Brunner, Daniel; Heiss, Florian; Romahn, André; Weiser, Constantin

Working Paper

Reliable estimation of random coefficient logit demand models

DICE Discussion Paper, No. 267

Provided in Cooperation with:
Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf

Suggested Citation: Brunner, Daniel; Heiss, Florian; Romahn, André; Weiser, Constantin (2017) : Reliable estimation of random coefficient logit demand models, DICE Discussion Paper, No. 267, ISBN 978-3-86304-266-0, Heinrich Heine University Düsseldorf, Düsseldorf Institute for Competition Economics (DICE), Düsseldorf

This Version is available at:
https://hdl.handle.net/10419/168359

DISCUSSION PAPER

No 267

Reliable Estimation of Random Coefficient Logit Demand Models

Daniel Brunner,
Florian Heiss,
André Romahn,
Constantin Weiser

September 2017

HEINRICH HEINE
UNIVERSITÄT DÜSSELDORF

d|u|p düsseldorf university press

IMPRINT

DISCUSSION PAPER

Published by

düsseldorf university press (dup) on behalf of
Heinrich-Heine-Universität Düsseldorf, Faculty of Economics,
Düsseldorf Institute for Competition Economics (DICE), Universitätsstraße 1,
40225 Düsseldorf, Germany
www.dice.hhu.de

Editor:

Prof. Dr. Hans-Theo Normann
Düsseldorf Institute for Competition Economics (DICE)
Phone: +49(0) 211-81-15125, e-mail: normann@dice.hhu.de

DISCUSSION PAPER

ISSN 2190-9938 (online) – ISBN 978-3-86304-266-0

Reliable Estimation of Random Coefficient Logit Demand Models *

Daniel Brunner[†] Florian Heiss[‡] André Romahn[§]
Constantin Weiser[¶]

September 2017

Abstract

The differentiated demand model of Berry, Levinsohn and Pakes (1995) is widely used in empirical economic research. Previous literature has demonstrated numerical instabilities of the corresponding GMM estimator that give a wide range of parameter estimates and economic implications depending on technical details such as the choice of optimization algorithm, starting values, and convergence criteria. We show that these instabilities are mainly driven by numerical approximation errors of the moment function which is not analytically available. With accurate approximation, the estimator is well-behaved. We also discuss approaches to mitigate the computational burden of accurate approximation and provide code for download.

*Financial support was provided by the DFG. Computational support and infrastructure were provided by the "Centre for Information and Media Technology" (ZIM) at the University of Düsseldorf (Germany) and the Flemish Supercomputing Center (VSC) at the University of Leuven (Belgium). We would like to thank Richard Friberg, Mathias Reynaert, Joel Stiebale, Yutec Sun, Frank Verboven and seminar participants at EARIE 2017, the 2017 China Meeting of the Econometric Society, DICE, Télécom ParisTech, Tilburg University and the University of Cologne for helpful comments and suggestions. All remaining errors are our own.

[†]daniel.brunner@hhu.de; Heinrich-Heine-Universität Düsseldorf
[‡]florian.heiss@hhu.de; Heinrich-Heine-Universität Düsseldorf
[§]romahn@hhu.de; Heinrich-Heine-Universität Düsseldorf
[¶]constantin.weiser@uni-mainz.de, Johannes Gutenberg Universität Mainz

1 Introduction

The seminal contribution of Berry, Levinsohn and Pakes (1995; henceforth BLP) has provided economists with an oligopoly model of differentiated product markets that is capable of producing realistic substitution patterns. The BLP model allows for partially or fully unobserved preference heterogeneity among economic agents and explicitly deals with the endogeneity of product attributes, typically price, and makes it possible to investigate counterfactual market outcomes. Apart from being applied to markets and questions that lie at the heart of Industrial Organization, the model's use has spread to the areas of environmental economics, insurance, voting preferences, and housing markets among others (see Table 1 in Berry and Haile (2014)). The BLP model is parsimonious - compared to the standard logit model only a limited number of additional coefficients must be estimated - and its flexible functional form allows for arbitrary correlations between prices and markups. Products with similar attributes can be closer substitutes than products with very different characteristics.[1]

Consistent identification of the preference parameters depends on the sample moments, which are the product of relevant and valid instrumental variables (IVs) and the BLP model's structural error term. To obtain estimates of the structural error, the observed aggregate market shares have to be inverted. In contrast to the standard logit or nested logit model, where this inversion can be performed analytically (Berry, 1994), in the BLP model it must be computed numerically.

Using the U.S. automobile market data from BLP, Knittel and Metaxoglou (2014; henceforth KM) re-examine the nested fixed point estimator's behavior and find a very wide set of parameter estimates and economic implications. The particular choice of starting guess, optimization algorithm and inner convergence threshold has a substantial effect on the estimation outcomes. Given the large spread of the estimates, the credibility of the approach is drawn into question. This Knittel-Metaxoglou critique has urged researchers to more carefully implement the BLP estimation framework and to more transparently report implementation details and estimation results (see the supplemental appendix of Goldberg and Hellerstein (2013) for an example).

KM use the same set of 50 standard Monte Carlo draws throughout all of their estimations and thereby ignore the impact of simulation error. With the exact same data, instruments and starting guesses we show that their findings are fully explained by the low

[1] In the standard logit model prices and markups are negatively related: high-priced products have lower markups than their low-priced rivals. Moreover, it is highly likely that the best substitute for any other product is the product with the largest market share.

number of Monte Carlo draws used, which causes substantial simulation error.[2] Our main findings are twofold. First, with a crude numerical integration approach the estimate of the structural error term is overwhelmed by simulation error. The biased structural error estimates enter the GMM-IV objective function, which causes many local minima with widely varying parameter estimates and model-implied economic predictions. With a large number of simulation draws we obtain tightly clustered model estimates and economic predictions. Second, inaccurate numerical integration gives biased parameter estimates and economic implications. With a relatively low number of simulation draws we find that the magnitude of own-price elasticities is systematically estimated to be too high. In a simulation of a merger between GM and Chrysler this leads to a substantial underestimation of the merger's effect on prices. The estimated demand model therefore systematically and erroneously predicts the merger to cause less harm to consumers than when a large number of draws is used. Our findings should therefore also be relevant for competition policy practitioners.

Broadly, our results contribute to the body of literature that shows how accurate computational methods can be crucial for obtaining reliable results from the estimation of nonlinear econometric models. Such models pose two major difficulties. First, a highly nonlinear objective function can produce many candidate extreme points. Depending on the identification approach, the econometrician must identify either the global maximum or minimum. Second, it can be numerically challenging to compute the estimator's objective function or the economic model's moments with sufficient accuracy to reliably pin down the sought after coefficients. McCullough and Vinod (2003) illustrate the importance of carefully verifying the candidate extreme points from a nonlinear solver. Petrosky-Nadeau and Zhang (2017) show how inaccurate computation of the Diamond-Mortensen-Pissarides model, the workhorse approach for general equilibrium labor market models, produces biased moments. The results of KM suggest that both issues are relevant for the estimation of BLP models, while we show that the number and spread of the estimator's minima fall into an increasingly narrowing interval when the objective function is computed accurately.

In this sense, our findings are related to weak identification in nonlinear GMM-IV estimation as described in Stock, Wright, and Yogo (2002). Here, the weak identification is not caused by weak instruments, however, but by random simulation error in the estimates of the structural error terms, which introduces many local minima in the estimator's GMM-IV objective function. Moreover, our findings are in line with the results from Berry, Linton, and Pakes (2004), who derive the properties of the BLP nested fixed point estimator when

[2]The replication files of KM are exemplary.

the number of products becomes large. Simulation error in the estimates is bounded if and only if the number of simulation draws grows proportionally with the square of the number of products. The impact of simulation error is therefore more pronounced in samples with many products, which applies to the U.S. automobile data with markets having between 72 and 150 products. For the U.S. automobile data, we find the estimator no longer fails to converge to a local minimum for any of the random starting guesses if we use at least 5,000 Monte Carlo simulation draws to compute the aggregate market share inversion. With our least precise numerical integration approach, which as in KM uses 50 Monte Carlo draws, we obtain convergence to a local minimum in less than 63 percent of the estimations and a coefficient of variation among the objective values of the identified minima of more than 30 percent. In contrast, with our most accurate integration approach, which uses 10,000 modified latin hypercube sampling draws (MLHS draws, Hess et al. (2006)) we obtain convergence to a local minimum for 100 percent of the estimations and a coefficient of variation of less than 3 percent. This tight clustering of the identified minima also carries over to the parameter estimates and the model-implied economic predictions. To illustrate, with 50 Monte Carlo simulation draws the 95 percent confidence interval for the average own-price elasticity across all observations in the automobile data ranges from roughly -24 to -3. When 10,000 MLHS draws are used instead, the confidence interval tightens to the range from roughly -9 to -8. Our results are based on a total of 40,000 BLP model estimations, where we use two numerical integration techniques, standard Monte Carlo and MLHS draws. For each of these approaches we consider eight different numbers of simulation draws that range from 50 to 10,000. For each number of draws, 50 independently sampled sets are generated and for each of these sets the BLP model is estimated 50 times using the same specification and random starting guesses as in KM.

The importance of simulation error has largely been abstracted from in the existing literature. Dubé, Fox, and Su (2012) also identify the contraction mapping as the major source of numerical instabilities, but focus on the convergence threshold of the contraction that is set by the researcher and explicitly shut down the effect of simulation error.[3] A loose threshold speeds up the estimation, but also introduces approximation error in the objective function. In qualitative terms, we can confirm their findings, but in our setting with real world data we find the impact of the convergence threshold to be of second order (see the

[3]See p. 2263 in Appendix A of Dubé, Fox, and Su (2012): "...Because our focus is not on numerical integration error, we use the same sample of 1000 draws to compute the market shares in the data-generation and estimation phases."

bottom panel of Table 5).[4] Reynaert and Verboven (2014) show that approximately optimal instruments can substantially reduce weak identification in BLP model estimation that is caused by weak instruments. Again, the impact of simulation error is shut down, because highly accurate numerical integration approaches are used throughout the simulations and only markets with at most 20 products are considered.

A brute-force reduction of the approximation errors by increasing the number of simulation draws can be computationally costly up to a point where it seems infeasible to implement. We therefore also discuss how to increase the approximation accuracy in a computationally efficient way and provide an implementations in the form of an R package and Matlab code that offer substantial speedups over KM's Matlab implementation of the nested fixed point estimator.

The remainder of the paper is organized as follows. Section 2 briefly reviews the BLP model and its identification using the nested fixed point algorithm. It also theoretically shows how simulation error propagates in the GMM-IV sample moment and objective function using results from Berry, Linton, and Pakes (2004). Section 3 presents the setup for our large-scale study of the BLP estimator using the U.S. automobile data. We trace out the effects of simulation error in our 40,000 BLP model estimations in Section 4 and discuss the trade-off between the computational burden and the accuracy of the estimation and suggest ways to improve the computational efficiency in Section 5. There we also benchmark our R and Matlab programs and show that they are several times faster than the routine in KM's replication files. Finally, we conclude.

2 The BLP model

This section briefly presents the BLP model and its estimation using the nested fixed point algorithm. We also discuss the propagation of the simulation error in the estimator's moment function.

2.1 Setup and model-implied market shares

Each consumer in a market for differentiated products faces the discrete choice between the alternatives labeled by $j = 0, ..., J$, where $j = 0$ indicates the outside good. Typically, we model several markets jointly. For notational simplicity, we suppress an index for the

[4] We consider convergence thresholds of 10^{-16}, 10^{-9} and 10^{-4}. Only the latter, extremely lax criterion, yields a noticeably wider spread in the model estimates.

market on all relevant variables. Since the level of utility is not separately identified, the indirect utility each consumer attaches to the outside option is normalized to zero. Consumer i's indirect utility from purchasing product j is specified as a function of its price p_j, K observed product characteristics collected in the vector x_j, and the valuation of unobserved characteristics ξ_j. The average utility over consumers is specified as

$$\delta_j = \delta_j(\xi_j) = x_j\beta - \alpha p_j + \xi_j. \tag{1}$$

The BLP model also allows for preference heterogeneity over the K characteristics in the population. It is captured by the vector $\nu_i = [\nu_{i1}, \ldots, \nu_{iK}]$. For the model's exposition and to offer a meaningful comparison with the results in Knittel and Metaxoglou (2014), we assume that the K dimensions of ν_i are independently distributed. This assumption can be relaxed and preference correlations between the K characteristics can be modeled to achieve more flexible substitution patterns. In our simpler case, the only parameters that capture heterogeneity are the standard deviations of the preference parameters $\theta = [\sigma_1, \ldots, \sigma_K]$.[5] The consumer-specific deviation from mean utility is defined as

$$\mu_{ij} = \mu_j(\theta, \nu_i) = p_j\sigma_p\nu_{ip} + \sum_{k=1}^{K} x_{jk}\nu_{ik}\sigma_k \tag{2}$$

The overall utility also includes consumer-product specific utility residuals ε_{ij} and can be written as

$$u_{ij} = \delta_j + \mu_{ij} + \varepsilon_{ij}. \tag{3}$$

Making the typical assumption that the ε_{ij} are independent and follow a Type I extreme value distribution gives the closed form expressions for the consumer-specific choice probabilities for product j

$$Pr_j(\delta, \theta, \nu_i) = \frac{\exp(\delta_j + \mu_j(\theta, \nu_i))}{1 + \sum_{\ell=1}^{J} \exp(\delta_\ell + \mu_\ell(\theta, \nu_i))}, \tag{4}$$

where $\delta = [\delta_1, \ldots, \delta_J]$. The model-implied aggregate market share function integrates over the consumer-specific choice probabilities, where we let $F(\nu)$ denote the population distribution of consumer heterogeneity.

$$s_j^*(\delta, \theta) = \int Pr_j(\delta, \theta, \nu) dF(\nu) = \int \frac{\exp(\delta_j + \mu_j(\theta, \nu))}{1 + \sum_{\ell=1}^{J} \exp(\delta_\ell + \mu_\ell(\theta, \nu))} dF(\nu) \tag{5}$$

[5]Note that in practice, often one or more of the preference parameters are restricted to be constant across the population which can be implemented in this notation by restricting the respective σ_k parameters to zero.

