Poor Substitutes? Counterfactual methods in IO and Trade compared*

Keith Head† Thierry Mayer‡

December 2, 2019

Abstract

Monopolistic competition with constant-elasticity-of-substitution preferences (CES-MC), is a widely-used tool in many fields of economics. A major reason for this success in International Trade and Macro is that it yields predictions that are simple to interpret, easy to implement, and particularly well-suited for counterfactual analysis of policy changes with minimal data requirements. Industrial Organization economists however view CES-MC as irredeemably unrealistic both on the demand and supply sides. They overwhelmingly prefer random coefficients demand structures that feature *rich substitution* patterns, combined with multiproduct oligopoly. The lack of account for oligopolistic markups and cannibalization effects in CES-MC can be corrected while maintaining most of the tractability advantages for estimation and counterfactual analysis, a framework we label CES-Oly. We present simulation-based evidence that even if the IO models are the true data generating process, CES-Oly and even CES-MC can be used with remarkable accuracy as an approximation to conduct counterfactuals for aggregations over multiple varieties, even when allowing for realistic amounts of consumer preference heterogeneity and market power.

*This research has received funding from the European Research Council (ERC) under the Grant Agreement No. 313522. We thank participants to the SMU-NUS-Paris Joint Trade Workshop (Merlion 2018 Workshop). Very helpful comments were received during seminar presentations at Princeton, Columbia and Harvard during the fall of 2018.

†Sauder School of Business, University of British Columbia, CEPR, and CEP, keith.head@sauder.ubc.ca

‡Sciences Po, Banque de France, CEPII, and CEPR, thierry.mayer@sciencespo.fr

1 Introduction

> "[W]hile the [CES] functional form is convenient it imposes a very strong restriction on the demand system. The simplicity of the model and its analytic tractability make it a popular choice in theory and it is also heavily used in trade and in macro, but *it is not appropriate to explain micro data* and is essentially never used in empirical IO." Nevo (2011), italics added

Monopolistic competition with constant-elasticity-of-substitution preferences (CES-MC), also referred to as Dixit-Stiglitz, is a widely-used tool in many fields of economics. As the quote above illustrates, industrial organization economists view it as irredeemably unrealistic. We will not argue with the premise that CES is an unreasonable way to represent variety-level data. Instead, we present simulation evidence that CES-MC can be used with remarkable accuracy as an approximation to conduct counterfactuals for aggregations over multiple varieties.

Anti-CES, anti-MC feeling runs particularly strong with respect to the car industry, an emblematic case studied by the very first papers of modern empirical IO (Berry et al. (1995), Goldberg (1995), Verboven (1996), Goldberg and Verboven (2001), Petrin (2002) among others). While our paper is primarily based on simulations we will often calibrate the models using parameters from papers focused on the car industry (Berry et al. (1999), Goldberg and Verboven (2001) and Coşar et al. (2018) in particular) and refer to varieties as models produced by firms/brands.

There are three main features that IO economists believe to be essential for realism but which are absent in CES-MC.

1. Rich substitution patterns.
2. Oligopolistic price setting with variable markups.
3. Internalization of cannibalization effects by multiproduct firms.

The first is a consequence of the restrictive substitution elasticities inherent in the CES functional form. The second issue arises because CES combined with mass-less firms yields a constant markup. The third is a consequence of assuming that individual varieties are too small for their prices to affect the price index, which is a sufficient statistic in both linear and CES monopolistic competition models.

A related critique of the Dixit-Stiglitz framework is that, by forcing log-linearity of the demand function, it imposes a unitary pass-through of costs into prices. Nakamura and Zerom (2010) argue that "a key advantage of [BLP] over ... Dixit-Stiglitz demand models:

implications for pass-through, as determined by the curvature of demand, depend on estimated parameters in the demand system as opposed to purely on functional form assumptions." A very active literature has recently shown that the curvature of demand and associated pass-through rate of a model has critical consequences on whether the model has pro-competitive or anti-competitive properties, with CES-MC being the limit case where equilibrium markups and prices do not depend on the size of the market and the mass of competitors (Zhelobodko et al. (2012), Fabinger and Weyl (2016) and Mrázová and Neary (2017) in particular).[1]

Notwithstanding these valid critiques, CES-MC has advantages that may have been under-emphasized and the models that can accommodate the three features above present serious challenges in computation (Knittel and Metaxoglou (2014)), identification (Gandhi and Houde (2016)), data requirements, and transparency of estimation. Moreover, functional forms may impose more on results than is obvious, e.g. pass-through predictions that seem contradictory with what is commonly estimated in the literature. In particular, increasing the degree of richness in the substitution patterns (through consumer heterogeneity), tends to raise the pass-through to make it complete or even more than complete, while a large macro literature has found empirical evidence largely in favor of incomplete pass-through (Burstein and Gopinath (2014)). A related effect of "extreme" rich substitution pattern is that, associated with more than complete pass-through, markups are higher for small and unproductive firms, a prediction at odds with recent empirical finding by De Loecker et al. (2016) or Atkin et al. (2015).[2]

In our paper, we investigate the implications of estimating CES-based equations when the true data generating process (DGP) is the random coefficients discrete choice model with oligopolistic multiproduct price setting. We think of these two features as the defining attributes of a class of models that we abbreviate as "BLP." Within the set of BLP models in the literature there are many variations, but one that is important to highlight here is the distinction of whether demand is for a single unit of consumption or a continuous quantity. By far the most common demand formulation is the "mixed logit," which stems back to the original work by Berry et al. (1995) . The problem with this form for our purposes is that it does not nest CES as a special case. Thus, no matter what set of parameters one uses in the DGP—even as the variance of the random coefficients goes to zero and the number of independent varieties becomes very large—the demand function is still wrong. We therefore focus our attention on a demand formulation inspired

[1] Our contribution to this literature (for now relegated in the appendix) is to characterize precisely how the Mixed Logit and Mixed CES models used in IO behave in terms of pass-through depending on the degree of heterogeneity exhibited by consumers.

[2] Both papers are estimating markups using methods that do not rely on demand estimation.

by Anderson et al. (1992), and used recently in empirical applications in IO (Björnerstedt and Verboven (2016)) and Trade (Adao et al. (2017), and Redding and Weinstein (2016)). The "Mixed CES" formulation introduces heterogeneity in consumer sensitivity to price in an otherwise CES demand system. A notabale advantage of MCES for our purposes is that it can be used very easily to assess the ability of the CES-MC approximation to forecast welfare outcomes from policy interventions. MCES is also attractive because several well-known models can be thought of as special cases. As the variance of price elasticity across households goes to zero, MCES can reach three different limiting cases. First, with many single-variety firms it becomes the Dixit-Stiglitz model which we have also referred to as CES-MC. Second, with a small number of single-variety firms, the limiting case is a version of Atkeson and Burstein (2008) with the upper level CES set to one. Finally with several large multiproduct firms, MCES converges on models used by Hottman et al. (2016) and Bernard et al. (2016). We therefore concentrate on MCES in the main text, while presenting results for the more traditional Mixed Logit case in the appendix.

Our counterfactuals simulate a trade liberalization and generate as key outcomes the decline in domestic firms's share of the domestic market, as well as the changes in each component of welfare. For a wide value of parameter settings, we find that counterfactuals based on CES estimates yield results that give good guidance on the "true" quantitative response of the aggregate market share of domestic firms. For some plausible settings, the declines in the domestic share associated with tariff reductions are can be remarkably accurate. This is true both with the Mixed CES and Mixed Logit DGPs.

The paper proceeds as follows. In Section 2, we start by describing the Mixed CES data generating process. This will be regarded as the "truth." The CES approximation we use is laid out in section 3 with the generalization to oligopoly in section 4. These approximations comprise a method for estimating the price elasticity and a "hat algebra" method for predicting counterfactual market shares, prices, and welfare components. We report the results of the counterfactual analysis of a tariff cut in section 5. The paper ends with some still preliminary conclusions and an outline of some potential extensions.

2 The Mixed CES (MCES) model

The Mixed CES model assumes individual consumers have CES utility but that their price elasticity is heterogeneous. It is micro-founded by starting with the variable consumption discrete-choice model of Anderson et al. (1992) (section 3.7) extending it to include heterogeneity in the price responsiveness parameter. In the spirit of BLP models more generally, the MCES also allows for random coefficients on the consumers' indirect utility derived from product attributes. The "mixed CES" terminology comes from Adao et al. (2017) and Redding and Weinstein (2016). Björnerstedt and Verboven (2016) refer to the model using the descriptive, but unwieldy "random coefficients specification of the constant expenditure logit." They conclude that "the constant expenditures specification entails a more plausible pattern of price elasticities across products."

On the supply side, the MCES DGP includes potentially large multi-product firms as is standard in the empirical IO literature. We start by presenting the analysis in a single market context, before extending it to the multi-market case, which will be particularly useful for estimation.

2.1 Mixed CES random-coefficients multiproduct oligopoly

Denoting y_h the income spent on the differentiated goods (the upper level utility is Cobb-Douglas), the (indirect) utility of household h is given by

$$U_{mh} = \ln y_h - \tilde{\alpha}_h \ln p_m + \tilde{\beta}_h x_m + \tilde{\xi}_m + \varepsilon_{mh}, \tag{1}$$

with x_m being a characteristic of the model and p_m being its price. We think of x_m as an observed component of quality, whereas the unobserved component is captured in ξ_m. In the CES approximation regression, ξ_m will be the major (sometimes the only) component of the error term. Assuming that the individual random term of households for specific models, ε_{mh}, is distributed Gumbel with scale parameter $1/\eta$, the choice probability of household h for model m takes the usual logit form:

$$\mathbb{P}_{mh} = \frac{\exp(\beta_h x_m - \alpha_h \ln p_m + \xi_m)}{\sum_i \exp(\beta_h x_i - \alpha_h \ln p_i + \xi_i)}, \tag{2}$$

with $\alpha_h = \eta\tilde{\alpha}_h$, $\beta_h = \eta\tilde{\beta}_h$, and $\xi_m = \eta\tilde{\xi}_m$. Note that the specification of the random coefficients does not impose a relationship between α_h and β_h but it does imply that all buyers view the unobserved quality ξ_m in the same way. We adopt this approach to parallel the one taken by IO economists in the mixed logit models. An alternative, considered by Redding and Weinstein (2016), places the household heterogeneity in the η parameter. This has the consequence of making consumers who are more price sensitive also more sensitive to differences in quality, both observed and unobserved. This approach is attractive in many respects but we have not pursued it in this version so as to limit the number of permutations to consider.

The richness of substitution patterns in the model will come entirely from the consumer-level heterogeneity captured by the randomness of coefficients α_h and β_h. The specification of those is therefore critical. Beginning with Berry et al. (1995) the IO literature has postulated that the price responsiveness term in the logit equation is inversely proportionate to income. Thus, higher income people react less to price increases. We incorporate this possibility in mixed CES by assuming $\tilde{\alpha}_h = y^{a_2}$. The IO literature does not appear to have considered an alternative, to us equally plausible, assumption that valuation of quality is increasing in income (Petrin (2002) may be an exception). As we will argue, this assumption has the ability to explain an important moment in the data, income sorting across products. We therefore specify β_h and α_h such that

$$\beta_h \sim \mathcal{N}(\bar{\beta} + b_2 \ln y_h, \sigma_\beta) \quad \text{and} \quad \ln \alpha_h \sim \mathcal{N}(\eta + a_2 \ln y_h, \sigma_\alpha).$$

Each individual spends y_h on their preferred variety. Total expenditures on m are therefore $s_m Y$, where $Y \equiv \sum_h y_h$ and s_m is the variety's market share—defined in value. This market share is given by the expenditure-weighed average of the individual probabilities from equation (34):

$$s_m = \frac{\sum_h \mathbb{P}_{mh} y_h}{Y}, \tag{3}$$

On the supply side, the model-specific primitives, x_m and ξ_m, are also assumed to be normally distributed. As with BLP, each firm f offers a portfolio of models, and unobserved quality ξ has both a model-level and a firm-level component (common to all varieties manufactured by firm f). We let the quality of the model influence marginal cost which is given by

$$\ln c_m = \gamma_0 + \gamma_1 x_m + \gamma_2 \xi_m + \nu_m + \ln \tau_m, \tag{4}$$

ν_m being normal as well to yield log-normal marginal costs.[3] Our assumption of constant returns to scale, which is the usual one in both IO and trade literature, allows to consider profit maximization on each market separately. The delivered cost c_m includes a trade cost τ_m which is 1 if the model is domestic, and > 1 for foreign models.

Model-level profits are

$$\pi_m = q_m(p_m - c_m) = \frac{s_m Y}{p_m}(p_m - c_m) = s_m L_m Y,$$

where $L_m \equiv (p_m - c_m)/p_m$ is the Lerner markup. The firm maximizes the sum of the π_m

[3]This distribution is a natural choice since it ensures positive costs, and under CES demand would lead to log-normal sales distributions, a dominant feature of the micro-level data in many countries (see Head et al. (2014) or Fernandes et al. (2015) for instance).

over the set of models it owns. Define the symmetric M-by-M co-ownership matrix as $\mathbf{\Omega}_{jm} = 1$ if models j and m have a common owner and zero otherwise.

The FOC for model m's price in the Bertrand-Nash equilibrium is

$$s_m \frac{\partial L_m}{\partial p_m} + L_m \frac{\partial s_m}{\partial p_m} + \sum_{j \neq m} \Omega_{jm} L_j \frac{\partial s_j}{\partial p_m} = 0.$$

Noting that $\partial L_m / \partial p_m = c_m / p_m^2$, we can write this as

$$-\frac{\partial s_m}{\partial p_m}(p_m - c_m) = \frac{s_m c_m}{p_m} + p_m \sum_{j \neq m} \Omega_{jm} \frac{p_j - c_j}{p_j} \frac{\partial s_j}{\partial p_m}.$$

This FOC can therefore also be written as

$$-\frac{\partial \ln s_m}{\partial \ln p_m}(p_m - c_m) = c_m + \left(\frac{p_m}{s_m} \sum_{j \neq m} \Omega_{jm} L_j \frac{\partial \ln s_j}{\partial \ln p_m} s_j \right).$$

or

$$p_m \left[-\sum_{j \neq m} \Omega_{jm} L_j \frac{\partial \ln s_j}{\partial \ln p_m} \frac{s_j}{s_m} - \frac{\partial \ln s_m}{\partial \ln p_m} \right] = c_m \left(1 - \frac{\partial \ln s_m}{\partial \ln p_m} \right).$$

finally, letting $\epsilon_m \equiv -\frac{\partial \ln s_m}{\partial \ln p_m}$ denote the own price elasticity,

$$p_m = c_m \times \frac{(\epsilon_m + 1)}{\left[\epsilon_m - \frac{1}{s_m} \sum_{j \neq m} \Omega_{jm} \frac{\partial \ln s_j}{\partial \ln p_m} L_j s_j \right]}. \tag{5}$$

This equation for p_m is not, of course, a closed form since ϵ and s both depend on p_m. A closed form is only achieved under the Dixit-Stiglitz assumptions that $\epsilon_m = \sigma - 1$ and $\Omega_{jm} = 0$, which lead to the familiar $p_m = c_m \frac{\sigma}{\sigma - 1}$ (σ denoting the CES common to all consumers in Dixit-Stiglitz). More generally, without cannibalization effects ($\Omega_{jm} = 0$) or consumer heterogeneity ($\alpha_h = \eta$ and $\mathbb{P}_{mh} = \mathbb{P}_m = s_m, \forall h$), the elasticity formulas derived in the appendix gives

$$p_m = c_m \left(\frac{\epsilon_m + 1}{\epsilon_m} \right) = c_m \left(1 + \frac{1}{\eta(1 - s_m)} \right).$$

Under the monopolistic competition case where all market shares converge to 0, we have a constant markup of $\frac{\eta+1}{\eta} = 1.20$, when using our calibrated value of $\eta = 5$.

It will prove useful to re-express equation (5) in vector form. To this end, we first write the summation in the denominator of equation (5) using a matrix multiplication. Define an M by M matrix of common ownership cross-price elasticities on market share as $\mathbf{D}$. Elements of $\mathbf{D}$ are given by $D_{jm} = \Omega_{jm} \frac{\partial \ln s_j}{\partial \ln p_m}$ for $j \neq m$ and zero otherwise. Using $\mathbf{D}$, we can express the impact of a price increase of one model on the profits earned on all the rest of the models in a firm's portfolio compactly as

$$\mathbf{r} = \frac{\mathbf{D}(\mathbf{L}\mathbf{s})}{\mathbf{s}},$$

where $(\mathbf{L}\mathbf{s})$ is an element-by-element vector multiplication, the result of which is matrix multiplied by $\mathbf{D}$ before element-by-element division by the market share vector $\mathbf{s}$. Using $\boldsymbol{\epsilon}$ to denote the vector of absolute own-price elasticities $-\frac{\partial \ln s_m}{\partial \ln p_m}$, the vector form of equation (5) is given by

$$\mathbf{p}^{\star} = \mathbf{c}\frac{\boldsymbol{\epsilon} + 1}{\boldsymbol{\epsilon} - \mathbf{r}}, \tag{6}$$

where $\mathbf{p}^{\star}$ is the vector of optimal response prices. The markups implied here are guaranteed to be greater than one as long as $\boldsymbol{\epsilon} > \mathbf{r} > 0$ for all elements. Iteration of equation (6) is not a contraction mapping but we can still solve for the equilibrium price vector via fixed point iteration on the equation

$$\mathbf{p}^{i+1} = \lambda \mathbf{p}^{\star} + (1 - \lambda)\mathbf{p}^{i},$$

where λ is the weight accorded to the new best response and $1 - \lambda$ is the weight on the previous vector of iterations.

This completes the description of the equilibrium, characterized by two vectors of optimal model-level prices and resulting market shares. We can calculate those under a setting of benchmark trade costs affecting foreign varieties, and one of counterfactual policy changes, which provides the "true" change in the distribution of market shares. Denoting with s_m^0 the initial market share and s_m^1 the one following trade liberalization, our first outcome to monitor is the aggregate change in market shares for domestic firms:

$$\Delta S = \sum_{m \in f_{\text{dom}}} (s_m^1 - s_m^0). \tag{7}$$

2.2 Measurement of welfare changes

The MCES setup allows for easy computation of welfare changes from our counterfactual experiment of a trade liberalization episode. Those can be decomposed in our setup into changes in the consumer surplus, to which are added the declines in domestic firms' profits and tariff revenues.

