

MANAGEMENT SCIENCE

Volume 63 Number 4 April 2017



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Dynamic Quality Ladder Model Predictions in Nonrandom Holdout Samples

Linli Xu, Jorge M. Silva-Risso, Kenneth C. Wilbur

To cite this article:

Linli Xu, Jorge M. Silva-Risso, Kenneth C. Wilbur (2017) Dynamic Quality Ladder Model Predictions in Nonrandom Holdout Samples. Management Science

Published online in Articles in Advance 28 Jul 2017

. <https://doi.org/10.1287/mnsc.2017.2780>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Dynamic Quality Ladder Model Predictions in Nonrandom Holdout Samples

Linli Xu,^a Jorge M. Silva-Risso,^b Kenneth C. Wilbur^c

^a Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455; ^b A. Gary Anderson Graduate School of Management, University of California, Riverside, Riverside, California 92521; ^c Rady School of Management, University of California, San Diego, La Jolla, California 92093

Contact: linlixu@umn.edu,  <http://orcid.org/0000-0001-6832-0981> (LX); jorge.silva-risso@ucr.edu (JMS-R); kennethcwilbur@gmail.com (KCW)

Received: April 12, 2015

Revised: October 6, 2015; December 7, 2016

Accepted: January 19, 2017

Published Online in *Articles in Advance*:
July 28, 2017

<https://doi.org/10.1287/mnsc.2017.2780>

Copyright: © 2017 INFORMS

Abstract. In light of recent calls for further validation of structural models, this paper evaluates the popular dynamic quality ladder (DQL) model using a nonrandom holdout approach. The model is used to predict data following a regime shift—that is, a change in the environment that produced the estimation data. The prediction performance is evaluated relative to a benchmark vector autoregression (VAR) model across three automotive categories and multiple prediction horizons. Whereas the VAR model performs better in all scenarios in the compact car category, the DQL model tends to perform better on multiple-year horizons in both the midsize car and full-size pickup categories. A supplementary data analysis suggests that DQL model performance in the nonrandom holdout prediction task is better in categories that are more affected by the regime shift, helping to validate the usefulness of the dynamic structural model for making predictions after policy changes.

History: Accepted by Matthew Shum, marketing.

Keywords: automobiles • product quality • dynamic oligopoly competition • product innovation • nonrandom holdout validation

1. Introduction

Structural dynamic oligopoly models have enabled researchers to provide theoretically grounded answers to such previously intractable questions as how competition spurs innovation (Goettler and Gordon 2011, 2014), how environmental regulation affects entry and market power (Ryan 2012), and how demand-smoothing fiscal policy influences market structure (Collard-Wexler 2013). This literature has mostly relied on the “dynamic quality ladder” (DQL) framework introduced by Ericson and Pakes (1995). The majority of papers that investigate firms’ dynamic investments have followed Ericson and Pakes (1995) closely.

The usefulness and appeal of structural dynamic oligopoly models are well established. However, there has been substantial recent debate on the general validity of structural empirical work: untested behavioral assumptions may produce misleading parameter estimates and mistaken policy experiments (e.g., Chintagunta et al. 2006). For example, Angrist and Pischke (2010) have called for further validation of structural models, focusing particularly on the importance of model uncertainty. Rust (2014) pointed out that a key element in building confidence in a structural modeling framework is to rigorously validate the model’s predictions.¹ As Keane (2010, p. 18) noted, the structural literature “[has] tended to pay little attention to the issue of model validation.” The structural dynamic

oligopoly literature, in particular, has not emphasized empirical model validation.²

The purpose of the current paper is to evaluate the dynamic quality ladder framework using the nonrandom holdout validation approach. Specifically, we follow Keane and Wolpin’s (2007) suggestion to evaluate the model’s ability to predict data following a “regime shift”—a change in the environment that produced the estimation data.³ The idea is simple: if a model provides good forecasts for data produced by a different regime than the data used for estimation, the model can be judged to reliably predict the impact of large changes in the environment.

We do not advocate evaluating any class of model based on holdout performance alone. However, holdout validation has played a rare but important role in building confidence in several classes of structural models. For example, McFadden et al. (1977) estimated a random utility model that produced famously accurate predictions of demand for a new transportation service. Ailawadi et al. (2005) estimated a structural manufacturer–retailer Stackelberg model using data prior to a change in manufacturer pricing policy and found that it predicted post-policy-change data better than two alternative models. Misra and Nair (2011) estimated a dynamic structural model of individual salesperson effort allocation and then shifted salesperson incentives in a field experiment. They showed

that, as predicted, the recommended compensation plan led to a 9% increase in overall revenue, providing face validity for dynamic agency theory. Pathak and Shi (2015) estimated a discrete choice model of Boston families' public school choices and predicted how household choices would change after a policy intervention. They have promised to report prediction accuracy after collecting postintervention data. Raval et al. (2016) evaluated the predictive accuracy of different discrete choice demand models using natural disasters that unexpectedly removed hospitals from consumers' choice sets. Their holdout predictions showed the importance of flexibly accounting for consumer heterogeneity and that no single model dominates all others in all cases.

All of these studies estimated structural models using data prior to a regime shift or a policy change, predicted behaviors for the sample after the regime shift, and compared predictions with the actual behavior observed following the regime shift. The regime shift provided a natural (i.e., "nonrandom") opportunity to separate the estimation sample from the validation sample. We contribute to this literature by evaluating structural dynamic oligopoly model predictions after a regime shift.

The context we consider is the U.S. automobile market in three categories: compact car, midsize car, and full-size pickup. The price of gas rose to a record high of more than \$4 a gallon in 2008, increasing consumers' postpurchase automobile costs and reducing automotive demand to varying degrees in accordance with how much each category was affected by the gas price change (Meinero and Rooney 2008). Full-size pickup sales fell 20% in 2008, midsize car sales fell 11%, and compact car sales fell 1%. The change in fuel cost provides a regime shift at which we separate estimation data and nonrandom holdout validation data.

We specify a dynamic quality ladder (DQL) model in which automakers invest in perceived product quality to maximize the net present value of current and future profits.⁴ More investment today increases the chance of realizing a quality improvement in the next period, but the outcome of investment is stochastic. We estimate the model using the mathematical programming with equilibrium constraints (MPEC) approach of Su and Judd (2012). We use data up to 2007 to estimate the model, conduct a counterfactual to predict the perceived quality and market shares for a four-year period from 2008 to 2011, and then compare the model predictions to the observed market outcomes. The nonrandom holdout performance of the DQL model is evaluated relative to a benchmark: the vector autoregression (VAR) model, a common approach to modeling the joint evolution of multiple related time series. All predictions are made by sampling from the

asymptotic distributions of the parameter estimates to account for estimation error.

Predictions are made over two different time horizons: one-year horizons in which year t data are used to predict the perceived quality and market shares in year $t + 1$ and multiple-year horizons in which all holdout predictions are based on the observed market outcomes in 2007. Multiple-year horizons are important for policy makers who are interested in understanding policy impacts that last longer than one year.

The empirical results suggest that the DQL model performs best, relative to a benchmark vector autoregression, in multiple-year horizon predictions in the full-size pickup and midsize car categories. A supplementary analysis of geospatial variation shows that new vehicle sales in these two automotive categories are more susceptible to changes in gasoline price than those in the compact car category. Taken together, the results suggest that the structural model is better suited for prediction purposes over multiple-year horizons in settings where the regime shift has larger effects.

2. Relationship to Prior Literature

The theoretical concept of Markov perfect equilibrium (MPE) was developed to understand dynamic interactions in oligopoly settings (Maskin and Tirole 1987, 1988a, b). From this concept, Ericson and Pakes (1995, hereafter referred to as EP) proposed an empirical framework to model an oligopolistic market over time: each competing firm can change its "state" (e.g., product quality) through investment, and uncertain outcomes of investments determine the evolution of the industry state. Most articles studying industry dynamics of firm investment have followed the EP framework closely.⁵ It has been adapted to a variety of settings, including inpatient hospital services (Gowrisankaran and Town 1997), mergers (Gowrisankaran 1999), capacity accumulation (Besanko and Doraszelski 2004, Besanko et al. 2010), advertising policies (Tan 2006, Qi 2013), research joint ventures (Song 2011), network effects (Markovich 2008), quality innovation in the PC microprocessor industry (Goettler and Gordon 2011), and innovation in the global automobile industry (Hashmi and Van Biesebroeck 2016).⁶

The primary challenge in estimating the EP framework is computational: it requires solving for MPE strategies at every point in the state space. The size of the state space grows exponentially with the number of firms and possible states, leading to a curse of dimensionality; see Doraszelski and Pakes (2007) for a review. This computational burden has limited the empirical work in this literature to mostly rely on computing MPE based on parameters calibrated using available data and simulating counterfactual experiments to understand the impacts of policy changes. To the best of our understanding of this literature,

only a limited number of published studies have estimated the EP framework while solving for the full equilibrium. Gowrisankaran and Town (1997) were the first to adopt the Pakes and McGuire (1994) algorithm to compute MPE while estimating dynamic structural parameters using a nested fixed point approach. They examined hospitals' quality provision strategies in a model with endogenous patient selection of hospitals and evaluated several counterfactual experiments such as changes in Medicare reimbursements and non-profit taxation policy. Goettler and Gordon (2011) used data from the PC microprocessor industry to estimate a DQL model under which both consumers and firms are forward looking. They found that competition lowered innovation rates relative to monopoly but still improved consumer surplus through lower prices. Borkovsky et al. (2017) solved for equilibrium in a DQL model extension in which firms use advertising to build brand equity.⁷ They found that the brand equity depreciation rate plays a key role in determining the value of the brand as well as the value of the firm. Like Borkovsky et al. (2017), we apply the MPEC estimation approach of Su and Judd (2012), which reduces the computational burden of estimation but still solves for the full equilibrium.