This integral does not have an analytic solution, but can be approximated numerically. We have to deal with the fact that the population distribution of consumer preferences is not directly observed by the econometrician.[6] We therefore have to assume a joint distribution of preferences over the K characteristics. A common assumption is a joint normal distribution. The most straightforward algorithm for this approximation is Monte Carlo simulation. We draw a sample $\tilde{\nu} = [\tilde{\nu}_1, \ldots, \tilde{\nu}_R]$ of size R from the joint distribution of ν. The approximated version of $s_j^*(\delta, \theta)$ is

$$s_j(\delta, \theta, \tilde{\nu}) = \frac{1}{R} \sum_{r=1}^{R} Pr_j(\delta, \theta, \tilde{\nu}_r) = \frac{1}{R} \sum_{r=1}^{R} \frac{\exp(\delta_j + \mu_j(\theta, \tilde{\nu}_r))}{1 + \sum_{\ell=1}^{J} \exp(\delta_\ell + \mu_\ell(\theta, \tilde{\nu}_r))}. \tag{6}$$

2.2 Instrumental variables and identification

The unobserved characteristics or structural error terms ξ_j are vertical product attributes. Consumer utility for product j is increasing in ξ_j, so that consumers always prefer more of it. Contrary to the econometrician, both firms and consumers observe all $\xi = [\xi_1, \ldots, \xi_J]$, which yields positive correlations between the error term and price. We obtain consistent estimates of the preference parameters by imposing a standard GMM-IV moment restriction. Let z_j denote a row vector of $L \geq K$ relevant and valid instrumental variables. The moment restriction is

$$E\left[\mathcal{G}(\theta)\right] = E\left[\frac{1}{J} \sum_{j=1}^{J} z_j \xi_j(\theta)\right] = 0. \tag{7}$$

In the typical case where we model several markets jointly, we average over all available products in all markets. Note that the ξ_j are by definition unobserved in the data. Here, $\xi_j(\theta)$ denotes the implied values as detailed in the next section.

Cost shifters that vary at the product level would be ideal candidates for the excluded instruments. The required data, however, is often not available. To construct suitable instruments we make the assumption that the ξ's are mean independent of the observed product characteristics.

$$E\left(\xi | x\right) = 0 \tag{8}$$

If this assumption holds, any function of the observed product characteristics qualifies as a potentially valid instrument for price. Such functions also give relevant instruments, because the observed characteristics enter each product's equilibrium pricing function. BLP use this

[6] Depending on data availability, consumer heterogeneity can be partially directly observed by using (relevant) consumer demographics. This introduces an additional term that enters utility additively separably and that interacts the observable product attributes and consumer demographics.

insight to derive a set of instruments that can be viewed as a first-order approximation of a pricing game played between firms: for each product j sold by firm f the observable characteristics of all products sold by the same firm are summed over and the observable characteristics of all products sold by rival firms are summed over, $z_{j,own} = \sum_{k \neq j, k \in \mathcal{F}_f} x_k$, $z_{j,other} = \sum_{k, k \notin \mathcal{F}_f} x_k$. This gives $2K_1 \geq K_2 + 1$ excluded instruments to identify the price coefficient α, and the standard deviations of the random coefficients θ.

Let θ^* denote the true population preference parameters. Given a suitable weighting matrix W, we obtain a consistent and, as Berry, Linton, and Pakes (2004) prove, asymptotically normally distributed estimator of θ^* by minimizing the GMM-IV objective function, which is a norm of the sample moment $\frac{1}{J}\sum_j z_j \xi_j(\theta)$.[7]

$$\widehat{\theta}_2 = \arg\min_{\theta} \mathcal{J}(\theta) = \arg\min_{\theta} \xi(\theta)' z W z' \xi(\theta) \tag{9}$$

Here, $\xi(\theta)$ and z are the vertically stacked market-specific structural error terms and instrument matrices, respectively.

2.3 The fixed-point algorithm for obtaining $\xi(\theta)$

The objective function (9) involves the vector of unobserved product characteristics $\xi(\theta)$ which need to be evaluated numerically for a given set of parameters. To simplify the notation in this section, we abstract from any sampling errors and assume throughout that market shares are observed without error.

For each candidate vector of the nonlinearly entering preference parameters θ, the BLP model chooses the values of the product-specific mean utility $\delta = [\delta_1, \ldots, \delta_J]$ such that for each product j, the model-implied approximated market share matches the observed share in the data $\mathcal{S}_j$.[8]

$$\mathcal{S}_j = s_j(\delta, \theta, \tilde{\nu}) \quad \text{for all } j \tag{10}$$

We cannot solve these equations for δ analytically but have to resort to numerical methods. BLP prove the existence of a fixed point that gives the unique solution δ_{match} to this system of equations for any candidate vector θ and set of simulation draws $\tilde{\nu}$. We iterate over the equation

$$\delta_{j,iter+1} = \delta_{j,iter} + \log(\mathcal{S}_j) - \log(s_j(\delta_{iter}, \tilde{\nu}, \theta)) \tag{11}$$

[7] Without loss of generality, we assume that the weighting matrix is homoscedastic, $W = (z'z)^{-1}$.

[8] The magnitude of the relative approximation error, $|(\mathcal{S}_j - s_j)/\mathcal{S}_j|$, is bounded from above by the inner convergence threshold, γ_{inner}.

until the distance between successive iterates falls below the chosen convergence threshold, $|\delta_{j,iter+1} - \delta_{j,iter}| \leq \gamma_{inner}$ for all j. If this inequality holds, the current update for the vector of mean utilities is accepted as the solution $\delta_{match}(\theta, \tilde{\nu})$ to (10). We obtain the corresponding vector of structural error terms $\xi_{match}(\theta, \tilde{\nu})$ as the residuals of a two-stage least squares regression of $\delta_{match}(\theta, \tilde{\nu})$ on the observed product characteristics. This step also delivers the estimates of the linearly entering parameters α and β.

$$\delta_{j,match}(\theta, \tilde{\nu}) = x_j \hat{\beta} - \hat{\alpha} p_j + \xi_{j,match}(\theta, \tilde{\nu}) \tag{12}$$

2.4 Simulation errors and their propagation

At each iteration of the contraction mapping, (11), the model-implied aggregate market shares must be computed using (6). Simulation error is introduced because of differences between the consumer population and the simulated sample of consumer preferences using R simulation draws.[9]

In order to define the simulation errors similar to Berry, Linton, and Pakes (2004), let $\delta^*_{match}(\theta)$ and $\xi^*_{match}(\theta)$ denote the solutions to $\mathcal{S}_j = s^*_j(\delta, \theta)$ in the absence of simulation errors in the market shares. Simulation error is defined as

$$e(\theta, \tilde{\nu}) \equiv s^*\big(\delta^*_{match}(\theta), \theta\big) - s\big(\delta_{match}(\theta, \tilde{\nu}), \theta, \tilde{\nu}\big). \tag{13}$$

By construction, the simulation errors would vanish for any candidate parameter vector θ if we were able to solve the market share integral (5) exactly. Berry, Linton, and Pakes (2004), show that $\xi_{match}(\theta, \tilde{\nu})$ can approximately be written as

$$\xi_{match}(\theta, \tilde{\nu}) \approx \underbrace{\xi^*_{match}(\theta)}_{\text{exact inversion}} - \underbrace{\left[\left. \frac{\partial s^*(\delta(\xi), \theta^*)}{\partial \xi'} \right|_{\xi^*} \right]^{-1} e(\theta, \tilde{\nu})}_{\text{effect of simulation error}}, \tag{14}$$

where θ^* and ξ^* denote the true population values of θ and ξ, respectively. The expression $[\partial s^*/\partial \xi']$ is the $J \times J$ matrix of market share derivatives with respect to the unobservable product characteristics.

The first term on the right-hand side is the estimate of the structural errors that we would obtain if we could match the population distribution of preference heterogeneity exactly in

[9] We ignore an additional potential error here. We impose that consumer preference heterogeneity follows a normal distribution. The true preference distribution could be non-normal, which would potentially lead to biased parameter estimates.

the numerical integration of the aggregate shares. We only use a sample of R simulation draws, however, which causes deviations of the model-implied shares from their observed sample counterparts at ξ^*. How these deviations affect our computations of ξ_{match} depends on how sensitive ξ_{match} is with respect to changes in the entries of the model-implied aggregate market share vector s^*. This sensitivity is measured by the inverse of the matrix $[\partial s^*/\partial \xi']$. The smaller the derivatives, the larger is the distortion of ξ_{match} that is caused by simulation error. Thus, the inversion of aggregate market shares magnifies simulation error in the estimates of the structural error term.

By distorting the estimates of ξ, simulation error propagates in the sample moments and thereby in the GMM-IV objective function. Plugging (14) into the sample moment gives

$$\mathcal{G}(\theta, \tilde{\nu}) \approx \frac{1}{J} \sum_{j=1}^{J} z_j \left(\xi^*_{match}(\theta) - \left[\frac{\partial s^*(\delta(\xi), \theta^*)}{\partial \xi'} \bigg|_{\xi^*} \right]^{-1} e(\theta, \tilde{\nu}) \right), \tag{15}$$

which stresses that the computed sample moment depends explicitly on the simulation error that is caused by the specific set of draws $\tilde{\nu}$. There is an analogy to the definition of weak identification in Stock, Wright, and Yogo (2002) for nonlinear GMM estimation. Given that e^R is random and propagates into the GMM-IV objective function, the shape and location of the objective function (9), are affected. There can be several values for $\theta \neq \theta^*$ for which the objective function attains a local minimum. This explains how many local minima with widely varying parameter estimates and economic implications are found with a crude numerical integration approach.

Berry, Linton, and Pakes (2004) also show that the extent of the magnification depends on the number of products in the market. In equilibrium, as more and more products enter a market, it must be the case that product-level market shares fall. This is because in the BLP model, each product is substitutable with every other product to some extent. Specifically, it is assumed that all shares move inversely proportional with J (Condition S/equation (20) in Berry, Linton, and Pakes (2004)). The derivatives of the shares with respect to ξ are proportional to market shares and therefore also decline with J. As simulation errors are scaled by the inverse of $[\partial s^*/\partial \xi']$, the magnification of simulation errors is greater in samples with many products. To bound the simulation errors as the number of products becomes large, the number of simulation draws must grow proportionally with the square of the number of products in the market.

3 Computational experiments: Setup

We study how numerical integration accuracy affects the behavior and outcomes of the BLP nested fixed point estimation algorithm using the original automobile market data from BLP. This data set covers 20 years of annually aggregated car model-level sales for the United States starting in 1971.[10] We think this choice presents two advantages. First, this is a real world data set where the number of products ranges from 72 to 150 and that is based on a large sample of individual consumer purchases. Sampling error, therefore, is likely to be negligible, while simulation error should play a substantial role in this setting. Second, the same data set has been used by KM to carefully document several numerical instabilities in the BLP estimation algorithm. The study is exemplary in terms of its replicability and transparency and has motivated researchers to more carefully implement and report the outcomes of their BLP model estimations (e.g. Goldberg and Hellerstein (2013)). We therefore base our large-scale study of the BLP estimation algorithm on KM's replication files to demonstrate that the reported numerical instabilities are tackled once the numerical integration of the model-implied aggregate shares is performed accurately. Specifically, we estimate exactly the same specification using the same set of instruments and random starting guesses for θ.

3.1 Model specification

In the automobile market application, we follow the literature and interpret the yearly data as separate markets $t = 1, \ldots, 20$. The indirect utility of consumer i in year t for car j is specified as

$$u_{ijt} = \beta_{i0} + hpwt_{jt}\beta_{i1} + space_{jt}\beta_2 + aircon_{jt}\beta_{i3} + mpg_{jt}\beta_{i4} - \alpha_i price_{jt} + \xi_{jt} + \varepsilon_{ijt}, \quad (16)$$

where *hpwt* is the horsepower-weight ratio, *space* is the length times the width of the car, *aircon* is a dummy indicating whether the car has air conditioning built in and *mpg* measures the car's miles per gallon. Except for *space*, all observable characteristics, including price and the constant term, have a random coefficient. The specification therefore involves 5 random coefficients in total. We assume that the random coefficients are distributed normally and independently. Thus, $\alpha_i = \alpha + \sigma_p \nu_{i,p}$ and $\beta_{i,k} = \beta_k + \sigma_k \nu_{i,k}$ with $\nu_{i,k}, \nu_{i,p} \sim N(0,1)$ for $k = 1, \ldots, K_2 = 5$.

[10] For a detailed description of the data set, see Berry, Levinsohn, and Pakes (1995).

3.2 Instruments

We use the instruments from the Knittel and Metaxoglou (2014) replication files. These are the standard characteristics-based or BLP-type instruments. Using all five non-price product characteristics including the constant, these instruments sum over the characteristics of all other cars produced by the same firm, and sum over the characteristics of all cars produced by rival firms. We therefore have 10 instruments for price and the 5 nonlinearly entering parameters. Given that the literature on approximately optimal instruments shows that these standard characteristics-based instruments can be weak and thereby yield weak identification of the random coefficients, it is important to show that for specification (16) this is not the case so that we are dealing with a sensible model.

We simply run the first-stage regression of price on the instruments for two cases. First, we only explain the variation in price using the excluded (BLP-type) instruments. This regression gives an F-statistic of 43.9. Second, we use the full instrumental variable matrix that also contains the observed non-price characteristics, which given their assumed exogeneity instrument for themselves. Not surprisingly, this gives a higher F-statistic, namely roughly 248. To assess whether the observed characteristics drive out the excluded instruments, we compute the F-statistic for the null that only the coefficients of the excluded instruments are zero. This F-statistic has a value of 43.7, almost unchanged from the first-stage regression without the observed characteristics. In both cases, we comfortably pass the rule of thumb that the F-statistic should be greater than 10. The excluded instruments also comfortably pass the critical values reported in table 1 of Stock, Wright, and Yogo (2002). We conclude that the example model is well identified and we don't have to worry about weak instruments.