We use Compensating Variation (CV) as our measure of consumer surplus. The mixed CES utility being based on a discrete choice setup, it can be analyzed using results by Anderson et al. (1992) and extending those to the case of random coefficients. Anderson et al. (1992) (p 45) show the expression for the expectation of maximum utility when the individual random term, ε_{mh}, is distributed Gumbel, and refer to Small and Rosen (1981) for the interpretation in terms of consumer surplus. Denoting the observed part of indirect utility (1) as v_{mh}, we can express consumer surplus as

$$\mathbb{E}\left(\max_{m=1..M} U_{mh}\right) = \frac{1}{\alpha_h}\ln\sum_m \exp(v_{mh}\alpha_h) = \ln y_h + \frac{1}{\alpha_h}\ln\sum_m \exp(\beta_h x_m - \alpha_h \ln p_m + \xi_{mh}). \quad (8)$$

The expenditure function $E(U, \mathbf{p})$ is obtained by inverting consumer surplus, to put y_h on the left hand side. We can write

$$\ln y_h = \mathbb{E}\left(\max_{m=1..M} U_{mh}\right) + \ln P_h \quad \text{with} \quad P_h \equiv \left(\sum_m \exp(\beta_h x_m + \xi_{mh}) p_m^{-\alpha_h}\right)^{-1/\alpha_h}. \quad (9)$$

Hence

$$y_h = \exp\left(\mathbb{E}\left(\max_{m=1..M} U_{mh}\right)\right) P_h \quad (10)$$

The expenditure functions are therefore just $E(U_h, \mathbf{p}) = \exp(U_h)P_h$. The compensating variation CV is the difference in expenditure for initial (0) and counterfactual (1) prices, holding utility constant at the initial level based on the consumer having chosen their preferred variety, U_h^{*0}. It writes as $\text{CV}_h = \exp(U_h^0)(P_h^0 - P_h^1)$. Substituting back $\exp(U_h^0) = E_h^0/P_h^0$, the consumer surplus change measure is $\text{CV}_h = E_h^0 \times \frac{P_h^0 - P_h^1}{P_h^0}$. The final step is to aggregate over the consumers to obtain

$$\text{CV} = \sum_h y_h \times \frac{P_h^0 - P_h^1}{P_h^0} \quad (11)$$

Using notation $L^0 = (p_m^0 - c_m)/p_m^0$ and $L^1 = (p_m^1 - c_m)/p_m^1$ for the Lerner index of domestic firms (note that the marginal cost of those is not affected by the policy experiment), the change in profits for domestic models in the Mixed CES case are $\pi_m^1 - \pi_m^0 = Y(s_m^1 L_m^1 - s_m^0 L_m^0)$. We then sum those profit changes over domestic firms:

$$\Delta\Pi = \sum_{m \in f_{\text{dom}}} \pi_m^1 - \pi_m^0 = Y \sum_{m \in f_{\text{dom}}} (s_m^1 L_m^1 - s_m^0 L_m^0). \quad (12)$$

To calculate changes in tariff revenues, we assume that tariffs are imposed on the marginal cost of the imported models. Our policy experiment is the removal of a 10%

ad valorem tariff, leaving in place a non-tariff barrier. Tariff revenues on imports of each foreign model are $0.10 \times (c_m/\tau_m)q_m$. Bringing tariffs down to 0, the aggregate change in tariff revenues is

$$\Delta\text{TR} = -0.10 \sum_{m \in f_{\text{for}}} \frac{c_m}{\tau_m} \times \frac{s_m Y}{p_m}. \tag{13}$$

The "true" total welfare change is summing over the three components:

$$\Delta\text{Welfare} = \text{CV} + \Delta\Pi + \Delta\text{TR}. \tag{14}$$

3 The CES-MC approximation

Having described all the relevant equilibrium true levels and changes in the MCES DGP, we proceed to their CES approximation. The critical primary step to this approximation is the estimation of a price elasticity denoted $\breve{\eta}$, which the CES approximation assumes to be constant across consumers. The typical dataset available to both IO and trade economists involves a combination of market shares and prices at the variety level. Recognizing the endogeneity of own price when explaining market shares (here made very clear from the fact that unobserved quality ξ_m enters both utility and marginal costs), IO economists use different sets of instruments for a variety's price (see Gandhi and Houde (2016) for a discussion of those). Here we follow the trade approach, which uses price-shifters rather than prices to obtain the price elasticity. Country-pair freight or tariff rates have been used in the literature (Hummels (1999), Romalis (2007), Fitzgerald and Haller (2015), Spearot (2013), Bas et al. (2017)), based on their plausible exogeneity to variety-level outcomes.

At this stage, we therefore need to account for the country of origin of model m, denoted $i(m)$, and introduce a bilateral dimension to trade costs, with ad-valorem equivalent denoted by $\tau_{i(m)n} > 1$, which is assumed to apply on marginal costs, for simplicity. We take logs of the market shares generated by the MCES DGP and use those as the dependent variable in the following regression:

$$\ln s_{mn} = -\eta \ln \tau_{i(m)n} + \text{FE}_m + \text{FE}_n + \upsilon_{mn}, \tag{15}$$

where FE_m are model-level fixed effects, which account for differences in all variables which makes m a successful variety (low c_m and/or high x_m and ξ_m). Destination-level FE_n account for the varying degrees of competition on each market faced by model m. The error term in this regression, υ_{mn} comes from the mis-specification between the DGP and the CES approximation.

This estimation provides an estimate of the trade cost elasticity, denoted $\breve{\eta}$, which is also the price elasticity when CES demand is combined with monopolistic competition. Thus, the CES-MC model predicts prices

$$\breve{p}_{mn} = \frac{\breve{\eta} + 1}{\breve{\eta}} c_m \tau_{i(m)n}.$$

Using this price and regression coefficients yields CES-predicted market shares, to be compared with the BLP-based market shares from equation (3):

$$\breve{s}_{mn} = \frac{\exp(\breve{\text{FE}}_m + \breve{\text{FE}}_n)\tau_{i(m)n}^{-\breve{\eta}}}{\sum_j \exp(\breve{\text{FE}}_j + \breve{\text{FE}}_n)\tau_{i(j)n}^{-\breve{\eta}}}. \tag{16}$$

There are two fundamental differences between "true" s_{mn} and CES-MC predicted $\breve{s}_{mn}$. First, s_{mn} sums over heterogeneous-coefficient consumers which all have different probabilities to buy each model m, as seen in (34). Second, the prices determining market shares are different. The s_{mn} are based on the equilibrium oligopoly prices that take into account cannibalization effects, whereas $\breve{s}_{mn}$ predicts prices by applying a constant markup (based on the–potentially biased–estimate of the trade cost elasticity) to costs. With multiproduct oligopolists, such markups would not be optimal even under CES demand. The question we explore next is whether the $\breve{s}_{mn}$, despite all their differences, might nevertheless succeed in capturing the cross-sectional variation in s_{mn} and more crucially for our purposes, the changes in aggregate domestic market shares ΔS.

The CES approximation lends itself to what Costinot and Rodriguez-Clare (2014) call the "Exact Hat Algebra" (EHA) method to counterfactual computations. We simulate a change in trade costs and compute the predicted change in model-level market shares by the CES-MC approximation, where $\hat{s}_{mn} = s'_{mn}/s_{mn}$, with s'_{mn} being the market share of m in n under the reduced trade costs. Assuming both observable and unobservable components of quality to be left unchanged, the EHA of market shares is simply

$$\hat{s}_{mn} = \frac{\hat{\tau}_{i(m)n}^{-\breve{\eta}}}{\sum_j s_{jn}\hat{\tau}_{i(j)n}^{-\breve{\eta}}}, \tag{17}$$

where $\hat{\tau}_{i(j)n} = \tau'_{i(j)n}/\tau_{i(j)n}$, and ℓ is the country of origin of model j. It is clear from that equation that the predicted percentage change in market share will be common across all models that face the same shock (the ones in each foreign country). Domestic firms will also be affected uniformly (and negatively), by the rise in competition captured in the denominator of (17). There are three sources of errors in the computation of the individual market share changes: i) the omission of any rich substitution patterns that emerge as soon as consumers are heterogeneous, ii) the lack of markup adjustment due to oligopoly, iii) potential bias in $\breve{\eta}$. It will depend on particular settings of the simulation whether those three sources of discrepancies between the true change and EHA approximation reinforce or compensate each other. Aggregated for all domestic firms, we therefore obtain

$$\Delta S^{\text{CES-MC}} = \sum_{m \in f_{\text{dom}}} s_{mn}(\hat{s}_{mn} - 1). \tag{18}$$

Let us now turn to the CES-MC approximation of welfare changes. The formula relevant for the consumer surplus (11) becomes common to all (homogenous) consumers:

$$\text{CV}^{\text{CES-MC}} = Y \times \frac{P - P'}{P} = Y \times \left(1 - \frac{P'}{P}\right) \quad \text{with} \quad P \equiv \left(\sum_j \exp(\breve{\text{FE}}_j + \breve{\text{FE}}_n)\breve{p}_{jn}^{-\breve{\eta}}\right)^{-1/\breve{\eta}}. \tag{19}$$

As with market shares, we can use the EHA method to evaluate the change in the price index brought by trade liberalization. It only requires knowledge of the true market share pre-liberalization, the estimated price elasticity approximation $\breve{\eta}$ together with the trade cost shock:

$$\hat{P} = P'/P = \left(\sum_j s_{jn}\hat{\tau}_{i(j)n}^{-\breve{\eta}}\right)^{-1/\breve{\eta}}. \tag{20}$$

Regarding profits, since in the CES-monopolistic competition case, markups are assumed constant, we have that $p'_{mn} = p_{mn}\hat{\tau}_{i(m)n}$, and $L'_{mn} = L_{mn}$. We can therefore write the change in profits as

$$\Delta\Pi^{\text{CES-MC}} = Y \sum_{m \in f_{\text{dom}}} (\hat{s}_{mn} - 1)s_{mn}L_{mn}. \tag{21}$$

Since tariffs are set to 0 in the policy experiment, the CES-MC change in tariff revenues is identical to the Mixed CES case. Total welfare change is then

$$\Delta\text{Welfare}^{\text{CES-MC}} = \text{CV}^{\text{CES}} + \Delta\Pi^{\text{CES}} + \Delta\text{TR}. \tag{22}$$

4 The CES-OLY approximation

The CES-MC model approximates the true Mixed-CES model in two ways: by assuming homogenous consumers, and by falsely imposing constant markups in the pricing strategy of firms. This generates two main sources of errors in the predicted changes in market share and associated aggregate outcomes: i) the trade elasticity $\breve{\eta}$ is a unique parameter in charge of describing a whole set of heterogeneous own and cross-price elasticities. ii) the EHA formula used to compute counterfactual market shares only holds under the (wrong) assumptions of constant markups and trade elasticities.

The second deviation from the true DGP can be amended, by letting firms being non-atomistic, and therefore adjust markups optimally. Let us consider a version of the model featuring homogeneous consumers but oligopolistic multi-product firms. Firms face CES demand curves such that

$$s_{mn} = \frac{\exp(\beta x_m + \xi_m) p_{mn}^{-\eta}}{\sum_j \exp(\beta x_j + \xi_j) p_{jn}^{-\eta}}, \tag{23}$$

and choose markups μ_{mn} for each of their models in each market optimally such that

$$p_{mn} = \mu_{mn} c_m \tau_{i(m)n}.$$

Unlike in CES-MC, the optimal markup varies over both models and markets. Hottman et al. (2016) have shown the useful result that in this setup, the firm chooses prices such that markups are equalized across its varieties. The formula for the optimal markup depends on whether firms play Cournot or Bertrand. We maintain the usual assumption in empirical IO and choose price-setting strategies. Therefore we have

$$\mu_{mn} = \mu_{fn} = \frac{\eta(1 - s_{fn}) + 1}{\eta(1 - s_{fn})}, \quad \forall m \in f, \quad \text{with} \quad s_{fn} = \sum_{m \in f} s_{mn}, \tag{24}$$

converging to $(\eta+1)/\eta$ as market shares of firms go to zero. Except in that limit case, there is no closed form solution to the market share equation and estimation of η is problematic. We can however follow an iterative procedure to estimate η even under oligopoly.

Start with a guess called η_0. Since we observe firm market share s_{fn}, we can compute the equilibrium markup μ_{fn0} using (24). This markup is passed on the left-hand-side, and combined with the log of market shares to yield the following regression:

$$\ln s_{mn} + \eta_0 \ln \mu_{fn0} = -\eta \ln \tau_{i(m)n} + \text{FE}_m + \text{FE}_n + \upsilon_{mn}, \tag{25}$$

The coefficient on trade costs provides a new estimate η_1, with which we can recalculate markups. The process iterates until obtaining an estimate $\breve{\eta}$, consistent with oligopoly pricing. The appendix contains a section documenting the fact that this procedure estimates perfectly η when it should (without consumer heterogeneity), and how it performs in cases where the DGP is a random coefficients one.

Once the estimate $\check{\eta}$ is obtained, one can also work with Exact Hat Algebra to compute counterfactual market shares. Indeed, changes in market shares are as follows

$$\hat{s}_{mn} = \frac{(\hat{\mu}_{mn}\hat{\tau}_{i(m)n})^{-\check{\eta}}}{\sum_j s_{jn}(\hat{\mu}_{jn}\hat{\tau}_{i(j)n})^{-\check{\eta}}}. \tag{26}$$

And the change in markup is also very easily computed using

$$\hat{\mu}_{mn} = \hat{\mu}_{fn} = \frac{1}{\mu_{fn}} \frac{\check{\eta}[1 - \hat{s}_{fn}s_{fn}] + 1}{\check{\eta}[1 - \hat{s}_{fn}s_{fn}]}, \quad \forall m \in f, \quad \text{with} \quad \hat{s}_{fn} = \frac{\sum_{m\in f}\hat{s}_{mn}s_{mn}}{s_{fn}}, \tag{27}$$

We have all the elements to iterate over the EHA predictions. Start with $\hat{s}_{mn} = 1$, aggregate to obtain the firm-level market shares $\hat{s}_{fn}$. Using initial markup (24), one can retrieve its change from (27). The new vector of market share changes is finally obtained with (26).

In terms of welfare analysis, the results of the CES-MC approximation apply with the following changes:

1. The change in the price index now incorporates the change in markups, with resulting CV as:

$$\text{CV}^{\text{CES-OLY}} = Y(1 - \hat{P}) \quad \text{with} \quad \hat{P} = \left(\sum_m s_{mn}(\hat{\mu}_{mn}\hat{\tau}_{i(m)n})^{-\check{\eta}}\right)^{-1/\check{\eta}}. \tag{28}$$

2. The changes in profits adjust market shares and prices post-liberalization to the be the CES-oligopoly ones. The formula becomes

$$\Delta\Pi^{\text{CES-OLY}} = Y \sum_{m\in f_{\text{dom}}} (\hat{s}_{mn}\hat{L}_{mn} - 1)s_{mn}L_{mn}. \tag{29}$$

Changes in tariff revenues are not affected, since those are brought to zero in any case. The CES-OLY approximation is expected to improve the fit compared to the CES-MC approximation, in ways that can be analyzed comparing equations (20) with (28) for consumer surplus, and (21) with (29) for profits. On the consumer side, the only difference is that the EHA for the oligopoly case takes account of the change in markups, it should therefore improve the fit with the true DGP. In the case of panel C setting 1, the fit should actually be perfect. On the profit side, the CES-OLY and CES-MC cases differ in two ways: the CES-MC approximation will tend to overestimate the fall in the market share of domestic models ($\hat{s}_{mn}$), since the adjustment through optimally lowered markups is not accounted for. However, in terms of profits, the lowered markups reduce $\hat{L}_{mn}$, which will tend to make CES-OLY and CES-MC predictions closer. Intuitively, an oligopolistic firm confronted with more competition lowers its price to lose less market share than a Dixit-Stiglitz firm. The difference in profits lost can however be small because of the two countervailing effects.

5 Monte Carlo simulations of trade liberalization

We now turn to describing the details of our Monte Carlo simulations, starting with how we calibrate parameters of the Mixed CES DGP under different settings, before turning to results in terms of aggregate market share levels and changes, together with welfare analysis.

5.1 Calibration of different settings

We consider an industry consisting of 100 different varieties, which belong to either 100 or 10 firms. The first case is called "monopolistic competition", since with such a large number of decision makers, the optimal prices by oligopolists are very proximate to the monopolistic competition asymptotia. With 10 models per firm, oligopolistic pricing and cannibalization effects have an actual influence, and we call this case "multi-product oligopoly".

The firms are producing in 5 different i countries (either 20 or 2 firms in each country depending on panels). The five countries are also markets (denoted n) for all of those models (we do not consider the extensive margin of export decision as the canonical empirical IO setup under investigation). Apart from the 10% tariff paid on all international trade flows, we calibrate a bilateral distance ad-valorem equivalent d_{AVE} which multiplies "distance" specified as abs$(n - i)$. The target is to obtain a domestic market share of 50%, which corresponds approximately to the share of cars consumed in the USA that are produced locally, according to 2016 data from Head and Mayer (2017).

On the supply side, we calibrate the variance of the model-specific cost (σ_ν) such that we replicate the levels of concentration exhibited in the data in the panels featuring multi-product oligopoly. Coşar et al. (2018) report a share of the top five firms (CR5) of 68% on the US market, the most concentrated market in their data being Brazil at 82%. We target concentration ratios in the mid-80s in the oligopoly panels to make sure we have enough concentration compared to the monopolistic competition case. Observed quality also affects costs, with an elasticity $\gamma_1 = 0.24$ calibrated from estimates by Crozet et al. (2012).

Following the approach of Edmond et al. (2015), we adjust the central parameter driving price responsiveness, (η)[4] such that the CES cost elasticity estimate ($-\breve{\eta}$) rounds to -5.0 in all settings when averaged over the 1000 replications.[5] The rationale for fixing the same trade cost elasticity across settings is that we consider $\breve{\eta} = 5$ as data in the spirit of Arkolakis et al. (2012). The average sensitivity of consumers to quality is normalized to $\bar{\beta} = 1$.

The most important part of the calibration regards rich substitution, which is driven by the different parameters governing consumer heterogeneity. The five settings differ in the heterogeneity of the consumers' α_h and β_h parameters affecting sensitivity to price and to quality respectively. The first setting imposes homogeneous coefficients in consumer indirect utility. With $\sigma_\alpha = \sigma_\beta = a_2 = b_2 = 0$, demand takes the CES form. Setting 1 is intended to benchmark results on a case with lack of consumer heterogeneity, where a lot of predictions (markups, pass-through rates etc.) are well known.