3. Models and Estimation

Firms make investment decisions to compete in a quality ladder. Perceived quality (or *quality*, for short) is defined as consumers' perceptions about the product, including preferences for objective product characteristics as well as subjective attitudes about the product. We define quality in this way because (1) product characteristics such as design, brand equity, and perceived reliability are difficult to measure but may be important predictors of demand; and (2) it is the consumer perception of product quality that determines preference, satisfaction, loyalty, and profitability of purchase (Mitra and Golder 2006).

Each product's perceived quality in each year is estimated to be the demand intercept that rationalizes its observed market share, conditional on all prices and products in the market, and is denoted by $\tilde{\omega}$. Note that $\tilde{\omega}$ is continuous, but the dynamic quality ladder model assumes that product quality is discrete; therefore, we partition $\tilde{\omega}$ into discrete levels denoted by ω . The estimation and discretization of quality levels are not the main focus of the paper and therefore are described briefly in Section 4 and fully in Appendix A. Next, we first specify the demand model that generates the quality estimates.

3.1. Consumer Demand

Consumer i gets utility u_{ijt} from purchasing product $j \in \{1, \dots, J\}$ in year t :⁸

$$u_{ijt} = -\alpha TVC_{jt} + \tilde{\omega}_{jt} + \varepsilon_{ijt}, \quad (1)$$

where $TVC_{jt} = p_{jt} + EVFC_{jt}$ is the total vehicle cost (TVC) of product j at time t , which includes the vehicle price, p_{jt} , and the expected vehicle fuel cost, $EVFC_{jt}$. To construct $EVFC_{jt}$, we first define the annual fuel cost, $FC_{jt} = gp_t \times (VMT_t / MPG_j)$, as a function of the gasoline price (gp), vehicle miles traveled (VMT) per year, and vehicle fuel efficiency (miles per gallon, or MPG).⁹ Then, we define $EVFC_{jt}$ as the net present difference between new vehicle j 's annual fuel cost and a reference vehicle j' 's fuel cost over some horizon:

$$EVFC_{jt} = \sum_{t=1}^3 \frac{FC_{jt}}{(1+r)^t} - \sum_{t=1}^3 \frac{FC_{j't}}{(1+r)^t}, \quad (2)$$

where r is the interest rate,¹⁰ the summation assumes a three-year horizon, and the reference vehicle j' is specific to model j . An analysis of millions of automotive purchase records showed that 60% of new auto purchases involve a trade-in, that the trade-in is typically an older version of the same model that is purchased new, and that median trade-in age varies across categories (five years for compact car and midsize car, and six years for full-size pickup).

The perceived quality of product j at time t , $\tilde{\omega}_{jt} = \theta_j + \xi_{jt}$, is decomposed into a set of product intercepts θ_j and an unobserved quality term ξ_{jt} . Assuming that consumer i 's idiosyncratic preferences ε_{ijt} are independently distributed extreme value, and normalizing the utility of the outside option to ε_{i0t} , the market share of product j at time t is

$$s_{jt}(TVC_t, \tilde{\omega}_t) = \frac{\exp(-\alpha TVC_{jt} + \tilde{\omega}_{jt})}{1 + \sum_k \exp(-\alpha TVC_{kt} + \tilde{\omega}_{kt})}, \quad (3)$$

where TVC_t and $\tilde{\omega}_t$ are the vectors of J total vehicle costs and qualities at time t .

3.2. Dynamic Quality Ladder Model

This section specifies a dynamic stochastic oligopoly game in product-quality investments. We think of the choice variable as investments in research and development (R&D) to update product features, but it also may reflect advertising expenditures that influence consumer perceptions of product quality.¹¹ The game is played by a fixed number of firms over an infinite horizon.¹² Each firm $j \in \{1, \dots, J\}$ is described by its product's quality level, or "state," $\omega_{jt} \in \Omega = \{1, 2, \dots, \bar{\omega}\}$ in each year $t \in \{1, 2, \dots, \infty\}$. At any point of time, the industry is completely characterized by an "industry state" vector $\omega_t = (\omega_{1t}, \dots, \omega_{Jt})$, which describes all products' quality levels.

Firms choose prices and investments to maximize their expected discounted profits. Investment is a dynamic decision because more investment today increases the chance of realizing a quality improvement in the next period. Quality is also subject to an exogenous depreciation shock that is common to all firms,

to allow for the possibility of improvement in the “outside option.” Prices are set conditional on the industry state vector. Hence, each firm’s per-period profit is determined by all firms’ current qualities and prices, and it is treated as a primitive of the dynamic stochastic game.

The timing of the game is as follows. At the start of each time period, firms observe the current industry state and simultaneously choose investments x_{jt} to improve next period’s quality. Next, firms compete in the product market and simultaneously set Nash equilibrium prices p_{jt} to maximize per-period profits. At the end of the period, the outcomes of the investment and the industry-wide shock are realized, and as a result, the industry state is updated.

3.2.1. States and Transitions. Improving product quality is modeled as a time-consuming and uncertain process, with a probability of success that increases with the investment x_{jt} . Following Pakes and McGuire (1994), we restrict the outcome of investments, v_{jt} , to be either 0 (if innovation fails) or 1 (if innovation succeeds) to ensure a closed-form solution for optimal investment. The discrete distribution of v_{jt} is given by

$$v_{jt} = \begin{cases} 1 & \text{with probability } \frac{\rho x_{jt}}{1 + \rho x_{jt}}, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where $\rho > 0$ represents the effectiveness of investment. The quality state evolves as follows:

$$\omega'_{jt} = \omega_{jt} + v_{jt} - \eta_t, \quad (5)$$

where ω'_{jt} is the realized state at the end of period t . Firm j ’s state cannot improve without investment and may decay with the realization of the common industry shock η_t . This shock may represent technological improvements or other variables that lead consumers to favor substitutable products. It induces positive correlation among competing firms’ profits, as is often observed in market data. Similar to most papers applying the dynamic quality ladder framework, we assume $\eta_t = 1$ with an exogenous probability δ and $\eta_t = 0$ with probability $1 - \delta$.

3.2.2. Per-Period Profits. Within each time period, conditional on the current state of the market ω , each firm competes for a mass M^{13} of consumers by setting its price, p_{jt} , to maximize its per-period profit:

$$\max_{p_{jt}} \pi_{jt} = (p_{jt} - mc)Ms_{jt}(p_t, \tilde{\omega}_t), \quad (6)$$

where $mc = \exp(\gamma)$ is the constant marginal cost of production across firms. Given that the per-period profits depend only on the current prices and quality states, we can compute all per-period profits that correspond to equilibrium prices at every possible industry state and take them as the primitives of the dynamic investment problem.

3.2.3. Dynamic Investment Decisions and Equilibrium.

At the beginning of each time period, firms observe the industry state ω_t and simultaneously choose investments x_{jt} to maximize the expected discounted value of net cash flows:

$$\max_{x_{jt}} E \left\{ \sum_{\tau=t}^{\infty} \beta^{\tau-t} (\pi_j(\omega_\tau) - x_{j\tau}) \mid \omega_t \right\}, \quad (7)$$

where β is the discount factor¹⁴ and $\pi_j(\omega_\tau)$ is firm j ’s per-period profit before deducting the investment cost at time τ . As in EP, we focus on stationary Markov perfect equilibria, which are only a function of the current state and do not depend on calendar time, so time subscripts are omitted in equations hereafter in this subsection. The maximization problem in (7) implies that the Bellman equation for firm j is

$$V(\omega_j, \omega_{-j}) = \max_{x_j} \{ \pi(\omega_j, \omega_{-j}) - x_j + \beta E[V(\omega'_j, \omega'_{-j}) \mid \omega] \}, \quad (8)$$

where ω_{-j} is a vector of all competitors’ quality levels. The expectation on the right-hand side of Equation (8) is taken over the probability distribution of both firm j ’s own state and its competitors’ states in the next time period. As shown in Appendix B, solving the first-order conditions of the maximization problem on the right-hand side of Equation (8) yields the closed-form equilibrium investment policy:

$$x_j^*(\omega_j, \omega_{-j}) = \max \left\{ 0, \frac{-1 + \sqrt{\beta \rho [W(1 \mid \omega) - W(0 \mid \omega)]}}{\rho} \right\}, \quad (9)$$

if $W(1 \mid \omega) \geq W(0 \mid \omega)$ and $x_j^*(\omega_j, \omega_{-j}) = 0$ otherwise. Note $W(v_j \mid \omega) \equiv \sum_{\omega'_{-j}} \sum_{\eta} V(\omega_j + v_j - \eta, \omega'_{-j}) q(\omega'_{-j} \mid \omega, \eta) p(\eta)$ is firm j ’s expected payoff conditional on the outcome of its investment $v_j = \{0, 1\}$, and $q(\omega'_{-j} \mid \omega, \eta)$ represents firm j ’s expectation of its competitors’ future states ω'_{-j} .

A pure-strategy MPE involves value functions $V(\omega)$ and policy functions $x(\omega)$ such that (i) given policy functions, value functions solve the Bellman equation (8) for every firm j , and (ii) given value functions, policy functions solve the maximization problem on the right-hand side of Equation (8) for every firm j . The functional form of the state transition probabilities (Equation (5)) satisfies the unique investment choice admissibility condition derived by Doraszelski and Satterthwaite (2010), guaranteeing that optimal investment strategies are unique.