3.3 Simulation of the market shares

Different algorithms for the approximation of the integral in the market share equation (5) have been proposed in the literature. Since this paper focuses on the effect of approximation errors rather than on ways to avoid them, we restrict ourselves to two popular simulation methods: the standard Monte Carlo approach and modified latin hypercube sampling draws (MLHS draws). We come back to alternative approximation methods in Section 5.3. Hess, Train, and Polak (2006) find that in finite samples MLHS draws perform roughly on par with Halton draws. For our study, MLHS draws offer the advantage that it is straightforward to obtain measures of how the number of simulation draws affects the spread of estimation outcomes. We can simply compute the variance of some estimation outcome for a given

number of simulation draws. With standard Halton draws or any quadrature method this is no longer the case, because for these approaches the simulation draws or nodes are based on deterministic number sequences. By construction, therefore, for a given number of draws or nodes there is no variation across different estimations. To obtain a measure of simulation error in the estimation outcomes, we would have to compute error bounds for these methods, which are model-specific and cumbersome to implement.

We use 8 different numbers of draws for both simulation approaches that range from 50 to 10,000.[11] To exclude the possibility that our findings are due to any specific set of draws, we generate 50 independently sampled sets of ν for each of the 8 different numbers of draws. Therefore, with the 50 starting guesses for θ from KM, each number of draws requires us to estimate specification (16) 2,500 times. With 8 different numbers of simulation draws and 2 simulation approaches, we estimate the BLP model 40,000 times.

3.4 Optimization algorithms and inner convergence threshold

An important part of the Knittel-Metaxoglou critique is that the choice of optimization algorithm can have a substantial effect on the estimation outcomes. Similarly, Dubé, Fox, and Su (2012) caution that a loose inner convergence threshold can produce many local minima with widely varying estimates. We investigate both of these aspects in our setting and with an accurate numerical integration approach we find the choice of optimization algorithm to be irrelevant (see the top panel of Table 5) and the impact of the inner convergence threshold to be of second order (see the bottom panel of Table 5). We therefore base all of our 40,000 estimations in the main part of our study on a trust region optimizer with an analytical gradient[12] and on a stringent inner convergence threshold of 10^{-16}.

3.5 Benchmark comparison and additional computational details

We deviate from the implementation of the nested fixed point algorithm in some aspects from KM. The changes that we implement make the algorithm more robust and enforce a uniform convergence threshold for the market share inversion throughout. Specifically, KM follow the original code of Nevo (2000), which assigns very high but computable values to the objective function and analytical gradient if a specific parameter value results in numerical overflow. This issue can be easily avoided by rescaling price. We simply divide price by its standard

[11] These numbers are 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000.

[12] Specifically, we use Matlab's `fminunc` optimizer algorithm. This corresponds to KM's DER1-QN1 optimizer.

Table 1: Estimated Random Coefficients Using KM's 50 Monte Carlo Draws

	Min 1	Min 2	Min 3	Min 4	Min 5
price	0.328**	0.182	0.162	0.107**	0.134**
constant	7.480**	2.720**	5.232**	2.001	1.598**
hpwt	2.565	1.063	0.165	5.781**	1.481
aircon	8.800**	0.484	3.629	0.425	4.231**
mpg	0.098	0.687	0.134	1.767**	1.163**
$\mathcal{J}(\widehat{\theta}_2)$	207.7	215.1	216.0	224.6	226.9
$\overline{\eta}_{jj}$	-10.53	-7.782	-5.787	-4.606	-5.387
$\boldsymbol{\eta}$	-1.007	-1.374	-0.946	-0.945	-1.263
Wald-statistic	23.26	87.49	72.69	112.4	93.74

Note: * and ** indicate statistical significance at the 95 and 99 percent confidence levels, respectively. Only the estimated random coefficients, θ, are shown. All inputs to the estimation, including the 50 simulated draws for consumer preference heterogeneity, ν_{KM}, are identical to those used by KM. We compute HAC standard errors. $\overline{\eta}_{jj}$ is the average own-price elasticity and $\boldsymbol{\eta}$ is the aggregate demand elasticity averaged over all 20 markets. The null hypothesis of the Wald test is $\theta = 0$.

deviation. With this rescaling we have never had to contend with overflow problems in our 40,000 BLP model estimations. Moreover, in KM's "loose" implementation of the estimation algorithm, the convergence tolerance in the nested fixed point is dynamically adjusted. When successive iterates of (11) are close to each other, the convergence threshold is set to 10^{-9}. If this is not the case, the threshold is set at 10^{-6}. This dynamic adjustment was originally implemented by Nevo (2000) to reduce the computational burden of the estimation. Given that Dubé, Fox, and Su (2012) show that a loose convergence threshold is an additional source of numerical error and given that computational power has increased dramatically over the last two decades, we enforce a uniform convergence threshold of 10^{-16} throughout.

Using the same set of 50 Monte Carlo draws as KM, we demonstrate that our changes do not fundamentally impact the Knittel-Metaxoglou critique at this level of numerical integration accuracy. Table 1 presents the results of estimating specification (16) for 50 different starting values.

We find that 44 of the 50 random starting guesses for θ yield a local minimum. Rounding the objective function values of these minima to two digits, we obtain 5 minima that range between 207.72 and 226.94. This is a more narrow range than that reported by KM. This indicates that at least some of the lack of robustness in their estimation results could stem from scaling issues, which we avoid.[13] Overall however, the Knittel-Metaxoglou critique is

[13]KM's high cutoff of 30 for the Euclidean norm of the gradient is likely to contribute to a wider range of

broadly reaffirmed. For each random coefficient, the ratio of its largest to smallest point estimate across the 50 starting guesses is at least 3 (price) and reaches up to 35 (hpwt). The model-implied average own-price elasticity and the aggregate demand elasticity vary by factors of roughly 2.3 and 1.4, respectively. Moreover, the statistical significance of individual random coefficients changes substantially across minima. In fact, for each coefficient, it is possible to select a minimum where that coefficient is either statistically significant or insignificant at the 95 percent confidence interval. Finally, the Wald statistic we obtain by testing the estimated BLP model against the simple logit model also ranges widely from 23 to 112.

3.6 Verifying Candidate Minima of the Objective Function

We use two criteria to assess whether the output of the optimization algorithm delivers a minimum. First, it must be the case that all the eigenvalues of the Hessian at the estimated coefficient vector, $\widehat{\theta}_2$, are strictly positive. Second, the gradient must be sufficiently close to zero. The definition of sufficiently close to zero is arbitrary to some extent. We adopt a cutoff of 0.1 for the Euclidean norm of the gradient at $\widehat{\theta}_2$. Our qualitative results are robust to either tightening or relaxing this cutoff. This cutoff is substantially more stringent than the cutoff of 30 that is adopted by KM.

4 Main results

We present the outcomes of the 40,000 BLP model estimations in two parts. First, we demonstrate how simulation error propagates in the GMM-IV sample moments and thereby in the objective function of the estimator. This propagation explains the numerical instabilities documented by KM. Moreover, simulation error can be reduced substantially by increasing the number of simulation draws and thereby raising the accuracy of numerical integration.

Second, we document how the mean and spread of the estimation outcomes and the corresponding economic predictions change with the number of simulation draws. Beyond 500 draws, the spread of the estimation outcomes is falling monotonically in the number of simulation draws for both integration approaches. The estimated parameters and economic implications fall into increasingly narrowing intervals. With regards to the mean of the estimation outcomes, our findings show that simulation error biases the estimation outcomes

outcomes, too.

Figure 1: Empirical Distribution of $\widehat{\xi}$ for Selected Products

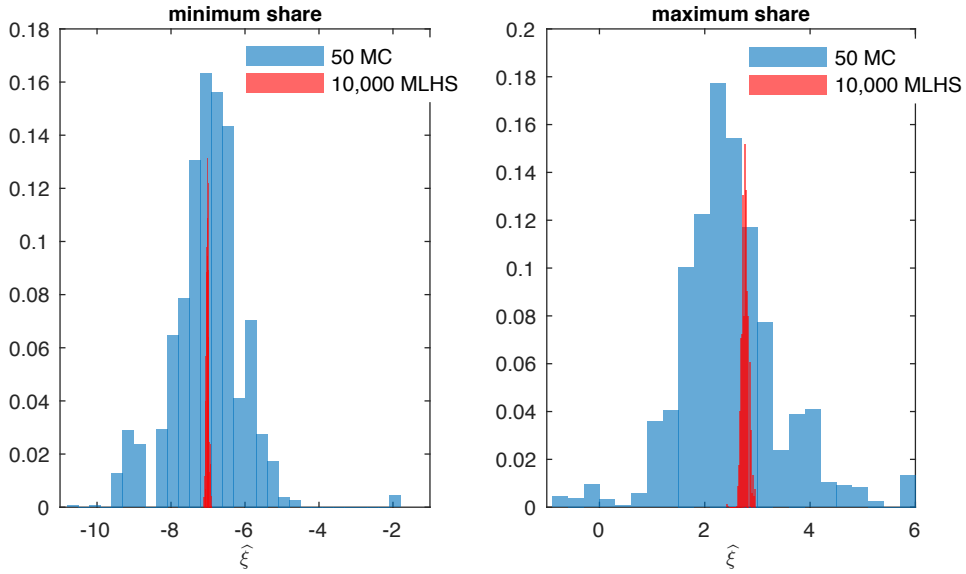

Note: Both panels show the empirical relative frequency plots for the estimated structural error term across 2,500 estimations of the BLP model for a given number of simulation draws. The 2,500 estimations are based on 50 independently sampled sets of preference heterogeneity for a given number of draws. For each of these 50 sets, we estimate the BLP model using 50 random starting guesses. With 50 MC draws, only 1,562 of the 2,500 estimation runs converge to a local minimum. For the 10,000 MLHS draws, all estimations converge. Estimations that fail to converge are not included in the plots.

in the sample of U.S. automobile market data. Thus, as the number of simulation draws changes, so do the means of the estimation outcomes.

4.1 Simulation error in the structural error term

Simulation error propagates in the estimates of the structural error term. Figure 1 shows how the number of simulation draws, which is inversely related to the magnitude of simulation error, affects the estimates of the structural error term. Both panels plot the empirical distributions of the estimated unobservable characteristic for the products with the smallest and largest market shares in the sample. These distributions are based on our least accurate numerical integration approach, namely 50 Monte Carlo draws (blue), and our most accurate approach, 10,000 MLHS draws (red). For each of these integration approaches the model is estimated 2,500 times and each estimation that converges to a local minimum gives us one estimate of the structural error.

The differences between the distributions are remarkable. Using only 50 Monte Carlo

draws, the variances of the estimated structural errors are 0.843 and 0.952 for the products with the smallest and greatest market shares in the sample, respectively. If we use 10,000 MLHS draws, instead, we obtain corresponding variances of only 0.001 and 0.004. In terms of 99 percent confidence intervals, with 50 Monte Carlo draws, the estimate of the unobservable attribute for the products with the smallest and greatest shares are, respectively, the ranges from -9.5 to -4.8 and -0.6 to 5.8. Using our most accurate numerical integration approach gives the corresponding confidence intervals of -7.1 to -6.9 and 2.6 to 3.0. Adopting a crude integration approach, therefore, produces simulation error that easily overwhelms the estimates of the error terms. This holds across the sample. We obtain qualitatively identical figures for the products with the mean and median market shares, for example. Thus, simulation error randomly perturbs the estimates of each product's unobserved characteristic and thereby it affects the shape of the GMM-IV objective function.

4.2 The level of the simulated objective function

The structural error term is a critical ingredient of the GMM-IV objective function, so the simulation errors in the former directly affect the latter. For a first indication of the magnitude of the problem, we first fix the nonlinearly entering parameters at our global minimum candidate, $\widetilde{\theta}_2 = (1.52, 5.84, 3.39, 0.41, 0.10)'$. It is not essential that we pick this specific point. We would obtain qualitatively identical results at other candidate values of θ. For each set, we therefore evaluate the objective function at exactly the same point and only vary the set of simulation draws. Without simulation error, there would be no variation across the objective function values that we obtain. To fix notation, let $\mathcal{J}(\xi(\widetilde{\theta}_2), \nu_i^m)$ denote the objective function value that we obtain at $\widetilde{\theta}_2$ using the particular set of simulation draws ν_i^m, where we use simulation approach $m = \{\text{MC}, \text{MLHS}\}$ and generate $i = 1, \ldots, 1000$ independent samples. We vary the number of draws between 50 and 100,000. Table 2 presents the results.

The spread in objective function values is striking. With only 50 Monte Carlo draws, we see a range of roughly 2,500 for the objective function values. As we hold everything else constant, the different random samples of ν are the sole driver of this effect. To assess how the variation in objective values changes with the number of simulation draws across the independently drawn samples, we report the coefficient of variation. For only 50 draws, we obtain coefficients of roughly 0.5 and 0.34 for Monte Carlo and MLHS integration, respectively. Given that these figures are based on evaluating the objective function at exactly the same point, this variation is indeed substantial. As we raise the number of simulation

Table 2: Objective Function Values Obtained Using Monte Carlo and MLHS Draws

	Monte Carlo draws				MLHS draws			
# draws	$\overline{\mathcal{J}}$	$\sigma_{\mathcal{J}}$	$\sigma_{\mathcal{J}}/\overline{\mathcal{J}}$	range of $\mathcal{J}$	$\overline{\mathcal{J}}$	$\sigma_{\mathcal{J}}$	$\sigma_{\mathcal{J}}/\overline{\mathcal{J}}$	range of $\mathcal{J}$
50	726.7	360.5	.496	2,471	381.6	130.8	.343	1,002
500	286.9	70.6	.246	455.3	251.9	37.9	.151	235.7
5,000	242.6	22.3	.092	141	238.1	11.1	.047	72.3
100,000	236.9	4.8	.020	30.8	236.8	2.2	.009	14.8

Note: $\overline{\mathcal{J}}$ and $\sigma_{\mathcal{J}}$ denote the mean and standard deviation of the objective function values for each number of simulation draws.

draws, however, we can observe a large drop in the coefficients of variation. For 10,000 draws the Monte Carlo and MLHS integration approaches deliver coefficients of around .07 and .03. Raising the number of draws further to 100,000 pushes the coefficient of variation for the MLHS approach below 1 percent, while its counterpart for Monte Carlo integration is 2 percent. Finally, we can see that the mean of the objective function tends towards the same value of roughly 237 for both simulation methods.