The second setting introduces heterogeneous valuations of quality, raising σ_β to one. This opens the possibility of richer substitution parameters than setting 1 because now high quality goods are closer substitutes for other high quality goods. This setting is hard to match to the literature because our x is a composite quality indicator whereas the IO implementations generally consider multiple product features. One way to do a units-free comparison is to take the ratio of the standard deviation of β to its mean value. In our setting 2 this is just one. Berry et al. (1999) consider four model attributes, with each having different relative importance for parameter heterogeneity. On average the ratio is 0.73, which suggests that our setting 2 has at least as much heterogeneity in the valuations of quality as in the car industry. Reynaert and Verboven (2014) set this ratio at 0.5 in their Monte-Carlo exercise, and their empirical estimates (on cars) show ratios that are under one for all but one of their characteristics.

The next two settings allow for heterogeneous price sensitivity. Setting 3 maintains $\sigma_\alpha = 0$, but introduces the relationship between price elasticity and income that is traditionally used in IO, with $a_2 = -1$, which creates a first source of consumer-based heterogeneity in price sensitivity (the consumer income distribution is calibrated on US 2010 GDP per capita and Gini index). In Berry et al. (1999), the income-dependence of heterogeneity parameters are set at $b_2 = 0$ and $a_2 = -1$. Coşar et al. (2018) estimate the equivalent of our parameter a_2 and find it to be strikingly close at -0.997.[6] While Berry

[4] $\mathbb{E}(\ln \alpha_h) = \eta + a_2 \ln y_h$.

[5] This corresponds to the average firm-level elasticity of trade with respect to tariffs found by Bas et al. (2017), pooling over 6-digit industries.

[6] With logit demand, α_h is the coefficient on price. However, $\alpha_h \propto y_h^{a_2}$ has the same implication for how demand elasticities responds to income.

et al. (1999) impose all the heterogeneity in price responsiveness to come from income variation, subsequent work allows for additional random heterogeneity, which we denote σ_α. Coşar et al. (2018) estimate this parameter to be quite small at 0.053. Nakamura and Zerom (2010) also find this additional dimension of consumer heterogeneity does not change results much after taking into account income-related heterogeneity in price responsiveness.

Setting 4 calibrates this important dimension of heterogeneity on a data moment that we can match in our simulations. A distinctive prediction of introducing $a_2 = -1$ is income sorting. Indeed, when rich consumers are less price sensitive, they buy high-price goods (foreign varieties in particular in our case) with larger probability, a prediction strongly confirmed empirically in Bils and Klenow (2001). The average income of buyers should therefore be higher for high price varieties. On the other hand, $\sigma_\alpha > 0$ will weaken the income-price relationship. In order to estimate the strength of that relationship in the actual data, we use the US consumer expenditure survey for the years 2014 to 2016, and compute, for passenger cars, the log average income of households buying a particular brand, which we regress on the log of the median price paid for the car. Figure 1 reports the scatter plot with regression coefficient associated (weighting by number of observations the plot and the regression). The Income Sorting Elasticity (ISE) is estimated at 0.85. It will turn out that with $a_2 = -1$, the ISE predicted by the MCES DGP is too high (around 1 for Mixed CES, even higher for Mixed logit). Setting 4 therefore adjusts the second source of heterogeneity in price sensitivity ($\sigma_\alpha > 0$) so as to obtain an ISE close to the 0.85 estimated for the US car industry.

Figure 1: Income Sorting Elasticity

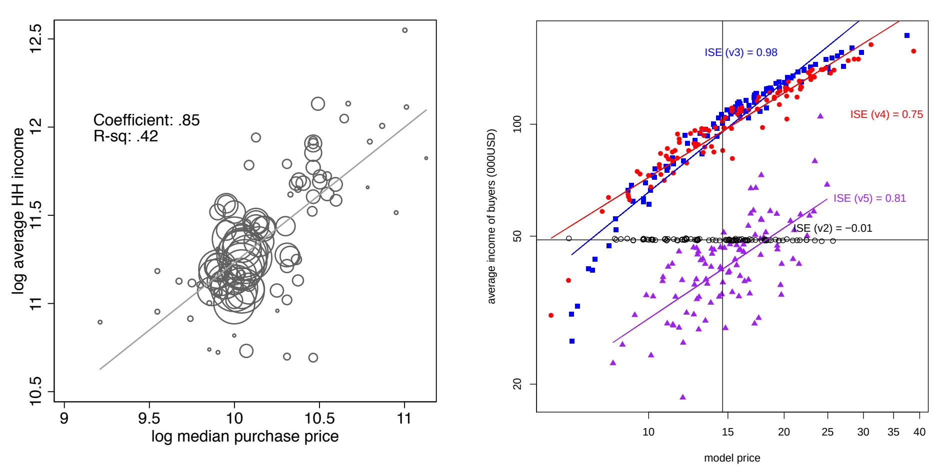

(a) cars/USA/2014–2016 (b) simulations

Last, setting 5 goes back to $a_2 = 0$ and $\sigma_\alpha = 0$ but now calibrates the heterogeneity in consumer sensitivity to quality with $b_2 > 0$ and $\sigma_\beta > 0$ so as to obtain the same ISE as in setting 4, while maintaining the other calibration targets. Note that this setting will exhibit rich substitution patterns because of $\sigma_\beta > 0$, and income sorting conforming to the data moment, because of $b_2 > 0$, meaning that rich households value quality more and therefore are more likely to buy expensive varieties (a high x_m means a high costs c_m since $\gamma_1 > 0$).

Table 1: Cross-price elasticities ($\partial \ln s_j / \partial \ln p_m$) of the MCES model

	Parameter Setting					Goldberg & Verboven		
$m \to j$	1	2	3	4	5	GBR	DEU	FRA
dom→ dom	0.097	0.099	0.090	0.090	0.099	0.076	0.015	0.170
dom→ for	0.097	0.099	0.061	0.061	0.099	0.025	0.036	0.024
for→ dom	0.030	0.029	0.026	0.025	0.029	0.030	0.043	0.015
for→ for	0.030	0.029	0.022	0.022	0.029	0.044	0.033	0.054
coef. on abs$(x_m - x_j)$	-0.000	-0.516	-0.580	-0.590	-0.583			

Note: Values of the cross-price elasticities are averaged over 100 replications for each setting. The coefficients reported in the last row are estimates of a regression where the log of the cross-price elasticity is explained by the absolute value of the difference in observed quality, and a set of fixed effects for each model-replication combination. All estimates are extremely precise, despite the use of (multi-way) model-replication clusters.

Whether substitution patterns are "rich enough" depends on the patterns of cross-price elasticities. Table 1 shows those elasticities (averaged over 100 repetitions, formulas for the elasticities are in the appendix) for the five settings of panel (c) in Table 2) which exhibits multi-product oligopoly. Setting 1 replicates the standard CES result that the cross elasticity depends only on the market share of the variety m whose price is being increased. Since foreign varieties are hampered by a trade cost, the rows where m is foreign have lower market shares and hence lower average cross-price elasticities. This pattern continues in setting 2, because at this point we assume that foreign varieties have the same mean quality as domestic ones.[7] In setting 3, the combination of heterogeneity in price elasticity with foreign firms having higher prices start to introduce segmentation: households buying foreign cars are less price-sensitive. Substitutability between foreign and domestic varieties is lower than between domestic varieties. This becomes even stronger in setting 4. Note that substitution between foreign varieties is the lowest

[7] Table 4 provides the cross-elasticities for the case where country 1 is high quality and countries 2–4 have decreasing mean levels of x. The dom-dom now differ from the dom-for in settings 2 and 5, and the for-dom also differs from for-for, as in the estimated data of Goldberg and Verboven (2001).

of all combinations, because the pool of consumers buying foreign varieties have a low average α_h. Setting 5 exhibits a pattern of cross-price elasticities that is similar to setting 2. The last three columns reproduce findings of Table 6 in Goldberg and Verboven (2001). The magnitude in cross elasticities are the same order of magnitude as in our simulations, although no country exhibits exactly the same pattern (probably because of systematic country-level differences in quality x that our simulations omit).[8]

Finally, the last row of Table 1 reports coefficients of a regression of logged cross-price elasticity on the absolute value of the observed quality of the two models. After taking logs, the cross-price elasticity is $\ln \eta + \ln s_m$ in CES without consumer heterogeneity. This model therefore predicts no relationship with the absolute quality differences once variety-level fixed effects are introduced to absorb variation in $\ln s_m$. This is exactly what we see in setting 1. Setting 2 allows for consumer heterogeneity in the taste for quality, which is enough to generate a very strong effect of the difference in model-level observed quality and cross-price elasticity. It is quite interesting that when introducing the second dimension of heterogeneity in price sensitivity in settings 3 and 4, the coefficient of the regression increases (in absolute value) as expected but moderately. Setting 5, which considers part of the heterogeneity in the taste for quality to come from rich consumers valuing quality more ($b_2 = 1$), but returns to homogeneous price sensitivity, has very comparable coefficient compared to settings with heterogeneity on price sensitivity.

5.2 Monte Carlo results

Based on the model described above, we simulate a tariff reduction scenario using the following data generating process:

1. For each of 100 models, draw quality (x_m and ξ_m) and marginal costs (ν_m). For each of 1000 households draw preferences (α_h and β_h).

2. Allocate models to firms equally, and then allocate firms to each of the five different countries. Impose a trade cost τ (specified as a higher marginal cost of delivering the product) on foreign varieties, consisting of a 10% ad-valorem tariff, and a bilateral Non-Tariff-Barrier that remains untouched.

3. Run the fixed point iteration to find the MCES equilibrium prices and market shares on each of the five markets.

4. Run equation (15), i.e. a regression of log share on log trade cost (controlling for model and market dummies) to estimate $\breve{\eta}$, the constant elasticity of market share

[8]The same table reports own elasticities that average to -5.5, a figure close to our calibration value.

with respect to τ, which also determines the markup under CES-MC. In CES-OLY case, run the fixed point iteration on (25) instead.

5. Substitute $\breve{\eta}$ into equation (16) to obtain CES analytical prices and market shares.

6. Reduce the trade cost on foreign varieties by cutting entirely the ad-valorem tariff, and redo step 3.

7. Calculate the CES-predicted change of (true) market shares from the pre-liberalization to the post-liberalization equilibrium together with welfare changes using Exact Hat Algebra. This is done using (17) for the CES-MC case, and the iterative procedure using (24), (26), and (27)in the CES-OLY case.

After completing these steps, we keep only market $n = 1$, whose income distribution is calibrated to match that of the USA in 2010. This means we can drop the n subscript at this stage. We then calculate the difference between the Mixed CES-generated market shares and the CES predictions at both the model and country level. We average the results of 1000 replications. Results are presented in three panels each containing the five different sets of MCES parameter values specified in the preceding subsection.

Panel (a) of Table 2 approximates monopolistic competition by setting the number of single-product firms to 100. Panel (b) reinstates oligopoly pricing of 10 multi-product firms aware of cannibalization effects between their 10 varieties. However, our approximation continues to treat the industry as if it were monopolistically competitive (CES-MC). In particular, the estimation and the counterfactual calculations ignore markup adjustments. Panel (c) employs the CES-OLY approximation on the same DGP as in Panel (b).

For each setting we report our targets ($\breve{\eta}$), CR5 and the domestic market share), together with a number of outcomes. We show three measures of pass-through of tariff changes into prices by foreign firms in order to illustrate the combined effect of demand curvature and oligopoly power. The first is the average derivative of prices with respect to costs (this "pass-through rate" is often used in the IO literature). The second is the average elasticity (often used in trade and international macro). Since most firms have low market share, we also show the elasticity of the foreign firm with the largest market share. The next column reports the "true" average own-price elasticity of demand implied by the DGP. The correlations in terms of levels and changes of s_m provide a metric for evaluating the accuracy of the approximation of the variety-level outcomes. The slopes of a regression of the CES approximation on true market shares (in levels and changes). The

Table 2: Monte Carlo simulation of the MCES data generating process

	Setting		CR5	Passthru			Mean	Corr		Slope			Agg. ΔS		Agg. bias	
#	$-\breve{\eta}$	ISE	(%)	rate	elas.	# 1	$\frac{\partial \ln s_m}{\partial \ln p_m}$	s_m	Δs_m	s_m	Δs_m	$\sum_{\text{dom}} s_m$	MCES	CES	Avg.	S.D.
Panel (a): 100 firms with 1 model each ("monopolistic competition"), CES-MC																
1	5.0	-0.00	38	1.20	1.00	0.97	-5.0	1.00	1.00	1.03	1.07	50	-9.1	-9.4	-0.3	0.2
2	5.0	-0.00	38	1.20	1.00	0.96	-5.0	1.00	1.00	1.05	1.10	50	-8.9	-9.4	-0.5	0.3
3	5.0	0.95	50	1.44	1.11	0.97	-3.6	0.95	0.94	0.84	0.79	49	-11.7	-9.7	2.0	1.9
4	5.0	0.85	53	1.47	1.12	0.96	-3.5	0.94	0.92	0.84	0.78	49	-12.0	-9.7	2.3	2.4
5	5.0	0.82	37	1.20	1.00	0.96	-5.0	1.00	1.00	1.04	1.10	51	-8.9	-9.4	-0.4	0.2
Panel (b): 10 firms with 10 models each (multiproduct oligopoly), CES-MC																
1	5.0	-0.00	83	1.16	0.98	0.94	-5.6	1.00	1.00	1.03	1.15	49	-8.3	-9.2	-1.0	0.4
2	5.0	-0.00	84	1.16	0.98	0.93	-5.6	0.99	0.99	1.05	1.20	51	-7.9	-9.1	-1.1	0.4
3	5.0	0.97	86	1.37	1.07	0.92	-4.1	0.94	0.92	0.84	0.83	51	-11.1	-9.6	1.5	2.2
4	5.0	0.87	86	1.40	1.08	0.90	-3.9	0.93	0.91	0.83	0.79	50	-11.6	-9.5	2.1	2.7
5	5.0	0.83	83	1.16	0.98	0.93	-5.6	1.00	0.99	1.04	1.18	50	-8.2	-9.2	-1.0	0.4
Panel (c): 10 firms with 10 models each (multiproduct oligopoly), CES-OLY																
1	5.0	-0.00	83	1.19	0.98	0.94	-5.0	1.00	1.00	1.00	1.00	49	-7.3	-7.3	0.0	0.0
2	5.0	-0.00	83	1.18	0.98	0.93	-5.0	1.00	1.00	1.02	1.03	49	-7.1	-7.3	-0.1	0.1
3	5.0	0.92	83	1.44	1.08	0.95	-3.4	0.94	0.94	0.82	0.75	48	-9.1	-7.6	1.5	1.4
4	5.0	0.86	84	1.46	1.09	0.94	-3.3	0.93	0.93	0.80	0.72	49	-9.3	-7.5	1.7	1.5
5	5.0	0.83	84	1.18	0.98	0.93	-5.0	1.00	1.00	1.02	1.05	50	-7.0	-7.2	-0.2	0.3

Note: CR5 is the 5-firm concentration ratio. "Corr" show Pearson correlations between levels and changes of MCES and CES market shares. "Agg. Δ" sums Δs_m (BLP) or $\Delta \breve{s}_m$ (CES) for all models produced in the home country. "Agg. Bias" shows the mean and standard deviations of the MCES aggregate change subtracted from the CES prediction.

aggregate decline in domestic market shares in the MCES and CES cases, with associated difference (average and s.d. of the 1000 repetitions).

Results on domestic market share

Setting 1 of panel (a), which eliminates rich substitution and oligopoly, replicates three results from CES-MC that were not imposed in the calibration. First, the average pass-through derivative and elasticity take the CES-MC results of 1.2 and 1 (although the top variety is non-negligible and therefore has a less-than-one pass-through elasticity). Second, the correlations of levels and changes in market shares round to 1.0. Finally, the own-price elasticity is -5, as it would be with infinitesimal firms, even though firms are sufficiently non-atomistic to yield a CR5 of 38% rather than the 5% that would be expected with 100 symmetric firms. There is a small bias in the CES-MC approximation, which predicts a larger reduction in the domestic share than what "truly" occurs under MCES. This error can be seen as an outcome of the two "Slopes" being greater than one. CES-MC predicts the biggest varieties will be larger than the true levels or changes. The reason traces back to the fact that the largest varieties are big enough to have less than unitary pass-through. In panel (b), the 100 varieties are owned by just 10 firms, with the 5 biggest accounting for over 80% of the market. This results in even smaller true pass-through. The slopes become steeper, and the aggregate prediction is more substantially biased even without random coefficients (setting 1). In panel (c) we adjust the approximation method to handle granular firms and find that it fully corrects the bias in the setting 1 predictions.

In setting 2, the richer substitution patterns brought about by $\sigma_\beta = 1$ lower the model-level correlations and raises the aggregate bias to half a percentage point. However, despite the presence of rich substitution patterns in setting 2, the outcomes are not very different. CES predicts domestic firms will go from taking 50% of the home market with the tariff to 40.6% of the market after it is removed. The "truth" is that domestic firms end up with a market share of 41.1%. The increase in the overestimate of the response to tariff reductions comes because of lower pass-through by the largest firm (0.96 instead of 0.97).

Setting 3 introduces variance in the price elasticity through the dependence on household income. Compared to the β heterogeneity in setting 2, α heterogeneity reduces the performance of the CES approximation more drastically. For the average firm, pass-through is now more than complete (a 10% cost increases leads to about an 11% price increase). Hence, faced with lower tariff costs, foreign firms *decrease* their markups, causing them to take even more market share from domestic firms. The CES approximation moves in the right direction but not enough, such that in setting 3 CES under-predicts the decline in the domestic market share. In order to match the "observed" moment of a trade elasticity of 5, the means of the random coefficients had to be recalibrated such that the own price elasticities of demand falls in absolute value from -5.0 to -3.6. Note that even though the average variety is in the "anti-competitive" region where costs increases lead to markup increases, the very largest variety remains in the "pro-competitive" region of incomplete pass-through. This is the consequence of the "battle" between convexification of demand versus non-granular firms. Setting 4 adds an additional moment to be targeted, an income-sorting elasticity of approximately 0.85. Matching that moment requires additional α heterogeneity, leading to slightly higher pass-through and greater reductions in the domestic share that the CES approximation fails to match. In this worst-case scenario for the CES approximation, the bias is about one standard deviation of the simulation and we would still emerge with the same basic takeaway, namely that the tariff cut will reduce the domestic firms' market shares by roughly 10 percentage points.

