Pass-Through and the Prediction of Merger Price Effects*

Nathan H. Miller†
Georgetown University

Marc Remer‡
Department of Justice

Conor Ryan‡
Department of Justice

Gloria Sheu‡
Department of Justice

December 31, 2015

Abstract

We develop how first order approximation can be used to make counterfactual price predictions in oligopoly models. We extend the theoretical results of **?** on mergers to any counterfactual scenario involving perturbations to firms' first order conditions. We then use Monte Carlo experiments to evaluate accuracy. We find that (i) first order approximation is more accurate than simulation in 91.7% of the mergers considered; (ii) it is more accurate than simulation in 98% of the cost shocks considered; and (iii) simple versions of approximation, of interest to antitrust practitioners, exist and systematically outperform merger simulation.

Keywords: first order approximation; cost pass-through; simulation; mergers
JEL classification: K21; L13; L41

*We thank Nicholas Hill, Liran Einav, Sonia Jaffe, Alexander Raskovich, Charles Taragin, Glen Weyl and Nathan Wilson, as well as seminar participants at the Department of Justice, Drexel University, Federal Trade Commission, Michigan State University, Stony Brook University, University of Virginia, and Williams College for valuable comments. The views expressed herein are entirely those of the authors and should not be purported to reflect those of the U.S. Department of Justice.

†Georgetown University, McDonough School of Business, 37th and O Streets NW, Washington DC 20057. Email: nhm27@georgetown.edu.

‡Department of Justice, Antitrust Division, Economic Analysis Group, 450 5th St. NW, Washington DC 20530. Email: marc.remer@usdoj.gov, conor.ryan@usdoj.gov, and gloria.sheu@usdoj.gov.

1 Introduction

The profit maximizing level of cost pass-through in many standard oligopoly models depends on both the first and second derivatives of the consumer demand schedules. This insight dates back at least to ? (?) and ? (?) and is extended and generalized in ? (?) and ? (?). The more recent literature emphasizes that observed pass-through can be used to answer questions that arise in the fields of industrial organization, international trade, and mechanism design. An emerging empirical literature uses pass-through to study trade costs (e.g., ? (?)), environmental regulation (e.g., ? (?)) and health insurance (e.g., ? (?)).

In this paper, we study how pass-through can inform predictions of merger price effects. Such counter-factual predictions are a persistent focus of empirical industrial organization, and the methodologies developed are of practical importance in antitrust enforcement. It is well established that simulation techniques can be sensitive to assumptions, made implicitly through the selection of function forms, regarding the curvature of demand (e.g. ? (?); ? (?)). Because pass-through is driven in part by demand curvature, it is natural to think that, when observed, pass-through could be used to reduce prediction error caused by functional form misspecification.

We focus in particular on two methodologies that incorporate pass-through information. The first is based on the theoretical finding of ? (?) that a first order approximation (FOA) to the post-merger price can be calculated given knowledge of the first and second derivatives of demand. Provided that elasticities can be estimated or calibrated, the second derivatives of demand can be inferred from pass-through. The first order effects of the merger then can be calculated without reliance on functional form assumptions.[1] ? (?) prove that FOA is precise for arbitrarily small price changes – here we extend inquiry to nontrivial mergers. The second methodology, which we refer to as "informed simulation," selects a demand system that generates pass-through that is close (by some distance metric) to observed pass-through. Simulation then is conducted with the selected demand system. Informed simulation is proposed in ? (?), and should mitigate misspecification error.

Our findings rely upon a series of Monte Carlo exercises. We generate a data set comprised of a large number of markets where the underlying demand system is either logit, linear, almost ideal, or log-linear. These four demand systems allow for a wide range of curvature and pass-through conditions and are commonly employed in antitrust analyses of mergers involving differentiated products (? (?), ? (?)). They have also been used in academic studies that examine the effect of demand curvature on the precision of counterfactual

[1] A minor caveat to this statement applies in many settings, and we discuss in Section 2.2.

simulations (e.g., ? (?), ? (?), ? (?)). In each market, we perfectly observe the price effects resulting from a hypothetical merger, and compare these effects to those predicted by FOA and informed simulation. We study scenarios in which pass-through is observed perfectly, and also scenarios in which pass-through is observed with measurement error or systematic bias.

A primary result is that FOA dominates standard merger simulation if pass-through is observed perfectly and there is some functional form misspecification in the simulation. The predictions of FOA are tightly distributed around true price effects, and the median absolute prediction error (MAPE) that arises with FOA typically is a fraction of the MAPE with standard merger simulation. The absolute prediction error (APE) of FOA is smaller than that of standard merger simulation in 93% of the merger scenarios considered. Further, when price effects are evaluated against a specific ten percent threshold, FOA produces far fewer false positives and false negatives than standard merger simulation.

Measurement error and bias in pass-through degrade the performance of FOA, both absolutely and relatively to standard merger simulation. For example, when pass-through is observed within 90% of its true value, the MAPE that arises with FOA and standard merger simulation are often of similar magnitudes, and when functional form specification is minor (e.g., linear demand when the true demand system is logit), merger simulation tends to be more accurate than FOA. Upward bias in pass-through causes FOA to over-predict price increases, just as downward bias leads to under-predictions. Taken together, our results on FOA underscore that having accurate information on pass-through is valuable when making counter-factual prediction, and should help motivate efforts to collect such information.

The accuracy gains from informed simulation are more modest. This is due in part to the experimental design. Given perfect knowledge of pass-through, as posited initially, and the design of our experiments, it is possible to identify the correct demand system with which to perform prediction. We view this as unrealistically optimistic because, in practice, consumer decisions need not align with *any* of the models used in our experiments. Thus, to implement informed simulation, we identify the misspecified demand system that produces pass-through that is closest to what is observed, and simulate using that demand system. By design, informed simulation should be less accurate than FOA in the presence of perfectly observed pass-through, and this is reflected in our results.

Informed simulation nonetheless improves upon standard merger simulation. For instance, with perfectly observed pass-through, informed simulation is at least as accurate as standard merger simulation in 90% of the merger scenarios and is more accurate in 57% of the merger scenarios. Further, informed simulation tends to be more robust than FOA to measurement error in pass-through. The MAPEs that arise with informed simulation and FOA are roughly equal when pass-through is observed within 60% of its true value, and informed simulation is clearly more accurate when pass-through is observed within 90% of its true value. This yields the insight that it can be appropriate to interpret pass-through with an economic model when pass-through is observed with significant measurement error, even if the model introduces functional form misspecification. Restated, in the presence of significant measurement error, pass-through is better employed as an object that informs functional forms than as a direct input to prediction.

The analysis is subject to the caveat that the data generating process used in the Monte Carlo experiments need not reflect the conditions of actual markets. The magnitude of prediction error that arises due to functional form misspecification, in particular, is driven by our reliance on the logit, linear, almost ideal, and log-linear demand systems. We nonetheless consider the results to be of value, insofar as they extend the theoretical insights of **?** (**?**) beyond arbitrarily small mergers, and inform how pass-through can best be used to improve counter-factual predictions. Additionally, we note that some of the accuracy gains we document can alternatively be achieved by using the random coefficients logit (RCL) demand system, as in **?** (**?**). The RCL is attractive because it is (theoretically) flexible enough to match the elasticities and curvature of the true underlying demand system, thereby mitigating substantially functional form misspecification. In most applications, however, the specification employed affords only limited flexibility. Further, the RCL is not often used antitrust enforcement due time constraints and the computation demands of estimation.