4.3 The shape of the simulated objective function

As the GMM estimator minimizes the objective function, its level is less important for parameter estimation than its shape. To provide an impression of the relevance of the simulation errors, we trace the objective function in one dimension. We take the parameter estimate from our global minimum candidate, which is based on numerical integration using 10,000 MLHS draws $\widetilde{\theta}_2$. We hold all σ values constant except for σ_{price}. Figure 2 plots the shape of the objective function along the σ_{price}-dimension for several sets of simulation draws. The red dashed, blue dashed and solid black lines are respectively based on three independently generated samples of 50, 500 and 5,000 MLHS draws. With 50 draws, the shape of the objective function changes markedly across the three sets. One of the three sets attains a local minimum at zero, which implies no preference heterogeneity along the price dimension. This outcome is strongly rejected by our full set of estimations. A second set produces a shape that yields two minima along the price dimension. The third set has only one local minimum, but gives a biased estimate of σ_{price} compared with our global minimum candidate. For 500 MLHS draws we can already see that the shape of the objective function stabilizes. There is only one local minimum for all sets, which is not located at zero. There

Figure 2: Shape of the Objective Function at the Global Minimum Candidate

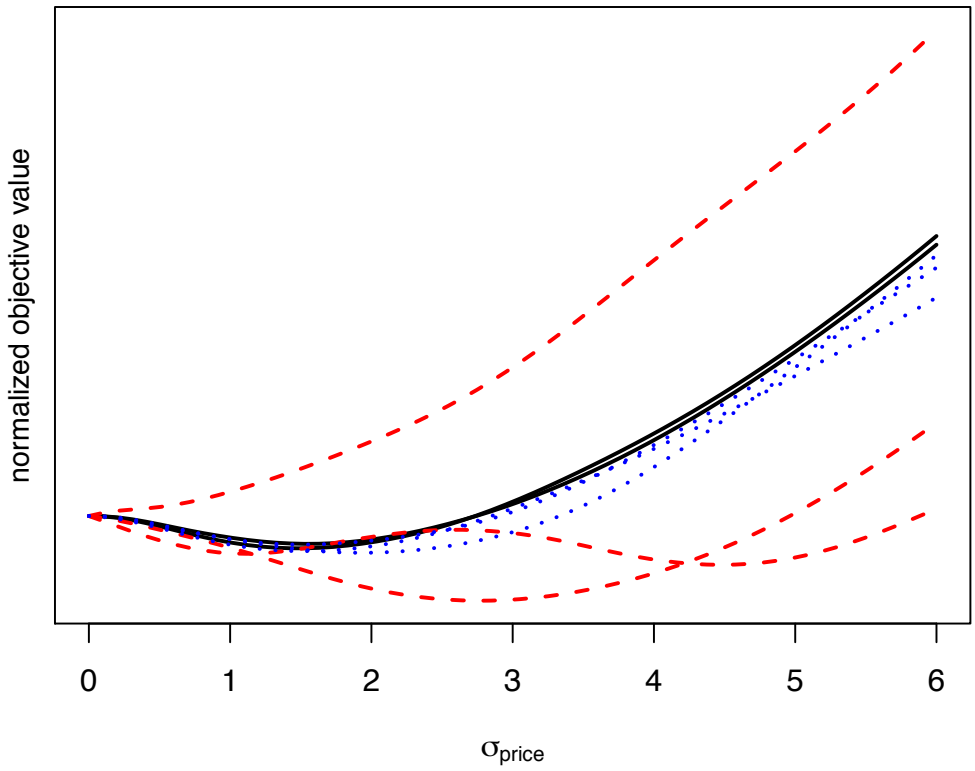

Note: The objective function is plotted along its σ_{price} dimension for 9 sets of Monte Carlo simulation draws: 3 sets using 50 draws (red dashed), 3 sets using 500 draws (blue dotted), and 3 sets using 5,000 draws (black solid). To ensure that all objective functions share the same value at $\sigma_{price} = 0$, we subtract the objective value at that point from each of the 9 objective function plots.

is, however, visible variation in the location of the minima. For 5,000 MLHS draws the three sets generate objective functions that appear to be congruent. The shape of the objective is stable, the minimum at zero is ruled out and the local minima across the three sets are located very close to our global minimum candidate.

4.4 Estimation results

The results presented so far show that simulation error randomly disturbs the point estimates of the structural error terms. These error terms directly enter the sample moment and thereby affect the level and shape of the GMM-IV objective function. Substantial simulation error can therefore produce ill-behaved objective functions with many local minima and widely ranging parameter estimates. We now present three sets of estimation outcomes.

Table 3: Range and Spread of the Identified Minima

	Monte Carlo draws				MLHS draws			
# draws	$\overline{\mathcal{J}}$	$\sigma_{\mathcal{J}}/\overline{\mathcal{J}}$	range of $\mathcal{J}$	# Minima	$\overline{\mathcal{J}}$	$\sigma_{\mathcal{J}}/\overline{\mathcal{J}}$	range of $\mathcal{J}$	# Minima
50	198.6	.323	283	126	179.6	.361	254	149
100	195.5	.311	265	129	183.9	.302	286	136
200	199.6	.298	278	128	188.0	.240	222	132
500	209.3	.201	222	123	204.4	.154	157	111
1,000	211.5	.146	166	99	207.5	.123	148	102
2,000	214.7	.135	159	93	221.7	.078	87	65
5,000	225.1	.090	101	80	229.3	.038	53	38
10,000	230.1	.065	82	57	232.1	.029	40	32

Note: MLHS stands for modified latin hypercube sampling. $\overline{\mathcal{J}}$ and $\sigma_{\mathcal{J}}$ denote the mean and standard deviation of the objective function values for each number of simulation draws. To count the number of unique minima we take all identified minima from the 2,500 estimations that are run for each number of draws and round the objective function values to whole numbers.

First, we trace out how an increasing number of simulation draws affects the behavior and robustness of the BLP estimator. Second, we turn to the point estimates of the 5 random coefficients and their statistical significance. Third, we examine the model-implied economic predictions by computing the own-price elasticities at the product level and the predicted price, profit, and consumer welfare effects of a counterfactual merger between Chrysler and GM.

4.4.1 Behavior and robustness of the nested fixed point estimator

We characterize the behavior of the estimator by examining the range and number of the identified minima. Table 3 shows that an increase in the number of simulation draws tightens the range and reduces the number of the identified local minima for both integration approaches. The pattern can be succinctly summarized using the coefficient of variation. Increasing the number of draws from 50 to 10,000 reduces the coefficient of variation for Monte Carlo integration from roughly 32 percent to 6.5 percent. For MLHS draws, the decrease is more substantial from 36.1 percent to 2.9 percent. Concomitantly, the number of unique minima is reduced by a factor exceeding 2 and close to 5 for Monte Carlo and MLHS draws, respectively. The 32 unique minima that are identified using our most accurate numerical integration approach are obtained across 50 independent samples of preference draws. Thus, there is less than one minimum per set of draws. Moreover, as we compute the model ag-

Table 4: Behavior of the Nested Fixed Point Estimator

	Monte Carlo draws			MLHS draws		
# draws	fraction minima	objective calls	inner iterations	fraction minima	objective calls	inner iterations
50	0.625	115.7	40.44	0.545	119.8	43.73
100	0.752	104.4	37.55	0.586	102.6	39.17
200	0.768	89.52	34.56	0.712	97.07	36.73
500	0.821	79.03	33.00	0.842	84.78	33.49
1,000	0.874	73.75	32.52	0.949	76.06	33.30
2,000	0.925	68.62	31.71	0.993	66.14	31.67
5,000	0.999	61.11	30.65	0.999	60.21	30.21
10,000	1.000	58.34	29.84	1.000	56.33	29.72

Note: MLHS stands for modified latin hypercube sampling. All statistics are computed as averages across all estimations for a given number of draws that converge to a local minimum. The number of objective calls is the number of GMM-IV objective function evaluations the optimization algorithm requires to converge to a candidate minimum.

gregate shares more accurately, the reduced number of minima fall into a narrowing range of values.

As a measure of the estimator's robustness we use the fraction of starting guesses that yield a local minimum. With 50 simulation draws, we see in Table 4 that for both simulation approaches, a large fraction of estimations fails to converge to a local minimum. For Monte Carlo simulation, this fraction is roughly 40 percent, while for MLHS draws, almost 45 percent of the attempts fail to converge to a local minimum. With 500 simulation draws, this fraction of failed estimation runs drops below 20 percent for both approaches and beyond 5,000 Monte Carlo draws and 2,000 MLHS draws almost every estimation run identifies a local minimum. With 10,000 draws, both approaches return a local minimum for all estimation runs. Thus, with high integration accuracy, the particular starting guess has no effect on whether the estimator converges to a minimum or not.

Additionally, to evaluate the computational complexity of identifying a candidate minimum, we trace out how the number of simulation draws affects the number of iterations in the estimator's inner loop, the nested contraction mapping, and the number of objective function evaluations, the outer loop, that are required for convergence. Table 4 shows that this measure of computational complexity is roughly identical across the two simulation approaches. In terms of the number of objective function evaluations we see a substantial reduction when raising the number of draws from 50 to 10,000. The latter requires around

Table 5: Choice of Optimizer and Convergence Threshold

Optimizer Effect

	Nelder-Mead		BFGS		Simulated Annealing	
draws	$\overline{\mathcal{J}}$	$\overline{\sigma}_{price}$	$\overline{\mathcal{J}}$	$\overline{\sigma}_{price}$	$\overline{\mathcal{J}}$	$\overline{\sigma}_{price}$
50	175.1	2.466	170.3	2.662	359.8	0.919
	[22.53, 251.7]	[1.337, 4.228]	[24.34, 251.7]	[1.334, 5.912]	[234.9, 628.5]	[.037, 2.466]
5,000	234.4	1.472	233.2	1.484	309.7	1.000
	[227.6, 244.4]	[1.327, 1.533]	[227.6, 242.0]	[1.428, 1.531]	[259.6, 405.7]	[.089, 1.866]
10,000	231.4	1.460	231.6	1.465	319.5	1.03
	[225.4, 237.9]	[1.423, 1.485]	[225.4, 236.9]	[1.446, 1.484]	[253.2, 481.6]	[.243, 1.888]

Convergence Threshold Effect

	50 MLHS draws			10,000 MLHS draws		
γ_{inner}	10^{-4}	10^{-9}	10^{-16}	10^{-4}	10^{-9}	10^{-16}
$\overline{\mathcal{J}}$	241.4	175.8	170.3	245.8	231.6	231.6
	[197.8, 311.2]	[43.88, 251.7]	[24.3, 251.7]	[224.6, 265.5]	[225.4, 236.9]	[225.4, 236.9]
$\overline{\sigma}_{price}$	1.488	2.453	2.662	1.394	1.466	1.465
	[.859, 2.462]	[1.335, 4.089]	[1.334, 5.911]	[1.169, 1.922]	[1.446, 1.486]	[1.446, 1.484]

Note: $\overline{\sigma}_{price}$ is the average of the estimated σ_{price} coefficients. The 2.5^{th} and 97.5^{th} quantiles of the outcome distributions for the objective function values and σ_{price} are shown in square brackets. To conserve space, we only report the outcomes from the MLHS simulation approach. The results are based on running the full 2,500 estimations each for a given number of draws. Thus, the top panel is based on a total of 22,500 estimations and the bottom panel is based on 15,000 estimations. The reported results for the simulated annealing optimizer are based on 1,000 iterations of the optimizer. As can be seen, the simulated annealing optimizer did not converge at this point.

57 iterations, while the former needs more than 115 evaluations to arrive at a local minimum candidate. We also obtain a sizable reduction in the number of iterations in the contraction mapping from more than 40 to less than 30.

Finally, we examine how sensitive the estimation outcomes are with respect to the choice of the optimization algorithm and the choice of the inner convergence threshold for the aggregate market share inversion. The top panel of Table 5 shows how the choice of optimization algorithm affects the outcomes of the estimation. The results are based on running the 2,500 estimations each for 50, 5,000 and 10,000 MLHS draws with different optimization algorithms. We select one representative algorithm from three classes of optimization approaches. The Nelder-Mead algorithm falls into the category of derivative-free optimizers, the BFGS optimizer is a quasi-Newton optimizer that is derivative-based and lastly, simulated annealing belongs to the class of stochastic optimizers. For the sake of brevity, we focus

on the average values of the objective function, σ_{price} and the estimates that are based on MLHS draws only. We obtain qualitatively identical outcomes for the remaining coefficients and the outcomes that we obtain using standard Monte Carlo draws. With 50 draws, we can see differences in the average outcomes and their empirical 95 percent confidence intervals across the optimization approaches. The estimates that we obtain with simulated annealing stand out in particular. Similar to KM, we have found that this optimization algorithm does not converge within a reasonable amount of time. For the Nelder-Mead and quasi-Newton approaches, the differences in estimation outcomes turn out to be negligible for both 5,000 and 10,000 MLHS draws. Thus, with a sufficiently accurate numerical integration of the aggregate market share function, the choice of optimization algorithm becomes irrelevant in our setting.

In the bottom panel of Table 5, we present evidence on the role of the inner convergence threshold, which Dubé, Fox, and Su (2012) demonstrate to have a major impact on the behavior of the BLP estimator. We run the 2,500 estimations each using 50 and 10,000 MLHS draws with three different inner convergence thresholds: 10^{-4}, which is the loose threshold defined by Dubé, Fox, and Su (2012), 10^{-9} and 10^{-16}. We impose the latter for all of our 40,000 estimations. With only 50 MLHS draws, we indeed find that the convergence threshold of the nested fixed point algorithm has a measurable impact on the estimation outcomes. The estimates of σ_{price} and the identified minima of the objective differ across the three different thresholds. With 10,000 MLHS draws, however, only the very lax criterion of 10^{-4} delivers results that differ markedly. The lax criterion yields a wider range for the identified minima and σ_{price}. The more stringent criteria of 10^{-9} and 10^{-16} are virtually identical in terms of the estimation outcomes. Thus, a sufficiently high simulation accuracy also substantially diminishes the impact of the nested fixed point's convergence threshold.