The quality state space must be bounded to numerically solve for the MPE. Therefore, we assume $\omega_j \leq \bar{\omega} \forall j$, where $\bar{\omega}$ is the highest level of quality that any firm could theoretically reach.¹⁵ We discuss the selection of $\bar{\omega}$ in Section 3.3.3.

3.3. Estimation

Parameters to be estimated include the static parameters of the demand model and the per-period profits $\{\alpha, \theta, \gamma, \hat{\omega}_t\}$ and the dynamic parameters $\{\rho, \delta\}$.

3.3.1. Estimation of Static Parameters. It is important to have a precise estimate of the price responsiveness parameter in the demand model in order to separate perceived quality from prices. We used the large, granular data set described in Appendix A.1 to estimate the price responsiveness parameter α . Estimates of θ and $\hat{\omega}$ are conditional on $\hat{\alpha}$ and estimated together with firms' marginal cost parameters γ using the generalized method of moments (GMM). First-order conditions of the maximization problem in Equation (6) are derived and used in estimating these parameters.

The estimated perceived quality $\hat{\omega}$ is discretized into partitioned levels ω following the state transition process (Equation (5)). Details of the discretization procedure are described in Appendix A.2. Taking the static parameter estimates $\{\hat{\alpha}, \hat{\theta}, \hat{\gamma}\}$ as inputs, we solve the systems of first-order conditions derived from Equation (6) for the Bertrand–Nash equilibrium prices at every possible industry state ω and then compute each firm's equilibrium per-period profits $\pi(\omega)$ that correspond to those equilibrium prices. Finally, the equilibrium prices and per-period profits are taken to the dynamic parameter estimation as inputs.

3.3.2. Estimation of Dynamic Parameters. Traditional estimation of the structural dynamic parameters ρ and δ requires solving the MPE investment policy function $x(\omega)$ and value function $V(\omega)$ for each firm at each possible industry state. The nested fixed point (NFXP) approach of Rust (1987) solves the MPE for each guess of ρ and δ , requiring substantial computational burden to estimate the model, especially when the number of possible states and the number of competing firms are both large. To alleviate this burden, we apply the MPEC approach proposed by Su and Judd (2012).¹⁶ MPEC does not require solving the MPE for every guess of the dynamic parameters; instead, it treats policy functions $x(\omega)$ and value functions $V(\omega)$ as parameters to be estimated, subject to equilibrium condition constraints.

More specifically, the MPEC constrained optimization problem is formulated as

$$\begin{aligned} & \max_{\rho, \delta, x, V} L(\rho, \delta, x, V) \\ & \text{subject to} \\ & (1) \quad V(\omega_j, \omega_{-j}) = \pi(\omega_j, \omega_{-j}) - x_j, \\ & \quad \quad \quad + \beta E[V(\omega'_j, \omega'_{-j}) | \omega], \quad \forall j, \quad (10) \\ & (2) \quad \beta \sum_{v_j} W(v_j | \omega) \frac{\partial p(v_j)}{\partial x_j} - 1 = 0, \quad \forall j, \\ & (3) \quad x_j \geq 0, \quad \forall j, \end{aligned}$$

where $L(\rho, \delta, x, V)$ is the full likelihood function, which is presented in Appendix C to conserve space. For a given value of the parameters ρ and δ , the model predicts quality changes for each firm in each time period according to the state transition process defined in Equation (5). This transition process depends on the realized firm-specific investment outcomes and industry-wide shocks. The likelihood function is formulated to match the predicted quality changes to the observed quality changes.

The policy functions $x(\omega)$ and value functions $V(\omega)$ are treated as parameters required to satisfy three sets of constraints imposed on the likelihood function. The first set is the Bellman equation defined by Equation (8). The second set is the first-order conditions derived from the right-hand side of the Bellman equation (as shown in Equation (B.4) in Appendix B), where $p(v_j)$ represents the probability distribution of the investment outcome v_j . The third set constrains investments to be nonnegative. Although estimating the value functions and policy functions increases the parameter space, the Jacobian matrix of the constraints is very sparse. This sparseness, along with analytical tractability, speeds optimization and allows (10) to be maximized using a Newton–Raphson method rather than a quasi-Newton algorithm.

The identification strategy is the following: the propensity of all firms' qualities to deteriorate or fail to improve in the same period identifies δ (i.e., the probability of the quality improvement of the outside good), because this industry-wide shock affects all firms in the same market. The effectiveness of investment, ρ , is identified by individual firms' tendency to realize quality improvements in the absence of a common industry shock.

An ideal data set might contain direct observations of firms' strategic R&D investments x . However, firms typically treat product-specific R&D expenditures as trade secrets, making it very difficult to obtain such data for multiple competing firms. As a result, we use the optimal investment policies x^* in place of the unobserved investment data, assuming rational investment behavior.

3.3.3. Estimation Exercise Based on Simulated Data.

We used the estimation algorithm to recover known values of dynamic parameters to investigate its properties and ensure the code was functional. We followed the Gauss–Jacobi algorithm described in Pakes and McGuire (1994) to compute the value functions and policy functions in equilibrium for a set of known parameter values. We did this using many different starting values to numerically verify equilibrium existence and uniqueness.

Starting from a given initial quality state, we obtained the corresponding optimal investment from the equilibrium policy functions and then computed the

Table 1. DQL Model Simulation Exercises to Recover Known Parameters

Parameters	No. of time periods	No. of firms		
		2	3	4
δ (std. dev.)	50	0.35 (0.06)	0.29 (0.06)	0.29 (0.08)
	100	0.32 (0.04)	0.28 (0.03)	0.29 (0.03)
ρ (std. dev.)	50	2.73 (0.68)	2.16 (0.27)	2.15 (0.17)
	100	1.91 (0.30)	2.01 (0.27)	1.99 (0.24)

Notes. True values for ρ and δ are 2 and 0.3, respectively. Marginal cost is assumed to be 5, discount factor β is assumed to be 0.925, and the number of quality states is 18.

next period's product quality by simulating investment outcomes and the realization of the industry shock. This simulation process was repeated to generate simulated data sets with 50 or 100 time periods and different numbers of firms ($J = 2, 3, 4$). For each synthetic data set, we used the estimation procedure described above to recover ρ and δ .

In the simulated data sets, we observe firms' investment levels. However, because we do not observe investment levels in the market data, we did not use the investment levels when recovering the known parameter values. Table 1 reports the means and standard deviations of the ρ and δ estimates. It shows that the MPEC approach recovers known parameter values fairly accurately with relatively small sample sizes. It also confirms that having longer time periods and/or more firms in the data facilitates accuracy in parameter estimation.

We also used simulations to gauge the sensitivity of the parameter estimates to the assumptions about the unobserved upper bound of the quality space, $\bar{\omega}$. Specifically, for all values of $\bar{\omega}$ up to 18, we generated synthetic data from known parameters and then estimated those parameters using the synthetic data, in the manner described above. We found that the parameter estimates displayed no apparent sensitivity to the choice of $\bar{\omega}$, so long as $\bar{\omega}$ was not reached within the simulated data set.¹⁷ Therefore, we set $\bar{\omega}$ to be at least one level above the maximum observed quality in the synthetic data.

3.4. Vector Autoregression Model

The dynamic quality ladder model is evaluated relative to a benchmark. The VAR model provides a familiar framework capable of describing joint evolution of competing firms' quality levels. The VAR allows for firm-specific autocorrelation parameters, cross-firm correlations in quality, and correlated errors across firms and time periods.¹⁸ The VAR model assumes a

first-order autoregressive process on the evolution of perceived quality to mimic the state transition process of the DQL model:

$$\begin{bmatrix} \tilde{\omega}_{1t} \\ \vdots \\ \tilde{\omega}_{Jt} \end{bmatrix} = A \begin{bmatrix} \tilde{\omega}_{1,t-1} \\ \vdots \\ \tilde{\omega}_{J,t-1} \end{bmatrix} + \begin{bmatrix} e_{1t} \\ \vdots \\ e_{Jt} \end{bmatrix}, \quad (11)$$

where $t = 2, \dots, T$ indexes time periods and T is the last period in the estimation sample; A is a $J \times J$ matrix of parameters measuring the correlation between current and past quality levels across J firms; and the errors $e_{jt} \sim N(0, \Sigma)$ are distributed Normal with a covariance matrix Σ to be estimated along with the parameters in A .

4. Data and Perceived Quality Estimates


The empirical context of the current paper is the automotive industry, which has historically contributed 3%–3.5% to gross domestic product. General Motors, Inc., introduced the practice of annual model-year design changes in the early 1920s. This regular cycle of innovation leads automakers to invest 4% of their U.S. revenues, about \$18 billion annually, on R&D (Hill et al. 2014).

The analysis focuses on three automotive categories: compact car, midsize car, and full-size pickup.¹⁹ These three categories were selected because they have relatively long histories in the industry and the primary competitors within these categories have been relatively stable. In each of the three categories, we identified the four largest manufacturers: in compact car and midsize car, they were Toyota, Honda, Ford, and Chevrolet; and in full-size pickup, the top four firms were Ford, Chevrolet, Dodge, and GMC (General Motors Truck Company).²⁰

We employ two automotive sales data sets to estimate model parameters. Category-specific price responsiveness parameters were estimated using 7.9 individual automotive sales transactions from 1997 to 2012. This exercise yielded highly precise estimates of $\hat{\alpha}$, allowing us to separate price response from perceived equality estimates.²¹ The detailed transaction data were not used to estimate all perceived quality levels because they only date back to 1996, providing insufficient degrees of freedom to estimate the dynamic parameters δ and ρ .