The paper proceeds as follows. We sketch the theoretical framework in Section 2. The focus is on mergers in differentiated-products Nash-Bertrand models, and we develop how pass-through can be used to inform prediction following **?** (**?**). Section 3 provides the details of the Monte Carlo experiments. Section 4 presents summary statistics on pass-through and the merger price effects that arise in the data. Section 5 presents results with and without measurement error in pass-through, and Section 6 concludes.

2 Theoretical Framework

2.1 Merger Price Effects In Bertrand-Nash Competition

We examine mergers in the context of a Bertrand-Nash oligopoly model of price competition among multi-product firms. Mergers change the unilateral pricing calculus of the merging firms and, provided products are substitutes and countervailing merger efficiencies are small, result in a new equilibrium characterized by higher prices. Assume that each firm faces a well-behaved, twice-differentiable demand function. The equilibrium prices of each firm $i \in I$ satisfy the following first order conditions:

$$f_i(P) \equiv -\left[\frac{\partial Q_i(P)}{\partial P_i}^T\right]^{-1} Q_i(P) - P_i + MC_i(Q_i(P)) = 0 \quad \forall i \in I \tag{1}$$

where P_i is a vector of firm i's prices, $Q_i(P)$ is a vector of firm i's unit sales, P is a vector containing the prices of every product, and MC_i is the marginal cost function. Consider a merger between firms j and k that, for simplicity, does not affect the marginal cost and demand functions. The first order condition changes such that:

$$h_i(P) \equiv f_i(P) + g_i(P) = 0 \quad \forall i \in I \tag{2}$$

where

$$g_j(P) = \underbrace{-\left(\frac{\partial Q_j(P)^T}{\partial P_j}\right)^{-1}\left(\frac{\partial Q_k(P)^T}{\partial P_j}\right)}_{\text{Matrix of Diversion from } j \text{ to } k} \underbrace{(P_k - MC_k^1)}_{\text{Markup of } k} \tag{3}$$

and $g_k(P)$ is defined analogously, while $g_i(P) = 0$ for all $i \neq j, k$. The g function is the product of firm k's markups and the matrix of diversion ratios between firms j and k, which depend upon the first derivatives of the demand functions. Equation (3) captures an opportunity cost created by the merger: each merging firm, when making a sale, possibly forgoes a sale of its merging partner (**?** (**?**)).[2]

The prices that satisfy equation (2) depend on how the demand and marginal cost functions change as prices move away from the pre-merger equilibrium. Nonetheless, the first order effects of the merger depend only on information that is local to the pre-merger equilibrium (**?** (**?**)). Specifically, a first order approximation ("FOA") to the price changes that arise from the merger is given by:

$$\Delta P = -\left(\frac{\partial h(P)}{\partial P}\right)^{-1}\Bigg|_{P=P^0} g(P^0) \tag{4}$$

where P^0 is the vector of pre-merger prices. The first order effects therefore depend upon the opposite inverse Jacobian of $h(P)$, which **?** refer to as the *merger pass-through* matrix. This matrix incorporates both the first and second derivatives of demand, and can be conceptu-

[2]The g function is referred to in the antitrust literature as upward pricing pressure (UPP).

alized as the rate at which the change in pricing incentives from the merger are transmitted to consumers. Therefore, when using equation (2) or (4) to infer the price changes that arise from a merger, the accuracy of the inference depends on how well the higher-order properties of real-world demand are captured.

2.2 Pass-Through and Prediction

Merger simulation is the predominant methodology in the industrial organization literature used to predict the price effects from a merger. This methodology requires functional forms for the demand and marginal cost functions to be selected and parameterized, which in turn allows post-merger prices to be computed as the solution to the post-merger first order conditions.[3] Because the assumed functional forms implicitly restrict the second derivatives of demand, misspecification bias can arise even if the demand function captures perfectly the elasticities (i.e. the first derivatives) that arise in the pre-merger equilibrium.[4]

We explore the extent to which pass-through can be used to inform the second derivatives of demand. **?** (**?**) demonstrate, through an application of the implicit function theorem, that the cost pass-through matrix in pre-merger equilibrium is given by:

$$\rho(P)|_{P=P_0} = -\left(\frac{\partial f(P)}{\partial P}\right)^{-1}\Bigg|_{P=P_0} \tag{5}$$

Thus, pass-through equals the opposite inverse of the pre-merger first order conditions, and it depends directly on both the first and second derivatives of demand. It follows that, given demand elasticities, equation (5) provides a mapping between pass-through and the second derivatives. This allows for the second derivatives of demand to be imputed from pass-through, and used to calculate FOA as proposed in **?** (**?**). Alternatively, equation (5) can be used to obtain the pass-through rates that arise under different candidate demand systems. Then, an informed simulation can be conducted using the functional form of demand that generates pass-through close to the observed pass-through rates (**?** (**?**)).[5] Either approach

[3]In practice, elasticities are typically obtained through demand estimation or calibration. A substantial literature focuses on the conditions under which regression analysis recovers consistent estimates of consumer substitution (e.g., **?** (**?**); **?** (**?**)). Elasticities alternatively could be calibrated to match price-cost margins and customer switching patterns, as is more common in merger enforcement (e.g., **?** (**?**)).

[4]For many common demand systems, the second derivatives are fully determined by the elasticities. This is the case for the linear, logit, nested logit, almost ideal, and log-linear demand systems. The random coefficient logit model is theoretically capable of divorcing the first and second derivatives, but in most applications the specification employed results in only a limited amount of flexibility.

[5]We define the distance metric that we use to evaluate "closeness" in the next section. Informed simulation

operates to mitigate misspecification error.

Because the number of second derivatives exceeds the number of pass-through terms, restrictions in addition to equation (5) are needed to identify the full set of second derivatives. This is relevant if one is attempting to calculate FOA based on cost pass-through. Slutsky symmetry is sufficient to identify all second derivatives for duopoly markets. If there are more than two firms, then second derivatives of the form $\partial^2 Q_i/(\partial P_j \partial P_k)$, for $i \neq j$, $i \neq k$ and $j \neq k$, remain unidentified without further restrictions. As suggested in **?** (**?**), the following assumption is sufficient:

$$\frac{\partial^2 Q_i}{\partial P_j \partial P_k} = \frac{\partial^2 Q_i}{\partial^2 P_i} \frac{\frac{\partial Q_i}{\partial P_j}\frac{\partial Q_i}{\partial P_k}}{\left(\frac{\partial Q_i}{\partial P_i}\right)^2} \qquad (i \neq j,\, i \neq k,\, j \neq k) \tag{6}$$

This restriction is exact only if demand adheres to a modified horizontality condition.[6] Thus, imputation based on equation (6) itself can introduce misspecification error. However, because error is only introduced for a limited subset of second derivative terms, one might expect this to be inconsequential relative to the misspecification error that may arise with simulation. Indeed, the Monte Carlo evidence we develop indicates that the loss of predictive accuracy that arises with this imputation tends to be small.[7]

3 Monte Carlo Experiments

3.1 Overview

In the remainder of the paper, we present numerical evidence on the extent to which cost pass-through information can improve the accuracy of counterfactual predictions, including instances in which pass-through is observed with measurement error or bias. All of the numerical experiments take as given the demand elasticities that arise in the pre-merger

still requires that a menu of candidate demand systems be selected, and the results can depend on which demand systems are included.