4.4.2 Estimated random coefficients and their statistical significance

To assess whether the estimated random coefficients are jointly statistically significant, we compute the Wald statistic for each local minimum. The null hypothesis is that the standard logit model is true, so that $H_0 : \theta = 0$.[14] We do not reject the null for one out of a total of 33,479 identified minima. The evidence in favor of consumer preference heterogeneity is therefore overwhelming. This is also in line with the evidence above that our findings are not driven by weak instruments, but by the propagation of simulation error.

[14]The test statistic follows a chi squared distribution with the degrees of freedom being equal to the number of entries in θ. At a 95 percent confidence level and with 5 random coefficients, the critical value for the Wald statistic is roughly 11.07.

Figure 3: Range of Random Coefficient Estimates and Their Joint Statistical Significance

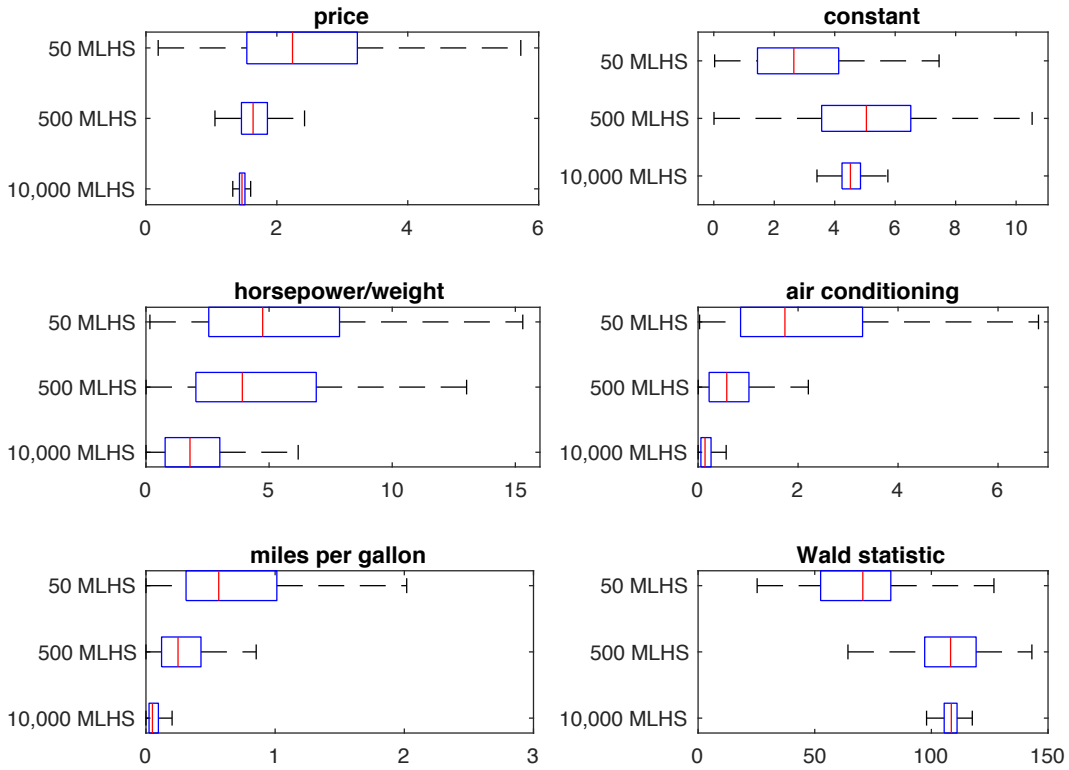

Note: To make changes in and around the medians of the point estimates easier, we do not plot the outliers. Moreover, to conserve space, we only show the box plots for the estimations using MLHS draws. The Monte Carlo counterparts are qualitatively identical. The Wald statistic is distributed chi squared with 5 degrees of freedom. The null hypothesis is that all random coefficients are zero, $\theta = 0$.

The boxplots in Figure 3 clearly show, however, that we obtain a lot of uncertainty in the random coefficient's point estimates when low numbers of draws are used to simulate ν. The range for the point estimates tightens drastically, however, as we move to 10,000 draws for both integration approaches. The random coefficient for price, for example, lies in a range between roughly 1.2 and 1.67 with a mean of 1.47 with 10,000 MLHS draws. The corresponding range for 50 MLHS draws is roughly 0.2 to 7.6 with a mean of 2.6. The random coefficient for the constant has a mean of 4.5 across all identified minima and also lies in a tight range. For air conditioning and miles per gallon, the point estimates strongly tend toward zero. We see a similar trend for the random coefficient that is placed on the horsepower-weight ratio. Compared to the other four coefficients, however, the range of the point estimates is still quite large for this random coefficient. We would need an even higher

Table 6: Range of Own-Price Elasticities Using Monte Carlo and MLHS Integration

	Monte Carlo draws				
	$\overline{\eta}_{jj}$				η_{jj}
# Draws	Min.	Mean	Max.	Std. Dev.	Std. Dev.
50	-43.8	-10.4	-2.31	5.32	9.97
500	-16.5	-9.13	-4.25	2.25	9.55
5,000	-11.1	-8.53	-6.14	0.84	9.12
10,000	-10.7	-8.49	-6.78	0.65	9.07
	MLHS draws				
50	-43.7	-13.4	-2.36	7.47	13.0
500	-14.9	-9.47	-5.71	1.64	10.1
5,000	-9.68	-8.54	-7.33	0.41	9.15
10,000	-9.64	-8.49	-6.97	0.32	9.07

Note: MLHS stands for modified latin hypercube sampling. $\overline{\eta}_{jj}$ denotes the average own-price elasticity. Each measure is computed across all local minima for a given number of draws. To arrive at the standard deviation of own-price elasticities for each number of draws, we average the standard deviations across all identified minima

number of draws to tighten this range further. This finding also illustrates that some random coefficients can be challenging to estimate. Nevertheless, when we examine the individual statistical significance of the random coefficients, a clear pattern emerges, which also applies to the estimated preference heterogeneity for the horsepower-weight ratio. The t-statistic for σ_{price} indicates that this coefficient is highly statistically significant. In fact, with 10,000 MLHS draws, there are only 8 cases, where the t-statistic drops below 2.[15] For $\sigma_{constant}$, we observe a similar pattern. Out of 2,500 estimations, only 207 yield a t-statistic below 1.65 and only 358 estimations produce t-statistics below 2. For each σ_{mpg} and σ_{air}, there is not a single case out of 2,500 estimated minima where the t-statistic exceeds 2. For σ_{hpwt}, there are only 19 such instances. Thus, with sufficient integration accuracy, it turns out that only the random coefficients on price and the constant are statistically significant. The average value of the Wald statistic, however, increases with the number of simulation draws even though we are left with only two statistically significant random coefficients. The range of the Wald statistic also tightens considerably.

4.4.3 Model-implied economic outcomes

We assess how sensitive the model-implied economic predictions are to numerical approximation error by characterizing the distribution of own-price elasticities. Table 6 presents statistics on the first and second moments of the distribution of own-price elasticities. We summarize the first moment of the distribution by showing the range and mean of the average own-price elasticity. With only 50 draws, we obtain the widest range, which reaches from roughly -44 to around -2.3. Moving to 10,000 draws reduces this dispersion substantially. The mean of the average own-price elasticity increases to -8.5 and the range covers only roughly -11 to -6.8 for Monte Carlo draws and -9.6 to -7 for MLHS draws. The reduction in the standard deviation of the estimated average own-price elasticity is impressive. For the Monte Carlo and MLHS integration approaches, it respectively falls from 5.3 to 0.65 and from 7.5 to only 0.32.

To examine the second moment of the distribution, we compute the standard deviation of the own-price elasticities for each local minimum and average the results over all minima. Both integration approaches tend towards the same measure of the distribution's spread. The estimated standard deviation is roughly 9 when using 10,000 draws to integrate the aggregate market share function. For lower numbers of draws, the spread is systematically higher.

4.5 Merger simulation

Lastly, we perform a simulation of the equilibrium that results from a merger between GM and Chrysler. These kind of counter-factual simulations are often the ultimate questions of interest in applied industrial organization studies. We simulate this scenario for each of the 20 years in the sample and average the results using units sold as weights. Figure 4 shows the distributions of the simulated change in consumer welfare following the merger for the two simulation approaches.

With only 50 draws, the estimation can deliver outcomes that range from hardly any detrimental effect to consumer welfare to an average annual welfare loss between 4 and 6 billion dollars. With 10,000 draws, the Monte Carlo and MLHS approaches deliver a mean annual consumer welfare loss of close to 1.5 billion dollars. With 50 draws, this estimate drops by roughly 40 percent to around 900 million dollars. The direction of this change is in line with how an increase in the number of simulation draws affects the estimates of

[15] We compute Eicker-Huber-White standard errors.

Figure 4: Change in Consumer Welfare

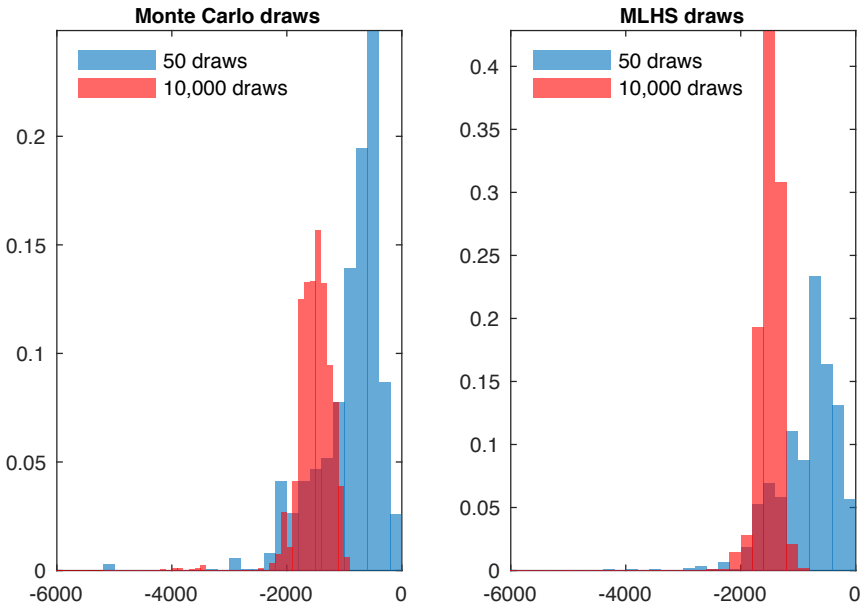

Note: The panels show the distribution of the average annual change in consumer welfare following a merger between GM and Chrysler.

own-price elasticities. We consistently obtain own-price elasticities of lower magnitude for a higher number of simulation draws. Thus, with fewer draws demand is estimated to be overly elastic. This immediately implies that the welfare losses and price changes following a merger in the market are smaller with a low integration accuracy. We surmise that this effect is driven by having sufficiently strong IVs. These effectively bound the estimates of own-price elasticities away from one. Simulation error produces a wider spread of the estimates. With a bound on own-price elasticities at one, this spread is likely to lead to an over-estimation of demand elasticities, which in turn affects the outcomes of our merger simulation.

Table 7 shows how this biases the estimates of post-merger price and profit changes. The relatively crude approximations to aggregate market shares deliver price and profit effects that are on average too low and yield substantially wider confidence intervals. In relative terms, this bias is substantial. For both simulation approaches, using 10,000 draws gives an average price effect that is roughly 45 percent greater than what we obtain with only 50 draws. For the profit effect, the bias is between 15 and 24 percent for the MLHS and Monte Carlo approach, respectively.

Table 7: Counterfactual Price and Profit Changes for the Merging Parties following a Chrysler-GM Merger

	Monte Carlo draws		MLHS draws	
draws	Δp (percent)	$\Delta\pi$ (mln 1983 dollars)	Δp (percent)	$\Delta\pi$ (mln 1983 dollars)
50	3.53	418	3.49	454
	[1.85, 6.85]	[191, 640]	[1.56, 6.37]	[297, 683]
500	4.78	503	5.03	526
	[2.61, 6.93]	[301, 660]	[3.21, 6.95]	[384, 666]
5,000	5.14	521	5.15	524
	[3.81, 6.40]	[397, 620]	[4.22, 5.84]	[457, 568]
10,000	5.09	519	5.10	521
	[4.17, 6.09]	[431, 604]	[4.56, 5.71]	[479, 568]

Note: The reported figures are based on simulating the GM-Chrysler merger for each of the 20 years in the sample and averaging the simulated outcomes by units sold. 95 percent confidence intervals are shown in square brackets.

5 Computational costs

Dubé, Fox, and Su (2012) point out that the desire to speed up the estimation of BLP models confronts the researcher with the temptation to introduce approximation or simulation error. Our results show that giving in to this temptation will backfire by undermining the replicability and reliability of the estimation results. On the other hand, increasing the number of Monte Carlo simulation draws to a sufficient level can induce burdensome or even infeasible computational costs.

In this section, we first highlight the trade-off between accuracy and speed in our simulation exercise (Section 5.1). Then we discuss different ways to achieve higher accuracy while reducing the computational costs by (1) tweaking the algorithm for the inversion of the market shares (Section 5.2) and (2) using different numerical integration algorithms (Section 5.3). Section 5.4 briefly introduces an R package and Matlab code for the estimation of BLP models that include these improvements and are a computationally more efficient implementation of the BLP estimation algorithm than the Matlab programs of KM, which are based on the code accompanying Nevo (2000). We attain speedups of 7 and 6 for the R package and our Matlab code, respectively.