To estimate perceived quality, we constructed a larger T data set of unit sales and list prices for each model-year from *Ward's Automotive Yearbook*. Data were available from 1975 to 2011 for compact car and from 1984 to 2011 for midsize car and full-size pickup categories. All list prices were converted to 1999 dollars using the Bureau of Labor Statistics Consumer Price Index. Data on gasoline price and MPG by vehicle model were collected from the U.S. Department of

Figure 1. Market Shares and Prices (in 1999 Dollars) by Category



Energy. VMT data were gathered from the Bureau of Transportation Statistics.

Figure 1 shows the market shares (left panel) and prices (right panel) of the four major manufacturers in each of the three categories. In compact car and midsize car, the market shares of Toyota and Honda display a slight upward trend, whereas the market shares of Ford and Chevrolet remained relatively stable. Similarly, the market shares of Ford and Chevrolet in the full-size pickup category remained stable over time, whereas the market shares of Dodge and GMC showed slight upward trends.²²


Figure 2 displays the perceived quality estimates $\tilde{\omega}$. The y axes represent the estimated quality in continuous space. The dotted horizontal lines separate the

ranges used to discretize the quality estimates ω . The discretization procedure is explained in Appendix A.2. The number of total discrete quality levels chosen in compact car, midsize car, and full-size pickup is eight, nine, and seven, respectively. In all three categories, we observe comovement of firms' product quality in some periods and divergence in quality movements in other periods.

5. Empirical Results

Our central research focus is to examine the performance of the DQL model in nonrandom holdout tests, to understand its ability to predict market changes after regime shifts. The regime shift analyzed is the

Figure 2. Estimated Perceived Quality with Discretization Boundaries



2008 rise in gas price, which allows us to conduct a counterfactual by solving equilibrium investment and market outcomes under the new regime. We estimate the DQL model in the three automotive categories using data up to the year 2007. The expected vehicle fuel cost (EVFC) enters the estimation of dynamic

parameters as an exogenous fixed-state variable. The dynamic parameters were estimated using the average EVFC across the in-sample periods. Given the parameter estimates, we then solve for the new equilibrium outcomes when the EVFC is changed to the average across the holdout periods—that is, the new regime.

Then, for each draw of the DQL parameter estimates from its asymptotic distribution, perceived qualities are predicted based on the new equilibrium outcomes, and then market shares are generated by solving the demand system based on the predicted quality levels and the corresponding equilibrium prices. Predictions of perceived quality and market shares are then compared to observations of market outcomes.²³

For the holdout years 2008–2011, we produce two types of predictions: one-year horizon predictions, in which year t data are used to predict perceived quality and market share in year $t + 1$, and multiple-year horizon predictions, in which we use observed market outcomes in 2007 to predict the perceived quality and market shares from 2008 until 2011. Multiple-year horizon predictions are useful because many policy changes have effects that are realized over multiple-year horizons, requiring policy makers to account for long-term impacts on agents' investments and subsequent investment payoffs.

The mean absolute error (MAE) is used to evaluate the accuracy of the predictions, and the DQL predictions are evaluated relative to VAR.²⁴ As the DQL model generates quality predictions in discrete levels, the prediction error is calculated as the difference between the midpoint of the range in which the predicted quality level resides and the actual quality in the holdout periods.²⁵

Table 2 presents ratios of DQL MAE to VAR MAE, based on quality-level predictions, for one-year horizon predictions and multiple-year horizon predictions. A ratio of less (more) than 1 means that the DQL model predictions exhibit a smaller (larger) MAE than the

Table 2. Quality Prediction MAE Ratio (DQL/VAR) in 2008–2011 Holdout Sample

	One-year horizon				Multiple-year horizon			
	2008	2009	2010	2011	2008	2009	2010	2011
Compact car	1.8	1.8	2.2	2.0	—	1.9	1.9	2.0
Midsize car	1.0	0.7	0.6	0.8	—	0.7	0.7	0.9
Full-size pickup	1.7	2.0	1.5	1.4	—	0.9	0.7	0.6

Note. When the ratio is less than 1 (as indicated in bold), VAR model predictions exhibit a larger MAE than DQL model predictions.

Table 3. Market Share Prediction MAE Ratio (DQL/VAR) in 2008–2011 Holdout Sample

	One-year horizon				Multiple-year horizon			
	2008	2009	2010	2011	2008	2009	2010	2011
Compact car	1.5	1.1	1.6	1.7	—	1.6	1.7	1.7
Midsize car	0.8	0.8	0.7	0.7	—	0.8	0.8	0.8
Full-size pickup	0.8	1.1	1.3	0.9	—	0.6	0.6	0.5

Note. When the ratio is less than 1 (as indicated in bold), VAR model predictions exhibit a larger MAE than DQL model predictions.

Table 4. Dynamic Quality Ladder Model Parameter Estimates

	Compact car	Midsize car	Full-size pickup
Price			
Est.	−0.52**	−0.18**	−0.13**
(Std. err.)	(0.02)	(0.01)	(0.01)
Mean price elasticity			
Est.	−6.93	−3.13	−2.63
(Std. dev.)	(0.27)	(0.26)	(0.24)
Marginal cost			
Est.	8.91**	10.36**	6.20*
(Std. err.)	(0.45)	(0.32)	(0.86)
Delta			
Est.	0.21**	0.29	0.50**
(Std. err.)	(0.08)	(0.19)	(0.17)
Mean probability of successful investment (%)			
Est.	54	66	71
(Std. dev.)	(18)	(16)	(16)

*Significant at the 95% level; **significant at the 99% level.

VAR model predictions. Table 3 presents similar information for market share predictions. Figure 3 depicts the information in Tables 2 and 3.

Section 5.1 discusses parameter estimates. Section 5.2 interprets the DQL model's prediction performance across categories and between the two time horizons. Section 5.3 presents a supplemental analysis to validate the results. Section 5.4 reports model fit statistics and welfare analyses.


5.1. Parameter Estimates

Table 4 presents the static demand parameter estimates as well as the parameter estimates governing firms' investment decisions in the dynamic quality ladder model. The mean price responsiveness parameters translate to price elasticities of about −7 in the compact car category, −3.1 in the midsize car category, and −2.6 in the full-size pickup category, showing that demand for larger automobiles is relatively less responsive to vehicle price.

Among the dynamic quality ladder model parameters, the estimated probability of an improvement in the outside option (δ) ranges from 21% to 50%. The investment effectiveness parameter ρ maps investment expenditure into the probability of successful investment outcomes; across the range of estimated investment levels, the mean probability of successful innovation is between 54% and 71%, which suggests that firms consistently invest in R&D and realize product-quality improvements.

Table 5 presents the parameter estimates from the VAR model. Most scenarios show significant correlations between the lagged qualities and the current qualities within each firm. Only a few cases showed significant correlations between a firm's current quality and its competitors' lagged qualities.

Figure 3. MAE Ratios (DQL/VAR) for 2008–2011 Holdout Sample



5.2. Nonrandom Holdout Results

Figure 3 illustrates the ratios of the quality level MAEs (Figure 3(a)) and market share MAEs (Figure 3(b)) across three categories in each of the holdout periods (2008–2011) for both one-year and multiple-year time horizons.²⁶

In the compact car category, MAE ratios are always above 1.0, for both one-year and multiple-year prediction horizons, and for predictions about both quality level and market shares. On the other hand, for the midsize car category, we see the reverse: in nearly all cases, the DQL model predicted market share and quality levels better than the VAR. The lone exception is the quality level prediction for 2008, in which the two models predicted about equally well.

In the full-size pickup category, the two prediction horizons diverged somewhat. In terms of quality levels, the VAR always predicted better than the DQL in one-year prediction horizons. However, for multiple-year horizons, the DQL model predicted better than the VAR. In terms of market shares, the DQL model's performance was slightly better, but the main result

persists that the relative prediction performance of the DQL was far better under multiple-year horizons than the one-year horizon.

Looking at the results between the two time horizons, the DQL model predicts better for two out of three categories, midsize car and full-size pickup, under multiple-year horizon, whereas the VAR model performs better for the compact car and full-size pickup categories under one-year horizon predictions. This pattern may imply that the DQL model is more suitable for predicting market outcomes over a longer time horizon. This is consistent with the assumptions of each model: the DQL considers the long-run stationary equilibrium, and the VAR model in the current study only included the first lag of the dependent variables.

The overall result, across categories and prediction horizons, suggests that the relative performance of the DQL model is somewhat situational. The next section investigates a possible explanation for the performance results: the importance of the gasoline price shock within each category.