[6]The condition, proposed in **?** (**?**), is that $Q_i(P) = \psi\left(P_i + \sum_{j \neq i} \mu_j(P_j)\right)$ for some $\psi : \mathbb{R} \rightarrow \mathbb{R}$ and $\mu : \mathbb{R} \rightarrow \mathbb{R}$. Among the four demand systems considered later in this paper, only linear demands satisfy the condition precisely.

[7]In Appendix Figure C.1, we use scatter-plots to compare the predictions of FOA calculated based on equation (6) to FOA predictions calculated with perfect knowledge of second derivatives. FOA predictions across the two approaches are nearly identical with logit, almost ideal, and linear demand, and remain similar with log-linear demand. It follows that imputation under the modified horizontality condition does not create meaningful misspecification error in our experiments.

equilibrium. We focus instead on perturbing what is known about demand curvature, as revealed through pass-through.[8]

We work with the logit, almost ideal, linear, and log-linear demand systems.[9] Because the curvature properties of these systems are fully determined by the elasticities, we can calibrate them such that the first derivatives are identical across the demand systems in the pre-merger equilibrium but the curvature (and pass-through) conditions differ. This conveys tractability to the data generating process and facilitates comparisons across demand systems. Given the theoretical relationship between demand curvature and the magnitude of merger price effects, a reasonable hypothesis is that FOA and informed simulation should outperform standard merger simulations if the observed pass-through information is of sufficiently high quality. Our experiments largely confirm this hypothesis. Importantly, we are able to quantify both how much pass-through can improve predictive accuracy and how quickly improvements diminish as measurement error and bias are introduced to the observed pass-through rates.

3.2 Data generating process

We generate simulated data that comport with the theoretical assumptions outlined previously. The markets feature four firms that produce differentiated products with a constant returns-to-scale production technology. Competition is in prices and equilibrium is Bertrand-Nash. Each draw of data is independent and characterizes the conditions of a single market, and the simulated data cover a wide range of competitive conditions that derive from the randomized draws. We normalize all prices to unity in the pre-merger equilibrium, which conveys the advantage that merger effects are the same in levels and percentages.[10] The details of the data generating process are as follows:

1. Randomly draw (i) market shares for four firms and an outside good, and (ii) the first firm's margin based on a uniform distribution bounded between 0.20 and 0.80.

2. Calibrate the parameters of a logit demand system based on the margin and market shares, and calculate the demand elasticities that arise in the pre-merger equilibrium. This entails selecting demand parameters that rationalize the random data. The parameters are exactly identified given market shares, prices, and a single margin.

[8] We use the term "demand curvature" interchangeably with second derivatives of demand.

[9] The four demand systems are commonly employed in academic research and antitrust analyses of mergers (e.g., **?** (?); **?** (?)); **?** (?).

[10] The loss of generality caused by the price normalization is limited, and we have confirmed that alternatives do not affect results.

3. Calibrate linear, almost ideal, and log-linear demand systems based on the logit demand elasticities. The parameters of these systems are exactly identified given market shares, prices, and the logit demand elasticities.[11]

4. Simulate the price effects of a merger between two firms under each of the demand systems.

5. Repeat steps (1) - (4) until 3,000 draws of data are obtained.

The algorithm generates 12,000 mergers to be examined, each defined by a draw of data and a demand system.[12] As discussed above, the data generating process imposes that pre-merger demand elasticities are identical across demand systems for a given draw of data. We provide mathematical details on the calibration process in Appendix A.

The data generating process allows us to isolate the role of demand curvature in driving merger price effects and to explore cleanly how curvature assumptions matter for simulation. For instance, consider a merger defined by a given draw of data and the logit demand system. The true price effect of the merger is obtained from a logit simulation, and this can be compared against simulation results obtained under alternative assumptions of almost ideal, linear, and log-linear demand. The existing literature indicates that prediction error due to functional form misspecification along these lines is substantial (e.g., **?** (**?**)), and this sensitivity is consistent with the theoretical results provided in Section 2.

To assess the extent to which pass-through improves predictive accuracy, we posit first that the cost pass-through matrix is available for use without measurement error or bias.[13] We then calculate FOA based on equation (4), using the horizontality restriction of equations (5) and (6) to impute second derivatives.[14]

[11] In the pre-merger equilibrium, consumer substitution between products is proportional to market share because all the systems are calibrated based on logit elasticities. This reduces the dimensionality of the random data that must be drawn. The substitution-by-share property is retained away from the pre-merger equilibrium only for logit demand.

[12] A small number of draws cannot be rationalized with logit demand – this arises if the first firm has both an unusually small market share and an unusually high price-cost margin. We replace these to obtain the 3,000 draws. The data generating process also produces some markets that exhibit extreme pass-through conditions, and others with no post-merger equilibria. We exclude those calibrations from the analysis, treating as extreme a pass-through rate that is negative or exceeds ten. The pass-through criterion eliminates 74 AIDS markets and 164 log-linear markets. We do not redraw these markets.

[13] We obtain the pre-merger cost pass-through matrix from equation (5), using second derivatives that are obtained analytically from the calibrated demand systems (see **?** (**?**)).

[14] To be clear, we use the actual second derivatives to obtain the pass-through matrices. Then we assume that pass-through, but not the second derivatives, is available for use in prediction. With FOA, the imputed second derivatives match the actual second derivatives exactly only for the case of linear demand, due to the misspecification bias that otherwise is introduced by the modified horizontality restriction.

