline{\mathbf{c}})$.

(b) calculate the profits of the merged firm with high synergy marginal costs ($\underline{\mathbf{c}}$) at both sets of CISSNE prices. Denote these profits, $\pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_M(\underline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}}))$ and $\pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_M(\overline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))$.

(c) compute the static profit-maximizing best response prices and profits of the merged firm with marginal costs $\underline{\mathbf{c}}$ when non-merging firms are setting prices $\mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})$. Denote these prices $\mathbf{p}^{\mathrm{BR}}_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))$ and the profits $\pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\mathrm{BR}}_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))$.

(d) calculate the critical discount factor ($\delta_M(\kappa)$) that would support a pooling MPBE given the calculated profits.

$$\delta_M(\kappa) = \frac{\pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\mathrm{BR}}_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})) - \pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_M(\overline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}}))}{\pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\mathrm{BR}}_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})), \mathbf{p}^{\dagger}_{-M}(\overline{\mathbf{c}})) - \pi_M(\underline{\mathbf{c}}, \mathbf{p}^{\dagger}_M(\underline{\mathbf{c}}), \mathbf{p}^{\dagger}_{-M}(\underline{\mathbf{c}}))}. \tag{5}$$

If $\delta_M(\kappa)$ is less than the assumed $\delta_M = 0.9^{\frac{1}{4}}$ then a pooling equilibrium can be supported.

Having identified the highest value of κ that supports pooling (κ^*) for a given market-merger, we calculate the profits of all firms at the low and high efficiency CISSNE prices, and we calculate the change in consumer surplus using the standard compensating variation formula for a random coefficients logit demand model.

Pooling Equilibrium When Synergies Can Lie on the Inside of the Interval. For the representative market, we also examine whether we could sustain a pooling equilibrium on the smallest possible synergy CISSNE prices when the merged firm's realized synergy can have any value on an interval. Specifically, we consider the case where marginal costs can take on values between $\mathbf{c}_M^{\hat{\text{PRE}}} - (1+\kappa^*)\mathbf{CMCR}_M$ and $\mathbf{c}_M^{\hat{\text{PRE}}} - (1-\kappa^*)\mathbf{CMCR}_M$, where κ^* is the value identified in the two-type exercise described above. By doing so, we are assuming that a change in the realized synergy moves the marginal costs of every product owned by the merging firms in the same way relative to the CI-CMCR marginal costs.

Taking 0.01 steps of κ, we find the merged firm's profits at $\mathbf{p}^\dagger(\mathbf{c}_M^{\hat{\text{PRE}}} - (1-\kappa^*)\mathbf{CMCR}_M)$, as well as the one-period best response price and profit of the merged firm if it was to deviate from the low synergy CISSNE (pooling) prices and when it best responds to non-merging firms setting $\mathbf{p}^\dagger_{-M}(\mathbf{c}_M^{\hat{\text{PRE}}} - (1+\kappa^*)\mathbf{CMCR}_M)$ prices. We use these profits to test whether the merged firm would want to deviate from the pooling equilibrium, under the assumption that non-merging firms believe that the merged firm has the lowest possible marginal costs after a deviation, even if, after the off-the-equilibrium path deviation, the merged firm sets prices that are inconsistent with this assumption. For all ten mergers and all of the gridpoints that we consider, we find that the critical discount factor for merged firms with marginal costs between $\mathbf{c}_M^{\hat{\text{PRE}}} - (1+\kappa^*)\mathbf{CMCR}_M$ and $\mathbf{c}_M^{\hat{\text{PRE}}} - (1-\kappa^*)\mathbf{CMCR}_M$ are less than our assumed $0.9^{\frac{1}{4}}$, implying that pooling can be sustained.

Varying Ownership Structures of Non-Merging Firms. As described in the text we can recompute pre-merger marginal costs, and the post-merger δ_M^*s and κ^*s assuming different ownership structures for the non-merging products. In the Section 3 simulations we observe a tight relationship between the critical discount factor and the HHI of non-merging firms when we hold the market shares of the merging firms and the share of the outside good fixed. Figure B.1 shows a scatterplot of the critical discount factors and rival HHI for a

Figure B.1: Critical Discount Factors and HHIs of Non-Merging Firms for a Hypothetical Coors/Miller Merger in the Representative Market. Variation in non-merging firm HHI is created by divesting products from existing brewers or merging them together. Points with critical discount factors above 1.6 are not shown.