Figure 5: Computational Burden versus Accuracy

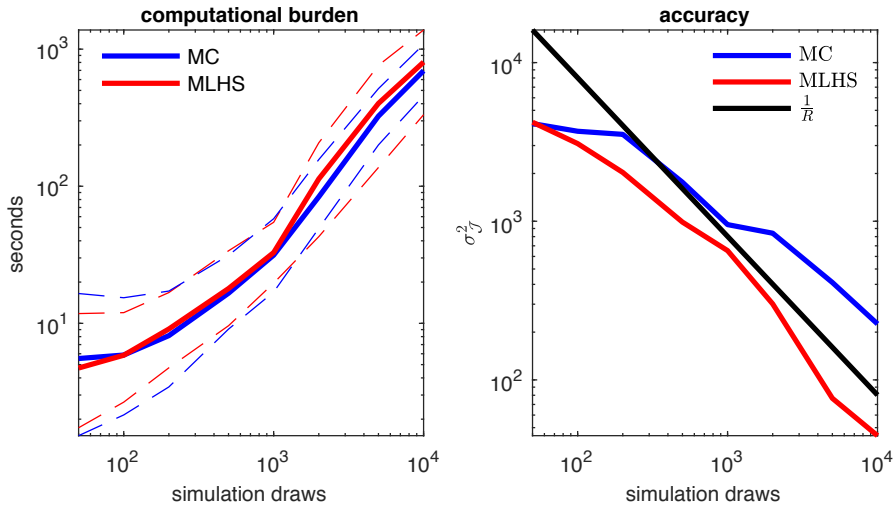

Note: The dashed lines in the left plot show the 95 percent empirical confidence intervals for the time required to complete a BLP model estimation with the corresponding number of draws on the x-axis. In the right panel, the solid black line shows the asymptotic convergence rate that applies to both numerical integration approaches.

5.1 The tradeoff between speed and reliability

How much precision is gained by an increase in the number of simulation draws and how much does this raise the computational burden of a single model estimation? The left panel of Figure (5) plots the average time it takes to estimate specification (16) for each number of simulation draws. Note that all axes are on a log-scale for the sake of readability. The computational burden increases linearly with the number of draws. In fact, we obtain an elasticity of close to one for the runtime of an estimation with respect to the number of simulation draws. Moreover, we also find that using MLHS draws requires a roughly 15 percent longer compute time on average.[16] MLH-sampling ensures that a uniform number of draws is generated for each quantile of the assumed distribution of ν. We therefore obtain more simulation draws in the tails of the distribution with this sampling scheme than with standard Monte Carlo sampling. This causes the market share inversion using MLHS draws to take a longer time to converge for values of θ_2 that are far away from θ^*.

The right panel shows how the variance of the identified minima from the 40,000 BLP model estimations decreases with the log-number of simulation draws. With only 50 simu-

[16]With 10,000 draws, the average BLP estimation that uses Monte Carlo draws takes 700 seconds, while estimation with MLHS draws takes roughly 800 seconds.

lation draws, we obtain a variance of roughly 4,000 for both integration approaches. With 10,000 draws the variance falls to 44 for MLHS integration and to 225 for Monte Carlo integration. Thus, roughly 5 times as many Monte Carlo draws are needed to attain the same integration accuracy of a given number of MLHS draws. We therefore find estimation using MLHS draws to be computationally more efficient even when each estimation run on average takes 15 percent longer than when using Monte Carlo draws. The solid black line plots the asymptotic convergence rate that applies to both integration approaches. We can see that the actual reduction in the variance of the identified local minima closely follows the asymptotic convergence rate.

Putting both panels together, we can state that doubling the number of simulation draws roughly reduces the variance of the objective function by half. This fits the behavior of the estimator well for a sufficiently large number of draws; in our case from roughly 500 draws onwards. The same convergence rate applies to MLHS draws. This sampling method can match the variance of the objective function obtained with standard Monte Carlo simulation with only a fifth of the number of Monte Carlo draws, however. The increase in accuracy requires a roughly one-for-one increase in compute times. It takes roughly 200 times longer to estimate specification (16) with 10,000 draws than with 50 draws. Even though this might sound dramatic, keep in mind that we are moving to 800 seconds or 13 minutes and 20 seconds on average for MLHS integration.

As our results above stress, performing relatively few estimations with many simulation draws that produce precise and reliable results are useful for answering economic questions of interest. Running many more estimations with few numbers of draws gives highly unstable and on average biased results. Following the implications of the Knittel-Metaxoglou critique in this situation requires re-estimating the model for different optimizers and convergence criteria, for example, and makes the transparent communication of the estimation results more difficult. A direct comparison of the compute times for a single estimation of (16) is therefore lopsided.

Accurate numerical integration comes at the price of a higher computational burden, but rewards the researcher with reliability and therefore a substantially lessened need for extensive robustness checks. The researcher's main concern should therefore be the reliability of the estimates and not the manageable computational burden of running the estimations. Having said that, we provide two approaches for decreasing the computational burden of BLP model estimations. First, the 40,000 BLP estimations are in fact an "embarrassingly parallel" computational task: each estimation is independent of all other estimations. The

speedup from parallelization therefore moves almost one-for-one with the number of compute cores that are used.[17] Second, we provide a simple reformulation of the BLP contraction mapping that gives a speedup of roughly 2. The nitty gritty of this reformulation is discussed next.

5.2 Speeding up the inversion of market shares

Iterating on the BLP contraction mapping, (11), until convergence is the most compute-intensive part of the estimation. We can speed up the inversion by a factor of roughly 2 by avoiding a large number of numerical divisions. Nevo (2000) notes in his Appendix that taking logs is a computationally costly operation and that the computational burden of repeatedly solving the fixed point during the estimation can be reduced by exponentiating the equation.

$$w_j^{iter+1} = w_j^{iter} \underbrace{\frac{\overbrace{S_j}^{J\times 1}}{R^{-1}\sum_r s_{rj}(w^{iter},\nu,\theta_2)}}_{J\times 1} \tag{17}$$

Here, $w \equiv \exp(\delta_j)$ is the exponential of the mean utility vector. This reformulation gives a substantial speedup in computing the contraction mapping by avoiding the need to re-compute the log shares during the iteration. In the denominator of the expression, the model-implied aggregate market share for each of the J products is computed. To do so, the full matrix of consumer-level choice probabilities must be computed.

$$[s_{rj}(\delta,\nu,\theta_2)]_{r=1,\ldots,R;j=1,\ldots,J} = \underbrace{\frac{\overbrace{\exp(\delta * \iota_R' + [\mu_j(\nu,\theta_2)])}^{J\times R}}{\iota_J * (1 + \sum_k \exp(\delta_k * \iota_R' + [\mu_k(\nu,\theta_2)]))}}_{J\times R} \tag{18}$$

ι_R is a vector of ones with R elements. It has the effect of stacking the vector of mean utilities horizontally R times. In the numerator ι_J stacks the denominator vertically J times. This makes the numerator and denominator conformable and the s_{rj}'s for a whole market can be computed in one matrix operation. We are dividing a $J \times R$ matrix by another $J \times R$ matrix, which requires $J * R$ divisions.

[17] The Matlab code that we have used will be made available online and we have made available the R package `BLPestimatoR` on the CRAN repository that uses the same speedup of the contraction mapping. Both programs implement a parallelization scheme that runs each estimation independently.

We can avoid a large number of these divisions by noting that the contraction mapping can be formulated in terms of consumer-specific choice probabilities for the outside option. Let $v_{rj} \equiv \exp(\mu_{rj})$.

$$w_j^{iter+1} = w_j^{iter} \frac{S_j}{R^{-1} \sum_r \frac{w_j^{iter} v_{rj}}{1+\sum_k w_k^{iter} v_{rk}}} = \underbrace{\frac{\overbrace{S_j}^{J \times 1}}{R^{-1} \sum_r v_{rj} s_{ro}(w^{iter}, \nu, \theta_2)}}_{J \times 1} \tag{19}$$

In the denominator on the rhs only the R choice probabilities for the outside option must be computed instead of the full matrix of choice probabilities for the inside products.

$$[s_{ro}(\delta, \nu, \theta_2)]_{r=1,\ldots,R} = 1/[(1 + \sum_k \exp(\delta_k * \iota_R' + \mu_k(\nu, \theta_2)))] \tag{20}$$

The alternative fixed point iteration (19) only requires a total of $J + R$ instead of $J + J * R$ numerical divisions. Unless there is only one inside product in the market, the computational burden for our reformulation in terms of the consumer-specific outside good choice probabilities has a strictly lower computational burden.

Figure 6 illustrates this. For the BLP automobile data, we solve the nested fixed point for each of KM's 50 starting guesses. We do this for numbers of draws between 50 and 100,000. The solid blue line plots the time to convergence required by the Nevo version of the fixed point, while the dashed red line corresponds to our version of the same fixed point problem. We want to emphasize that both versions need exactly the same number of iterations to reach the convergence threshold and give exactly the same δ^* for all of the starting guesses. Formulating the contraction mapping in terms of the outside good shares yields a speedup of at least 2. This matches the speedup of the approximate BLP estimator of Lee and Seo (2015), which uses a linear approximation of the market share equation to solve the fixed point problem analytically. We attain roughly the same speedup, but solve the fixed point problem exactly and thereby retain all the properties of the original BLP estimator without introducing an additional source of approximation error that propagates in the estimator's objective function.[18]

[18]The precision with which the fixed point is solved is of course limited by the convergence threshold.

Figure 6: Computational Burden of Solving the Fixed Point

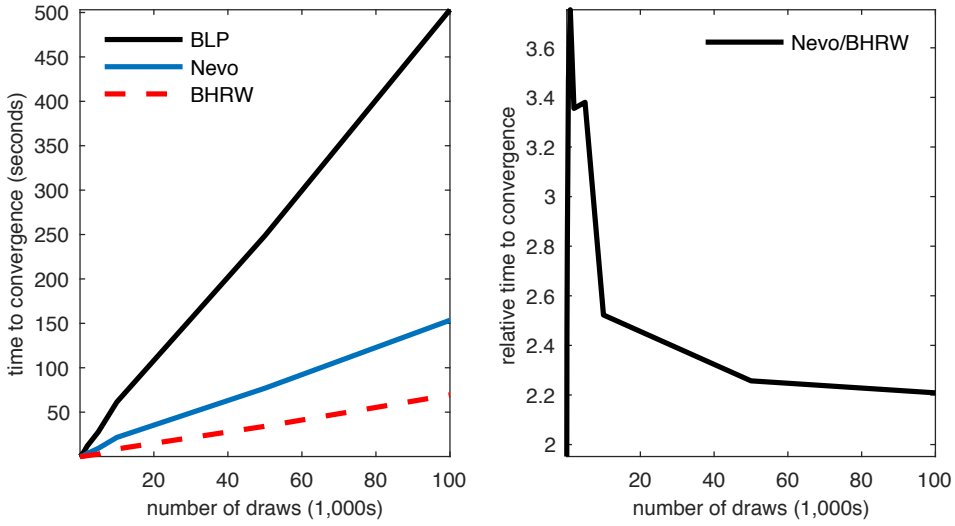

Note: The left panel shows the average time to reach the convergence threshold 10^{-14} in the contraction mapping over KM's 50 random starting guesses for the random coefficients. The black line shows the outcomes for the original BLP contraction mapping, while the solid blue line corresponds to the fixed point formulation of Nevo (2000) and the dashed red line shows our reformulation. We evaluate 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 50,000 and 100,000 draws using 10 independently generated samples and average the time until convergence across these 10 sets. The right panel plots the ratio of the time to convergence for the Nevo and our fixed point formulation.

5.3 Numerical integration algorithms

The critical integral in the market share equation, (5), can be approximated with different numerical algorithms. The choice of the algorithm can have large effects on the accuracy at a given computational cost (or on the required computational burden to attain a given accuracy). In our main analyses, we have seen that plain-vanilla Monte Carlo simulation requires roughly 4 to 5 times as many draws as MLHS-based approximation to attain the same integration accuracy. Given that the computational burden increases roughly one-for-one with the number of simulation draws, it follows that Monte Carlo draws present 4 to 5 times the computational burden of MLHS draws.

There are a number of other alternative approaches that may be more efficient than MLHS. For example, Nevo (2001) uses Halton draws, which Hess, Train, and Polak (2006) find to perform roughly on par with MLHS draws. Sovinsky Goeree (2008) applies antithetic sampling to increase simulation efficiency. This simulation approach has the additional advantage that it yields an objective function that is symmetric around the origin.

In low dimensions, Gaussian quadrature can be very efficient. In higher dimensions, sparse-grids quadrature is a potentially powerful approach, see Heiss and Winschel (2008). It has been successfully applied for the estimation of BLP models for example by Björnerstedt and Verboven (2016). In our own experience with our setup for this paper, sparse grids quadrature works very well in the majority of cases.[19]

Importance sampling and adaptive integration algorithms can greatly improve the approximation quality, see Heiss (2010). Berry, Levinsohn, and Pakes (1995) use importance sampling to increase the accuracy of numerically integrating the model-implied aggregate shares. Brunner (2017) shows that this is a very promising approach in the context of BLP model estimation. Due to space constraints, we cannot go into more detail here. As our results clearly indicate that accurate approximations are critical, and since the computational burden can be overwhelming, there is a large potential for further improving the approximation algorithms. To reduce the computational burden of BLP model estimation we also provide software that is several times faster than the replication code of KM, which in turn is based on the Matlab code of Nevo (2000).

5.4 The BLPestimatoR package and Matlab code for efficient estimation

We have implemented our suggestions to speed up the BLP estimation in the R package BLPestimatoR, available at the Comprehensive R Archive Network.[20] In addition to the algorithmic refinements from Section 5.2 and different approximation algorithms, the core of the package is written in the C language. We also provide our Matlab code that we have used for the 40,000 BLP model estimations in this paper. Both programs use parallelization to reduce the computational time required to run several independent estimations of BLP model specifications. To illustrate the speedups that can be expected we use a horse race between the code provided by KM, which in turn is based on the Matlab code accompanying Nevo (2000), the R package BLPestimatoR and our Matlab code.[21]

The horse race is structured as follows. We generate 1,000 independent sets of preference draws, ν, with 1,000 simulation draws each. We then fix the parameter vector

[19] In extreme cases with poor parameter starting guesses, sparse grids approximation can cause problems with negative market shares. These cause problems if they aren't dealt with appropriately.

[20] See `https://cran.r-project.org/web/packages/BLPestimatoR/`

[21] Aviv Nevo's code is available at `http://faculty.wcas.northwestern.edu/~ane686/supplements/rc_dc_code.htm`. However, it has not been updated for a while and some references to built-in Matlab functions are no longer valid. The code provided in the replication files of KM is directly based on Nevo's code and can be used as a replacement.