Settings 3 and 4 matched the observed tendency of rich consumers to buy expensive varieties through the assumption that rich consumers are less price sensitive. We think it is equally plausible that richer consumers value quality more highly. When building that assumption into setting 5, we can match the same amount of income sorting (ISE) but now the CES approximation works as well as it did in setting 2. In other words, the errors that CES makes come from combination of getting pass through wrong for the average firm and getting passthrough wrong for the leading firm(s). The former is a much bigger effect in settings 3 and 4 but negligible in settings 1, 2, and 5. The natural question is whether the latter effect would remain small if we let firms become bigger by combining varieties and pricing accordingly.

While panel (a) allows us to see the effects of random coefficients in a market structure with many (100) single-firms, panel (b) adds the second essential feature of the BLP data generating process: multi-product oligopoly. The most striking result in settings 1, 2, and 5 (the cases without α heterogeneity) is that counterfactual bias reaches a full percentage point, and the ratio of average bias to its standard deviation approximately *doubles*. On the other hand, the problems in panel (a) that were associated with super high pass-through (due to α heterogeneity) are now mitigated by the presence of large oligopolistic firms. This is actually intuitive. The presence of oligopoly depresses pass-through across the board with all three measures (rates, average and leading elasticities) lower that they would be under CES. This means the foreign firms do not increase market share as much and CES does a better job in settings 3 and 4. The poor performance in settings 1, 2, and 5 is linked to another notable difference between panels (a) and (b): The own-price elasticity needed to match the observed trade elasticity of 5 is now about -5.6. This further shifts the slopes away from one.

The takeaway from panels (a) and (b) is that CES does not have much trouble approximating counterfactuals even in a DGP exhibiting rich substitution—so long as the source is heterogeneity in consumer valuations of product attributes. There are two things that make CES performance deteriorate. The first is that firms are granular and potentially (as in panel (b)) large and multi-product. The second is that heterogeneous price responsiveness generates pass-through for the average firm that is far from the CES prediction. There is not much we can do about the latter problem except to note that it may not be a big issue in real data sets since such high pass-through rates have never been estimated (to our knowledge). With regard to the first problem, granular multi-product firms *are* features of the real world DGP so we have to come to terms with them if we are going to continue to use CES. Fortunately, section 4 has shown there are simple tools we can use to take into account oligopoly pricing in CES demand models. We deploy those tools in Panel (c) of Table 2.

The CES oligopoly estimation and approximation method solves entirely the bias in setting 1, as it should, since setting 1 corresponds to CES *without* randomness in the coefficients. The promising result of panel (c) is that the CES-OLY approximates counterfactuals very closely in settings 2 and 5, both of which feature rich substitution. It even mitigates some of the bias in the problematic settings 3 and 4. Since we do not need α heterogeneity to explain income-sorting and we do not need it to have substitution be stronger between varieties with similar observables, it is not clear that settings 3 and 4 are of first-order importance in real product markets. These settings have problematic implications for pass-through. Therefore, our tentative conclusion at this point in the analysis is that CES—*augmented* with oligopoly tools—may provide a much stronger than expected tool to approximate markets where substitution is in fact more complex.

Figure 2 illustrates the main outcome of that simulation graphing the average values of domestic sales changes together with its 95% confidence interval, over the three panels and five settings. The true change is the grey bar, the approximation is in blue. The bias apparent in settings 3 and 4 remains reasonably small across the three approximations.

Figure 3 replicates those figures, when the demand system is specified as Mixed Logit rather than Mixed CES. Full results, in appendix, show that one of the major differences between the two demand systems (calibrated on same data moments) is that the incompleteness of passthrough is larger with the logit than with CES. For instance, in setting 1 of panel (a), the pass-through elasticity is .84 with logit, vs 1 with CES. Therefore, we see in panel (c) of figure 3, that accounting for oligopolistic pricing is not enough to obtain a perfect prediction even without heterogeneity. The fall in domestic share is overes-

Figure 2: Summary of changes in domestic market share: Mixed CES

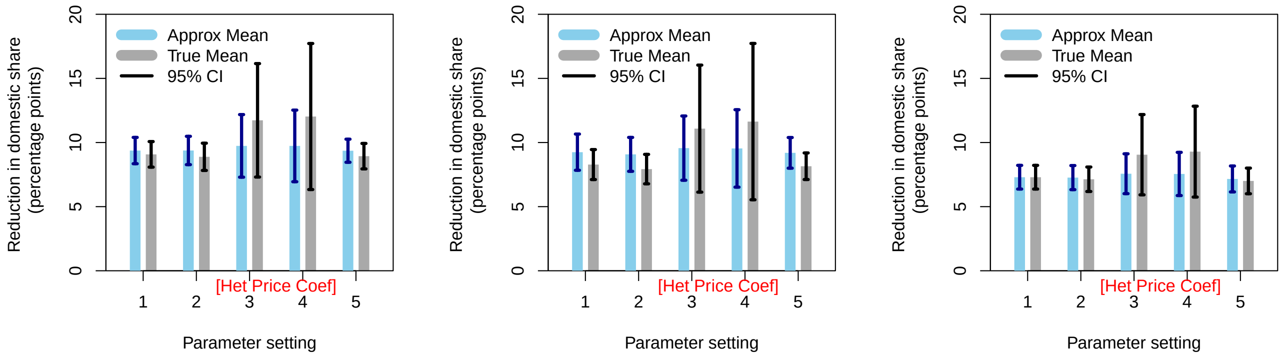

(a) MCES-MC vs CES-MC (b) MCES-OLY vs CES-MC (c) MCES-OLY vs CES-OLY

Figure 3: Summary of changes in domestic market share: Mixed Logit

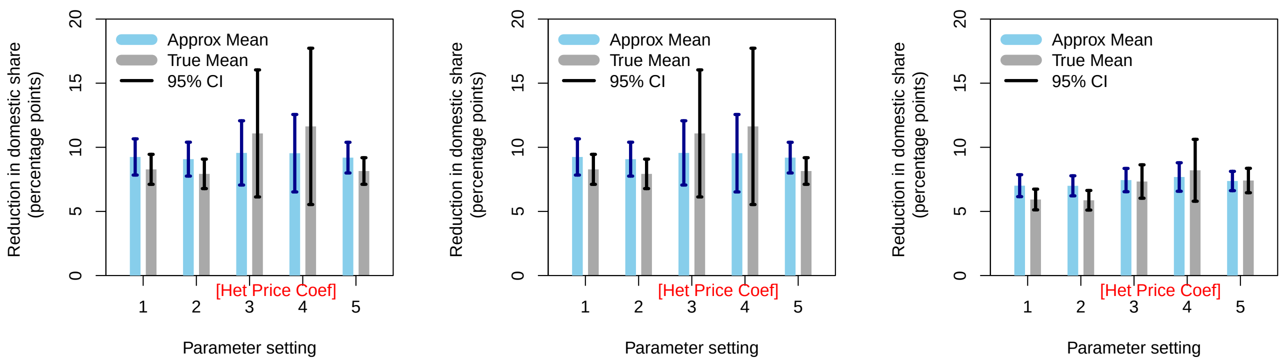

(a) MLOG-MC vs CES-MC (b) MLOG-OLY vs CES-MC (c) MLOG-OLY vs CES-OLY timated, because the fall in prices of foreign competitors is over-estimated. In settings 3 and 4, this weakness becomes an advantage, since the convexification of the demand system through income sorting now makes the average true pass-through rise, ending very close to the one predicted by the CES approximation. Under Mixed logit, adding the price dimension of consumer heterogeneity actually helps the CES prediction, since the misspecification on the oligopolistic forces and income sorting in the pass-through are compensating each other. Our findings seem therefore quite robust to the demand specification of the random coefficient model.

Figure 4: Initial model market shares: CES vs MCES

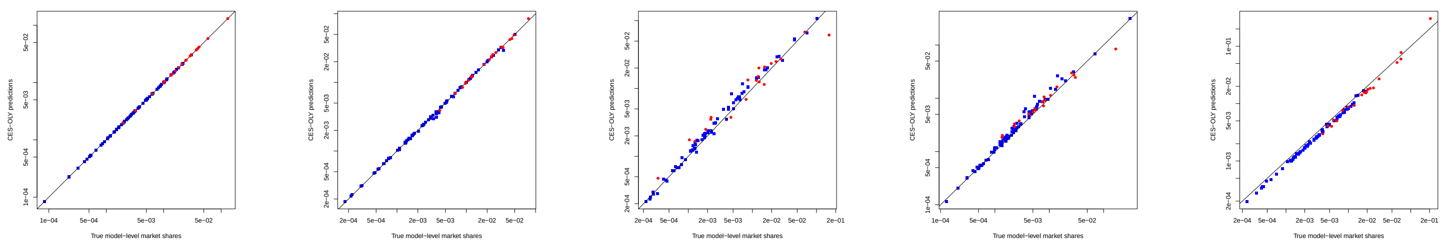

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 5: Changes in model-level market shares: CES vs MCES

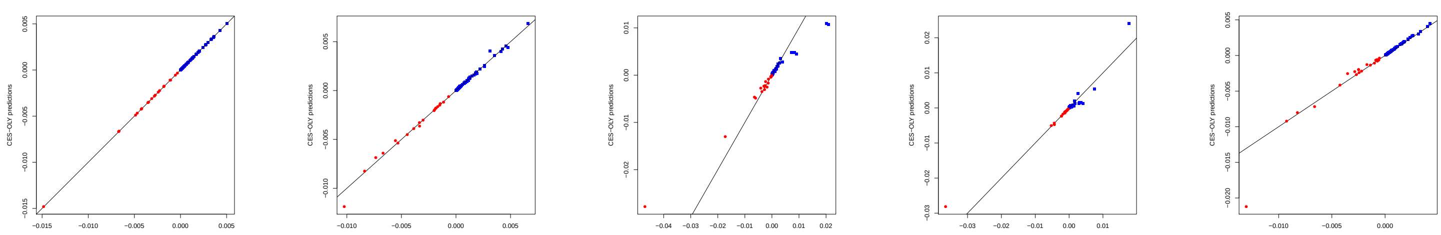

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figures 4 and 5 display the performance of the CES prediction at the micro level for one repetition of the DGP in Panel (c), with the MCES-generated data on the horizontal axis and the CES-OLY prediction on the vertical axis. The foreign models are represented with blue squares and the domestic models with red dots. Both figures show that most models are close to the 45-degree line where CES-OLY predictions match the generated data. Under Setting 1, the fit is perfect. Settings 2 and 5 are also near perfect in levels, with the main deviation being the error in the prediction regarding the drop in market share of the top national firm, which is overestimated.

Figure 6: Percent changes in domestic market share: CES-OLY vs MCES

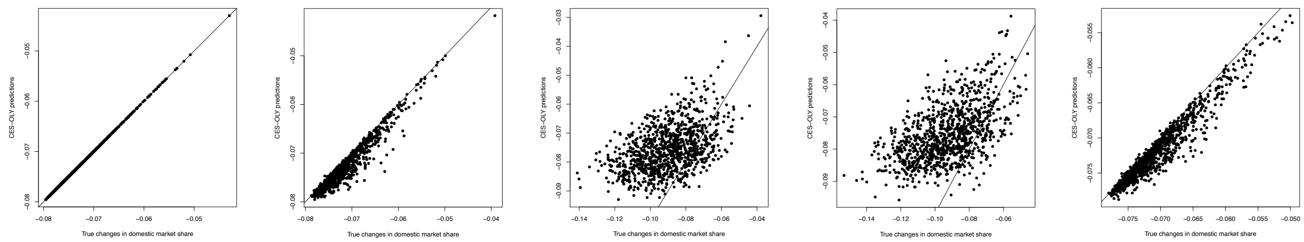

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (d) Setting 5

We illustrate graphically the change in the shares of all domestically produced models in response to trade liberalization in Figure 6. Going over the different settings illustrates the remarkable predictive performance of the CES approximation across 1000 repetitions, with the points clustering close to the 45 degree line in settings 2 and 5. Panels (c) and (d) of Figure 6 show that when consumers feature large heterogeneity in price elasticities the CES-MC approximation tends to under-predict the absolute change.

Welfare results

Table 3 summarizes the welfare consequences of our counterfactual unilateral trade liberalization exercise (the removal of a 10% tariff). We show the ability of CES to approximate the three components to domestic welfare changes individually and also evaluate the sum. Standard trade theory predicts that consumer surplus will rise while domestic profits and the home government's tariff revenues should decline. Costinot and Rodriguez-Clare (2014) quantify the optimal tariff for the class of models that setting 1 approximates (without random coefficients) to be approximately 20%. Since we start from a tariff that is below the optimal level, domestic welfare should decline in response to a unilateral tariff reduction. There is no theoretical result that we are aware of for the case of demand with random coefficients, but we indeed find (small) overall welfare losses to this experiment of trade liberalization.

Beginning with setting 1 in panel (a) of Table 3 we see that the CES approximation predicts changes in consumer surplus and tariff revenues very accurately.[9] There is slight

[9]Note that the change in tariff revenues is perfectly predicted by construction here because we reduce them from the observed level in the initial equilibrium to 0.

Table 3: Monte Carlo simulation: welfare effects

	Δ Cons. Surp.		Δ Profits		Δ Tar. rev.		Δ Welfare			
#	MCES	CES	MCES	CES	MCES	CES	MCES	CES	Diff.	
									avg	s.d.
Panel (a): 100 firms with 1 model each ("monopolistic competition"), CES-MC										
1	4.30	4.24	-1.69	-1.66	-3.16	-3.16	-0.55	-0.58	-0.03	0.07
2	4.32	4.27	-1.68	-1.68	-3.17	-3.17	-0.54	-0.59	-0.05	0.09
3	5.09	4.61	-1.79	-1.48	-3.39	-3.39	-0.09	-0.27	-0.18	0.42
4	5.11	4.59	-1.80	-1.46	-3.38	-3.38	-0.07	-0.25	-0.18	0.49
5	4.18	4.13	-1.67	-1.67	-3.06	-3.06	-0.56	-0.60	-0.04	0.09
Panel (b): 10 firms with 10 models each (multiproduct oligopoly), CES-MC										
1	4.60	4.26	-2.06	-1.86	-3.19	-3.19	-0.66	-0.78	-0.12	0.21
2	4.44	4.11	-2.04	-1.86	-3.06	-3.06	-0.66	-0.82	-0.16	0.22
3	5.14	4.38	-2.23	-1.57	-3.28	-3.28	-0.37	-0.48	-0.11	0.63
4	5.25	4.53	-2.22	-1.53	-3.41	-3.41	-0.38	-0.40	-0.03	0.69
5	4.55	4.22	-2.05	-1.86	-3.15	-3.15	-0.65	-0.79	-0.14	0.20
Panel (b): 10 firms with 10 models each (multiproduct oligopoly), CES-OLY										
1	4.44	4.44	-1.97	-1.97	-3.05	-3.05	-0.59	-0.59	0.00	0.00
2	4.42	4.42	-1.95	-1.97	-3.05	-3.05	-0.57	-0.60	-0.02	0.04
3	5.17	4.74	-2.10	-1.87	-3.20	-3.20	-0.13	-0.34	-0.21	0.34
4	5.11	4.69	-2.12	-1.88	-3.14	-3.14	-0.16	-0.33	-0.18	0.35
5	4.31	4.31	-1.95	-1.98	-2.94	-2.94	-0.58	-0.61	-0.03	0.07

Note: Columns titled MCES and CES compute averages over the 1000 replications for changes in the components of welfare. Figures reported are expressed in percentage of total expenditure value.

underestimation of the decline in domestic profits, probably because even with 100 independent varieties, the top 5 still have non-negligible market power (CR5 is 37%). The introduction of consumer heterogeneity in quality valuation in settings 2 and 5 hardly changes the overall very good predictive power of the CES-MC approximation.

In settings 3 and 4, the MCES model exhibits largely smaller welfare reductions, while CES is more stable. The essential source of the discrepancy is to be found in the fact that CES-MC understates both the gains of the consumers and the profit losses, but to different extents. Because of super pass-through, foreign firms reduce markups so that delivered prices fall even more than costs. This is missed by CES-MC, which underestimates (vastly) the gains in consumer surplus. For the same reason, foreign firms gain even more market share than expected if markups has stayed unchanged, hence the fall in domestic profits is also underestimated, but the bias is smaller than for consumer surplus, hence CES thinks the domestic welfare outcome is particularly bad. An important note is that the discrepancy of the prediction is contained to the settings with heterogeneity on price sensitivity: setting 5 is very close to settings 2 and 1.

Panel (b) of Table 3 illustrates the effect of oligopoly and joint optimization of profits across the varieties offered by a firm. The CES approximation is essentially unchanged as a result of these parameter changes. The MCES response to tariff reduction changes substantially but only for certain components. Compared to panel (a) the tariff revenue changes are very stable. The multiproduct oligopoly settings make tariff reductions look better for consumers and worse for firms. The latter effect dominates, leading to larger "true" domestic welfare losses. As in panel (a), CES-MC over-predicts welfare losses across all settings. The quality of the prediction however improves for settings 3 and 4. This paradoxically comes from a deterioration of the prediction on changes in domestic profits, which now compensates more closely the error on consumer surplus.

In panel (c), settings 1,2 and 5 are almost perfect as expected, settings 3 and 4 show again a paradox: both consumer surplus and profits predictions improve, but because those are conflicting effects, we still see an overall deterioration compared to panel (b). Overall, the income sorting and super passthrough induced by large amounts of consumer heterogeneity of the α_h type yields complex welfare effects for which is is hard to predict if the CES approximation will be severely biased or not.