Table 5. VAR Model Parameter Estimates

	Compact car			
	<i>Honda</i> _{<i>t</i>−1}	<i>Toyota</i> _{<i>t</i>−1}	<i>Ford</i> _{<i>t</i>−1}	<i>Chevrolet</i> _{<i>t</i>−1}
<i>Honda</i>	0.75** (0.19)	0.34 (0.19)	−0.04 (0.17)	−0.03 (0.14)
<i>Toyota</i>	0.14 (0.17)	0.78** (0.17)	0.09 (0.16)	0.00 (0.13)
<i>Ford</i>	0.26 (0.20)	0.08 (0.20)	0.65** (0.19)	0.00 (0.15)
<i>Chevrolet</i>	0.48 (0.20)	−0.31 (0.19)	0.16 (0.18)	0.66** (0.15)

	Midsize car			
	<i>Honda</i> _{<i>t</i>−1}	<i>Toyota</i> _{<i>t</i>−1}	<i>Ford</i> _{<i>t</i>−1}	<i>Chevrolet</i> _{<i>t</i>−1}
<i>Honda</i>	0.46 (0.30)	0.33 (0.21)	−0.39 (0.23)	0.26 (0.19)
<i>Toyota</i>	0.51 (0.39)	0.02 (0.27)	0.06 (0.30)	0.14 (0.24)
<i>Ford</i>	−0.10 (0.55)	0.22 (0.38)	0.24 (0.42)	0.13 (0.34)
<i>Chevrolet</i>	−0.05 (0.40)	0.21 (0.28)	−0.15 (0.31)	0.06 (0.25)

	Full-size pickup			
	<i>Ford</i> _{<i>t</i>−1}	<i>Chevrolet</i> _{<i>t</i>−1}	<i>Dodge</i> _{<i>t</i>−1}	<i>GMC</i> _{<i>t</i>−1}
<i>Ford</i>	0.91** (0.20)	0.22 (0.25)	0.04 (0.06)	−0.50* (0.23)
<i>Chevrolet</i>	0.11 (0.20)	0.90** (0.25)	0.01 (0.06)	−0.06 (0.23)
<i>Dodge</i>	−0.64 (0.33)	0.84 (0.40)	0.97** (0.10)	−0.26 (0.38)
<i>GMC</i>	−0.10 (0.20)	0.18 (0.24)	0.08 (0.06)	0.73** (0.23)

*Significant at the 95% level; **significant at the 99% level.

5.3. Supplemental Analysis of Increase in Gas Price

What explains the differential nonrandom holdout performance between the two models across three categories? It is possible that the difference corresponds to the extent to which each category is exposed to the rise of gasoline price in 2008.

To investigate the explanation further, we gathered additional data to estimate category-specific impacts of the increase in gas price on automotive demand. In 2008, the sales of full-size pickup and midsize car in the U.S. market fell by 20% and 11%, respectively; sales of compact cars only fell by 1%. These results, along with contemporaneous press accounts, suggest that the latter category is less sensitive to gas price overall (Associated Press 2008). It stands to reason that demand for large automobiles is more sensitive to gasoline price shocks than demand for more fuel-efficient cars. To examine whether this is true, we matched the detailed automotive transaction sales data described in Appendix A.1 to weekly gas prices for 10

Table 6. The Effects of Gas Price on Sales by Category

Variable	Estimate (std. err.)
<i>Compact car</i> × <i>Lagged gas price</i>	−0.07* (0.03)
<i>Midsize car</i> × <i>Lagged gas price</i>	−0.10** (0.03)
<i>Full-size pickup</i> × <i>Lagged gas price</i>	−0.65** (0.03)

Notes. The estimates of brand and city fixed effects were excluded for brevity. The R-square is 0.52.

*Significant at the 95% level; **significant at the 99% level.

major cities from 2003 to 2012. The 10 cities were chosen according to the availability of geographic gas price data.²⁷ We regress weekly model sales on brand dummies, city dummies, and an interaction between automotive category and log of gas price.²⁸ This regression uses geospatial variation as well as time-series variation to estimate the impact of gas price on sales of new automobiles.

Table 6 shows the category-specific effects of gas price on new auto sales. As expected, gas prices tend to reduce automotive demand in all three categories. However, the point estimates of lagged gas price on new vehicle sales are larger for midsize car and full-size pickup than in the compact car category.

Overall, this separate longitudinal analysis supports the idea that the DQL model's relative prediction performance in nonrandom holdout exercises was better in categories that were more affected by the regime shift. It does not explain the prediction results perfectly; the VAR model was relatively better in the full-size pickup category in one-year horizon prediction exercises. However, it does offer limited evidence for the idea that the DQL model performs better in multiple-year predictions when regime shifts are larger, as it did in midsize car and full-size pickup categories.²⁹

5.4. Goodness of Fit and Welfare Analysis

A common practice to evaluate a particular model is to look at how well the model fits the estimation sample. On the basis of the parameter estimates, perceived quality and market shares were predicted for the estimation sample periods, and the average MAE across the holdout periods was computed to measure the in-sample fit in each scenario. Table 7 presents the findings, showing that the VAR model fits the estimation data better than the dynamic quality ladder model across three categories.³⁰ These results suggest that a model that fits the data better does not always predict better out of sample, consistent with general concerns about overfitting and calls for cross-validation (e.g., Picard and Cook 1984, Shao 1993, Chintagunta et al. 2006, Greene 2012).

Table 7. Model In-Sample Fit Statistics (MAE)

	VAR	DQL
Market share predictions		
Compact car	0.04	0.06
Midsize car	0.02	0.03
Full-size pickup	0.02	0.04
Quality level predictions		
Compact car	0.28	0.74
Midsize car	0.13	0.37
Full-size pickup	0.12	0.44

We can also use the dynamic quality ladder model to predict consumer welfare and per-period profits for the nonrandom holdout sample periods and compare them to the VAR predictions. For both models, we relied on the demand model and the per-period profit function specified in Section 3 to calculate consumer surplus and firm profits in each of the out-of-sample periods. We solved for the firms' optimal prices by inserting the predicted quality into the demand model, computed the per-period profits at the implied prices, and converted consumers' utilities to monetary values using the price responsiveness parameter α . Finally, we computed the average consumer surplus for the holdout sample 2008–2011 and the average per-period profits across firms and the holdout sample periods.

Table 8 shows the average consumer surplus and the average firm per-period profits under each model along with the actual values calculated based on the observed qualities and prices.³¹ Two findings are worth discussion: first, the full-size pickup category generates the highest per-period profits and consumer surplus, followed by midsize car and compact car in descending order. Second, profits based on both models are reasonably close to the per-period profits based on the actual quality levels. However, we want to avoid

Table 8. Average Consumer Surplus and Average Firm Per-Period Profits

	Actual	VAR (std. dev.)	DQL (std. dev.)
Average consumer surplus (\$ billion)			
Compact car	3.8	6.0 (3.1)	10.2 (4.1)
Midsize car	88.4	105.1 (19.0)	113.0 (14.1)
Full-size pickup	114.2	112.0 (9.4)	120.7 (17.3)
Average firm static profits (\$ billion)			
Compact car	2.1	1.5 (0.2)	1.9 (0.2)
Midsize car	10.8	8.7 (0.6)	8.3 (0.3)
Full-size pickup	11.3	9.8 (0.5)	9.3 (0.3)

excessive interpretation of these comparisons as the demand model and the per-period profit function used to calculate equilibrium prices, per-period profits, and consumer welfare are not mechanisms built into the VAR model.

6. Discussion

In light of recent calls for further validation of structural models, the present article evaluates the dynamic quality ladder model using the nonrandom holdout approach proposed by Keane and Wolpin (2007). The DQL model's ability to predict data after a regime shift is evaluated relative to a benchmark VAR model. Looking across three automotive categories and two time horizons, we found that the predictive performance of the DQL model performs better in multiple-year predictions when the size of regime shifts experienced in the category is larger.

This paper has a number of limitations which suggest avenues for further investigation. The most important caveat is that only a single, stylized dynamic quality ladder model has been evaluated. We chose to do this because of the overwhelming prevalence of direct application of the EP framework in the literature. However, one could reasonably question whether the DQL model tested in the current paper could be adapted to fit the automotive industry more closely. For example, both advertising and R&D investment could influence the evolution of perceived product quality. If data on advertising were available, a model could accommodate two dynamic decisions affecting quality transitions. Or if major automotive model revisions take more than a year to complete, a model could allow for time lags between periods when investment decisions are made and when quality improvements are realized. However, although these modifications may bring the model closer to reality, they are not theoretically straightforward as they may influence the conditions that guarantee the existence of equilibrium. They remain as interesting directions for future research to extend the EP framework.

Another possible direction for future research is to test the DQL model in nonautomotive settings or to test extensions of the DQL model that have been proposed by other authors. For example, Borkovsky et al. (2017) allow for multistep product-quality movements, Goettler and Gordon (2011) model firm-specific innovation parameters, Gowrisankaran and Town (1997) model investment cost and account for endogenous entry/exit decisions, and Markovich (2008) considers consumers' forward-looking behavior and its consequences for durable goods.

Overall, we believe we have provided some initial results about the reliability of policy experiments using the popular dynamic quality ladder model framework. Broadly speaking, our results show that the

DQL model performs well in multiple-year nonrandom holdout exercises when regime shifts are larger but that a benchmark VAR model performs better in one-year predictions. We suspect that further investigations of model validity will increase confidence in the modeling paradigm and continue to prove that dynamic oligopoly models are valid tools to estimate theoretical parameters and answer policy questions.

Acknowledgments

The authors thank Hai Che, Andrew Ching, Anthony Dukes, Brett Gordon, Karsten Hansen, Masakazu Ishihara, Ahmed Khwaja, Hongju Liu, Lan Luo, S. Siddarth, Che-Lin Su, Botao Yang, Song Yao, Yi Zhu, and numerous seminar audiences for helpful discussions and suggestions.

Appendix A. Price Responsiveness Parameter Estimation and Quality Discretization

It is important to have a precise estimate of the price responsiveness parameter in the demand model in order to separate perceived quality from prices. This led us to use a separate and more granular data set to estimate consumers' price responsiveness, as described in Section A.1. The perceived-quality estimates are discretized in accordance with the specification of the DQL model. The discretization procedure is discussed in Section A.2.