Table 8: Benchmarking BLP Estimation Routines

Matlab	R - BLPestimatoR			Matlab	
KM	BHRW - pure R	Rcpp	2 cores	BHRW - pure Matlab	2 cores
2.57	5.20	0.36	0.18	0.43	0.22

Note: All runtimes are reported in seconds. The results use the automobile data from BLP and implement specification (16) above. The runtimes are based on 1,000 evaluations of the BLP estimator's GMM-IV objective function at independently drawn sets of preference heterogeneity, ν. Each set has 1,000 draws per characteristic and resamples by market. The reported runtime is the time it took to evaluate the objective once averaged over the 1,000 evaluations. The parameter vector at which we evaluate the objective is $\theta_2 = (1.52, 5.84, 3.39, 0.41, 0.10)$. KM stands for the implementation of Knittel and Metaxoglou (2014), while BHRW is the implementation used for the 40,000 BLP estimations in this paper. It uses the fixed point formulation (19). The Rcpp entry for BLPestimtoR implements the contration mapping in C within R using the `Rcpp` package (Eddelbuettel and François (2011)). The entries "2 cores" use parallelization at the level of each combination of starting guess and set of preference draws. The R implementation uses `mclapply`, while our Matlab code uses `parfor`. All benchmarks were run on the same Linux workstation with an Intel Xeon E5-2640 v3 CPU.

$\theta_2 = (1.52, 5.84, 3.39, 0.41, 0.10)$ and evaluate the GMM-IV objective function value for each of the 1,000 sets of draws. This requires the numerical inversion of the aggregate market shares and therefore covers the bulk of the computational burden for a full BLP model estimation. To obtain measures of the computational cost of each, we then average the runtimes of each objective evaluation across the 1,000 sets of draws.

Table 8 presents the outcomes of this benchmarking exercise. The entry in the first column uses the replication code of KM. The code is modified to only perform the objective function evaluation and nothing else. We also make sure that all programs use the same sets of draws, convergence criteria and input variables. Thus, all programs perform mathematically equivalent evaluations of the objective function. Last but not least, all evaluations have been run on the same computer. On average the Matlab code provided by KM requires 2.57 seconds per function evaluation. All entries for the BLPestimatoR package implement our more computationally efficient reformulation of the contraction mapping, (19). An implementation in pure R requires roughly twice the runtime of KM's Matlab routine. When we implement the contraction mapping in the C programming language using the Rcpp package, however, the runtime of the R package drops to a mere 0.36 seconds, which gives roughly a 7 times speedup over the KM code. Parallelizing the objective evaluations on 2 cores is very efficient, as it yields a speedup of almost exactly 2. The runtime drops to 0.18 seconds, a 14 times speedup over KM's program. Our Matlab routine is written purely in Matlab and does not outsource the computation of the contraction mapping to an external C func-

tion.[22] Nevertheless, it performs almost on par with the BLPestimatoR package. Without parallelization the runtime is 0.43 seconds, which is roughly 20 percent slower than our R package and gives a 6 times speedup over KM's Matlab routine. Parallelizing the objective evaluations is again highly efficient and roughly halves the runtime to 0.22 seconds.[23]

Both our R package and Matlab code are available for researchers who want to estimate BLP models reliably and quickly. The R package also contains a convenient interface.

6 Conclusions

The BLP model's nested fixed point estimator is susceptible to numerical instabilities if simulation error in the model's aggregate market share function is large. By substantially raising the number of simulation draws, however, the sample moments are computed accurately and the estimator's sensitivity to the specific combination of starting guess, optimization algorithm and the convergence threshold of the nested fixed point disappears. Instead, the estimator delivers an increasingly narrowing set of minima of its objective function, which also brings with it tighter sets of parameter estimates and implied economic predictions. Given a suitable set of instrumental variables, the main concern for the reliable numerical implementation of the BLP model should therefore be to reduce the approximation error in the market share integral.

Berry, Linton, and Pakes (2004) show that in a single cross-section the BLP model's nested fixed point estimator satisfies asymptotic normality if the ratio of the number of products squared over the number of simulation draws, J^2/R, is bounded as the number of products becomes large. This asymptotic result clearly resonates with our findings and in this sense the estimator behaves as advertised. In the automobile data the number of products varies between 72 and 150 products per market with on average roughly 111 cars per year. For 10,000 draws the estimator delivers a local minimum of the objective function for every combination of starting guess and set of simulation draws.

The reduction of simulation errors simplifies the implementation, verification and communication of BLP model estimates relative to the guidance offered by KM and Goldberg and Hellerstein (2013). We find it unnecessary to re-estimate the model with multiple optimization algorithms once simulation error is taken seriously. This also highlights potential

[22] This is possible in Matlab using a mex file.

[23] We have also implemented the contraction mapping in pure Fortran. The maximum speedup over pure Matlab was around 2 and requires switching on several compiler optimizations, which make the code less safe than Matlab. The program, for example, does no longer check whether arrays are accessed inbounds.

gains in computational efficiency. The Simplex or Nelder-Mead optimization algorithm is frequently used, because it is seen as particularly robust. We could not find any measurable difference in terms of estimation outcomes between the Nelder-Mead algorithm and a trust-region method that uses an analytical gradient. The latter approach, however, is computationally much more efficient. Similarly, the impact of the nested fixed point's convergence threshold is substantially reduced with an accurate approximation of the model's aggregate market share function. A maximum threshold of 10^{-9} seems to work well for the automobile data. The loose threshold of 10^{-4} should simply not be used in any setting.

We caution, however, to push our findings regarding the use of different starting guesses too far. KM have selected these 50 starting guesses after having evaluated the objective function for many more values. Thus, these guesses are likely to cover the potential parameter space well. The higher the dimensionality of the estimation problem, the more difficult it becomes to provide a good coverage of the parameter space. Therefore, all else equal, more guesses should be used for BLP models with a larger number of random coefficients. We therefore do not recommend a reduction in the number of starting guesses.[24] Moreover, each candidate minimum should be carefully verified. At the estimated parameter vector, the Hessian matrix must be positive definite and the norm of the gradient must be close to zero.

We re-emphasize that consistent identification requires strong and valid IVs. Our results show, however, that simulation error can easily overwhelm the estimates of the structural error terms even when the IVs are not weak. The error propagates in the GMM-IV objective function and produces many local minima with a wide range of parameter estimates and corresponding economic implications. The accurate numerical integration of the BLP model's aggregate market share function is therefore necessary to attain reliable identification. A high degree of numerical integration accuracy and relevant and valid IVs are therefore complements, not substitutes.

One way to reduce the approximation errors is to increase the number of Monte Carlo simulation draws. As this can be challenging or prohibitive in terms of computational costs, we finally discuss ways to improve the computational efficiency of numerically integrating the aggregate shares and provide efficient R and Matlab code. In comparison with the replication programs of KM, our routines provide speedups of roughly 7 and 6 for our R package and Matlab program, respectively. We can therefore substantially alleviate the burden of estimating BLP models reliably.

[24] Selecting 50 starting guesses from thousands of evaluations, however, is likely unnecessary.

References

Berry, Steven T. and Philip A. Haile (2014). "Identification in Differentiated Products Markets Using Market Level Data". In: *Econometrica* 82.5, pp. 1749–1797.

Berry, Steven (1994). "Estimating Discrete-Choice Models of Product Differentiation". In: *The RAND Journal of Economics* 25.2, pp. 242–262.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium". In: *Econometrica* 63.4, pp. 841–890.

Berry, Steven, Oliver B. Linton, and Ariel Pakes (2004). "Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems". In: *Review of Economic Studies* 71, pp. 613–654.

Björnerstedt, Jonas and Frank Verboven (2016). "Does Merger Simulation Work? Evidence from the Swedish Analgesics Market". In: *American Economic Journal: Applied Economics* 8, pp. 125–164.

Brunner, Daniel (2017). "Implications of Adaptive Integration Rules for the Performance of Random Coefficient Models of Demand". mimeo, University of Duesseldorf.

Dubé, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su (2012). "Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation". In: *Econometrica* 80.5, pp. 2231–2267.

Eddelbuettel, Dirk and Romain François (2011). "Rcpp: Seamless R and C++ Integration". In: *Journal of Statistical Software* 40.8, pp. 1–18.

Goldberg, Pinelopi and Rebecca Hellerstein (2013). "A Structural Approach to Identifying the Sources of Local Currency Price Stability". In: *Review of Economic Studies* 80.1, pp. 175–210.

Heiss, Florian (2010). "The Panel Probit Model: Adaptive Integration on Sparse Grids". In: *Advances in Econometrics, Vol. 26: Maximum Simulated Likelihood Methods and Applications*. Ed. by William Greene and R. Carter Hill. Emerald, pp. 41–64.

Heiss, Florian and Viktor Winschel (2008). "Likelihood aroximation by numerical integration on sparse grids". In: *Journal of Econometrics* 144.1, pp. 62–80.

Hess, Stephane, Kenneth E. Train, and John W. Polak (2006). "On the use of a Modified Latin Hypercube Sampling (MLHS) method in the estimation of a Mixed Logit Model for vehicle choice". In: *Transportation Research Part B: Methodological* 40.2, pp. 147 –163.

Knittel, Christopher R. and Konstantinos Metaxoglou (2014). "Estimation of Random-Coefficient Demand Models: Two Empiricists' Perspective". In: *Review of Economics and Statistics* 96.1, pp. 34–59.

Lee, Jinhyuk and Kyoungwon Seo (2015). "A computationally fast estimator for random coefficients logit demand models using aggregate data". In: *The RAND Journal of Economics* 46.1, pp. 86–102.

McCullough, B. D. and H. D. Vinod (2003). "Verifying the Solution from a Nonlinear Solver: A Case Study". In: *American Economic Review* 93.3, pp. 873–892.

Nevo, Aviv (2000). "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand". In: *Journal of Economics & Management Strategy* 9.4, pp. 513–548.

— (2001). "Measuring Market Power in the Ready-to-Eat Cereal Industry". In: *Econometrica* 69.2, pp. 307–342.

Petrosky-Nadeau, Nicolas and Lu Zhang (2017). "Solving the Diamond-Mortensen-Pissarides Model Accurately". In: *Quantitative Economics* 8.2, pp. 611–650.

Reynaert, Mathias and Frank Verboven (2014). "Improving the performance of random coefficients demand models: The role of optimal instruments". In: *Journal of Econometrics* 179.1, pp. 83–98.

Sovinsky Goeree, Michelle (2008). "Limited Information and Advertising in the U.S. Personal Computer Industry". In: *Econometrica* 76.5, pp. 1017–1074.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo (2002). "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments". In: *Journal of Business & Economic Statistics* 20.4, pp. 518–529.

PREVIOUS DISCUSSION PAPERS

267 Brunner, Daniel, Heiss, Florian, Romahn, André and Weiser, Constantin, Reliable Estimation of Random Coefficient Logit Demand Models, September 2017.

266 Kosse, Fabian, Deckers, Thomas, Schildberg-Hörisch, Hannah and Falk, Armin, The Formation of Prosociality: Causal Evidence on the Role of Social Environment, July 2017.

265 Friehe, Tim and Schildberg-Hörisch, Hannah, Predicting Norm Enforcement: The Individual and Joint Predictive Power of Economic Preferences, Personality, and Self-Control, July 2017.

264 Friehe, Tim and Schildberg-Hörisch, Hannah, Self-Control and Crime Revisited: Disentangling the Effect of Self-Control on Risk Taking and Antisocial Behavior, July 2017.

263 Golsteyn, Bart and Schildberg-Hörisch, Hannah, Challenges in Research on Preferences and Personality Traits: Measurement, Stability, and Inference, July 2017.

262 Lange, Mirjam R.J., Tariff Diversity and Competition Policy – Drivers for Broadband Adoption in the European Union, July 2017.

261 Reisinger, Markus and Thomes, Tim Paul, Manufacturer Collusion: Strategic Implications of the Channel Structure, July 2017.

260 Shekhar, Shiva and Wey, Christian, Uncertain Merger Synergies, Passive Partial Ownership, and Merger Control, July 2017.

259 Link, Thomas and Neyer, Ulrike, Friction-Induced Interbank Rate Volatility under Alternative Interest Corridor Systems, July 2017.

258 Diermeier, Matthias, Goecke, Henry, Niehues, Judith and Thomas, Tobias, Impact of Inequality-Related Media Coverage on the Concerns of the Citizens, July 2017.

257 Stiebale, Joel and Wößner, Nicole, M&As, Investment and Financing Constraints, July 2017.

256 Wellmann, Nicolas, OTT-Messaging and Mobile Telecommunication: A Joint Market? – An Empirical Approach, July 2017.

255 Ciani, Andrea and Imbruno, Michele, Microeconomic Mechanisms Behind Export Spillovers from FDI: Evidence from Bulgaria, June 2017.
Forthcoming in: Review of World Economics.

254 Hunold, Matthias and Muthers, Johannes, Capacity Constraints, Price Discrimination, Inefficient Competition and Subcontracting, June 2017.

253 Dertwinkel-Kalt, Markus and Köster, Mats, Salient Compromises in the Newsvendor Game, June 2017.
Forthcoming in: Journal of Economic Behavior and Organization.

252 Siekmann, Manuel, Characteristics, Causes, and Price Effects: Empirical Evidence of Intraday Edgeworth Cycles, May, 2017.

251 Benndorf, Volker, Moellers, Claudia and Normann, Hans-Theo, Experienced vs. Inexperienced Participants in the Lab: Do they Behave Differently?, May 2017.
Forthcoming in: Journal of the Economic Science Association.

250 Hunold, Matthias, Backward Ownership, Uniform Pricing and Entry Deterrence, May 2017.

249 Kreickemeier, Udo and Wrona, Jens, Industrialisation and the Big Push in a Global Economy, May 2017.

248 Dertwinkel-Kalt, Markus and Köster, Mats, Local Thinking and Skewness Preferences, April 2017.

247 Shekhar, Shiva, Homing Choice and Platform Pricing Strategy, March 2017.

246 Manasakis, Constantine, Mitrokostas, Evangelos and Petrakis, Emmanuel, Strategic Corporate Social Responsibility by a Multinational Firm, March 2017.

245 Ciani, Andrea, Income Inequality and the Quality of Imports, March 2017.

244 Bonnet, Céline and Schain, Jan Philip, An Empirical Analysis of Mergers: Efficiency Gains and Impact on Consumer Prices, February 2017.