5.3 Pass-through and Markups

With heterogeneity in price sensitivity, Burstein and Gopinath (2014) point out that pass-through elasticities can exceed 100%, which we confirm in our Table 2. Furthermore,

Figure 7: Changes in domestic consumer surplus: CES vs MCES

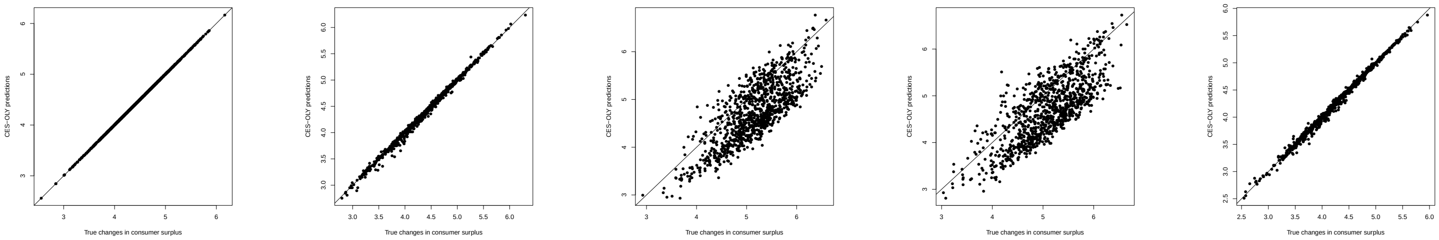

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 8: Changes in domestic firms' profits: CES vs MCES

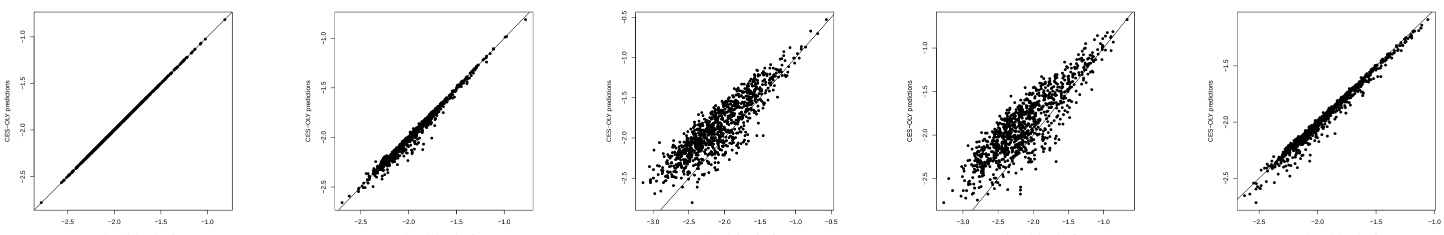

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 9: Changes in domestic welfare: CES vs MCES

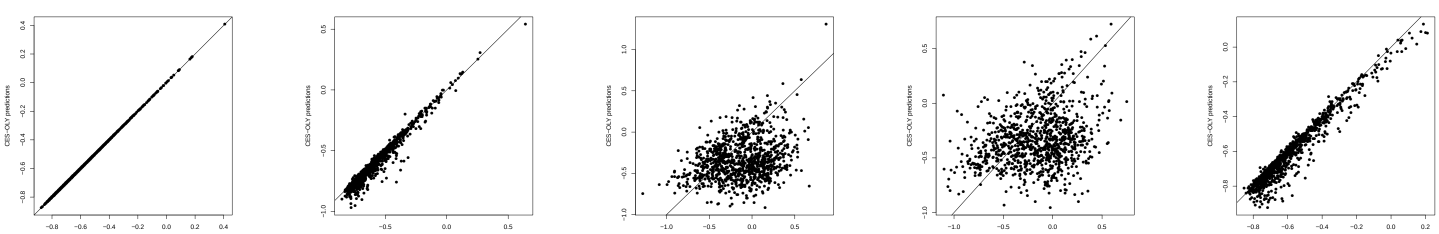

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5 we find that it is very difficult to obtain pass-through derivatives under one combined with heterogeneity in the price responsiveness terms. This is problematic, since demand-side estimation in BLP studies typically estimate large heterogeneity in price sensitivity, while it is consensual that the empirical pass-through is quite incomplete. For instance, Goldberg and Verboven (2001) report exchange rate pass through elasticities for cars that are considerably smaller than one (54% is their main estimate). The literature on pass-through for tariffs is considerably smaller than for exchange rate changes. Fontagne et al. (2017) use French firm-level exports over all goods to estimate a tariff passthrough elasticity around 0.65. Lud (2016) estimate a passthrough for US firms that is actually negative (a result they attribute to quality changes).

A related issue is that BLP with large enough heterogeneity in consumer preferences may violate the stylized fact that markups are increasing with market shares and decreasing with marginal costs. We derive in the appendix the conditions under which the pass-through elasticity is above or below one, and show that higher heterogeneity in consumer sensitivity to price will increase the elasticity from below to over one through an effect on the super-elasticity of demand. In order to illustrate the consequences of this finding for how markups vary with size and costs, we show two figures in the spirit of De Loecker et al. (2016), where we use the output of one repetition for the five settings of Panel (c) Table 2.

Figure 10: Markups and market share - Mixed CES

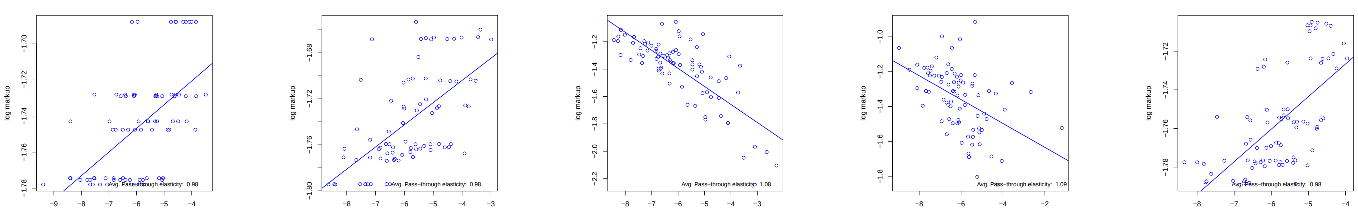

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 10(a) takes the extreme case with no consumer heterogeneity, and plots model-level markups against size (note that without consumer heterogeneity, markups are constant across varieties for a given firm in this DGP as emphasized by Hottman et al. (2016)). The figure exhibits the expected positive relationship between markups and size, as found in De Loecker et al. (2016) for the Indian economy (note that the production function approach used by the quoted paper does not impose the functional forms of BLP to estimate

Figure 11: Markups and market share - Mixed logit

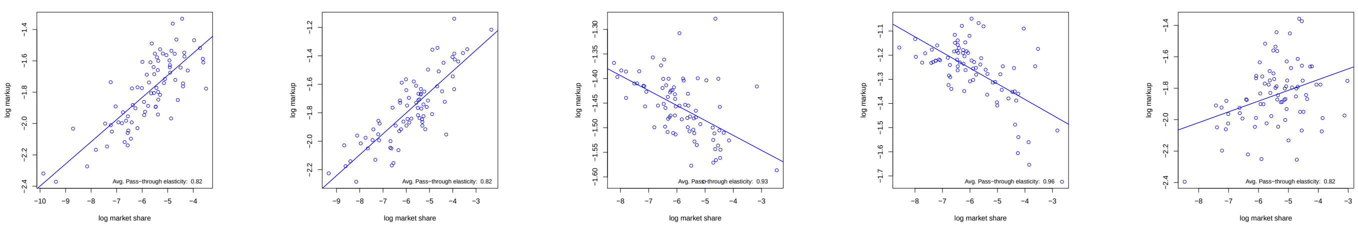

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

the markups). This positive relationship is maintained with heterogeneous consumers in terms of quality. Setting 2 and 5 keep average pass-through elasticities under one, which is consistent with most empirical findings, and yield the positive markup-size slope. On the contrary, pass-through elasticities above one, due to random coefficients on prices in settings 3 and 4 generate a negative relationship.

Figure 11 repeats the analysis for the Mixed logit case. Notice that with logit demand, markups now vary across varieties for a given firm, even for setting 1. Without random coefficients, oligopoly combined with logit demand creates a pass-through elasticity of .82 (still higher than what is commonly found), and confirms the strong positive relationship. Adding heterogeneity across consumers in price sensitivity largely increases the pass-through and reverses the sign of the relationship. Overall, the pattern is the same as with Mixed CES: expected negative relationship with CES or with rich substitution patterns created by heterogeneity in quality valuation, and counter-intuitively high markups for small firms with high enough levels of consumer heterogeneity in price sensitivity.

Mrázová and Neary (2017) point to an additional prediction related to cases where pass-through rates and elasticities are larger than one. Not only does the cross-sectional distribution of markups becomes inversely related to market share, but the markups of domestic firms are also predicted to *rise* following trade liberalization. The model then becomes "anti-competitive" in the sense that reducing the delivered costs of foreign competitors induce domestic firms to raise their price. This is a theoretical result for the monopolistic competition case with homogeneous consumers (also highlighted in Zhelobodko et al. (2012)), and arises when the demand system violates Marshall's Second Law of Demand (MSLD) under which own price elasticity should increase with price. The prediction is much more complex under BLP, and we are not aware of unambiguous theoretical result in that case. We can however measure the change in markup for domes-

Figure 12: Markup changes after trade liberalization - Mixed CES

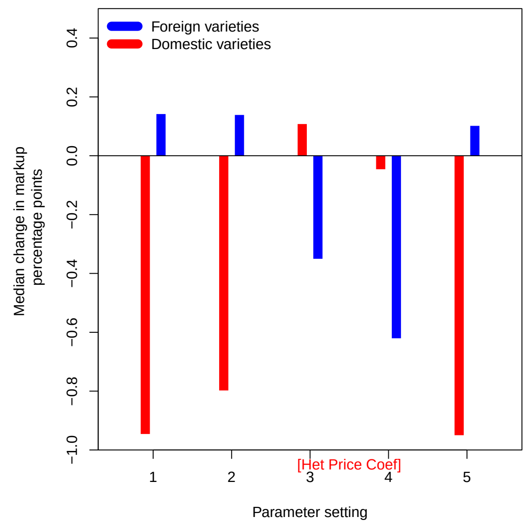

tic and foreign models in our numerical simulation (for the same draw as in figure 10). This is done in figure 12, where we show the change in markups of domestic firms with a red bar and the change for foreign firms with a blue one (for each of the five settings). Settings 1, 2 and 5 are "pro-competitive": foreign varieties increase their markups following the improvement in their access to the domestic market, while domestic varieties find it optimal to strongly cut their markups, and therefore prices. In settings 3 and 4, the foreign firms now reduce markups (more than complete pass-through), while domestic ones now increase it in setting 3, or leave it almost unchanged in setting 4.

Overall, our conclusion sheds light to an important aspect of the quality of the CES approximation. Table 2 shows that the approximation might be problematic when heterogeneity in price sensitivity is pronounced. However, this way of generating rich substitution patterns is not without issues. When calibrated using common practice and estimates, both the MCES and MLOG versions have a counterfactual prediction on how markups relate to size. In addition, the model essentially loses its pro-competitive features, domestic firms leaving markups unchanged or even raising those as competition from foreign varieties gets fiercer.

5.4 Retrieving quality with the CES approximation

Khandelwal (2010) initiated a residual-based method for retrieving quality as a demand shifter using a logit demand system with monopolistic competition. Khandelwal et al.

(2013) implemented a version with CES demand which has been since then applied in many contexts. The method involves passing prices on the left-hand-side of the (logged) CES demand equation and retrieving a residual intended to capture unobserved quality of the variety. In a multi-market context, the demand equation writes

$$s_{mn} = \frac{\exp(\beta x_m + \xi_m) p_{mn}^{-\eta}}{\sum_j \exp(\beta x_i + \xi_i) p_{jn}^{-\eta}}, \tag{30}$$

and the KSW approach involves running the following regression:

$$\ln s_{mn} + \breve{\eta} \ln p_{mn} = \mathrm{FE}_n + \breve{\beta} x_m + \psi_{mn}, \tag{31}$$

where the destination fixed effect captures the competition faced in country n (the denominator of equation 30), and x is the observed part of m's quality. One can then demean the observable variables over the n dimension to obtain an estimate of the model-destination appeal term $\breve{\psi}_{mn}$, which is then averaged to produce an estimate of the unobservable dimension of quality:

$$\breve{\xi}_m = \frac{\sum_n \breve{\psi}_{mn}}{N} = \frac{\sum_n (\ln s_{mn} - \overline{\ln s_{mn}}) + \breve{\eta}(\ln p_{mn} - \overline{\ln p_{mn}}) - \breve{\beta}(x_m - \overline{x_m})}{N}. \tag{32}$$

Equation (32) is equivalent to a logged version of Redding and Weinstein (2018) equation (17), which we follow by normalizing the mean of ξ_m to be zero. We expect $\breve{\xi}_m$ to be very close to true unobserved quality ξ_m if the model does not feature consumer heterogeneity, since in that case $\breve{\eta}$ is exactly reflecting the true parameter governing the demand-side response to a price change, which implies that the coefficient on observed quality $\breve{\beta}$ should be quite close to the true one. An interesting aspect of that approach is that it uses delivered price, and therefore is valid under both the MC and the OLY approach to CES, as long as $\breve{\eta}$ is estimated without bias. However, under consumer heterogeneity, the regression is mis-specified. We will here quantify how bad is the resulting measurement of unobservable quality in both cases where the true DGP is Mixed CES and Mixed logit.

The Mixed CES case should be the one with the best fit since the approximation maintains the functional form of the true demand function. Figure 13 plots the estimated unobserved model quality against the true one, for one repetition of each of the five settings of panel C in Table 2. The correlation between true and estimated levels of unobserved quality is 1 in the first setting, and closely aligned on the 45 degree line. The quality of the prediction deteriorates slightly as consumers are assumed to be more heterogeneous, but remains overall very high.

Figure 13: Unobserved quality and the residual approach - Mixed CES

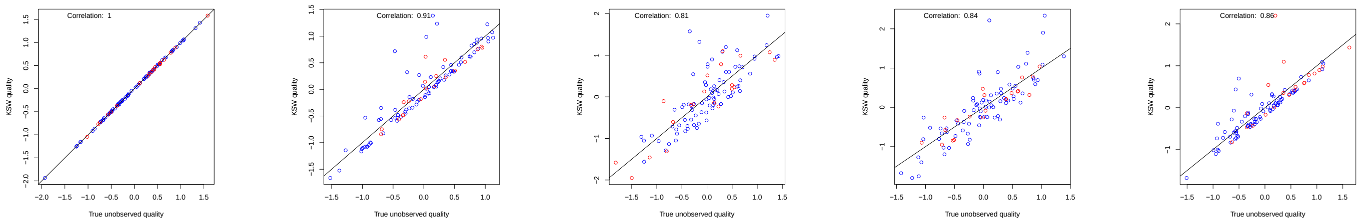

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 14 replicates the scatter plots of Figure 13 for the five settings of panel C in Table 6, where the true data generating process is Mixed logit rather than Mixed CES. This case of logit demand (the original case considered by Khandelwal (2010)) introduces a new source of discordance between the DGP and the CES approximation. As a consequence, the scatter plot is more noisy, even in setting 1 where consumers are assumed homogeneous. However, the CES approximation seems globally reliable in its capacity to recover the pattern of unobserved quality, even with large amounts of consumer heterogeneity, in settings 3 to 5.

Figure 14: Unobserved quality and the residual approach - Mixed Logit

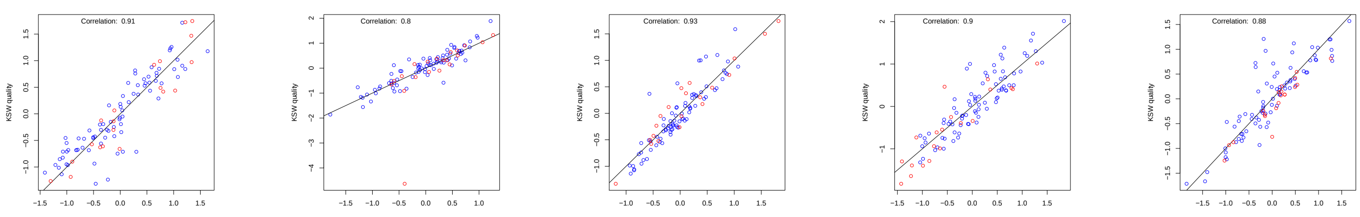

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

To summarize, the popular residual-based methods using the CES approximation to recover variety-level unobservable quality seem to deliver results that are robust to the demand-side function forms (logit vs CES), to the existence of rich substitution patterns in the true data, and to potential supply-side mis-specification multi-product oligopoly vs monopolistic competition).

6 Conclusion and Next Steps

For counterfactuals about aggregate responses to trade liberalization, CES makes good predictions even if BLP is the true model. Of course it remains true that if you want to know how specific models will be impacted by a change, BLP is important. An important *caveat* to this paper is that we have not put forward a general theorem. On the other hand, we view our simulation results as a disproof of the folk theorem that CES is doomed to perform poorly for all purposes if BLP is the "right" model.

There are several directions we plan to work on next.

1. **Merger Analysis:** The random coefficients models with multiproduct oligopoly is often used to evaluate the welfare consequences of mergers. The CES - OLY approximation we proposed is able to produce such evaluation with very low data requirement, in particular because of the use of Exact Hat Algebra. It is therefore an attractive alternative to full-fledge BLP if the consumers are not too heterogeneous. We will evaluate how wrong does CES-OLY get it under reasonable degrees of heterogeneity.

2. **Estimation:** Our approach compares BLP "truth" with the estimated CES model. The practical alternatives are *estimated* BLP versus estimated CES. This would recognize that BLP, even if it is the correct specification, will be estimated with error. Depending on the strength/validity of the instruments, the errors could be large.

3. **Cost elasticities:** The trade literature is interested in how changes in different measures of cost shifters affect patterns of trade flows. The cost elasticity is the product of the price elasticity and the pass-through elasticity. At this stage we have implemented an "empirical" approach to the second. We need to implement the analytical versions contained in the appendix in our simulations to quantify the bias in the CES approximation of the cost elasticity estimated in the trade literature.

References

(2016). Tariff pass-through, firm heterogeneity and product quality. *Journal of International Economics 103*, 234–249.

Adao, R., A. Costinot, and D. Donaldson (2017). Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review 107*(3), 633–89.

Anderson, S., A. De Palma, and J. Thisse (1992). *Discrete choice theory of product differentiation*. MIT Press.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *The American Economic Review 98*(5), 1998–2031.

Atkin, D., A. Chaudhry, S. Chaudhry, A. K. Khandelwal, and E. Verhoogen (2015). Markup and cost dispersion across firms: Direct evidence from producer surveys in pakistan. *American Economic Review 105*(5), 537–44.

Bas, M., T. Mayer, and M. Thoenig (2017). From micro to macro: Demand, supply, and heterogeneity in the trade elasticity. *Journal of International Economics 108*, 1–19.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2016). Global firms. Working Paper 22727, National Bureau of Economic Research.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.

Berry, S., J. Levinsohn, and A. Pakes (1999). Voluntary export restraints on automobiles: evaluating a trade policy. *American Economic Review 89*(3), 400–430.

Bils, M. and P. J. Klenow (2001, September). Quantifying quality growth. *American Economic Review 91*(4), 1006–1030.