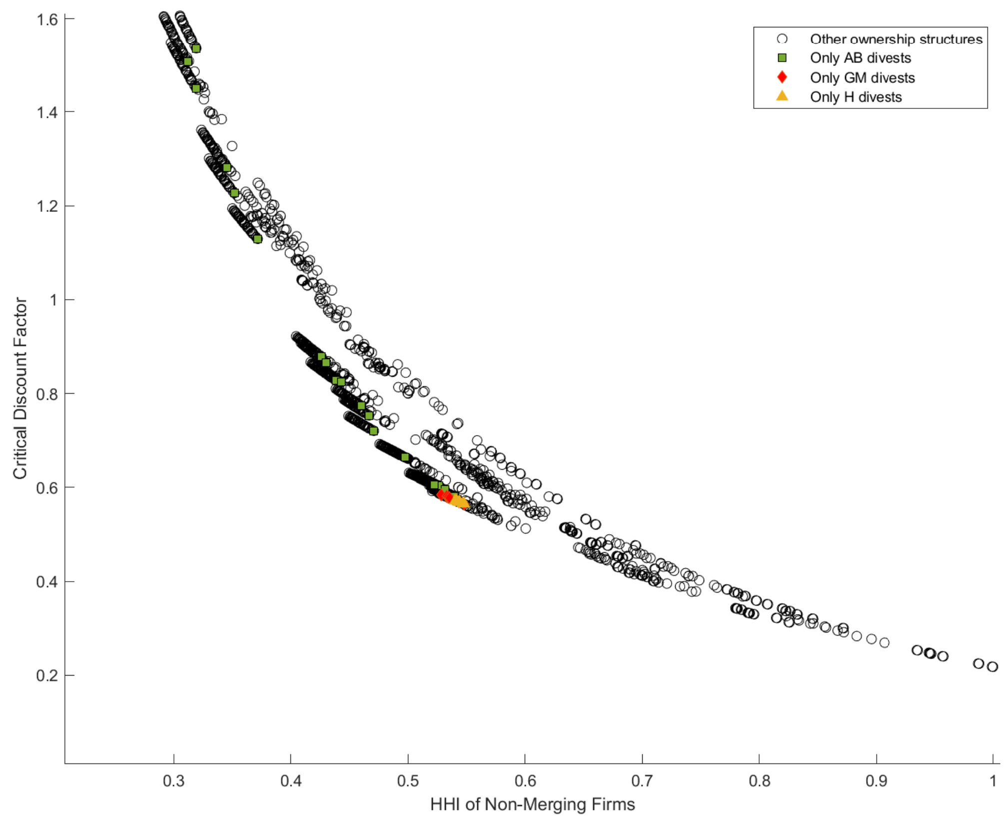

Coors/Miller merger in the representative market. The alternative values of rival HHI come from either breaking up rival brewers or merging them together, and the figure identifies some examples where particular brewers are broken up. The taste heterogeneity implied by the demand system leads to more heterogeneity than the multinomial logit, linear and AIDS demand examples, but the overall negative correlation is clear.

C Changes in Pricing After the MillerCoors Joint Venture

This Appendix presents new evidence on changes in pricing after the MillerCoors joint venture, which the Department of Justice decided not to challenge in June 2008, on the grounds that the joint venture was expected to lower distribution costs and the Department anticipated that these benefits would lower prices to consumers.[4]

A key component of the marginal costs of selling beer are transportation costs. We can think of marginal transportation costs from a brewery to a market as the distance in thousands of miles (distance) multiplied by a cost per unit per thousands of miles (cost coefficient). A lower cost coefficient would represent a distribution network operating more efficiently. The JV did reduce transportation distance, as several brands, such as Miller Lite and Coors Light, could be brewed at both parties' breweries around the country. Ashenfelter, Hosken and Weinberg (2015) show that there was a relative price decrease of MillerCoors products in markets where the reduction in transportation costs was greatest. However, this does not necessarily tell us anything about what happened to the cost coefficient, which is likely to be an aspect of efficiency that is not directly observed by rivals.