We also evaluate an "informed simulation" in which pass-through is used to select an appropriate demand system. Given perfect knowledge of pass-through, as posited initially, and the design of our experiments, it is possible to identify the correct demand system with which to perform prediction. We view this as unrealistically optimistic because, in practice, consumer decisions need not align with *any* of the models used in our experiments. It follows that testing the accuracy of informed simulation should involve the mitigation, but not the complete elimination, of functional form misspecification. Thus, in our implementation, we identify the misspecified demand system that produces pass-through that is closest to what is observed, and simulate using that demand system.[15]

These results in hand, we next incorporate measurement error and bias into the observed pass-through data, and evaluate how the predictive accuracy of FOA and informed simulation changes. To add noise, we add a uniformly distributed error to each element of the pass-through matrix. Mathematically, we define the observed pass-through element (j,k) to be

$$\widetilde{\rho}_{jk} = \rho_{jk} + \epsilon \quad \text{where} \quad \epsilon \sim U(\rho_{jk} - t\rho_{jk}, \rho_{jk} + t\rho_{jk}) \tag{7}$$

The support of the error is element-specific and depends on t. We use three different levels for t, such that pass-through is observed alternately within 30, 60, and 90 percent of its true value. To add bias, we suppose that what is observed for each element (j,k) is $\widetilde{\rho}_{jk} = \rho_{jk}(1+s)$, where we set $s = \pm 0.15$ to reflect some degree of upward or downward bias.[16]

4 Summary Statistics

4.1 Empirical distributions

In Table 1, we summarize the empirical distributions that arise in the data. The market shares and margins of firm 1 are obtained from random draws. Because shares are allocated among the four products and the outside good, the distribution of firm 1's share is centered around 20 percent. The margin distribution reflects uniform draws with support over (0.20, 0.80). The own-price elasticity of demand, which equals the inverse margin, has a distribution centered around 2.08, and 90 percent of the elasticities fall between 1.32 and 4.38.

[15] We use mean squared error as the distance measure. Let ρ_{jk} be the (j,k) element of the observed pass-through matrix, and let $\widehat{\rho}^i_{jk}$ be the analog for demand system i. The mean squared error for demand system i is given by $MSE_i = \sum_{j,k}(\rho_{jk} - \widehat{\rho}^i_{jk})^2$.

[16] Recent research demonstrates that standard orthogonality conditions are insufficient to ensure that reduced-form regressions of prices on cost shifters yield unbiased estimates of pass-through (**?** (**?**)). Bias arises, for example, if pass-through is not constant in prices and the cost distribution is asymmetric.

These statistics are invariant to the posited demand system because the demand systems are calibrated to reproduce the same first-order characteristics in the pre-merger equilibria.

Cost pass-through depends on demand curvature and varies across the four demand systems. We define *own pass-through* as the effect an individual firm's costs on its equilibrium price. The own pass-through terms fall along the diagonal of the pass-through matrix. Median own pass-through equals 0.80, 1.19, 0.53, and 1.87 for the logit, almost ideal, linear and log-linear demand systems, respectively. Own-cost pass-through has wide support for the almost ideal and log-linear demand systems but is more tightly distributed for the logit and (especially) the linear demand systems. We define *cross pass-through* as the effect of a specific competitor's cost on an equilibrium price – cross pass-through is isomorphic to strategic complementarity in prices (**?** (**?**)). The cross pass-through terms are the off-diagonal elements of the pass-through matrix. Median cross pass-through equals 0.04, 0.22, 0.09, and 0.00 across the four demand systems. Thus, while the almost ideal and log-linear demand systems both tend to generate large own pass-through, only the AIDS generates large cross pass-through because prices are not strategic complements (or substitutes) with log-linear demand.

We also report statistics for *industry pass-through*, which we define as the effect on equilibrium prices of a cost increase that is experienced by all firms. While knowledge of industry pass-through alone is insufficient to obtain a FOA to a merger price effect, it can inform counter-factual prediction in some simpler settings. Further, much of the existing empirical literature relies on industry-wide cost changes for identification, such as exchange rate fluctuations (e.g., **?** (**?**)), sales taxes (e.g., **?** (**?**)), and input prices (e.g., **?** (**?**)). Our data inform the levels of industry pass-through that one might expect to estimate with reduced-form techniques, absent trade costs and other market frictions. As shown, median industry pass-through equals 0.95, 1.90, 0.79, and 1.87 for the logit, almost ideal, linear, and log-linear demand systems, respectively. Industry pass-through will always exceed own pass-through if prices are strategic complements, as is the case for three of our demand systems.

The median merger price effects are 0.09, 0.18, 0.08, and 0.30 for the logit, almost ideal, linear, and log-linear demand systems, respectively. Because pre-merger prices are normalized to one, these statistics reflects both the median level change and median percentage change. Dispersion within demand systems mainly reflects the range of upward pricing pressure that arises from the data generating process. Dispersion across demand systems reflects the specific pass-through properties of the systems, with greater own pass-through associated with larger price effects. This relationship, first observed in **?** (**?**), is explained by the

Table 1: Order Statistics

	Median	5%.	10%	25%	75%	90%	95%
Characteristics Invariant to Demand Form							
Market share	0.21	0.03	0.06	0.13	0.28	0.35	0.40
Margin	0.48	0.23	0.26	0.34	0.62	0.72	0.76
Elasticity	2.08	1.32	1.38	1.60	2.94	3.91	4.38
Own-Cost Pass-Through							
Logit	0.80	0.63	0.67	0.73	0.88	0.94	0.97
AIDS	1.19	0.75	0.78	0.90	1.72	2.36	2.82
Linear	0.53	0.51	0.51	0.52	0.55	0.57	0.58
Log-Linear	1.87	1.29	1.34	1.50	2.52	3.39	3.98
Cross-Cost Pass-Through							
Logit	0.04	0.00	0.01	0.02	0.06	0.09	0.11
AIDS	0.22	0.03	0.06	0.12	0.39	0.70	0.98
Linear	0.09	0.01	0.02	0.05	0.12	0.15	0.17
Log-Linear	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Industry Pass-Through							
Logit	0.95	0.84	0.87	0.91	0.97	0.99	0.99
AIDS	1.90	1.10	1.18	1.39	2.92	4.32	5.27
Linear	0.79	0.67	0.70	0.74	0.87	0.94	0.97
Log-linear	1.87	1.29	1.34	1.50	2.52	3.39	3.98
Merger Price Effects							
Logit	0.09	0.01	0.02	0.05	0.16	0.24	0.30
AIDS	0.18	0.01	0.03	0.08	0.46	1.09	1.88
Linear	0.08	0.01	0.02	0.04	0.14	0.21	0.28
Log-Linear	0.30	0.02	0.05	0.12	0.77	2.08	4.11

Notes: Summary statistics are based on 3,000 randomly-drawn sets of data on the pre-merger equilibria. The market share, margin and elasticity are for the first firm. Market share and margin are drawn randomly in the data generating process while the elasticity is the own-price elasticity of demand and equals the inverse margin. Pass-through is calculated, following calibration, based on the curvature properties of the respective demand systems. Own-cost pass-through is the derivative of firm 1's equilibrium price with respect to its own marginal cost. The cross-cost pass-through statistics are based on the derivative of firm 1's equilibrium price with respect to firm 2's marginal cost. The merger price effects are the change in firm 1's equilibrium price.

theoretical results of ? (?).