A.1. Price Responsiveness Parameter Estimation

The price responsiveness parameter could not be precisely estimated using the annual data described in Section 4. Instead, we use detailed transaction-level sales data of the three categories (compact car, midsize car, and full-size pickup) in the U.S. market from 1996 to 2012, which was collected by the Power Information Network (PIN), a division of J.D. Power and Associates. For each transaction, the data report the transaction date, type (lease, finance, or cash), pricing terms (e.g., down payment, rebate, annual percentage rate), and vehicle characteristics (make, model, and model-year). Within each category, we analyzed transactions of the most dominant model for each of the four major manufacturers, as discussed in Section 4. This resulted in sample sizes of 2.5 million observations for compact car, 2.7 million observations for midsize car, and 2.7 million observations for full-size pickup.

Following Xu et al. (2014), we first constructed the transaction price as the net present value of each transaction across three types: lease, dealer finance, and "cash," in which customers finance through an outside lender or pay in full. We then aggregated the transaction-level data on sales for each model in each category to the weekly level and constructed a panel of average prices at the firm-category-week level to estimate the price responsiveness parameter.

Specifically, consumer i purchasing from product j in week w gets utility

$$u_{ijw} = \beta_j + \alpha TVC_{jw} + \zeta_{jw} + e_{ijw}, \quad (\text{A.1})$$

where $TVC_{jw} = p_{jw} + EVFC_{jt}$ is the total vehicle cost, which includes the vehicle price, p_{jw} , and the expected vehicle fuel cost, $EVFC_{jt}$. The cost $EVFC_{jt}$ is constructed in the same way

as described in Equation (2) of Section 3. The only difference is that the annual fuel cost, $FC_{jt} = gp_w \times (VMT_t/MPG_j)$, is computed based on gas price in week w of year t .

Product intercept β_j captures consumers' mean preferences for product j ; ζ_{jw} represents any unobserved weekly departures from the mean product preferences, which is assumed to be normally distributed with a variance to be estimated. Similar to the way we conceptualize the term $\tilde{\omega}_{jt} = \theta_j + \xi_{jt}$ in Equation (1) as the perceived quality of product j in year t , the term $\beta_j + \zeta_{jw}$ can be considered as the perceived quality of product j in week w . Given the assumption that the consumer idiosyncratic preference e_{ijw} is independently and identically distributed of Type I extreme value,³² the market share for model j in week w is given by³³

$$s_{jw} = \frac{\exp(\beta_j + \alpha TVC_{jw} + \zeta_{jw})}{1 + \sum_k \beta_k + \alpha TVC_{kw} + \zeta_{kw}}. \quad (\text{A.2})$$

We log-transform the market shares relative to the outside option and estimate the parameters for each of the three categories following Berry (1994). The price responsiveness parameter estimates $\hat{\alpha}$ are all significant and precise, with t -statistics ranging from 13 to about 26, as reported in Section 5.

A.2. Quality Discretization

The estimated perceived quality $\tilde{\omega}_t$ is discretized into partitioned levels (denoted by ω_t), in a manner consistent with the state transition process described by Equation (5). This transition process implies two requirements needed for logical consistency with the model: (i) a firm's quality cannot move more than one level from one period to the next, and (ii) if one firm's quality increases (decreases), none of the other firms' quality levels can decrease (increase) simultaneously.

To understand these two requirements, recall that the values of v_j can be either 0 (the investment fails) or 1 (the investment succeeds). The same is true for η : if $\eta = 0$, the industry shock is not realized and firms' quality only depends on their own investment; if $\eta = 1$, the outside option improves and all firms' qualities are reduced by one level. Therefore, all possible values of $v_j - \eta$ are in the set $\{-1, 0, 1\}$. Requirement (i) says that no firm may move more than one quality level in one period; this is implied by the maximum change in v_j . Requirement (ii) says that if any firm's quality is observed to rise in one period, no other firm's quality may be observed to fall in the same period. If one firm's quality rises, then $v_j - \eta = 1$, so it must be that $\eta = 0$. Conversely, if any firm's quality is observed to fall in some period, it must be that $\eta = 1$, so $v_j - \eta$ must equal either -1 or 0 for every other firm, depending on the outcome of that firm's investment. Therefore, for logical consistency, the quality levels must be partitioned such that firms' quality levels never move in opposite directions in the same period, consistent with requirement (ii).

To discretize the estimated qualities, we found that for two categories (compact car and midsize car), partitions chosen at intervals of 0.5 in continuous quality space produced quality discretizations that satisfy these two requirements. Minor adjustments in either direction (smaller or larger) led to violations of either requirement (i) or requirement (ii). For full-size pickup, we found that slightly larger intervals were needed to avoid violating requirement (i), so an interval of 0.7 was chosen. These partitions are shown graphically in Figure 2.

It is possible that the discretization may affect the models' ability to predict product quality and market shares. For this reason, the benchmark VAR is estimated in continuous units of quality, to prevent the quality discretization from harming the performance of the benchmark model.

Appendix B. Derivation of the Optimal Investment

To compute the expectation term in Equation (8), each firm must form an expectation of future market states, conditional on its competitors' optimal strategies and possible outcomes of the shock η . We can rewrite the expectation as

$$E[V(\omega'_j, \omega'_{-j}) | \omega] = \sum_{v_j} W(v_j | \omega) p(v_j), \quad (\text{B.1})$$

where

$$W(v_j | \omega) \equiv \sum_{\omega'_j} \sum_{\eta} V(\omega_j + v_j - \eta, \omega'_{-j}) q(\omega'_j | \omega, \eta) p(\eta) \quad (\text{B.2})$$

is firm j 's expected payoff conditional on the outcome of its investment and $p(v_j)$ is the distribution of v_j as given in Equation (4); $q(\omega'_j | \omega, \eta)$ represents firm j 's expectation of its competitors' next-period states ω'_{-j} conditional on the common shock η . These beliefs are rational in the sense that they are consistent with competitors' equilibrium investment strategies.

Plugging Equations (B.1) and (B.2) into (8) gives

$$\begin{aligned} V(\omega_j, \omega_{-j}) &= \max_{x_j} \left\{ \pi(\omega_j, \omega_{-j}) - x_j + \beta \sum_{v_j} W(v_j | \omega) p(v_j) \right\} \\ &= \max_{x_j} \left\{ \begin{aligned} &\pi(\omega_j, \omega_{-j}) - x_j \\ &+ \beta \sum_{v_j} \left[\sum_{\omega'_{-j}} \sum_{\eta} V(\omega_j + v_j - \eta, \omega'_{-j}) \right. \\ &\quad \left. \cdot q(\omega'_{-j} | \omega, \eta) p(\eta) \right] p(v_j) \end{aligned} \right\}. \end{aligned} \quad (\text{B.3})$$

Note that $W(v_j | \omega)$ is not a function of x_j , so differentiating Equation (B.3) with respect to investment x_j leads to the first-order condition for the investment policy function:

$$-1 + \beta \sum_{v_j} W(v_j | \omega) \frac{\partial p(v_j)}{\partial x_j} = 0. \quad (\text{B.4})$$

Equation (4) implies equalities $\partial p(v_j = 1) / \partial x_j = \rho / (1 + \rho x_j)^2$ and $\partial p(v_j = 0) / \partial x_j = -\rho / (1 + \rho x_j)^2$. Plugging these into Equation (B.4),

$$-1 + \beta \frac{\rho}{(1 + \rho x_j)^2} \{W(v_j = 1 | \omega) - W(v_j = 0 | \omega)\} = 0. \quad (\text{B.5})$$

Rearranging (B.5) and imposing the nonnegativity constraint $x_j \geq 0$, the analytical solution of the optimal investment policy is

$$x_j^*(\omega_j, \omega_{-j}) = \max \left\{ 0, \frac{-1 + \sqrt{\beta \rho (W(1 | \omega) - W(0 | \omega))}}{\rho} \right\} \quad (\text{B.6})$$

if $W(1 | \omega) \geq W(0 | \omega)$ and $x_j^*(\omega_j, \omega_{-j}) = 0$ otherwise.

Appendix C. Likelihood Function Formulation and Alternative Estimation Methods

Two-step estimation approaches have been developed to reduce the computational burden of solving for optimal policies at every point in the state space (e.g., Aguirregabiria and Mira 2007, Bajari et al. 2007, Pakes et al. 2007).³⁴ Doraszelski and Pakes (2007) and Borkovsky et al. (2012) reviewed this literature. The essence of the two-step approach is to avoid the computation of the fixed point by flexibly estimating state transition probabilities and policy functions in the first step, assuming the observed data are generated by equilibrium play. However, a general concern about this approach is that the first-step estimates might not be compatible with the equilibrium implied by the underlying model, resulting in serious finite sample biases of the second-stage estimates of the structural parameters (Aguirregabiria and Mira 2007). Given the small number of quality observations typically available in technology-related product categories, these potential finite sample problems are considered to be a first-order concern.