We estimate several regressions where the dependent variable is the deflated average retail price per 12-pack, either in levels or logs. Observations are brand-pack size-market-weeks, and the sample is the 13 brands included in the Miller and Weinberg (2017) demand estimation in six common pack sizes (equivalent to 6, 12, 18, 24, 30 and 36 standard-sized

[4]The Department of Justice press release (`https://www.justice.gov/archive/atr/public/press_releases/2008/233845.htm`, accessed May 17, 2023) explained that the Antitrust Division's assessment was that the JV was "likely to produce substantial and credible savings that will significantly reduce the companies' costs of producing and distributing beer. These savings meet the Division's criteria of being verifiable and specifically related to the transaction and include large reductions in variable costs of the type that are likely to have a beneficial effect on prices."

cans or bottles). The data cover the period January 2001 to December 2011, but we exclude the period for one year following the JV, so that the results should not reflect short-term reorganization effects. All regressions include date (i.e., week) fixed effects, and the reported specifications vary in how distances are measured and the combination of product and market fixed effects that are included.

Table C.1 reports estimates. In columns (1) and (5) the main regressors are brewery-to-market trucking distance, measured in 1,000 miles and taken from Miller and Weinberg (2017), for each brewer (we estimate a single importer coefficient for Heineken and Grupo Modelo) and the interaction of these variables with a post-JV indicator. Our main interest is in the post-JV$\times$ distance coefficients, with positive coefficients interpreted as meaning that a brewer priced in a way consistent with shipping beer a given distance becoming more expensive.

We include fixed effects for each brand-pack size (product) before and after the JV, as well as market fixed effects. The inclusion of post-JV product fixed effects is used to control for the possibility that tacit collusion between domestic brewers (Miller and Weinberg (2017)) may have caused the domestic product prices to rise, as well as possible production efficiencies that may have lowered price levels in every market. The distance coefficients are therefore identified from within-product variation in prices across markets, controlling for the level of average prices in the market, which may be influenced by taxes or retail market local structure for selling beer. The coefficients show that, after the JV, there was a statistically significant increase in the slope of the price-trucking distance relationship for Miller and Coors branded products, and no change in the slope for imported brands (we will discuss the Anheuser-Busch coefficients below). The Miller and Coors post-JV coefficients imply quite large price effects: for example, the Miller coefficient implies a 21 cent price increase per 12-pack in a market that is one standard deviation (269 miles) away from its brewery after a JV. The log price coefficient from the similar specification in column (5) implies a 2.1% price increase. In specification (2), the distance variable is the trucking distance multiplied by the real price of diesel fuel, which should also be observable. While the scale of the coefficients changes, the pattern remains the same.

The specifications in columns (3), (4), (7) and (8) include market-product fixed effects to control for time-invariant differences in retail prices across market-products, so that, conceptually, only changes in how prices vary with distance are identified.[5] The post-JV interaction coefficients change relatively little.

While our regression specifications are not structural pricing best response functions, we interpret the fact that importers' post-JV distance coefficients are not statistically significant (and are small in magnitude, especially in the log price specifications) as being consistent with the idea that the pricing behavior of these firms did not change, which is consistent with our modeling assumption that rivals would continue to set best response prices after the merger, as well as indicating that there was not an industry-wide decrease in distribution efficiency. The interpretation of the Anheuser-Busch coefficients, which are positive and statistically significant in some specifications, is more ambiguous. For example, the calibrated simultaneous dynamic signaling model of Sweeting, Tao and Yao (2023), where non-merging firms also invest more in signaling when rivals merge, can explain why the post-JV Anheuser-Busch coefficients should not be statistically different to those of Miller and

[5] In the Miller and Weinberg (2017) data there are handful of small distance changes for markets before the JV. The brand post-JV distance coefficients are almost unchanged if additional pre-JV distance coefficients are estimated.