4.2 Merger price effects

Figure 1 shows how functional form assumptions affect the predictions of simulation. The scatter plots characterize the accuracy of merger simulations when the underlying demand system is logit (column 1), almost ideal (column 2), linear (column 3), and log-linear (column 4). Merger simulations are conducted assuming demand is logit (row 1), almost ideal (row 2), linear (row 3), and log-linear (row 4). Each dot represents the predicted and true changes in firm 1's price for a given draw of data; its vertical position is the prediction of simulation and its horizontal position is the true price effect. Dots that fall along the 45-degree line represent exact predictions while dots that fall above (below) the line represent over (under) predictions. Prediction error is zero when the functional form used in simulation matches that of the underlying demand system.[17]

Prediction error is substantial and systematic when using misspecified simulations. Logit and linear simulations under-predict the price effects of mergers when the underlying demand system is almost ideal or log-linear. AIDS simulation over-predicts prices increases when the underlying demand system is logit or linear but under-predicts when it is log-linear. Log-linear simulation over-predicts price increases in all cases. This sensitivity of prediction to functional form assumptions is well known (e.g., ? (?)) and, in antitrust settings, it is standard practice to generate predictions under multiple different assumptions as a way to evaluate the scope for price changes. We explore next the extent to which cost pass-through can be used to improve the precision of merger predictions.

5 Results

5.1 Perfect information on pass-through

Figure 2 provides scatter plots of the prediction error that arises with FOA and informed simulation, for the cases in which pass-through is observed precisely. As shown, FOA yields accurate predictions when the underlying demand system is logit or almost ideal, as demon-

[17] Two clarifications may assist in the interpretation of Figure 1. First, the post-merger prices are censored at 1.25 and, in some instances, the simulated price increases are well above this level. This may lead the figure to optically understate the degree of prediction error. Second, the figure is symmetric by construction. For example, the scatterplot for logit merger simulation when underlying demand is AIDS is the inverse of the scatterplot for AIDS merger simulation when underlying demand is logit.

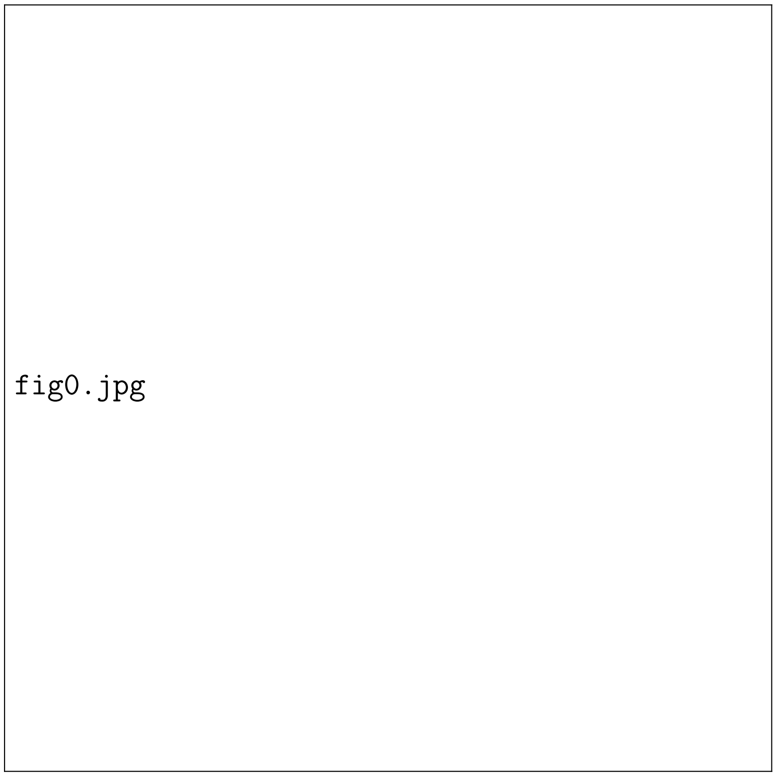

Figure 1: Prediction Error from Standard Merger Simulations

Notes: The scatter plots characterize the accuracy of merger simulutions when the underlying demand system is logit (column 1), almost ideal (column 2), linear (column 3) and log-linear (column 4). Merger simulations are conducted assuming demand is logit (row 1), almost ideal (row 2), linear (row 3) and log-linear (row 4). Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data.

strated by the clustering of dots around the 45° line. It is exact with linear demand, as it is in any setting that produces a quadratic profit function. Prediction error is somewhat larger with the log-linear demand system. Informed simulation provides noisier estimates than FOA, but the biases are reduced when compared with standard merger simulations. These results are consistent with expectations: FOA provides high quality predictions when pass-through is perfectly observed, while (by design) informed simulation only partially mitigates functional form misspecification.[18]

Table 2 provides the median absolute prediction errors (MAPEs) generated by the different methodologies. As shown, the MAPEs of FOA tends to be an order of magnitude smaller than those of standard simulations, provided there is some functional misspecifica-

[18] If, in our setting, one allowed the true underlying demand system to be identified from pass-through then informed simulation would predict merger effects with zero prediction error. As discussed above, we consider this possibility to be unlikely outside our numerical experiments, as there is no reason that consumer decisions should be expected to conform to any of the standard (tractable) models.

basescatter.jpg

Figure 2: Prediction Error from FOA and Informed Simulation

Notes: The scatter plots characterize the accuracy of FOA and Informed Simulation when the underlying demand system is logit, almost ideal, linear and log-linear. Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data.

Table 2: Median Absolute Prediction Error

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
FOA	0.002	0.018	0.000	0.101
Informed Simulation	0.020	0.078	0.019	0.133
Logit Simulation	0.000	0.088	0.016	0.207
AIDS Simulation	0.090	0.000	0.103	0.122
Linear Simulation	0.016	0.102	0.000	0.220
Log-Linear Simulation	0.215	0.122	0.228	0.000

Notes: The table provides the median absolute prediction error of FOA, informed simulation, and standard simulations. Pass-through is assumed to be observed perfectly.

tion in the simulation. The improvements in accuracy with informed simulation are more modest – the MAPEs are close to what arises under the least consequential functional form misspecification. The similarity exists because informed simulation largely serves to select these among the misspecified demand systems in a way that minimizes prediction error.

Table 3: Frequency that Pass-Through Improves Accuracy

Panel A: First Order Approximation

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
Logit Simulation	·	90.2%	100%	95.3%
AIDS Simulation	99.4%	·	100%	53.1%
Linear Simulation	89.2%	92.0%	·	93.9%
Log-Linear Simulation	100%	97.7%	100%	·

Panel B: Informed Simulation

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
Logit Simulation	·	88.1%	86.3%	97.8%
AIDS Simulation	94.4%	·	98.4%	81.3%
Linear Simulation	80.9%	70.4%	·	90.8%
Log-Linear Simulation	100%	92.4%	100%	·

Notes: Panel A shows the fraction of data draws for which FOA has a smaller absolute prediction error than standard merger simulations in predicting firm 1's price change. Panel B shows the same statistic for informed simulation.

Panel A of Table 3 provides the frequencies with which FOA has smaller absolute prediction error (APE) than standard merger simulations. As shown, FOA is more accurate than standard AIDS, linear and log-linear simulations for 99%, 89% and 100% of the mergers, respectively, when the underlying demand system is logit. Similarly high frequencies arise with the other demand systems. Aggregating across the four demand systems, FOA is more accurate than standard simulations for 93% of the mergers considered, provided there is some functional form misspecification. Panel B shows that informed simulation also yields more accurate predictions for the bulk of mergers when compared to standard simulation. Ties occur here, by design, because informed simulation is always identical to one of the misspecified simulations. We report the fraction of mergers for which informed simulation has an APE that is at least as small as standard merger simulations. Aggregating across the systems, informed simulation is at least as accurate as the standard merger simulations in 90% of the mergers, provided some misspecification exists, and more accurate in 57% of the mergers.