The current paper adopts the MPEC approach to estimate the parameters of interest. The objective function of the constrained optimization problem is a likelihood function that matches the predicted quality changes to the observed quality changes. Let $\hat{d}_{jt} \in \{-1, 0, 1\}$ represent the possible quality change for firm j at the end of time period t , as a result of the realization of its investment x_{jt} and the realization of the industry-wide shock η_t . As explained in Appendix A.2, three mutually exclusive events might happen in each time period: (a) *at least one* firm's quality improves, denoted by $I1_t = 1$; (b) *at least one* firm's quality deteriorates, denoted by $I0_t = 1$; or (c) no firm's quality changes (i.e., $I1_t = 0$ and $I0_t = 0$). The probability that each of these three events occurs is given by

- (a) $\Pr(I1_t = 1)$
 $= (1 - \delta) \prod_j \Pr(v_{jt} = 1)^{I(\hat{d}_{jt} = 1)} \Pr(v_{jt} = 0)^{1 - I(\hat{d}_{jt} = 1)},$
- (b) $\Pr(I0_t = 1)$
 $= \delta \prod_j \Pr(v_{jt} = 0)^{I(\hat{d}_{jt} = -1)} \Pr(v_{jt} = 1)^{1 - I(\hat{d}_{jt} = -1)},$ and
- (c) $\Pr(I1_t = 0 \text{ and } I0_t = 0)$
 $= \delta \prod_j \Pr(v_{jt} = 1)^{I(\hat{d}_{jt} = 0)} + (1 - \delta) \prod_j \Pr(v_{jt} = 0)^{I(\hat{d}_{jt} = 0)},$

where $\Pr(I1_t = 1)$ is the joint probability that event (a) happens, $\Pr(I0_t = 1)$ is the joint probability that event (b) occurs, and $\Pr(I1_t = 0 \text{ and } I0_t = 0)$ is the joint probability that event (c) occurs. Here, $\Pr(v_{jt} = 1) = \rho x_{jt} / (1 + \rho x_{jt})$ and $\Pr(v_{jt} = 0) = 1 / (1 + \rho x_{jt})$ are the probability that firm j 's investment succeeds and fails, respectively, provided in Equation (4) of Section 3. Given an industry state ω_t , we could compute these three probabilities based on the optimal investment policies $x_t^* = (x_{1t}^*, \dots, x_{jt}^*)$.

Now let $\hat{l}_t \equiv \omega_{t+1} - \omega_t$ represent the observed vector of firms' quality changes. We can write out the likelihood function of observing the quality changes given the probability of each event defined above:³⁵

$$\begin{aligned} L(\cdot; \rho, \delta) &= \prod_t \{ (\hat{I}1_t = 1)^{\Pr(I1_t = 1)} (\hat{I}0_t = 1)^{\Pr(I0_t = 1)} \\ &\quad \cdot (\hat{I}1_t = 0 \text{ and } \hat{I}0_t = 0)^{\Pr(I1_t = 0 \text{ and } I0_t = 0)} \}. \end{aligned} \quad (\text{C.1})$$

The MPEC approach is then to maximize this likelihood function subject to the three sets of equilibrium constraints

defined in Equation (10) by searching for optimal values of the structural parameters ρ and δ as well as solving for the value functions $V(\omega)$ and policy functions $x(\omega)$.

Appendix D. Quality Predictions

This section describes how we use the parameter estimates of the dynamic quality ladder model as well as the benchmark VAR model to predict product qualities and market shares for the out-of-sample periods 2008–2011. Both sets of predictions account for estimation error in the parameters.

D.1. Dynamic Quality Ladder Model Predictions

The DQL model predictions are done via a counterfactual by computing equilibrium under the new regime (i.e., the average fuel cost in 2008–2011) based on the model parameter estimates. That is, we solve the Bellman equation (8) to obtain the equilibrium investments at each possible industry state under the new regime.

To predict product-quality levels after 2007, we need to simulate the outcomes of firms' investments based on the state transition process, Equation (5) in Section 3.2. We make predictions in the following steps for the one-year prediction horizon:

Step 1. Solve for the optimal investment policies x^* in year t using the observed quality levels in year $t - 1$ as the initial state vector.

Step 2. Draw from the asymptotic distributions of ρ and δ , obtained by the estimation of the DQL model. We then calculate the probability that each firm's investment succeeds as $\rho x_j^*/(1 + \rho x_j^*)$.

Step 3. Simulate the values of v_j and η by drawing from the standard uniform distribution. If the draw is greater than $\rho x_j^*/(1 + \rho x_j^*)$, then $v_j = 0$; otherwise, $v_j = 1$. Similarly, if a separate uniform random draw is greater than δ (recall that δ represents the probability that the negative demand shock is realized), $\eta = 0$; otherwise, $\eta = 1$.

Step 4. Based on the simulated values of v_j and η , year t quality levels are predicted using Equation (5). The predicted qualities are in discrete levels and are then converted to continuous measures by taking the midpoint of the quality range that the discrete quality belongs to.

Step 5. Given the predicted quality in year t , we solve the per-period profit maximization problem (Equation (6)) simultaneously for all firms to obtain the optimal prices. The predicted quality and the corresponding optimal prices are then taken back into the demand system (Equation (3)) to calculate the predicted market shares.

Step 6. Repeat Steps 2–5 1,000 times. The MAE is calculated as the mean of the 1,000 absolute prediction errors.

This completes the description of the one-year prediction horizon, in which year t prediction errors are always calculated as a function of the observed state in year $t - 1$.

The multiple-year prediction horizon uses a very similar algorithm. The 2008 predictions are based on the observed state in 2007. Then, for a draw of ρ and δ , Steps 2–5 are repeated in all subsequent years, with each year taking the previous year's predicted state vector as the initial state. Finally, we take the average MAE across firms and holdout periods to compute the MAE ratio presented in Section 5.

Table D.1. Prediction Performance Measure Comparison

	One-year horizon				Multiple-year horizon			
	2008	2009	2010	2011	2008	2009	2010	2011
a: Perceived quality predictions								
MAE								
Compact car	1.8	1.8	2.2	2.0	—	1.9	1.9	2.0
Midsize car	1.0	0.7	0.6	0.8	—	0.7	0.7	0.9
Full-size pickup	1.7	2.0	1.5	1.4	—	0.9	0.7	0.6
RMSE								
Compact car	1.6	1.6	2.0	2.1	—	1.7	1.8	1.9
Midsize car	0.9	0.6	0.6	0.7	—	0.7	0.7	0.8
Full-size pickup	1.7	1.9	1.5	1.4	—	0.9	0.8	0.7
b: Market share predictions								
MAE								
Compact car	1.5	1.1	1.6	1.7	—	1.6	1.7	1.7
Midsize car	0.8	0.8	0.7	0.7	—	0.8	0.8	0.8
Full-size pickup	0.8	1.1	1.3	0.9	—	0.6	0.6	0.5
RMSE								
Compact car	1.4	1.1	1.5	1.7	—	1.6	1.7	1.7
Midsize car	0.8	0.8	0.7	0.7	—	0.8	0.7	0.7
Full-size pickup	0.8	1.1	1.3	1.0	—	0.7	0.6	0.5

Note. When the ratio is less than 1 (as indicated in bold), VAR model predictions exhibit a larger error than DQL model predictions.

We also compute other prediction accuracy metrics such as root mean squared error (RMSE). Table D.1 shows that different prediction performance metrics lead to similar qualitative conclusions about the nonrandom holdout performance of the DQL model.

D.2. The VAR Model

The benchmark VAR model predictions are done by taking draws from the asymptotic distribution of parameter estimates in Equation (11), provided by the model estimation. We make the predictions in the following steps for the one-year prediction horizon:

Step 1. Draw from the asymptotic distribution of the parameter matrix A and Σ .

Step 2. Taking the observed quality in year $t - 1$ as the initial state vector, compute the prediction quality of year t by plugging the numbers from Step 1 into the right-hand side of Equation (11).

Step 3. Given the predicted quality in year t , we solve the per-period profit maximization problem (Equation (6)) simultaneously for all firms to obtain the optimal prices. The predicted quality and the corresponding optimal prices are then taken back into the demand system (Equation (3)) to calculate the predicted market shares.

Step 4. Repeat Steps 1–3 1,000 times and compute the MAE as the mean of the 1,000 absolute prediction errors.

For the multiple-year prediction horizon, Steps 1–4 are repeated for 2009–2011, each conditional on the predicted values of the previous year as the initial state. Similarly, we take the average MAE across firms and holdout periods and compute the MAE ratio reported in Section 5.

Figure D.1. Predicted and Observed Market Shares for Holdout Periods





Figure D.1. (Continued)



Note. The dark grey area presents the 90% confidence bound of a brand's market share predictions made by the DQL model, the light grey area shows the 90% confidence bound of the VAR model market share predictions, and the solid line in black represents the observed market shares.

Figure D.2. Predicted and Observed Market Shares for In-Sample Periods





Figure D.2. (Continued)



Note. The dark grey area presents the 90% confidence bound of a brand's market share predictions made by the DQL model, the light grey area shows the 90% confidence bound of the VAR model market share predictions, and the solid line in black represents the observed market shares.

Endnotes

¹ From a purely practical perspective, further validation of structural dynamic oligopoly models may help to show that their benefits merit the nontrivial coding and computational costs of estimating them.

² For example, we systematically reviewed 125 published papers that cite Pakes and McGuire (1994). To the best of our understanding, none of them performed model validation using holdout samples.

³ Keane and Wolpin (2007) estimated a structural dynamic model of individuals' welfare program participation and showed that it performed better in nonrandom holdout samples than several benchmarks.

⁴ We estimate products' "perceived quality" as the brand intercepts in a discrete choice model that rationalize observed market shares at observed prices. More details are provided in Section 3.

⁵ A few published studies that have gone beyond the EP framework to study other dynamic decisions, such as dynamic pricing decisions under a learning environment (Ching 2010) and dynamic production decisions (Benkard 2004, Esteban and Shum 2007, Chen et al. 2013).