Table C.1: Changes in Distance Pass-Through After the MillerCoors Joint Venture

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Variable	Real Price	Real Price	Real Price	Real Price	Log	Log	Log	Log
Price Measure					Real Price	Real Price	Real Price	Real Price
Distance Measure	Distance	Diesel × Distance	Distance	Diesel × Distance	Distance	Diesel × Distance	Distance	Diesel × Distance
Distance Measure ×								
Anheuser-Busch	-0.276	-0.040			-0.022	-0.003		
	(0.185)	(0.067)			(0.016)	(0.006)		
Imports	0.451	0.088			0.039	0.011		
	(0.181)	(0.040)			(0.015)	(0.003)		
Coors	0.245	0.131			0.026	0.013		
	(0.069)	(0.020)			(0.006)	(0.002)		
Miller	-0.083	0.032			0.003	0.006		
	(0.093)	(0.035)			(0.010)	(0.004)		
Post JV × Distance Measure ×								
Anheuser-Busch	0.619	0.189	0.537	0.182	0.052	0.015	0.044	0.015
	(0.283)	(0.091)	(0.302)	(0.094)	(0.029)	(0.009)	(0.030)	(0.009)
Imports	-0.166	-0.040	-0.107	-0.070	-0.010	-0.005	-0.007	-0.004
	(0.183)	(0.041)	(0.154)	(0.038)	(0.013)	(0.003)	(0.011)	(0.003)
Coors	0.451	0.116	0.692	0.226	0.054	0.014	0.069	0.022
	(0.145)	(0.046)	(0.133)	(0.041)	(0.015)	(0.005)	(0.014)	(0.004)
Miller	0.694	0.194	0.704	0.235	0.076	0.021	0.076	0.025
	(0.137)	(0.038)	(0.145)	(0.044)	(0.015)	(0.004)	(0.016)	(0.005)
Fixed Effects	Pre/Post × Product, Market, Week		Pre/Post × Product, Product × Market, Week		Pre/Post × Product, Market, Week		Pre/Post × Product, Product × Market, Week	
Observations	869,018	869,018	869,007	869,007	869,018	869,018	869,007	869,007
R-squared	0.923	0.922	0.947	0.947	0.915	0.915	0.942	0.942

Notes: Robust standard errors, clustered on the market, in parentheses.

Coors, whereas the model in the current paper predicts that they would not be significantly different from zero. Therefore, these specification are not able to clearly indicate which one of these alternative theories is correct.

As stressed in the text, there are two possible interpretations of finding that the merged firms priced as if the efficiency of their distribution network decreased. One interpretation is that distribution efficiency actually decreased, and the other is that efficiency did not decrease but MillerCoors priced in this way because of how rivals would respond. Without detailed data on costs we cannot tell these explanations apart (and indeed the possibility that efficiency could have decreased is an essential part of a pooling equilibrium), but we would note that MillerCoors's financial reports claimed that the joint venture realized overall efficiencies (presumably a mix of fixed and variable cost reductions) more quickly than was anticipated and news reports documented efficiencies directly related to transportation.[6]

[6] For example, in their annual report MolsonCoors noted "The expected timing to achieve the original goal of $50 million in synergies in the first twelve months of operations has accelerated, and the MillerCoors team now expects to realize $128 million of synergies by June 30, 2009. By the end of calendar year 2009, MillerCoors expects to achieve a total of $238 million in synergies, surpassing the original forecast of $225 million." A number of trade press sources (e.g., `https://www.logisticsmgmt.com/article/millercoors_taps_dunnage_innovation`) note that the merged firm introduced new systems that substantially reduced worker injuries and product losses in its distribution systems.

D Vertical Merger Example: EDM Pass-Through

This Appendix details our vertical merger example. As noted in the text, we do not fully develop the assumptions that would lead the elimination of double marginalization (EDM) not to be (partly or fully) realized in equilibrium. Instead we show, by example, that such an outcome could be profitable for a merging firm.

Our description of the example is motivated by mergers between health insurers and providers. For example, UnitedHealth Group has been expanding its Optum provider network (for example, Optum purchased the DaVita Medical Group in 2019). UnitedHealth and Optum continue to deal with independent providers and insurers respectively, which creates the possibility of raising rivals' costs effects, which, in a standard framework might offset some of the benefits of EDM. These mergers are often also justified on the grounds that the merged firm may be able to compete more effectively with fully integrated providers, such as Kaiser Permanente, which are major competitors in some regions of the country.

Specification. The pre-merger structure of the industry has two independent upstream providers (U1 and U2) and two independent insurers (D1 and D2), and a pre-existing fully vertically integrated insurer-provider (VI3) who does not contract with other providers or insurers. VI3 could represent Kaiser Permanente in some large urban markets. Di's network (N_{Di}) could contain one or both of U1 and U2. An enrollee c's indirect utilities from insurance plans are

$$u_{Di,c} = a_i + \log\left(\sum_{j \in N_{Di}} \exp(v_{Uj})\right) - \alpha p_{Di} + \theta_c + \varepsilon_{Di,c} \tag{6}$$

$$u_{VI3,c} = a_{VI3} - \alpha p_{VI3} + \theta_c + \varepsilon_{VI3,c},\ u_{0,c} = \varepsilon_{0,c}, \tag{7}$$

where the ps are downstream prices, $\theta_c \sim \mathrm{N}(0, \sigma^2)$ and the εs are logit payoff shocks. An enrollee in plans D1 or D2 expects to consume one unit of in-network care, with the provider chosen in a separate logit choice problem where v_{Uj} is the expected valuation of Uj, and the probability of using U1 when both providers are in-network is $\frac{\exp(v_{U1})}{\sum_{j=1,2} \exp(v_{Uj})}$.