One measure of whether these improvements in accuracy are economically meaningful is whether they would improve enforcement decisions made on the basis of the predicted price effects. To explore this, we examine the propensity of the prediction methodologies to produce "false positives" and "false negatives." We define false positives as price increase predictions that exceed ten percent when the true effect that is less than ten percent. We define false negatives analogously.[19] The results are summarized in Table 4. FOA generates both few false positive and few false negatives, while standard merger simulations yield either many false positives or many false negatives, provided there is some misspecification in functional form. Thus, for example, it is possible to generate conservative predictions of merger price effects with linear and logit simulations, but if such simulations receive weight in enforcement decision-making then a nontrivial number of anticompetitive mergers would proceed. Informed simulation also tends to improve the balance of false positives and negatives, albeit to a lesser extent than FOA.

One might suspect – based on economic theory and the scatterplots above – that prediction error decreases for smaller mergers, for each of the methodologies examined. This is indeed the case. To illustrate, we regress absolute prediction error on the price change using nonparametric techniques, and examine the obtained fitted values.[20] In all cases, precision improves as merger price effects become small. Figure 3 shows the results when the true underlying demand system is almost ideal, focusing on mergers that increase price by less than 10 percent. Interestingly, the relative accuracy of FOA is maintained throughout this range. The result is similar under the other demand systems, as well. Thus, the data support that FOA improves the accuracy of prediction, relative to simulation, even for smaller price effects.

5.2 Measurement error and bias in pass-through

Table 5 shows the MAPEs that arise with FOA and informed simulation when pass-through is observed with measurement error or bias. The accuracy of FOA deteriorates with the magnitude of measurement error regardless of the underlying demand system. This is consistent with expectation, as pass-through determines the extent to which the opportunity costs created by the merger affect prices, and therefore measurement error affects the accuracy with which merger pass-through can be recovered. However, even when the pass-through data are quite noisy, prediction error usually is smaller than what arises under standard merger

[19]We select ten percent threshold solely based on the empirical distribution of true prices changes: in each demand system, many mergers produce true price effects both above and below this threshold. We have examined other thresholds, and the qualitative results are unaffected.

[20]We run kernel-weighted local polynomial regressions using the standard Epanechnikov kernel.

Table 4: Type I and II Prediction Error

Panel A: Frequency of False Positives (Type I)

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
FOA	1.6%	1.5%	0.0%	0.9%
Informed Simulation	6.7%	2.4%	11.3%	0.5%
Logit Simulation	·	0.3%	9.6%	0.0%
AIDS Simulation	23.2%	·	30.4%	0.0%
Linear Simulation	2.2%	0.0%	·	0.0%
Log-Linear Simulation	34.4%	12.6%	41.8%	·

Panel B: Frequency of False Negatives (Type II)

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
FOA	0.0%	1.0%	0.0%	14.5%
Informed Simulation	8.2%	19.0%	1.5%	16.9%
Logit Simulation	·	25.0%	2.2%	39.1%
AIDS Simulation	0.2%	·	0.0%	13.3%
Linear Simulation	9.6%	32.4%	·	46.8%
Log-Linear Simulation	0.0%	0.0%	0.0%	·

Notes: Panel A shows the fraction of data draws for which the true price change in firm 1's price is less than 10 percent but the prediction exceeds 10 percent. Panel B the fraction of data draws for which the true price change exceeds 10 percent but the prediction is less than 10 percent. FOA is calculated using the pass-through that arises in the pre-merger equilibrium.

SmallPriceDensities.pdf

Figure 3: Absolute Prediction Error for Small Price Effects
Notes: The panels plot fitted value for absolute prediction error (APE) obtained with nonparametric regressions of APE on the true price effect. The underlying demand system is almost ideal.

simulation provided that some misspecification exists (see Table 2). The small amount of bias introduced also increases MAPE in most cases. The effect of bias is more easily seen graphically, and we return to this shortly.

The accuracy of informed simulation also deteriorates with measurement error in pass-through, but less quickly relative to FOA. This is because measurement error only rarely causes a change in the misspecified model selected for use in the informed simulation. For example, when underlying demand is logit, the linear demand system is selected for use in simulation for 74.4% of mergers if there is no measurement error, and for 73.7% of mergers if pass-through is observed within 90% of its true value. Due to the robustness of this selection routine, the accuracy of informed simulation exceeds that of FOA when measurement error in pass-through is large.[21] Bias at the level considered also does not affect the selection routine substantially, and so the MAPE of informed simulation is mostly unaffected.

Figures 4 provides scatter plots of the prediction error with FOA.[22] The presence of measurement error in pass-through leads to greater spread of FOA predictions, and the spread increases in the magnitude of the measurement error. Predictions remain centered

[21] Note that our results are in part dependent on the menu of demand systems we have chosen for our exercise.

[22] The figure shows the cases in which cost pass-through is observed within 30% and 90% of its true value. We omit the case of 60% due to space considerations.

Table 5: MAPE with Imperfect Pass-Through Data

Panel A: First Order Approximation

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
30% Measurement Error	0.013	0.032	0.009	0.102
60% Measurement Error	0.023	0.071	0.019	0.120
90% Measurement Error	0.038	0.118	0.034	0.159
15% Downward Bias	0.015	0.021	0.014	0.142
15% Upward Bias	0.019	0.067	0.015	0.069

Panel B: Informed Simulation

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
30% Measurement Error	0.020	0.078	0.019	0.131
60% Measurement Error	0.021	0.081	0.019	0.132
90% Measurement Error	0.022	0.088	0.019	0.138
15% Downward Bias	0.016	0.084	0.018	0.134
15% Upward Bias	0.026	0.073	0.020	0.129

Notes: The table shows the median absolute prediction error that arises with (i) FOA supported by cost pass-through observed within 30%, 60%, and 90% of its true value; (ii) FOA supported by cost pass-through with 15% downward and upward biases; and (iii) informed simulation, as defined by the most accurate misspecified simulation model.

around zero, however, so measurement error does not lead to systematic over-prediction or under-prediction. The predicted price effects of FOA are muted when cost pass-through is observed with downward bias, and amplified when pass-through is observed with upward bias. Again this is consistent with expectations, based on theory. Because FOA tends to under-predict price effect with log-linear demand, it also explains why MAPE could be diminished by an upward bias in pass-through for that case.

Figures 5 provides the same scatter plots for informed simulation. Predictions are not centered around the true price effects. Nonetheless, the extent is visibly reduced relative to standard merger simulations (see Figure 1), and the magnitude of measurement error does not lead to a greater spread of predictions. The presence of bias at the level examined does not affect much the predictions with informed simulation, again because the selection routine that selects the demand model proves to be robust.

6 Concluding Remarks

have to conclude with shortcomings of FOA – when it does poorly.