Björnerstedt, J. and F. Verboven (2016, July). Does merger simulation work? evidence from the swedish analgesics market. *American Economic Journal: Applied Economics 8*(3), 125–64.

Burstein, A. and G. Gopinath (2014). *International Prices and Exchange Rates*, Volume 4, pp. 391–451. Elsevier.

Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics 110*, 135–150.

Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In E. Helpman, G. Gopinath, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Chapter 4, pp. 197–261. Elsevier.

Crozet, M., K. Head, and T. Mayer (2012). Quality sorting and trade: Firm-level evidence for french wine. *Review of Economic Studies 79*(2), 609 – 644.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica 84*(2), 445–510.

Edmond, C., V. Midrigan, and D. Y. Xu (2015). Competition, markups, and the gains from international trade. *The American Economic Review 105*(10), 3183–3221.

Fabinger, M. and E. G. Weyl (2016). The average-marginal relationship and tractable equilibrium forms. *arXiv preprint arXiv:1611.02270*.

Fernandes, A., P. Klenow, S. Meleshchuk, M. D. Pierola, and U. A. Rodriguez-Clare (2015). The intensive margin in trade: Moving beyond pareto. mimeo.

Fitzgerald, D. and S. Haller (2015). Exporters and shocks: Dissecting the international elasticity puzzle. manuscript.

Fontagne, L., P. Martin, and G. Orefice (2017). The international elasticity puzzle is worse than you think.

Gandhi, A. and J.-F. Houde (2016). Measuring substitution patterns in differentiated products industries. *University of Wisconsin-Madison and Wharton School*.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the us automobile industry. *Econometrica*, 891–951.

Goldberg, P. K. and R. Hellerstein (2008). A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market. *American Economic Review 98*(2), 423–29.

Goldberg, P. K. and F. Verboven (2001). The evolution of price dispersion in the european car market. *The Review of Economic Studies 68*(4), 811–848.

Head, K. and T. Mayer (2017). Brands in motion: How frictions shape multinational production.

Head, K., T. Mayer, and M. Thoenig (2014). Welfare and trade without pareto. *American Economic Review 104*(5), 310–16.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics 131*(3), 1291–1364.

Hummels, D. (1999). Towards a geography of trade costs. Technical Report 17, GTAP Working Paper.

Khandelwal, A. (2010). The long and short (of) quality ladders. *The Review of Economic Studies 77*(4), 1450–1476.

Khandelwal, A. K., P. K. Schott, and S.-J. Wei (2013). Trade liberalization and embedded institutional reform: evidence from chinese exporters. *American Economic Review 103*(6), 2169–95.

Knittel, C. R. and K. Metaxoglou (2014). Estimation of random-coefficient demand models: Two empiricists' perspective. *Review of Economics and Statistics 96*(1), 34–59.

Mrázová, M. and J. P. Neary (2017, December). Not so demanding: Demand structure and firm behavior. *American Economic Review 107*(12), 3835–74.

Nakamura, E. and D. Zerom (2010). Accounting for incomplete pass-through. *The Review of Economic Studies 77*(3), 1192–1230.

Nevo, A. (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy 9*(4), 513–548.

Nevo, A. (2011). Empirical models of consumer behavior. *Annual Review in Economics 3*(1), 51–75.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of political Economy 110*(4), 705–729.

Redding, S. J. and D. E. Weinstein (2016, August). Measuring aggregate price indexes with demand shocks: Theory and evidence for ces preferences. Working Paper 22479, National Bureau of Economic Research.

Redding, S. J. and D. E. Weinstein (2018). Accounting for trade patterns. Discussion Paper 12446, Center for Economic Policy Research.

Reynaert, M. and F. Verboven (2014). Improving the performance of random coefficients demand models: the role of optimal instruments. *Journal of Econometrics 179*(1), 83–98.

Romalis, J. (2007). Nafta's and cusfta's impact on international trade. *Review of Economics and Statistics 89*(3), 416 – 435.

Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica* (49(1)), 105–130.

Spearot, A. C. (2013). Variable demand elasticities and tariff liberalization. *Journal of International Economics 89*(1), 26–41.

Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, 240–268.

Zhelobodko, E., S. Kokovin, M. Parenti, and J.-F. Thisse (2012). Monopolistic competition: Beyond the constant elasticity of substitution. *Econometrica 80*(6), 2765–2784.

A Distribution of income and price elasticities

In traditional BLP, which uses mixed logit utility, the price sensitivity parameter α is specified as a function of income such that rich households are assumed to be less price sensitive:

$$\ln \alpha_h \sim \mathcal{N}(a_1 + a_2 \ln y_h, \sigma_a).$$

Since households are not observed, their income (and choice probabilities) are simulated. Income of h is assumed to be drawn from a log-normal distribution, with

$$\ln y_h \sim \mathcal{N}(\mu_y, \sigma_y).$$

Following CGLT, the means and standard deviations of log income, μ_y and σ_y, are set to match moments of the income distribution, specifically, average income $\bar{y}$, and the Gini coefficient, G. In our single-market analysis we use 2010 US values of $\bar{y} = 48374$ (USD) and $G = 0.41$, with both figures obtained for 2010 from the WDI. Knowing G one can invert Lubrano's formula for the Gini to obtain

$$\sigma_y = \sqrt{2}\Phi^{-1}((1+G)/2)$$

where $\Phi()$ is the cumulative of the standard normal. Then inverting the formula for the expected value of a log-normal variable, we obtain

$$\mu_y = \ln(\bar{y}) - 0.5\sigma_y^2.$$

B Cross-country quality differences

C DEV counterfactuals

We replicate panels a and b of the benchmark table, with counterfactuals using what Head and Mayer (2017) call the Difference in Expected Values method. After estimating the cost elasticity, this method calculates the predicted CES market shares (that differ from the true BLP ones). It then recomputes the CES market shares after the policy change, and computes the change between the pre and post CES market shares for domestic firms.

The important difference with EHA is that it starts from CES prediction rather than from "observed" shares. In settings 1, 2 and 5, it means starting from an overestimated market share of the top domestic firms in particular. Those are charging higher markups

Table 4: Cross-price elasticities ($\partial \ln s_j / \partial \ln p_m$) of the MCES model for panel (d)

	Parameter Setting					Goldberg & Verboven		
$m \rightarrow j$	1	2	3	4	5	GBR	DEU	FRA
dom→ dom	0.096	0.131	0.111	0.109	0.114	0.076	0.015	0.170
dom→ for	0.096	0.115	0.078	0.076	0.088	0.025	0.036	0.024
for→ dom	0.030	0.028	0.028	0.027	0.036	0.030	0.043	0.015
for→ for	0.030	0.035	0.030	0.028	0.047	0.044	0.033	0.054
coef. on abs$(x_m - x_j)$	-0.000	-0.571	-0.551	-0.575	-0.763			

Note: Values of the cross-price elasticities are averaged over 100 replications for each setting. The coefficients reported in the last row are estimates of a regression where the log of the cross-price elasticity is explained by the absolute value of the difference in observed quality, and a set of fixed effects for each model-replication combination. All estimates are extremely precise, despite the use of (multi-way) model-replication clusters.

and therefore have a smaller market share than CES-MC predicts. The same adjustment to the trade policy change is then applied, resulting in larger deviations in DEV when predicting changes in market shares (the slope on Δs_m is larger in DEV than in the EHA equivalent). The difference is however extremely small. For settings 3 and 4, where the starting point of CES-MC in levels is more drastically different, we see a more noticeable deterioration in overall prediction of ΔS.

Table 5: Monte Carlo simulation of the MCES data generating process-DEV approach

#	Setting $-\breve{\eta}$	ISE	CR5 (%)	Passthru rate	elas.	# 1	$\frac{\partial \ln s_m}{\partial \ln p_m}$	Corr s_m	Δs_m	Slope s_m	Δs_m	$\sum_{\text{dom}} s_m$	Agg. ΔS MCES	CES	Agg. bias Avg.	S.D.
							Panel (a): 100 firms with 1 model each ("monopolistic competition"), CES-MC									
1	5.0	0.00	37	1.20	1.00	0.97	-5.0	1.00	1.00	1.03	1.08	50	-9.1	-9.4	-0.3	0.1
2	5.0	-0.00	38	1.20	1.00	0.96	-5.0	1.00	0.99	1.05	1.12	50	-8.9	-9.4	-0.5	0.2
3	5.0	0.95	50	1.44	1.11	0.96	-3.6	0.95	0.90	0.86	0.70	49	-11.7	-9.2	2.5	2.3
4	5.0	0.84	53	1.47	1.12	0.95	-3.5	0.94	0.88	0.84	0.66	50	-12.4	-9.2	3.1	2.7
5	5.0	0.83	37	1.20	1.00	0.96	-5.0	1.00	0.99	1.04	1.12	50	-8.9	-9.3	-0.4	0.2
							Panel (b): 10 firms with 10 models each (multiproduct oligopoly), CES-MC									
1	5.0	-0.00	83	1.16	0.98	0.94	-5.6	1.00	0.99	1.03	1.15	49	-8.3	-9.2	-0.9	0.4
2	5.0	-0.00	84	1.16	0.98	0.92	-5.6	0.99	0.98	1.05	1.20	50	-7.9	-9.0	-1.0	0.4
3	5.0	0.97	86	1.38	1.07	0.92	-4.0	0.94	0.88	0.84	0.72	51	-11.0	-8.8	2.2	2.9
4	5.0	0.85	87	1.40	1.08	0.91	-3.9	0.93	0.87	0.82	0.65	51	-11.8	-8.9	3.0	3.4
5	5.0	0.83	82	1.16	0.98	0.93	-5.5	1.00	0.99	1.05	1.19	48	-8.1	-9.1	-1.0	0.4

Note: CR5 is the 5-firm concentration ratio. "Corr" show Pearson correlations between levels and changes of MCES and CES market shares. "Agg. Δ" sums Δs_m (BLP) or $\Delta \breve{s}_m$ (CES) for all models produced in the home country. "Agg. Bias" shows the mean and standard deviations of the MCES aggregate change subtracted from the CES prediction.

D Appendix for Mixed Logit vs CES

D.1 Mixed logit (original BLP) compared with CES

In this section, we describe the data-generating process when the demand system is mixed logit rather than mixed CES. This is the most common usage in empirical IO following the lead of Berry et al. (1995) in particular. We follow the set of demand-side assumptions used in the "user guide" paper by Nevo (2000), which has been widely influential in that literature.

D.2 The random-coefficients multiproduct oligopoly model (BLP)

There are M models subscripted with m and N households subscripted with h. The (indirect) utility of household h is given by

$$U_{mh} = \tilde{\alpha}_h(y_h - p_m) + \tilde{\beta}_h x_m + \tilde{\xi}_m + \varepsilon_{mh}, \tag{33}$$

with being the price of model m. As for Mixed CES in the text, x_m is an observed component of quality, whereas the unobserved component is captured in ξ_m. Assuming that the individual random term of households for specific models, ε_{mh}, is distributed Gumbel with scale parameter $1/\eta$, the choice probability of household h for model m takes the usual logit form, with $\beta_h = \eta\tilde{\beta}_h$, $\alpha_h = \eta\tilde{\alpha}_h$, and $\xi_m = \eta\tilde{\xi}_m$:

$$\mathbb{P}_{mh} = \frac{\exp(\beta_h x_m - \alpha_h p_m + \xi_m)}{\sum_i \exp(\beta_h x_i - \alpha_h p_i + \xi_i)}. \tag{34}$$

The main difference with the corresponding equation for Mixed CES is that the price enters linearly rather than in logs.

As for Mixed CES, β_h and α_h capture household heterogeneity with

$$\beta_h \sim \mathcal{N}(\bar{\beta} + b_2 \ln y_h, \sigma_\beta) \quad \text{and} \quad \ln \alpha_h \sim \mathcal{N}(\eta + a_2 \ln y_h, \sigma_\alpha).$$

On the supply side, the model-specific primitives, x_m, ξ_m and ν_m, are also assumed to be normally distributed, with

$$\ln c_m = \gamma_0 + \gamma_1 x_m + \gamma_2 \xi_m + \nu_m.$$

With all individuals buying one car, the market share of model m simply averages individual probabilities from equation (34) over the N consumers:

$$s_m = \frac{\sum_h \mathbb{P}_{mh}}{N} = \frac{1}{N}\sum_h \frac{\exp(\beta_h x_m - \alpha_h p_m + \xi_m)}{\sum_i \exp(\beta_h x_i - \alpha_h p_i + \xi_i)}. \quad (35)$$

Note that, keeping with the tradition in empirical IO, the market share s_m is now the market share in quantities, unlike with mixed CES, where it is market share in terms of values. Profits for firm f are given by model-level profits $\pi_m = Ns_m(p_m - c_m)$ for the models owned by the firm. Define the symmetric M-by-M co-ownership matrix as $\mathbf{\Omega}_{jm} = 1$ if models j and m have a common owner and zero otherwise.

The FOC for model m's price in the Bertrand-Nash equilibrium is

$$s_m + (p_m - c_m)\frac{\partial s_m}{\partial p_m} = -\sum_{j\neq m}\Omega_{jm}(p_j - c_j)\frac{\partial s_j}{\partial p_m}.$$

Isolating p_m on the left-hand side, while noting p_m is implicit in the s_m on the right-hand side, we have

$$p_m = c_m - \frac{s_m + \sum_{j\neq m}\Omega_{jm}(p_j - c_j)\frac{\partial s_j}{\partial p_m}}{\frac{\partial s_m}{\partial p_m}}$$

Since matrix formulations can drastically improve computation time, we express this equation in terms of the "optimal" response vector.

$$\mathbf{p}^\star = \mathbf{c} + \frac{\mathbf{s} + \mathbf{r}}{\mathbf{d}}, \quad (36)$$

where $\mathbf{r}$ is the effect on profits earned by the rest (r being a mnemonic for this rest) of the firm's models caused by raising model m's price and $\mathbf{d}$ is minus the derivative of market share with respect to own price (elements of $\mathbf{d}$ are given by $-\frac{\partial s_m}{\partial p_m}$). Define an M by M matrix of cross-price derivatives on market share as $\mathbf{D}$. Elements of $\mathbf{D}$ are given by $D_{jm} = \Omega_{jm}\frac{\partial s_j}{\partial p_m}$ for $j \neq m$ and zero otherwise. With the aid of $\mathbf{D}$ we can express the cross-variety profit impact of price rises compactly as

$$\mathbf{r} = \boldsymbol{D}(\mathbf{p} - \mathbf{c}).$$

Numerical experimentation indicates that this system is not a contraction mapping but it can be solved via fixed point iteration on the equation

$$\mathbf{p}^{i+1} = \omega\mathbf{p}^\star + (1-\omega)\mathbf{p}^i,$$

where ω is the weight accorded to the new best response and $1 - \omega$ is the weight on the previous vector of iterations. This closes the description of the demand and supply sides of the BLP data generating process.

D.3 The CES-MC approximation

As for Mixed CES, the CES approximation takes logs of the market shares generated by the multi-market version of the BLP DGP equation (35) and use it as the dependent variable in the following regression equation:

$$\ln s_{mn} = -\eta \ln \tau_{i(m)n} + \text{FE}_m + \text{FE}_n + \varepsilon_{mn}, \tag{37}$$

where FE_m are model-level fixed effects, which account for differences in all variables which makes m a successful variety (low c_m and/or high x_m and ξ_m). Destination-level FE_n account for the varying degrees of competition on each market faced by model m. The estimate of the trade cost elasticity $\breve{\eta}$ is is used to predicts prices

$$\breve{p}_{mn} = \frac{\breve{\eta}}{\breve{\eta} - 1} c_m \tau_{i(m)n},$$

and CES-MC predicted market shares, to be compared with the BLP-based market shares:

$$\breve{s}_{mn} = \frac{\exp(\breve{\text{FE}}_m + \breve{\text{FE}}_n)\breve{p}_{mn}^{-\breve{\eta}}}{\sum_j \exp(\breve{\text{FE}}_j + \breve{\text{FE}}_n)\breve{p}_{jn}^{-\breve{\eta}}}. \tag{38}$$

There are two fundamental differences between "true" s_{mn} and CES-MC predicted $\breve{s}_{mn}$. First, the two functional forms are radically different with s_{mn} constructed from an *inner* logit (which is not CES) and an *outer* summation over heterogeneous-coefficient consumers. Second, as for Mixed CES, the prices determining market shares are different. The true market shares are based on the equilibrium oligopoly prices that take into account cannibalization effects, that the CES-MC approximation neglects. We do implement the CES-OLY approximation with Mixed logit, while it should be noted that (unlike with Mixed CES), CES-OLY is not the natural estimation / approximation. One of the most important differences is that the CES-OLY estimation and approximation uses the result that markups are constant within a firm, when this is not the case with the logit demand system. The fit of CES-OLY with the true DGP will not be perfect even in cases without consumer heterogeneity.

D.4 Monte Carlo results for tariff reductions

Based on the model described above, we simulate a tariff reduction scenario using the same calibration and steps as in section 5. In particular, we set the key parameters of the BLP DGP so as to meet our targets: a trade cost elasticity $-\breve{\eta} = -5.0$, a total market share of domestic firms around 50%, a CR5 in the mid 80s for the multi-product oligopoly panel, and a income sorting elasticity around 0.85.

Figure 15: Income Sorting Elasticity in the MLOG case

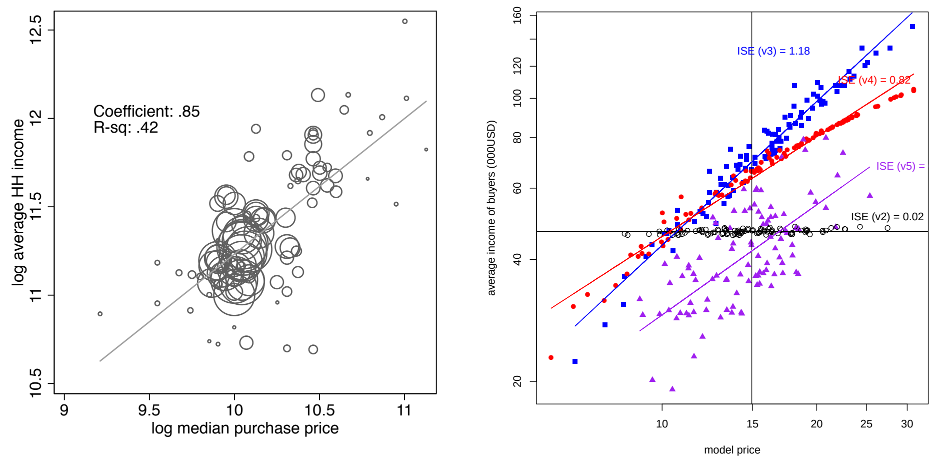

(a) cars/USA/2014–2016

(b) simulations

Our first implementation, shown as panel (a) of Table 6, retains a market structure approximating monopolistic competition, as in panel (a) of Table 2. This allows us to focus on the effects of changing the functional form to random coefficients logit under the familiar set of assumptions, where a number of analytical results are known for both functional forms. The collective share of the top 5 variety-firms is 33% rather than the 37% in the MCES case, both quite a lot larger than the 5% that would be expected with 100 symmetric firms. This does not seem to affect the results much in Mixed logit either as we see a pass-through rate of one, exactly what Anderson et al. (1992) obtain analytically for the case of symmetric firms.