We solve the model assuming that all upstream marginal costs (c) are 4, downstream retail costs (r) are 0.1, $\sigma = 2$, and $\alpha = 0.3$. We set $a_{D1} = a_{D2} = 8$, $v_{U1} = v_{U2} = 0$ and $a_{VI3} = 8 + \log(2)$, so that, when D1 and D2 contract with both providers, all downstream firms have symmetric demand, although with different costs. These parameters imply almost all consumers will buy from one insurer in the pre-merger equilibrium.

Bargaining and Pre-Merger Outcomes. We solve for pre-merger outcomes assuming that there are two stages and complete information. In the first stage, independent provider-insurer pairs simultaneously engage in Nash-in-Nash bargaining over the wholesale price of patient care, where the outside option is that the provider is not in-network. We assume expected surplus is split equally. In the second stage, insurers simultaneously set retail prices to maximize their static profits taking the first-stage prices and networks as given.

Solving for the pre-merger complete information equilibrium requires us to find the following:

- five sets of downstream (retail) prices: the retail prices when networks are complete, and retail prices where one of the four possible U-D bargains is not completed. This is a total of 15 (5×3) retail prices, for which there are 15 associated pricing first order conditions that treat the negotiated ws as given. For example, for downstream firm i, the first-order condition, when we normalize market size to 1, is

$$\frac{\partial s_{Di}}{\partial p_{Di}}(p_{Di} - \tilde{c}_i - r_i) + s_{Di} = 0$$

where $\tilde{c}_i$ is Di's cost of a customer given its negotiated wholesale prices (or the upstream marginal cost of VI3).

- four wholesale prices. The Nash-in-Nash assumption implies these are determined as the solutions to maximization problems of the form

$$w_{Uj,Di} = \arg\max_{w_{Uj,Di}} \left(\pi_{Uj}(c_j, \mathbf{w}) - \pi_{Uj}(c_j, \mathbf{w}\backslash w_{Uj,Di})\right)^{\tau} \times \left(\pi_{Di}(\mathbf{w}) - \pi_{Di}(\mathbf{w}\backslash w_{Uj,Di})\right)^{1-\tau}$$

where $\pi_{Uj}(c_j, \mathbf{w})$ is the profit of Uj given patient flows when both providers are in both insurer networks and D1 and D2 set optimal retail prices given the ws. τ is the upstream share of bargaining surplus, assumed to be one-half in our example. $\pi_{Uj}(c_j, \mathbf{w}\backslash w_{Uj,Di})$ is the profit of Uj when it is not in the network of Di, but it is in the other insurer's network, the other bargained prices are held fixed, and, in the second stage, the Ds set retail prices that reflect this incomplete network configuration. The associated first order conditions have the form:

$$\tau\left(\pi_{Di}(\mathbf{w}) - \pi_{Di}(\mathbf{w}\backslash w_{Uj,Di})\right)\frac{\partial \pi_{Uj}(\mathbf{w})}{\partial w_{Uj,Di}} + (1-\tau)\left((\pi_{Uj}(c_j, \mathbf{w}) - \pi_{Uj}(c_j, \mathbf{w}\backslash w_{Uj,Di})\right)\frac{\partial \pi_{Di}(\mathbf{w})}{\partial w_{Uj,Di}} = 0$$

where the derivative terms reflect how the retail prices of all firms will change with the wholesale price $w_{Uj,Di}$, with the other wholesale prices held fixed.

Column (1) of Table D.1 shows the pre-merger equilibrium prices and shares. Upstream and downstream margins are substantial (i.e., $7.98 - 4$ upstream, and $12.56 - 8.08$ downstream), which leads to the possibility that EDM effects could be large. The vertically integrated firm, VI3, charges a significantly lower price. The choice of parameters was partly motivated by the urban market examples in Ho and Lee (2017), although the assumption that there are only two independent upstream providers, and that they are symmetric is a significant simplification. For example, in Ho and Lee's Los Angeles data, Kaiser has just under 50% of the public-sector insurees in their sample, its premiums are somewhat lower than Blue Cross and Blue Shield, and its demand is less elastic.