[23]

Our results demonstrate that first order approximation can be a powerful tool with which to make counterfactual predictions while remaining agnostic about the functional forms of the underlying economic model. While grounded in the oligopoly theory of industrial organization, its usefulness extends into other fields of economics including macroeconomics and international trade. A full application of the methodology requires the researcher to bring more information to bear – namely information on pass-through or demand curvature in the observed equilibria – than do most simulation-based prediction methodologies. This points to a potentially valuable research agenda that we sketch here.

First, the prospect of first order approximation accentuates the value of empirical research that examines the pass-through behavior of firms. Recent progress has been made on that front. For instance, in addition to the research cited herein, **?** and **?** explore pass-through in the gasoline retail markets and wholesale electricity markets, respectively. Related is the econometric question taken up by **?** about whether, and under which conditions, reduced-form regressions obtain consistent estimates of pass-through when pass-through is non-constant or when costs are partially observed.

[23]Pass-through also can be used to estimate the curvature properties of flexible demand systems (e.g., **?** (**?**); **?** (**?**)). The consideration of this methodology would complicate the data generating process substantially, and we leave it for future research.

fig4a.jpg

Figure 4: Prediction Error from FOA with Imperfect Pass-Through Data
Notes: The scatter plots characterize the accuracy of FOA when the underlying demand system is logit, almost ideal, linear and log-linear. Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data. FOA is calculated based on pass-through observed within 30% and 90% of its true value (rows 1 and 2), and observed with 15% downward and upward biases (rows 3 and 4).

Second, placing weight on estimated pass-through begs the question of whether the derived theoretical relationship between local demand curvature and cost pass-through extends to real-world settings, or whether menu costs and rule-of-thumb pricing obstruct the connection. Such issues may create a relevant distinction between long run and short run pass-through rates. This distinction is emphasized in the sizable literature on asymmetric fig4b.jpg

Figure 5: Prediction Error from Informed Simulation with Imperfect Pass-Through Data
Notes: The scatter plots characterize the accuracy of informed simulation when the underlying demand system is logit, almost ideal, linear and log-linear. Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data. Informed simulation is calculated based on pass-through observed within 30% and 90% of its true value (rows 1 and 2), and observed with 15% downward and upward biases (rows 3 and 4).

pass-through (e.g., **?**, **?**) and increasingly is modeled explicitly (e.g., **?**, **?**, **?**, **?**), but more work on this subject would be valuable.

Lastly, first order approximation is not the only way to make reasonable counterfactual predictions without imposing functional form assumptions on an economic model. Recent contributions demonstrate the random coefficients logit model is non-parametrically identified (**?**, **?**) and therefore capable of providing estimates of demand curvature that are independent of the elasticity estimates. Research that ascertains the empirical variation required in practice to identify the second order properties of the model would have value, as would research that examines the accuracy of simulations based on flexibly-estimated random coefficients logit model when the true underlying model is not logit.

Appendix for Online Publication

A Mathematical Details of the Calibration Process

We provide mathematical details on the calibration process in this appendix. To distinguish the notation from that of Section 2, we move to lower cases and let, for example, s_i and p_i be the market share and price of firm i's product, respectively.[24] Recall that in the data generating process we randomly assign market shares among the four single-product firms and the outside good, draw the price-cost margin of the first firm's product from a uniform distribution with support over (0.2, 0.8), and normalize all prices to unity. The calibration process then obtains parameters for the logit, almost ideal, linear and log-linear demand systems that reproduce these draws of data.

Calibration starts with multinomial logit demand, the basic workhorse model of the discrete choice literature. The system is defined by the share equation

$$s_i = \frac{e^{(\delta_i - \alpha p_i)}}{1 + \sum_{j=1}^{N} e^{(\delta_j - \alpha p_j)}} \tag{A.1}$$

The parameters to be calibrated include the price coefficient α and the product-specific quality terms δ_i. We recover the price coefficient by combining the data with the first order conditions of the first firm. Under the assumption of Nash-Bertrand competition this yields:

$$\alpha = \frac{1}{m_1 p_1 (1 - s_1)} \tag{A.2}$$

where m_1 is the price-cost margin of firm 1. We then identify the quality terms that reproduce the market shares:

$$\delta_i = \log(s_i) - \log(s_0) + \alpha p_i \tag{A.3}$$

for $i = 1 \ldots N$. We follow convention with the normalization $\delta_0 = 0$. Occasionally, a set of randomly-drawn data cannot be rationalized with logit demand and we replace it with a set that can be rationalized. This tends to occurs when the first firm has both an unusually small market share and an unusually high price-cost margin.

The logit demand system often is criticized for its inflexible demand elasticities. Here, the restrictions on substitution are advantageous and allow us to obtain a full matrix of elasticities with a tractable amount of randomly drawn data. The derivatives of demand

[24] We define market share $s_i = q_i / \sum_{j=1}^{N} q_j$, where q_i represents unit sales.

with respect to prices, as is well known, take the form

$$\frac{\partial q_i}{\partial p_j} = \begin{cases} \alpha s_i(1 - s_i) & \text{if } i = j \\ -\alpha s_i s_j & \text{if } i \neq j \end{cases} \tag{A.4}$$

We use the logit derivatives to calibrate the more flexible almost ideal, linear and log-linear demand systems. This ensures that each demand system has the same first order properties in the pre-merger equilibrium, for a given draw of data.

The AIDS is written in terms of expenditure shares instead of quantity shares (**?**). The expenditure share of product i takes the form

$$w_i = \alpha_i + \sum_{j=0}^{N} \gamma_{ij} \log p_j + \beta_i \log(x/P) \tag{A.5}$$

where x is total expenditure and P is a price index. We incorporate the outside good as product $i = 0$ and normalize its price to one; this reduces to N^2 the number of price coefficients in the system that must be identified (i.e., γ_{ij} for $i, j \neq 0$). We further set $\beta_i = 0$ for all i, a restriction that imposes in income elasticity of unity. Under this restriction, total expenditures are given by

$$\log(x) = (\widetilde{\alpha} + u\widetilde{\beta}) + \sum_{k=1}^{N} \alpha_k \log(p_k) + \frac{1}{2} \sum_{k=1}^{N} \sum_{j=1}^{N} \gamma_{kj} \log(p_k) \log(p_j) \tag{A.6}$$

for some utility u. We identify the sum $\widetilde{\alpha} + u\widetilde{\beta}$ rather than $\widetilde{\alpha}$, u and $\widetilde{\beta}$ individually.[25]

Given this structure, product i's unit sales are given by $q_i = x w_i / p_i$ and the first derivatives of demand take the form

$$\frac{\partial q_i}{\partial p_j} = \begin{cases} \frac{x}{p_i^2}(\gamma_{ii} - w_i + w_i^2) & \text{if } i = j \\ \frac{x}{p_i p_j}(\gamma_{ij} + w_i w_j) & \text{if } i \neq j \end{cases} \tag{A.7}$$

The calibration process for the AIDS then takes the following four steps:

1. Calculate x and w_i from the randomly drawn data on market shares, using a market size of one to translate market shares into quantities.