The very strong correlations between market share in levels and changes (both 0.96) show that CES is capable of closely approximating logit market shares. High correlations at the model level are neither a sufficient nor a necessary condition for good fit of aggregates. On the one hand, there may be outliers that drive a substantial macro response without substantially weakening the correlation, and on the other hand, offsetting deviations could mute the aggregate response. Summing changes at the national level, we

Table 6: Monte Carlo simulation of the MLOG data generating process

	Setting		CR5	Passthru				Corr		Slope			Agg. ΔS		Agg. bias	
#	$-\breve{\eta}$	ISE	(%)	rate	elas.	# 1	$\frac{\partial \ln s_m}{\partial \ln p_m}$	s_m	Δs_m	s_m	Δs_m	$\sum_{\text{dom}} s_m$	MLOG	CES	Avg.	S.D.
							Panel (a): 100 firms with 1 model each ("monopolistic competition"), CES-MC									
1	5.0	0.00	33	1.00	0.84	0.77	-6.9	0.96	0.96	1.17	1.21	49	-7.9	-9.2	-1.4	0.4
2	5.0	-0.00	33	1.00	0.84	0.75	-6.9	0.95	0.95	1.17	1.24	50	-7.7	-9.2	-1.5	0.4
3	5.0	1.18	40	1.26	0.97	0.91	-4.4	0.98	0.96	0.99	1.00	50	-9.9	-9.7	0.2	0.9
4	5.0	0.84	44	1.38	1.01	0.91	-3.8	0.96	0.92	0.93	0.88	50	-10.8	-9.8	1.0	1.6
5	5.0	0.84	32	1.00	0.85	0.80	-6.9	0.97	0.97	0.97	1.01	50	-9.6	-9.6	-0.0	0.4
							Panel (b): 10 firms with 10 models each (multiproduct oligopoly), CES-MC									
1	5.0	-0.00	83	0.97	0.83	0.71	-7.5	0.94	0.92	1.21	1.38	48	-6.4	-8.7	-2.3	0.6
2	4.9	-0.00	83	0.97	0.83	0.70	-7.5	0.93	0.91	1.20	1.37	48	-6.3	-8.6	-2.3	0.5
3	5.0	1.25	84	1.20	0.95	0.83	-5.0	0.98	0.93	1.01	1.12	50	-8.6	-9.5	-0.9	0.9
4	5.1	0.86	85	1.30	0.98	0.84	-4.4	0.95	0.89	0.89	0.93	50	-10.1	-9.8	0.3	2.0
5	4.9	0.86	82	0.97	0.83	0.76	-7.5	0.96	0.95	0.98	1.10	48	-8.1	-9.1	-1.0	0.5
							Panel (c): 10 firms with 10 models each (multiproduct oligopoly), CES-OLY									
1	5.0	-0.00	83	0.97	0.82	0.70	-7.2	0.94	0.92	1.17	1.17	49	-5.9	-7.0	-1.1	0.3
2	5.0	-0.00	84	0.97	0.82	0.69	-7.2	0.93	0.92	1.17	1.18	48	-5.9	-7.0	-1.1	0.3
3	5.0	1.19	83	1.21	0.93	0.82	-4.5	0.98	0.95	1.00	1.02	50	-7.3	-7.4	-0.1	0.5
4	5.0	0.85	83	1.33	0.96	0.84	-3.9	0.96	0.92	0.90	0.87	50	-8.2	-7.7	0.5	1.1
5	5.0	0.85	83	0.97	0.82	0.75	-7.1	0.96	0.96	0.94	0.96	49	-7.4	-7.4	0.0	0.3

Note: CR5 is the 5-firm concentration ratio. "Corr" show Pearson correlations between levels and changes of MCES and CES market shares. "Agg. Δ" sums Δs_m (BLP) or $\Delta \breve{s}_m$ (CES) for all models produced in the home country. "Agg. Bias" shows the mean and standard deviations of the MCES aggregate change subtracted from the CES prediction.

find a less than one percentage point deviation between the BLP change and the change predicted by the CES model.

In setting 2 we see that the richer substitution patterns brought about by $\sigma_\beta = 1$ lower the model-level correlations and raise the aggregate bias. However the gap between a 15.7% change and 16.8% is still quite small—less than the standard deviation across replications. Setting 3 exhibits negligible bias and we suspect that this arises because both models have identical pass-through rates and nearly the same price-elasticities of demand (-5.0 in CES vs -4.7 in BLP). Setting 4 shows that even with the heterogeneity of price sensitivity set to matching empirical ISE, the performance of the CES approximation remains acceptable, and the average bias stays smaller than its standard deviation. The high σ_α leads to extraordinarily high pass-through rates—a \$1 increase in costs leads to a \$1.38 increase in prices—of tariffs into higher import prices. However, the restriction that cost elasticities of market share remain at -5 leads to a much lower value of μ_α which shrinks the absolute price elasticity to 3.8. Thus, foreign firms raise their price more, but lower substitutability between varieties dampens the response of domestic market shares.

Panel (a) shows that neither the logit form nor the rich substitution allowed by random coefficients prevents the CES approximation from matching counterfactuals. The next step, considered in panels (b) and (c), is to bring in oligopoly with multiple varieties sold by each firm.

Figure 16: Initial model market shares: CES vs BLP

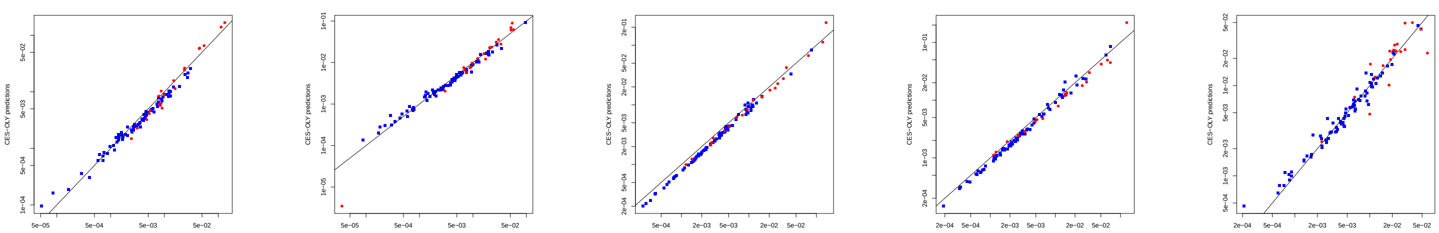

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Setting 1 now differs from CES-MC in two new dimensions that were excluded from this setting in panel (a): oligopoly, and cannibalization. As results of these changes, correlations at the variety level fall slightly. Figures 16(a) and 17(a) display the performance of the CES prediction at the micro level, with the BLP-generated data on the horizontal axis and the CES prediction on the vertical axis. The foreign models are represented with blue squares and the domestic models with red dots. Both figures show that most models

Figure 17: Changes in model-level market shares: CES vs BLP

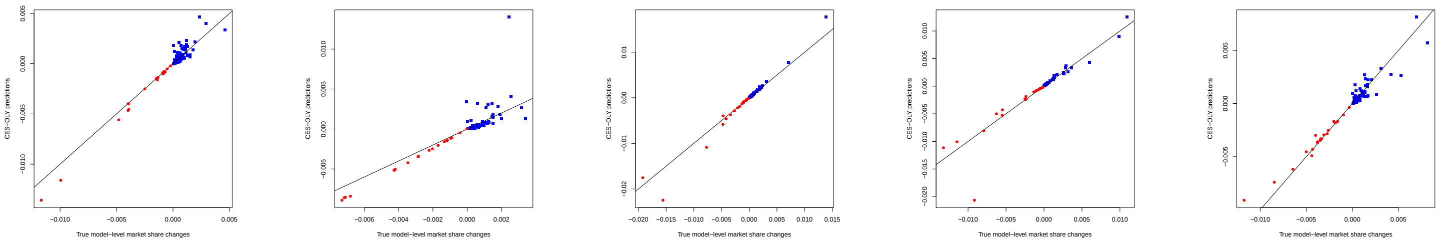

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5

Figure 18: Percent changes in domestic market share: CES vs BLP

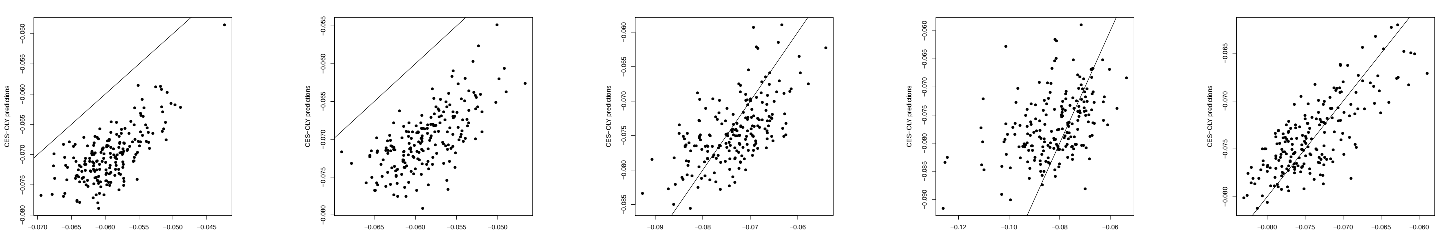

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4 (e) Setting 5 are close to the 45-degree line where CES predictions match the BLP generated data.

Aggregating across models, we graph the change in the shares of all domestically produced models in response to the 10% tariff reduction in Figure 18. There is a small downward bias in the CES prediction but it is only slightly worse than what was obtained under monopolistic competition. Adding richer substitution patterns in settings 2 lowers correlations marginally and hardly changes the macro level fit. Setting 3, with pass-through rates corresponding to CES-MC, also continues to fit the BLP data very well. The prediction and the data both round to a 17% domestic output reduction. Figure 18(c) illustrates the remarkable predictive performance of the CES approximation across 1000 repetitions, with the points clustering close to the 45 degree line. This fit is especially striking in light of the individual models that have large deviations in Figure 17(c). This illustrates the idea that rich substitution can induce the CES approximation to fail badly for individual models, while nevertheless maintaining strong predictive power at the macro level.

Setting 4 of panel (b) shows the largest case of CES bias in Table **??**. While the CES prediction does not change relative to panel (a), the BLP domestic output reduction rises by a half a percentage point. Figure 18(d) shows that in repetitions featuring large aggregate reductions under BLP, CES tends to under-predict the absolute change (the opposite of the pattern exhibited in figure 18(b)). The average bias, however, is about one third of the standard deviation across replications.

The CES-MC model performs very well in panel (b) despite the myopic pricing policy which omits adjustments to avoid cannibalization between the varieties that a firm offers. Equation (36) allows us to decompose the final price into three components: cost, a markup term given by $\mathbf{s}/\mathbf{d}$, and the cannibalization adjustment given by $\mathbf{r}/\mathbf{d}$. The simulation results show that on average just 1.6% to 1.8% of the price is attributable to the cannibalization adjustment. Under CES-MC, the markup share obtained by dividing $\mathbf{s}/\mathbf{d}$ by the price vector is just a constant given $1/\eta = 0.2$. In setting 3 the simulations show this markup term averages 0.2, with an interquartile range across models in a given replication of 0.19 to 0.21. This close similarity may account for why CES is so successful in this setting. In setting 4, where we observed deteriorating performance of CES, the markup term is 0.26 on average, with much more dispersion: the IQR is 0.22–0.30.

D.5 Elasticities

D.5.1 Mixed Logit (original BLP) elasticities

We are interested in reactions of a model's market share to changes in either its own price p_m or the price of a competing model p_j. With all individuals buying one car, the market share of model m simply averages individual probabilities over the N consumers:

$$s_m = \frac{\sum_h \mathbb{P}_{mh}}{N},$$

which enables to calculate the own price elasticity as

$$\frac{\partial \ln s_m}{\partial \ln p_m} = \frac{\sum_h \frac{\partial \mathbb{P}_{mh}}{\partial p_m}}{N} \frac{p_m}{s_m}.$$

Since the individual partial effect of a change in p_m is

$$\frac{\partial \mathbb{P}_{mh}}{\partial p_m} = -\alpha_h \mathbb{P}_{mh}(1 - \mathbb{P}_{mh}),$$

we can rewrite the own price elasticity as

$$\frac{\partial \ln s_m}{\partial \ln p_m} = -p_m \sum_h \omega_{mh} \alpha_h (1 - \mathbb{P}_{mh}), \quad \text{with} \quad \omega_{mh} \equiv \frac{\mathbb{P}_{mh}}{\sum_h \mathbb{P}_{mh}}.$$

Model m's own elasticity therefore is a weighted average of the individual household elasticities, which write

$$\frac{\partial \ln \mathbb{P}_{mh}}{\partial \ln p_m} = -\alpha_h (1 - \mathbb{P}_{mh}) p_m.$$

The weight ω_{mh} applied to each of those elasticities is the share of each household in total sales of the model. Note that in the individual elasticity, a low p_m will be associated with a high purchasing probability $\mathbb{P}_{mh}$, both contributing to a lowering of $\frac{\partial \ln \mathbb{P}_{mh}}{\partial \ln p_m}$. The individual response to price increases is therefore unambiguously concave, getting more and more pronounced as the price goes up. At the model level, however, a composition effect enters the picture. Low price models are preferred by low income individuals which are assumed to have a larger sensitivity for prices (a high α_h). Those low price models therefore face high α_h households with larger weight ω_{mh}, which raises the overall price elasticity. This introduces an element of convexity, which can dominate the individual-level concavity.

Let us turn to cross-price elasticities: the impact of an increase in the price of model j on demand for m. Those are given by

$$\frac{\partial \ln s_m}{\partial \ln p_j} = \frac{\sum_h \frac{\partial \mathbb{P}_{mh}}{\partial p_j}}{N} \frac{p_j}{s_m}.$$

The partial effect of j's price on $\mathbb{P}_{mh}$ is

$$\frac{\partial \mathbb{P}_{mh}}{\partial p_j} = \alpha_h \mathbb{P}_{mh} \mathbb{P}_{jh},$$

which yields

$$\frac{\partial \ln s_m}{\partial \ln p_j} = p_j \sum_h \omega_{mh} \alpha_h \mathbb{P}_{jh}.$$

Again, this is a weighted average of the individual choice probability cross elasticities,

$$\frac{\partial \ln \mathbb{P}_{mh}}{\partial \ln p_j} = \alpha_h \mathbb{P}_{jh} p_j.$$

D.5.2 Non-Mixed Logit elasticities

Setting variance of α and β across consumers to 0, we can obtain a very simple expression for the impact of a price change on market share:

$$\frac{\partial s_m}{\partial p_m} = \frac{\partial \mathbb{P}_m}{\partial p_m} = -\alpha \mathbb{P}_m (1 - \mathbb{P}_m) = -\alpha s_m (1 - s_m),$$

and for the own price elasticity:

$$\frac{\partial \ln s_m}{\partial \ln p_m} = -\alpha p_m (1 - \mathbb{P}_m) = -\alpha p_m (1 - s_m),$$

D.6 Partial effects of frictions

We are interested in observing effects of changes in frictions, for example how tariffs affect quantities.

$$\frac{\partial \ln s_m}{\partial \ln \tau_m} = \frac{\partial \ln s_m}{\partial \ln p_m} \times \frac{\partial \ln p_m}{\partial \ln \tau_m} \quad \text{and} \quad \frac{\partial \ln s_m}{\partial \ln \tau_j} = \frac{\partial \ln s_m}{\partial \ln p_j} \times \frac{\partial \ln p_j}{\partial \ln \tau_j}$$

The first factor is the demand own or cross price elasticity considered in the previous section. The second factor is exactly one under CES Dixit-Stiglitz iceberg assumptions. With BLP demand combined with oligopoly, firms will adjust their markups in response to trade cost shocks and pass-through will not be unitary. Under the kinds of demand curves considered by Brander and Krugman, pass-through is less than one. We are not aware of an analytical result for BLP oligopoly. However, given the convexity we see it exhibiting, it is possible that pass-through could be greater than one.[10] In general, pass-through will not only be non-unitary but it will also depend on the variety in question.

D.6.1 Single product case

The derivation of theoretical pass-through starts from FOC for model m:

$$s_m + (p_m - c_m)\frac{\partial \ln s_m}{\partial \ln p_m} = p_m + (p_m - c_m)\epsilon_{mm} = 0,$$

$\epsilon_{mm} < 0$ being the own price elasticity. Implicit differentiation gives

$$\frac{\partial p_m}{\partial c_m} = \frac{\epsilon_{mm}}{\epsilon_{mm} + 1 + (p_m - c_m)\frac{\partial \epsilon_{mm}}{\partial p_m}}.$$

Using the first order condition to replace $(p_m - c_m) = -p_m/\epsilon_{mm}$, one obtains

$$\frac{\partial p_m}{\partial c_m} = \frac{\epsilon_{mm}}{\epsilon_{mm} + 1 - E_{mm}}, \quad \text{where} \quad E_{mm} \equiv \frac{\partial \ln \epsilon_{mm}}{\partial \ln p_m}.$$

E_{mm} is often referred to as the superelasticity of demand, i.e. the elasticity of own price elasticity with respect to a change in own price (see Mrázová and Neary (2017) for a recent treatment). Under CES demand and monopolistic competition, ϵ_{mm} is a constant. Hence, $E_{mm} = 0$, and the pass-through derivative is a constant equal to $\epsilon/(\epsilon + 1)$. Writing pass-through as an elasticity (unlike Fabinger and Weyl),

$$\frac{\partial \ln p_m}{\partial \ln c_m} = \frac{\epsilon_{mm}}{\epsilon_{mm} + 1 - E_{mm}} \times \frac{c_m}{p_m} = \frac{\epsilon_{mm} + 1}{\epsilon_{mm} + 1 - E_{mm}}. \tag{39}$$

The sign of E_{mm} is therefore the determinant of whether the pass-through elasticity is greater or smaller than one. In the Dixit-Stiglitz case, $E_{mm} = 0$ implies a unitary pass-through elasticity.