Standard Analysis of a Vertical Merger. We consider a U1-D1 merger, forming U1D1, with demand and marginal costs unaffected. Consistent with our motivating example, where D1 is UnitedHealth and U1 is a provider, we assume that U1D1 continues to negotiate with U2 and D2. The standard analysis assumes that, after the merger,

- D1 will set its retail price recognizing that its marginal cost is $\frac{exp(v_{U1})}{\sum_{j=1,2}\exp\left(v_{Uj}\right)} \times c_1 + \frac{\exp(v_{U2})}{\sum_{j=1,2}\exp\left(v_{Uj}\right)} \times w_{U2,D1}$. The replacement of $w_{U1,D1}$ with c_1 reflects EDM.
- when U1 is bargaining with D2 it will recognize that, in the event of a failure to agree a deal, demand will switch from D2 to D1, and that this will imply some profits for U1D1. In bargaining, this will tend to increase the U1-D2 wholesale price, reflecting the merged firm's increased bargaining leverage (its profit when D2 does not have U1 in its network).

Table D.1: Alternative Outcomes After a Vertical U1-D1 Merger.

	(1)	(2)	(3)	(4)	(5)
Pre-/Post-Merger	Pre	Post	Post	Post	Post
p_{VI3}	Endogenous	Endogenous	Endogenous	Fixed	Fixed
EDM		EDM	No EDM	EDM	No EDM
$w_{U1,D1}$	7.98	4 ($\equiv$ mc)	7.98 (fixed)	4 ($\equiv$ mc)	7.98 (fixed)
$w_{U1,D2}$	7.98	9.99	10.41	9.39	9.25
$w_{U2,D1}$	7.98	6.31	7.34	5.88	6.91
$w_{U2,D2}$	7.98	8.04	8.44	7.53	7.64
Retail Prices					
p_{D1}	12.56	11.20	13.20	11.09	12.82
p_{D2}	12.56	13.22	13.79	12.79	12.93
p_{VI3}	10.48	10.25	10.93	10.81	12.41
Downstream Shares					
s_{D1}	0.256	0.345	0.259	0.354	0.249
s_{D2}	0.256	0.188	0.217	0.212	0.240
s_{VI3}	0.478	0.458	0.512	0.424	0.501
Profits Per Potential Consumer					
U1+D1	2.167	2.613	2.620	2.714	2.436
U2	1.019	0.778	0.915	0.707	0.798
D2	1.148	0.773	0.925	0.899	1.054
VI3	3.049	2.819	3.495	2.709	3.196
Consumer Surplus per Potential Consumer	20.958	21.327	20.282	21.351	20.801

Notes: Authors' calculations.

- when D1 is bargaining with U2 it will recognize that, in the event of a failure to agree a deal, D1's customers will use U1 and that the cost to D1 will be c_1 (lower than the wholesale price it had to pay prior to the merger). This will increase the bargaining leverage of D1, causing downwards pressure on the U2-D1 wholesale price. D1's leverage may also be increased by its greater incentive, when it benefits from EDM, to lower its retail price to increase its demand if a U2-D1 deal is not agreed.

The equations are the same as before the merger except that: (i) the bargaining first-order condition for $w_{U1,D1}$ is eliminated; (ii) the equation for $w_{U2,D1}$ reflects the profit that the merged U1D1 will make on patients who use U1, including when U2 is not in D1's network; (iii) the equation for $w_{U1,D2}$ will reflect the profit that the merged U1D1 will make on patients who use D1, including when U1 is not in D2's network; and, (iv) D1's marginal cost in its retail pricing equation reflects c_1 for the share of its customers who get care from U1. This approach matches how we have seen agencies evaluate vertical mergers, even if this type of modelling might not be presented as evidence in litigation, and it is also consistent with the standard approach in the structural academic literature.

Column (2) of Table D.1 shows the outcome when the game is resolved with exactly the same order of moves, and all players understand that D1's retail price will be set recognizing that the cost of care of a D1 enrollee at U1 is 4. The elimination of the upstream margin causes U1D1 to lower its downstream price. The merger also implies changes in bargaining leverage that increase $w_{U1,D2}$ and lower $w_{U2,D1}$. The merger is profitable for the merging firms (i.e., U1D1 profits are higher than combined U1 and D1 pre-merger profits) and good for consumers, but its profitability is reduced by p_{VI3} falling significantly.