2. Recover the price coefficients γ_{ij} for $i, j \neq 0$ that equate the AIDS derivatives given in

[25] The price index P is defined implicitly by equation (A.6) as the combination of prices that obtains utility u given expenditure x. A formulation is provided in **?**.

equation (A.7) and the logit derivatives given in equation (A.4). Symmetry is satisfied because consumer substitution is proportional to share in the logit model. The outside good price coefficients, γ_{i0} and γ_{0i} for all i, are not identified and do not affect outcomes under the normalization the $p_0 = 1$. Nonetheless, they can be conceptualized as taking values such that the adding up restrictions $\sum_{i=0}^{N} \gamma_{ij} = 0$ hold for all j.

3. Recover the expenditure share intercepts α_i from equation (A.5), leveraging the normalization that $\beta_i = 0$. The outside good intercept α_0 is not identified and does not affect outcomes, but can be conceptualized as taking a value such that the adding up restriction $\sum_{i=0}^{N} \alpha_i = 1$ holds.

4. Recover the composite term $(\widetilde{\alpha} + u\widetilde{\beta})$ from equation (A.6).

This process creates an AIDS that, for any given set of data, has quantities and elasticities that are identical in the pre-merger equilibrium to those that arise under logit demand. The system possess all the desirable properties defined in **?**. Our approach to calibration differs from **?**, which does not model the price index as a function of the parameters, and from **?**, which assumes total expenditures are fixed.

We turn now to the linear and log-linear demand systems. The first of these takes the form

$$q_i = \alpha_i + \sum_j \beta_{ij} p_j, \tag{A.8}$$

The parameters to be calibrated include the firm specific intercepts α_i and the price coefficients β_{ij}. We recover the price coefficients directly from the logit derivatives in equation (A.4). We then recover the intercepts to equate the implied quantities in equation (A.8) with the randomly drawn market shares, again using a market size of one. Of similar form is the log-linear demand system:

$$\log(q_i) = \gamma_i + \sum_j \epsilon_{ij} \log p_j \tag{A.9}$$

where the parameters to be calibrated are the intercepts γ_i and the price coefficients ϵ_{ij}. Again we recover the price coefficients from the logit derivatives (converting first the derivatives into elasticities). We then recover the intercepts to equate the implied quantities with the market share data. This process creates linear and log-linear demand systems that, for any given set of data, has quantities and elasticities that are identical to those of the calibrated logit and almost ideal demand systems, in the pre-merger equilibrium.

B Derivation of Merger Pass-Through

In this appendix, we provide an expression for the Jacobian of $h(P)$, which can be used to construct the merger pass-through matrix as used in Theorem 1. Using the definition $h(P) \equiv f(P) + g(P)$, we have

$$\frac{\partial h(P)}{\partial P} = \frac{\partial f(P)}{\partial P} + \frac{\partial g(P)}{\partial P} \tag{B.1}$$

The Jacobian of $f(P)$ can be written as:

$$\frac{\partial f(P)}{\partial P} = \begin{bmatrix} \frac{\partial f_1(P)}{\partial p_1} & \cdots & \frac{\partial f_1(P)}{\partial p_N} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_J(P)}{\partial p_1} & \cdots & \frac{\partial f_J(P)}{\partial p_N} \end{bmatrix} \tag{B.2}$$

where N is the total number of products and J is the number of firms. The vector P includes all prices; we use lower case to refer to the prices of individual products, so that p_n represents the price of product n. In the case that product n is sold by firm i,

$$\frac{\partial f_i(P)}{\partial p_n} = -\begin{bmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{bmatrix} + \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} \left[\frac{\partial^2 Q_i}{\partial P_i \partial p_n}^T\right] \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} Q_i - \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} \left[\frac{\partial Q_i}{\partial p_n}\right] \tag{B.3}$$

where Q_i and P_i are vectors representing the quantities and prices respectively of the products owned by firm i, and the initial vector of constants has a 1 in the firm-specific index of the product n. For example, if product 5 is the third product of firm 2, then the 1 will be in the 3^{rd} index position when calculating $\partial f_2(P)/\partial p_5$. If product n is not sold by firm i, the vector of constants is $\vec{0}$, and thus

$$\frac{\partial f_i(P)}{\partial p_n} = \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} \left[\frac{\partial^2 Q_i}{\partial P_i \partial p_n}^T\right] \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} Q_i - \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1} \left[\frac{\partial Q_i}{\partial p_n}\right] \tag{B.4}$$

The matrix $\partial g(P)/\partial P$ can be decomposed in a similar manner:

$$\frac{\partial g(P)}{\partial P} = \begin{bmatrix} \frac{\partial g_1(P)}{\partial p_1} & \cdots & \frac{\partial g_1(P)}{\partial p_N} \\ \vdots & \ddots & \vdots \\ \frac{\partial g_K(P)}{\partial p_1} & \cdots & \frac{\partial g_K(P)}{\partial p_N} \\ & & \\ 0 & \cdots & 0 \\ \downarrow & & \downarrow \end{bmatrix} \tag{B.5}$$

where N is the number of products and K is the number of merging firms. Notice that $\partial g(P)/\partial P$ includes zeros for non-merging firms, because the merger does not create opportunity costs for these firms. In the case that product n is sold by a firm merging with firm i (this does not include firm i itself),

$$\begin{aligned} \frac{\partial g_i(P)}{\partial p_n} &= -\left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial Q_j}{\partial P_i}^T\right]\begin{bmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{bmatrix} \\ &+ \left(\left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial^2 Q_i}{\partial P_i \partial p_n}^T\right]\left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial Q_j}{\partial P_i}^T\right] - \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial^2 Q_j}{\partial P_i \partial p_n}^T\right]\right)(P_j - C_j) \end{aligned} \tag{B.6}$$

where Q_j, P_j, and C_j are vectors of the quantities, prices, and marginal costs respectively of products sold by firms merging with firm i, and the vector of 1s and 0s has a 1 in the merging firm's firm-specific index of the product n. For example, if product 5 is the third product of firm 2, and firm 2 is merging with firm 1, then the 1 will be in the 3^{rd} index position when calculating $\partial g_1(P)/\partial p_5$. It is an important distinction that – supposing there are more than two merging parties – the index j refers to all of the merging parties' products, excluding firm i's products. If product n is not sold by any firm merging with firm i (including a product sold by firm i),

$$\frac{\partial g_i(P)}{\partial p_n} = \left(\left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial^2 Q_i}{\partial P_i \partial p_n}^T\right]\left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial Q_j}{\partial P_i}^T\right] - \left[\frac{\partial Q_i}{\partial P_i}^T\right]^{-1}\left[\frac{\partial^2 Q_j}{\partial P_i \partial p_n}^T\right]\right)(P_j - C_j) \tag{B.7}$$

C Additional Figures

FullbsHoriPanel.jpg

Figure C.1: FOA under the Modified Horizontality Condition

Notes: The panels plot the predictions of FOA obtained with full knowledge of second derivatives (on the vertical axis) versus those supported by the modified horizontality condition (on the horizontal axis).