- Under non-mixed logit

$$\epsilon_{mm} = \frac{\partial \ln \mathbb{P}_m}{\partial \ln p_m} = -\alpha p_m(1 - \mathbb{P}_m),$$

[10] The convexity occurs, because as prices rise the high a consumers are deterred from buying, which means the relevant consumers for the high price firm are those with low a (on average high income consumers). A similar intuition was given by Goldberg and Hellerstein (2008), Nakamura and Zerom (2010), and Burstein and Gopinath (2014).

$$E_{mm} = \frac{\partial \ln \epsilon_{mm}}{\partial \ln p_m} = [1 + \alpha p_m \mathbb{P}_m] .$$

Since $\alpha > 0$, the super-elasticity is positive (greater than one, its value when the market share of m approaches 0) and pass-through is incomplete. Using the penultimate equation to substitute out αp_m, we have

$$E_{mm} = \frac{\partial \ln \epsilon_{mm}}{\partial \ln p_m} = \left[1 - \epsilon_{mm} \frac{\mathbb{P}_m}{1 - \mathbb{P}_m}\right] .$$

Since the last fraction is the odds of choosing product m it can be infinitely large (for sufficiently low price).

- The mixed logit case is more complex. Recall that BLP demand at the household-model level implies $\frac{\partial \mathbb{P}_{mh}}{\partial p_m} = -\alpha_h \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})$, and therefore the following own elasticity:

$$\epsilon_{mm} = -\frac{p_m}{s_m} X_m, \quad \text{with} \quad X_m \equiv \frac{\sum_h \alpha_h \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})}{N} = -\frac{\partial s_m}{\partial p_m}.$$

The superelasticity is using

$$E_{mm} = \frac{\partial \epsilon_{mm}}{\partial p_m} \frac{p_m}{\epsilon_{mm}}, \quad \text{with} \quad \frac{p_m}{\epsilon_{mm}} = -\frac{s_m}{X_m}.$$

Taking the derivative of the own elasticity wrt price,

$$\frac{\partial \epsilon_{mm}}{\partial p_m} = -\frac{X_m}{s_m} - \frac{\partial X_m}{\partial p_m} \frac{p_m}{s_m} + \frac{p_m X_m}{s_m^2} \frac{\partial s_m}{\partial p_m}.$$

Using $\frac{\partial s_m}{\partial p_m} = -X_m$, one can re-write

$$\frac{\partial \epsilon_{mm}}{\partial p_m} = -\frac{X_m}{s_m} \left[1 + \frac{\partial X_m}{\partial p_m} \frac{p_m}{X_m} + \frac{p_m X_m}{s_m}\right] = -\frac{X_m}{s_m} \left[1 - \epsilon_{mm} + \frac{\partial \ln X_m}{\partial \ln p_m}\right] .$$

This expression will probably be useful for the multiproduct case, where is seems difficult to express things in terms of superelasticities. Hence the superelasticity is

$$E_{mm} = \frac{\partial \epsilon_{mm}}{\partial p_m} \frac{p_m}{\epsilon_{mm}} = \left[1 - \epsilon_{mm} + \frac{\partial \ln X_m}{\partial \ln p_m}\right] ,$$

where $\frac{\partial \ln X_m}{\partial \ln p_m}$ is the elasticity of the slope of demand to a change in price. One therefore needs to study how X_m varies with p_m

$$\frac{\partial \ln X_m}{\partial \ln p_m} = \frac{\sum_h \frac{\partial \mathbb{P}_m}{\partial p_m} \alpha_h (1 - 2\mathbb{P}_{mh})}{N} \frac{p_m}{X_m} = -\frac{p_m}{X_m} \frac{\sum_h \alpha_h^2 \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})(1 - 2\mathbb{P}_{mh})}{N}$$

hence,

$$\frac{\partial \ln X_m}{\partial \ln p_m} = -p_m \frac{\sum_h \alpha_h^2 \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})(1 - 2\mathbb{P}_{mh})}{\sum_h \alpha_h \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})}$$

D.6.2 Multiple product case

The FOC for a brand holding several models indexed j (including m) is

$$s_m + (p_m - c_m) \frac{\partial s_m}{\partial p_m} = -\sum_{j \neq m} (p_j - c_j) \frac{\partial s_j}{\partial p_m}.$$

It is now useful to use X notation with a double subscript such that

$$\frac{\partial s_m}{\partial p_m} = -X_{mm} = -\frac{\sum_h \alpha_h \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})}{N}.$$

$$\frac{\partial s_j}{\partial p_m} = X_{jm} = \frac{\sum_h \alpha_h \mathbb{P}_{mh} \mathbb{P}_{jh}}{N}.$$

The FOC can be re-written

$$\text{FOC} = s_m - (p_m - c_m) X_{mm} + \sum_{j \neq m} (p_j - c_j) X_{jm} = 0.$$

Implicitly differentiating

$$\frac{\partial p_m}{\partial c_m} = -\frac{\partial \text{FOC}/\partial c_m}{\partial \text{FOC}/\partial p_m}.$$

One then simply has to find expressions for those two derivatives of the FOC

$$\frac{\partial \text{FOC}}{\partial c_m} = X_{mm} - \sum_{j \neq m} X_{jm},$$

and

$$\frac{\partial \text{FOC}}{\partial p_m} = -2X_{mm} - (p_m - c_m) \frac{\partial X_{mm}}{\partial p_m} + \sum_{j \neq m} (p_j - c_j) \frac{\partial X_{jm}}{\partial p_m}.$$

$$\frac{\partial X_{mm}}{\partial p_m} = \frac{\sum_h \frac{\partial \mathbb{P}_{mh}}{\partial p_m} \alpha_h (1 - 2\mathbb{P}_{mh})}{N} = -\frac{\sum_h \alpha_h^2 \mathbb{P}_{mh}(1 - \mathbb{P}_{mh})(1 - 2\mathbb{P}_{mh})}{N},$$

and

$$\frac{\partial X_{jm}}{\partial p_m} = \frac{\sum_h \alpha_h \mathbb{P}_{jh} \frac{\partial \mathbb{P}_{mh}}{\partial p_m} + \alpha_h \mathbb{P}_{mh} \frac{\partial \mathbb{P}_{jh}}{\partial p_m}}{N} = \frac{\sum_h \alpha_h^2 \mathbb{P}_{mh} \mathbb{P}_{jh} (2\mathbb{P}_{mh} - 1)}{N},$$

E Mixed CES appendix

E.1 Elasticities

E.1.1 Own-price elasticities

We are interested in reactions of a model's market share to changes in either its own price p_m or the price of a competing model p_j. Under mixed CES demand, the market share of model m simply sums individual probabilities times income over the heterogeneous consumers:

$$s_m = \frac{\sum_h \mathbb{P}_{mh} y_h}{Y}$$

which enables to calculate the own price elasticity as

$$\frac{\partial \ln s_m}{\partial \ln p_m} = \frac{\sum_h \frac{\partial \mathbb{P}_{mh}}{\partial p_m} y_h}{Y} \frac{p_m}{s_m} = \frac{\sum_h \frac{\partial \mathbb{P}_{mh}}{\partial p_m} y_h}{\sum_h \mathbb{P}_{mh} y_h} p_m$$

Since the individual partial effect of a change in p_m is

$$\frac{\partial \mathbb{P}_{mh}}{\partial p_m} = -\frac{\alpha_h}{p_m} \mathbb{P}_{mh}(1 - \mathbb{P}_{mh}),$$

we can rewrite the own price elasticity using notation $\epsilon_m \equiv -\frac{\partial \ln s_m}{\partial \ln p_m}$ (it will prove convenient to write such that own elasticity is a positive number):

$$\epsilon_m = \sum_h \omega_{mh} \alpha_h (1 - \mathbb{P}_{mh}) = \sum_h \omega_{mh} \epsilon_{mh}, \quad \text{with} \quad \omega_{mh} \equiv \frac{\mathbb{P}_{mh} y_h}{\sum_h \mathbb{P}_{mh} y_h}. \tag{40}$$

Model m's own elasticity therefore is a weighted average of the individual elasticities, which write $\epsilon_{mh} = -\frac{\partial \ln \mathbb{P}_{mh}}{\partial \ln p_m} = \alpha_h(1 - \mathbb{P}_{mh})$, where the weight ω_{mh} is the share of each household in total sales of the model. Without consumer heterogeneity $\epsilon_m = \eta(1 - \mathbb{P}_m)$, which goes to η as the market share of individual models goes to zero (the monopolistic competition assumption).

One can study the concavity/convexity of the mixed CES demand function, which will be important for studying how pass-through and markups vary in the cross-section of firms, or following a counterfactual:

$$\frac{\partial \epsilon_m}{\partial p_m} = \sum_h \frac{\partial(\omega_{mh}\epsilon_{mh})}{\partial p_m} = \sum_h \omega_{mh} \frac{\partial \epsilon_{mh}}{\partial p_m} + \epsilon_{mh} \frac{\partial \omega_{mh}}{\partial p_m}. \tag{41}$$

The first term is simply $\omega_{mh} \frac{\partial \epsilon_{mh}}{\partial p_m} = \frac{\omega_{mh}}{p_m} \alpha_h \mathbb{P}_{mh} \epsilon_{mh}$. The second term is a little more complex and involves the derivative of household h share of m sales when raising price p_m:

$$\frac{\partial \omega_{mh}}{\partial p_m} = \frac{\omega_{mh}}{p_m}\left[\sum_k \omega_{mk}\epsilon_{mk} - \epsilon_{mh}\right] = \frac{\omega_{mh}}{p_m}\left[\epsilon_m - \epsilon_{mh}\right].$$

Notice that $\sum_k \omega_{mk}\epsilon_{mk} - \epsilon_{mh}$ is the difference between the weighted average of consumers' elasticities (weight reflecting their importance for m's sales) and the elasticity of h. Without consumer heterogeneity, all elasticities and weights are the same, which brings this term to 0. Rewriting the derivative of own elasticity

$$\frac{\partial \epsilon_m}{\partial p_m} = \sum_h \frac{\omega_{mh}\epsilon_{mh}}{p_m}\left[\alpha_h \mathbb{P}_{mh} + (\epsilon_m - \epsilon_{mh})\right]. \tag{42}$$

Without consumer heterogeneity, we therefore have

$$\frac{\partial \epsilon_m}{\partial p_m} = \frac{\eta^2}{p_m}(1 - \mathbb{P}_m)\mathbb{P}_m > 0.$$

The elasticity of demand rises as the price increase, which means that the demand function is concave. When adding the atomistic firm assumption of monopolistic competition, demand becomes linear in logs, with $\frac{\partial \epsilon_m}{\partial p_m} = 0$.

The mixed CES demand system can become convex when consumers get sufficiently heterogeneous. A household h contributes to convexity when adding a negative term to the sum over consumers. This happens when

$$\alpha_h(1 - 2\mathbb{P}_{mh}) > \epsilon_m.$$

This is more likely for firms with small market shares on very elastic consumers.

We can also express the super elasticity (noted E) in an interesting way:

$$E_m \equiv \frac{\partial \epsilon_m}{\partial p_m}\frac{p_m}{\epsilon_m} = \frac{\sum_h \omega_{mh}\epsilon_{mh}\left[E_{mh} + (\epsilon_m - \epsilon_{mh})\right]}{\epsilon_m}. \tag{43}$$

This is again ambiguous depending on whether for each household the deviation from weighted average elasticity on model m ($\epsilon_m - \epsilon_{mh}$) is larger or smaller than the household-level superelasticity $E_{mh} = \alpha_h \mathbb{P}_{mh}$. With homogeneous consumers, $E_m = \eta \mathbb{P}_m > 0$ as long as firms have more than zero mass.

We can simplify E_m by defining a new weight, $\tilde{\omega}_{mh} \equiv \omega_{mh}\epsilon_{mh}/\epsilon_m$, and using it into equation (43):

$$E_m = \epsilon_m + \sum_h \tilde{\omega}_{mh}(E_{mh} - \epsilon_{mh}) = \epsilon_m + \sum_h \tilde{\omega}_{mh}(2\mathbb{P}_{mh} - 1)\alpha_h. \tag{44}$$

E.1.2 Cross-price elasticities

Let us turn to cross-price elasticities: the impact of an increase in the price of model m on demand for j. Those are given by

$$\frac{\partial \ln s_j}{\partial \ln p_m} = \frac{\sum_h \frac{\partial \mathbb{P}_{jh}}{\partial p_m} y_h}{Y} \frac{p_m}{s_j} = \frac{\sum_h \frac{\partial \mathbb{P}_{jh}}{\partial p_m} y_h}{\sum_h \mathbb{P}_{jh} y_h} p_m$$

The partial effect of j's price on $\mathbb{P}_{jh}$ is

$$\frac{\partial \mathbb{P}_{jh}}{\partial p_m} = \frac{\alpha_h}{p_m} \mathbb{P}_{jh} \mathbb{P}_{mh},$$

which yields

$$\frac{\partial \ln s_j}{\partial \ln p_m} = \sum_h \omega_{jh} \alpha_h \mathbb{P}_{mh} \quad \text{with} \quad \omega_{jh} \equiv \frac{\mathbb{P}_{jh} y_h}{\sum_h \mathbb{P}_{jh} y_h}. \tag{45}$$

Again, this is a weighted average of the individual choice probability cross-elasticities, $\frac{\partial \ln \mathbb{P}_{jh}}{\partial \ln p_m} = \alpha_h \mathbb{P}_{mh}$.

E.2 Pro/anticompetitive effects

At this stage, we can turn to the pro-competitive / anti-competitive nature of the mixed CES model. Recall from equation (5) that the Lerner index expresses markups as

$$L_m = \frac{p_m - c_m}{p_m} = \frac{r_m + 1}{\epsilon_m + 1}, \quad \text{with} \quad r_m = \frac{1}{s_m} \sum_{j \neq m} \Omega_{jm} \frac{\partial \ln s_j}{\partial \ln p_m} L_j s_j. \tag{46}$$

With single product firms, $r_m = 0$ and it is sufficient to know the curvature of demand, i.e. the sign of $\frac{\partial \epsilon_m}{\partial p_m}$ to know if the model is pro or anti-competitive. With MSLD satisfied, the elasticity of demand increases with price, and the high-price, low market share varieties have low markups.

With multiproduct firms, the curvature is not sufficient since

$$\frac{\partial L_m}{\partial p_m} = \frac{\frac{\partial r_m}{\partial p_m}(\epsilon_m + 1) - (r_m + 1)\frac{\partial \epsilon_m}{\partial p_m}}{(\epsilon_m + 1)^2}.$$

E.2.1 CES Elasticities

In the CES version of discrete choice, the Gumbel parameter is common to all households. It is then very easy to calculate the models' market share response, which is equal to the uniform response of the different households:

$$\epsilon_m = -\frac{\partial \ln s_m}{\partial \ln p_m} = \eta(1 - s_m),$$

and

$$\frac{\partial \ln s_m}{\partial \ln p_j} = \eta s_j.$$

The super elasticity is unambiguously positive:

$$E_m = \frac{\partial \epsilon_m}{\partial p_m}\frac{p_m}{\epsilon_m} = \eta s_m > 0$$

E.3 Single product pass-through

The derivation of theoretical pass-through starts from FOC for model m, equation (5) simplified to apply to the single product case:

$$\epsilon_{mm} p_m - (\epsilon_{mm} + 1)c_m = 0 \quad \Leftrightarrow \quad \epsilon_{mm}(p_m - c_m) - c_m = 0,$$

$\epsilon_{mm} > 0$ being the own price elasticity. Implicit differentiation gives

$$\frac{\partial p_m}{\partial c_m} = \frac{\epsilon_{mm} + 1}{\epsilon_{mm} + (p_m - c_m)\frac{\partial \epsilon_{mm}}{\partial p_m}}.$$

Using the first order condition to replace $(p_m - c_m) = c_m/\epsilon_{mm}$, one obtains

$$\frac{\partial p_m}{\partial c_m} = \frac{\epsilon_{mm} + 1}{\epsilon_{mm} + E_{mm}\frac{c_m}{p_m}}, \quad \text{where} \quad E_{mm} \equiv \frac{\partial \ln \epsilon_{mm}}{\partial \ln p_m}.$$

E_{mm} is often referred to as the superelasticity of demand, i.e. the elasticity of own price elasticity with respect to a change in own price. Under homogeneous consumers with CES demand (with elasticity η), combined with monopolistic competition, ϵ_{mm} is a constant. Hence, $E_{mm} = 0$, and the pass-through derivative is a constant equal to $(\eta + 1)/\eta$. Writing pass-through as an elasticity, and using again the FOC to replace $\frac{c_m}{p_m} = \frac{\epsilon_{mm}}{\epsilon_{mm}+1}$

$$\frac{\partial \ln p_m}{\partial \ln c_m} = \frac{\partial p_m}{\partial c_m}\frac{c_m}{p_m} = \frac{\epsilon_{mm}}{\epsilon_{mm} + E_{mm}\frac{\epsilon_{mm}}{\epsilon_{mm}+1}} = \frac{\epsilon_{mm} + 1}{\epsilon_{mm} + 1 + E_{mm}}. \tag{47}$$

The sign of E_{mm} is therefore the determinant of whether pass-through is greater or smaller than one. The superelasticity of demand in the mixed CES case is of ambiguous sign, with expression given by (43).

In the homogenous consumers case of CES demand, however, the sign of the superelasticity is clearly positive. Using the formulas

$$\epsilon_{mm} = \eta(1 - s_m), \quad \text{and} \quad E_{mm} = \eta s_m,$$

and the fact that $\epsilon_{mm} > 0$, we have $E_{mm} > 0$ for the case of oligopoly where market shares are not negligible, which causes the pass-through elasticity to be less than unitary. In the Dixit-Stiglitz case, $E_{mm} = 0$ implies a unitary pass-through elasticity. The formula for CES oligopoly pass-through is

$$\frac{\partial \ln p_m}{\partial \ln c_m} = 1 - \frac{\eta}{\eta + 1} s_m. \tag{48}$$

With $\eta = 5$ for a model with a 10% market share, $E_{mm} = 0.5$, and the pass-through is $1 - \frac{0.5}{6} = 0.92$. A 50% market share with same η gives a pass-through of 0.58.