243 Benndorf, Volker and Martinez-Martinez, Ismael, Perturbed Best Response Dynamics in a Hawk-Dove Game, January 2017.
Published in: Economics Letters, 153 (2017), pp. 61-64.

242 Dauth, Wolfgang, Findeisen, Sebastian and Suedekum, Jens, Trade and Manufacturing Jobs in Germany, January 2017.
Forthcoming in: American Economic Review, Papers & Proceedings.

241 Borrs, Linda and Knauth, Florian, The Impact of Trade and Technology on Wage Components, December 2016.

240 Haucap, Justus, Heimeshoff, Ulrich and Siekmann, Manuel, Selling Gasoline as a By-Product: The Impact of Market Structure on Local Prices, December 2016.

239 Herr, Annika and Normann, Hans-Theo, How Much Priority Bonus Should be Given to Registered Organ Donors? An Experimental Analysis, November 2016.

238 Steffen, Nico, Optimal Tariffs and Firm Technology Choice: An Environmental Approach, November 2016.

237 Behrens, Kristian, Mion, Giordano, Murata, Yasusada and Suedekum, Jens, Distorted Monopolistic Competition, November 2016.

236 Beckmann, Klaus, Dewenter, Ralf and Thomas, Tobias, Can News Draw Blood? The Impact of Media Coverage on the Number and Severity of Terror Attacks, November 2016.
Forthcoming in: Peace Economics, Peace Science and Public Policy.

235 Dewenter, Ralf, Dulleck, Uwe and Thomas, Tobias, Does the 4th Estate Deliver? Towars a More Direct Measure of Political Media Bias, November 2016.

234 Egger, Hartmut, Kreickemeier, Udo, Moser, Christoph and Wrona, Jens, Offshoring and Job Polarisation Between Firms, November 2016.

233 Moellers, Claudia, Stühmeier, Torben and Wenzel, Tobias, Search Costs in Concentrated Markets – An Experimental Analysis, October 2016.

232 Moellers, Claudia, Reputation and Foreclosure with Vertical Integration – Experimental Evidence, October 2016.

231 Alipranti, Maria, Mitrokostas, Evangelos and Petrakis, Emmanuel, Non-comparative and Comparative Advertising in Oligopolistic Markets, October 2016.
Forthcoming in: The Manchester School.

230 Jeitschko, Thomas D., Liu, Ting and Wang, Tao, Information Acquisition, Signaling and Learning in Duopoly, October 2016.

229 Stiebale, Joel and Vencappa, Dev, Acquisitions, Markups, Efficiency, and Product Quality: Evidende from India, October 2016.

228 Dewenter, Ralf and Heimeshoff, Ulrich, Predicting Advertising Volumes: A Structural Time Series Approach, October 2016.
Published in: Economics Bulletin, 37 (2017), Volume 3.

227 Wagner, Valentin, Seeking Risk or Answering Smart? Framing in Elementary Schools, October 2016.

226 Moellers, Claudia, Normann, Hans-Theo and Snyder, Christopher M., Communication in Vertical Markets: Experimental Evidence, July 2016.
Published in: International Journal of Industrial Organization, 50 (2017), pp. 214-258.

225 Argentesi, Elena, Buccirossi, Paolo, Cervone, Roberto, Duso, Tomaso and Marrazzo, Alessia, The Effect of Retail Mergers on Prices and Variety: An Ex-post Evaluation, June 2016.

224 Aghadadashli, Hamid, Dertwinkel-Kalt, Markus and Wey, Christian, The Nash Bargaining Solution in Vertical Relations With Linear Input Prices, June 2016.
Published in: Economics Letters, 145 (2016), pp. 291-294.

223 Fan, Ying, Kühn, Kai-Uwe and Lafontaine, Francine, Financial Constraints and Moral Hazard: The Case of Franchising, June 2016.
Forthcoming in: Journal of Political Economy.

222 Benndorf, Volker, Martinez-Martinez, Ismael and Normann, Hans-Theo, Equilibrium Selection with Coupled Populations in Hawk-Dove Games: Theory and Experiment in Continuous Time, June 2016.
Published in: Journal of Economic Theory, 165 (2016), pp. 472-486.

221 Lange, Mirjam R. J. and Saric, Amela, Substitution between Fixed, Mobile, and Voice over IP Telephony – Evidence from the European Union, May 2016.
Published in: Telecommunications Policy, 40 (2016), pp. 1007-1019.

220 Dewenter, Ralf, Heimeshoff, Ulrich and Lüth, Hendrik, The Impact of the Market Transparency Unit for Fuels on Gasoline Prices in Germany, May 2016.
Published in: Applied Economics Letters, 24 (2017), pp. 302-305.

219 Schain, Jan Philip and Stiebale, Joel, Innovation, Institutional Ownership, and Financial Constraints, April 2016.

218 Haucap, Justus and Stiebale, Joel, How Mergers Affect Innovation: Theory and Evidence from the Pharmaceutical Industry, April 2016.

217 Dertwinkel-Kalt, Markus and Wey, Christian, Evidence Production in Merger Control: The Role of Remedies, March 2016.

216 Dertwinkel-Kalt, Markus, Köhler, Katrin, Lange, Mirjam R. J. and Wenzel, Tobias, Demand Shifts Due to Salience Effects: Experimental Evidence, March 2016.
Published in: Journal of the European Economic Association, 15 (2017), pp. 626-653.

215 Dewenter, Ralf, Heimeshoff, Ulrich and Thomas, Tobias, Media Coverage and Car Manufacturers' Sales, March 2016.
Published in: Economics Bulletin, 36 (2016), pp. 976-982.

214 Dertwinkel-Kalt, Markus and Riener, Gerhard, A First Test of Focusing Theory, February 2016.

213 Heinz, Matthias, Normann, Hans-Theo and Rau, Holger A., How Competitiveness May Cause a Gender Wage Gap: Experimental Evidence, February 2016.
Forthcoming in: European Economic Review, 90 (2016), pp. 336-349.

212 Fudickar, Roman, Hottenrott, Hanna and Lawson, Cornelia, What's the Price of Consulting? Effects of Public and Private Sector Consulting on Academic Research, February 2016.

211 Stühmeier, Torben, Competition and Corporate Control in Partial Ownership Acquisitions, February 2016.
Published in: Journal of Industry, Competition and Trade, 16 (2016), pp. 297-308.

210 Muck, Johannes, Tariff-Mediated Network Effects with Incompletely Informed Consumers, January 2016.

209 Dertwinkel-Kalt, Markus and Wey, Christian, Structural Remedies as a Signalling Device, January 2016.
Published in: Information Economics and Policy, 35 (2016), pp. 1-6.

208 Herr, Annika and Hottenrott, Hanna, Higher Prices, Higher Quality? Evidence From German Nursing Homes, January 2016.
Published in: Health Policy, 120 (2016), pp. 179-189.

207 Gaudin, Germain and Mantzari, Despoina, Margin Squeeze: An Above-Cost Predatory Pricing Approach, January 2016.
Published in: Journal of Competition Law & Economics, 12 (2016), pp. 151-179.

206 Hottenrott, Hanna, Rexhäuser, Sascha and Veugelers, Reinhilde, Organisational Change and the Productivity Effects of Green Technology Adoption, January 2016.
Published in: Energy and Ressource Economics, 43 (2016), pp. 172–194.

205 Dauth, Wolfgang, Findeisen, Sebastian and Suedekum, Jens, Adjusting to Globalization – Evidence from Worker-Establishment Matches in Germany, January 2016.

204 Banerjee, Debosree, Ibañez, Marcela, Riener, Gerhard and Wollni, Meike, Volunteering to Take on Power: Experimental Evidence from Matrilineal and Patriarchal Societies in India, November 2015.

203 Wagner, Valentin and Riener, Gerhard, Peers or Parents? On Non-Monetary Incentives in Schools, November 2015.

202 Gaudin, Germain, Pass-Through, Vertical Contracts, and Bargains, November 2015.
Published in: Economics Letters, 139 (2016), pp. 1-4.

201 Demeulemeester, Sarah and Hottenrott, Hanna, R&D Subsidies and Firms' Cost of Debt, November 2015.

200 Kreickemeier, Udo and Wrona, Jens, Two-Way Migration Between Similar Countries, October 2015.
Forthcoming in: World Economy.

199 Haucap, Justus and Stühmeier, Torben, Competition and Antitrust in Internet Markets, October 2015.
Published in: Bauer, J. and M. Latzer (Eds.), Handbook on the Economics of the Internet, Edward Elgar: Cheltenham 2016, pp. 183-210.

198 Alipranti, Maria, Milliou, Chrysovalantou and Petrakis, Emmanuel, On Vertical Relations and the Timing of Technology, October 2015.
Published in: Journal of Economic Behavior and Organization, 120 (2015), pp. 117-129.

197 Kellner, Christian, Reinstein, David and Riener, Gerhard, Stochastic Income and Conditional Generosity, October 2015.

196 Chlaß, Nadine and Riener, Gerhard, Lying, Spying, Sabotaging: Procedures and Consequences, September 2015.

195 Gaudin, Germain, Vertical Bargaining and Retail Competition: What Drives Countervailing Power?, May 2017 (First Version September 2015).
Forthcoming in: The Economic Journal.

194 Baumann, Florian and Friehe, Tim, Learning-by-Doing in Torts: Liability and Information About Accident Technology, September 2015.

193 Defever, Fabrice, Fischer, Christian and Suedekum, Jens, Relational Contracts and Supplier Turnover in the Global Economy, August 2015.
Published in: Journal of International Economics, 103 (2016), pp. 147-165.

192 Gu, Yiquan and Wenzel, Tobias, Putting on a Tight Leash and Levelling Playing Field: An Experiment in Strategic Obfuscation and Consumer Protection, July 2015.
Published in: International Journal of Industrial Organization, 42 (2015), pp. 120-128.

191 Ciani, Andrea and Bartoli, Francesca, Export Quality Upgrading under Credit Constraints, July 2015.

190 Hasnas, Irina and Wey, Christian, Full Versus Partial Collusion among Brands and Private Label Producers, July 2015.

189 Dertwinkel-Kalt, Markus and Köster, Mats, Violations of First-Order Stochastic Dominance as Salience Effects, June 2015.
Published in: Journal of Behavioral and Experimental Economics, 59 (2015), pp. 42-46.

188 Kholodilin, Konstantin, Kolmer, Christian, Thomas, Tobias and Ulbricht, Dirk, Asymmetric Perceptions of the Economy: Media, Firms, Consumers, and Experts, June 2015.

187 Dertwinkel-Kalt, Markus and Wey, Christian, Merger Remedies in Oligopoly under a Consumer Welfare Standard, June 2015.
Published in: Journal of Law, Economics, & Organization, 32 (2016), pp. 150-179.

186 Dertwinkel-Kalt, Markus, Salience and Health Campaigns, May 2015.
Published in: Forum for Health Economics & Policy, 19 (2016), pp. 1-22.

185 Wrona, Jens, Border Effects without Borders: What Divides Japan's Internal Trade? May 2015.

184 Amess, Kevin, Stiebale, Joel and Wright, Mike, The Impact of Private Equity on Firms' Innovation Activity, April 2015.
Published in: European Economic Review, 86 (2016), pp. 147-160.

183 Ibañez, Marcela, Rai, Ashok and Riener, Gerhard, Sorting Through Affirmative Action: Three Field Experiments in Colombia, April 2015.

182 Baumann, Florian, Friehe, Tim and Rasch, Alexander, The Influence of Product Liability on Vertical Product Differentiation, April 2015.
Published in: Economics Letters, 147 (2016), pp. 55-58 under the title "Why Product Liability May Lower Product Safety".

181 Baumann, Florian and Friehe, Tim, Proof beyond a Reasonable Doubt: Laboratory Evidence, March 2015.

180 Rasch, Alexander and Waibel, Christian, What Drives Fraud in a Credence Goods Market? – Evidence from a Field Study, March 2015.

179 Jeitschko, Thomas D., Incongruities of Real and Intellectual Property: Economic Concerns in Patent Policy and Practice, February 2015.
Forthcoming in: Michigan State Law Review.

178 Buchwald, Achim and Hottenrott, Hanna, Women on the Board and Executive Duration – Evidence for European Listed Firms, February 2015.

177 Heblich, Stephan, Lameli, Alfred and Riener, Gerhard, Regional Accents on Individual Economic Behavior: A Lab Experiment on Linguistic Performance, Cognitive Ratings and Economic Decisions, February 2015.
Published in: PLoS ONE, 10 (2015), e0113475.

176 Herr, Annika, Nguyen, Thu-Van and Schmitz, Hendrik, Does Quality Disclosure Improve Quality? Responses to the Introduction of Nursing Home Report Cards in Germany, February 2015.
Published in: Health Policy, 120 (2016), pp.1162-1170.

175 Herr, Annika and Normann, Hans-Theo, Organ Donation in the Lab: Preferences and Votes on the Priority Rule, February 2015.
Published in: Journal of Economic Behavior and Organization, 131 Part B (2016), pp. 139-149.

174 Buchwald, Achim, Competition, Outside Directors and Executive Turnover: Implications for Corporate Governance in the EU, February 2015.

173 Buchwald, Achim and Thorwarth, Susanne, Outside Directors on the Board, Competition and Innovation, February 2015.

172 Dewenter, Ralf and Giessing, Leonie, The Effects of Elite Sports Participation on Later Job Success, February 2015.

171 Haucap, Justus, Heimeshoff, Ulrich and Siekmann, Manuel, Price Dispersion and Station Heterogeneity on German Retail Gasoline Markets, January 2015.
Forthcoming in: The Energy Journal.

170 Schweinberger, Albert G. and Suedekum, Jens, De-Industrialisation and Entrepreneurship under Monopolistic Competition, January 2015.
Published in: Oxford Economic Papers, 67 (2015), pp. 1174-1185.

Older discussion papers can be found online at:
http://ideas.repec.org/s/zbw/dicedp.html

Heinrich-Heine-University of Düsseldorf

Department of Economics

Universitätsstraße 1 _ 40225 Düsseldorf
www.dice.hhu.de

ISBN 978-3-86304-266-0