Vertical Merger with No EDM. Column (3) shows the outcome where all players know that, when U1D1 is setting the D1 retail price, it will act as if its cost of care at U1 is the pre-merger $w_{U1,D1}$, although it will use the wholesale prices negotiated in a post-merger first stage for the other contracts. Note that here we are not micro-founding why D1 has not changed the wholesale price that it is using, but simply looking at how making this ad hoc assumption affects outcomes.

All prices and the profits of all of the firms increase relative to column (2), the standard prediction of the post-merger outcome, even though U1D1's downstream market share falls by around 30%.[7] Reflecting the continued existence of bargaining leverage that increases the costs of D2, and encourages VI3 to raise its prices, consumer surplus is lower than pre-merger. VI3's prices and profits are also higher than before the merger. U2 and D2's profits are smaller than before the merger, so this outcome is still consistent with the rivals that will need to deal with the merged firm complaining about the transaction.

The column (3) calculations assume that the merged firm does not achieve any of the benefits of EDM. However, as in many horizontal mergers, the most profitable outcome for the merged firm might be to partially realize the efficiency. We calculate that U1D1's profits would be maximized if it could commit that D1 would set the retail prices as if U1D1's cost of care at U1 was 6.07. This is roughly halfway between the actual marginal cost and the pre-merger wholesale price, so one might think of this as roughly 50% EDM pass-through. While consumers would be better off than in column (3), they would be worse off than before the merger as the prices of VI3 and D2 would increase, even though D1's price falls slightly.

[7] As pointed out by Salinger (1991), models where downstream firms sell multiple products can also lead to results where vertical integration raises all prices.

The Role of the Endogenous VI3 Price. Our horizontal merger analysis highlighted the key role of the responsiveness of rivals' prices to the prices set by the merging firm, and the role of rival market structure in determining the level of responsiveness. Unfortunately, changing the market structure of VI3 (for example, by breaking it up into three firms) changes strategic incentives in the first-stage bargaining game so that it is not possible to do an experiment where competition is increased while the size of the merged firm's pre-merger margins, market shares and costs remain the same.

Therefore columns (4) and (5) consider the simpler counterfactual experiment where we calculate post-merger outcomes when U1D1 fully internalizes EDM and when it does not internalize it at all (i.e., it continues to use 7.98 as the cost of care at U1) if VI3's price is known to be fixed at its pre-merger level.[8] We assume that D2's price is endogenous, so the raising rivals' costs incentive remains. With p_{VI3} fixed, U1D1's profits are higher when EDM is internalized than when it is not, which is consistent with our intuition that the incentive to realize efficiencies tends to be muted when rivals would respond by lowering their prices.

[8]We also assume that VI3's prices would be at the levels they would have been in the pre-merger game if there is no agreement between U1D1 and U2 or D2, or between U2 and D2.

References

Ashenfelter, Orley C, Daniel S Hosken, and Matthew C Weinberg. 2015. "Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry." *RAND Journal of Economics*, 46(2): 328–361.

Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela. 2008. "The IRI Marketing Dataset." *Marketing Science*, 27(4): 745–748.

Deaton, Angus, and John Muellbauer. 1980. "An Almost Ideal Demand System." *American Economic Review*, 70(3): 312–326.

Ho, Katherine, and Robin S Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica*, 85(2): 379–417.

Miller, Nathan H, and Matthew C Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica*, 85(6): 1763–1791.

Miller, Nathan H, Marc Remer, Conor Ryan, and Gloria Sheu. 2016. "Pass-Through and the Prediction of Merger Price Effects." *Journal of Industrial Economics*, 64(4): 683–709.

Miller, Nathan H, Marc Remer, Conor Ryan, and Gloria Sheu. 2017. "Upward Pricing Pressure as a Predictor of Merger Price Effects." *International Journal of Industrial Organization*, 52: 216–247.

Panhans, Matthew T, and Charles Taragin. 2023. "Consequences of Model Choice in Predicting Horizontal Merger Effects." *International Journal of Industrial Organization*, 89: 102986.

Salinger, Michael A. 1991. "Vertical Mergers in Multi-Product Industries and Edgeworth's Paradox of Taxation." *Journal of Industrial Economics*, 545–556.

Sweeting, Andrew, Xuezhen Tao, and Xinlu Yao. 2023. "Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers." University of Maryland.