⁶ Hashmi and Van Biesebroeck (2016) and the current paper are similar in the sense that they both apply the EP framework to the automotive industry. However, Hashmi and Van Biesebroeck (2016) focused on discovering the relationship between innovation and market structure, whereas the purpose of this study is to evaluate the dynamic quality ladder model using nonrandom holdout samples.

⁷ Other studies have also modeled how advertising builds goodwill stock in a dynamic oligopoly framework, including Dubé et al. (2004), Tan (2006), Qi (2013), and Liu et al. (2016).

⁸ Given the limited number of years in the sample, it was determined that there were insufficient degrees of freedom to estimate unobserved heterogeneity in demand for automobiles.

⁹ Ideally, one would want to include MPG as a state variable that is determined by firms' product strategies. However, MPG is treated as exogenously given in the current paper because of the lack of variation over time within a category. Time trend only explains 8%, 6%, and 0.2% of the total variation in MPG in compact car, midsize car, and full-size pickup categories, respectively.

¹⁰ We specify the interest rate as the national average interest rate on certificates of deposit.

¹¹ Automakers typically spend about \$18 billion annually on R&D and another \$5.5 billion on advertising (Kantar Media AdSpender report 2014). A substantial fraction of advertising is used to communicate pricing terms and other product information to consumers (Xu et al. 2014).

¹² By holding the number of firms constant throughout the course of the game, we have excluded endogenous entry and exit decisions from the model. This assumption may affect the model's nonrandom holdout performance and the generalizability of the model, but estimating a model with entry and exit requires sufficient observations of entry and exit, which are very rare in our empirical context.

¹³ Market size is defined as the maximum total annual sales for each category (compact car, midsize car, and full-size pickup) in the U.S. market.

¹⁴ The discount factor is chosen to be 0.925, which corresponds to an annual interest rate of 8%. Perturbing this assumption to 0.9, 0.95, or 0.975 does not change the qualitative results reported in Section 5.

¹⁵ To enforce this constraint, we assign $\eta = 1$ with probability 1 if any firm with $\omega_j = \bar{\omega}$ succeeds in its innovation.

¹⁶ The estimation typically converged in minutes for two-firm scenarios, a few days for three-firm scenarios, and one to three weeks for four-firm scenarios. Appendix C discusses alternative estimation methods.

¹⁷ We also investigated the sensitivity of the parameter estimates reported in Section 5 to the choice of $\bar{\omega}$ and found that the same

caveat applies: the point estimates are not sensitive to $\bar{\omega}$, so long as $\bar{\omega}$ is never observed within the estimation data.

¹⁸ We have also considered firm-specific AR(k) models as alternative benchmarks, but they do not present better prediction performance than the VAR model.

¹⁹ We estimate the model separately for three automotive categories as in other studies in the literature. This implicitly rules out consumer substitutions across categories.

²⁰ Typically, each manufacturer has a single dominant model within each category at a given time, and that is what we focus on when gathering data to estimate product qualities. However, a manufacturer occasionally "refreshed" an existing model by relaunching under a new name; at those points, we shift the focus from the prior model name to the new model name. We also note that GMC Sierra and Chevrolet Silverado were sold by separate divisions of General Motors. They are treated as two different firms in our empirical context because they often introduce new models at different times, charge different prices, have completely separate dealership networks, and target different segments of consumers.

²¹ Appendix A.1 describes the detailed transaction data, the demand system used to estimate the price responsiveness parameters and how they were estimated. In this framework, the demand intercept is indistinguishable from persuasive advertising or distribution strategies that might contribute to the demand intercept. In additional analyses, we found that the estimation error of the price coefficient does not affect the qualitative results presented in Section 5.

²² The combined sales of the top four firms, on average, account for 38% of the compact car market, 37% of the midsize car market, and 65% of the full-size pickup market. The combined sales of the 5th, 6th, and 7th firm, on average, account for 16% of the compact car market, 13% of the midsize car market, and 10% of the full-size pickup market. The compact car and midsize car markets seem to be less concentrated than the full-size pickup market, as the four largest manufacturers account for less than 50% of the total market. Our simulation exercises indicate that having more firms in the data could facilitate accuracy of the parameter estimates, but four firms is the largest number that our computational resource allows.

²³ The predictions are made using 1,000 draws from the joint distribution of the parameter estimates of each model in each scenario. Therefore, the holdout comparisons fully account for estimation error. A detailed description of the prediction procedure can be found in Appendix D.

²⁴ The comparison results are qualitatively similar if the predictions are evaluated using root mean squared error (RMSE), as shown in Appendix D.

²⁵ We also evaluated the comparisons in levels of quality (a discrete measure). The results were qualitatively similar and are therefore excluded.

²⁶ The 90% confidence bounds of the predicted market shares in 2008–2011 by both models along with the actual observed market shares are presented in Figure D.1 in the appendix.

²⁷ Regional gas prices were available from the U.S. Energy Information Administration for Boston, Chicago, Cleveland, Denver, Houston, Los Angeles, Miami, New York, San Francisco, and Seattle.

²⁸ The regression included the top four models in each of the three automotive categories. The fourth weekly lag of gas price was used, consistent with the typical 29-day automotive purchase cycle reported in J.D. Power's "2008 Auto Buyer Clickstream Study" (J.D. Power and Associates 2008), but the results are similar using other lags of gas price. Time fixed effects were excluded because they are highly collinear with the gas price data (R -square of 0.66).

²⁹ We considered and rejected several alternative explanations for the category-specific results, including regional economic growth and unemployment rate, concentration in category market shares,

the range and dispersion of quality estimates across products and time, correlations in products' quality levels, and the total number of quality changes observed in the data.

³⁰We plot the 90% confidence range of the predicted market shares by both models along with the actual observed market shares for the in-sample periods in Figure D.2 in the appendix.

³¹Note that, under each model, the quality predictions were made based on 1,000 simulations as discussed in Appendix D. For this exercise, we focus on the multiple-year horizon prediction approach. We calculate the consumer surplus and firm per-period profits for each of the 1,000 draws and report the average and the standard deviation across these draws in Table 6.

³²Note that the annual demand idiosyncratic preferences ε_{ijt} presented in Equation (1) are assumed to follow the same distribution. We acknowledge this as a modeling limitation to assume that consumers' weekly and annual demand for automobiles have the same distribution.

³³Market size is defined as the maximum total weekly unit sales within each year.

³⁴Weintraub et al. (2008) introduced a new equilibrium concept, oblivious equilibrium, through which dynamic investments can be solved as single-agent problems when the number of competing firms is large. However, this approach assumes that no single firm has influence over the overall market's evolution. For an application, see Qi (2013).

³⁵When any firm's quality is at either bound of the state space, we adjust the likelihood contribution of that particular observation as described in Section 3.2.1.

References

- Aguirregabiria V, Mira P (2007) Sequential estimation of dynamic discrete games. *Econometrica* 75(1):1–53.
- Ailawadi KL, Kopalle PK, Neslin SA (2005) Predicting competitive response to a major policy change: Combining game-theoretic and empirical analyses. *Marketing Sci.* 24(1):12–24.
- Angrist JD, Pischke J-S (2010) The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *J. Econom. Perspect.* 24(2):3–30.
- Associated Press (2008) Gas prices put Detroit Big Three in crisis mode. *NBCNews.com* (June 1), <http://www.nbcnews.com/id/24896359/ns/business-autos/t/gas-prices-put-detroit-big-three-crisis-mode/#.WVWouYjyIV>.
- Bajari P, Benkard CL, Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75(5):1331–1370.
- Benkard CL (2004) A dynamic analysis of the market for wide-bodied commercial aircraft. *Rev. Econom. Stud.* 71(3):581–611.
- Berry TS (1994) Estimating discrete-choice models of product differentiation. *RAND J. Econom.* 25(2):242–262.
- Besanko D, Doraszelski U (2004) Capacity dynamics and endogenous asymmetries in firm size. *RAND J. Econom.* 35(1):23–49.
- Besanko D, Doraszelski U, Lu LX, Satterthwaite M (2010) Lumpy capacity investment and disinvestment dynamics. *Oper. Res.* 58(3):1178–1193.
- Borkovsky RN, Doraszelski U, Kryukov Y (2012) A dynamic quality ladder model with entry and exit: Exploring the equilibrium correspondence using the homotopy method. *Quant. Marketing Econom.* 10(2):197–229.
- Borkovsky RN, Goldfarb A, Haviv A, Moorthy S (2017) Measuring and understanding brand value in a dynamic model of brand management. *Marketing Sci.*, ePub ahead of print May 11, <https://doi.org/10.1287/mksc.2016.1020>.
- Chen J, Esteban S, Shum M (2013) When do secondary markets harm firms? *Amer. Econom. Rev.* 103(7):2911–2934.
- Ching AT (2010) A dynamic oligopoly structural model for the prescription drug market after patent expiration. *Internat. Econom. Rev.* 51(4):1175–1027.
- Chintagunta P, Erdem T, Rossi PE, Wedel M (2006) Structural modeling in marketing: Review and assessment. *Marketing Sci.* 25(6):604–616.
- Collard-Wexler A (2013) Demand fluctuations in the ready-mix concrete industry. *Econometrica* 81(3):1003–1037.
- Doraszelski U, Pakes A (2007) A framework for applied dynamic analysis in IO. Armstrong M, Porter R, eds. *Handbook of Industrial Organization*, Vol. 3 (North-Holland, Amsterdam), 1887–1966.
- Doraszelski U, Satterthwaite M (2010) Computable Markov-perfect industry dynamics. *RAND J. Econom.* 41(2):215–243.
- Dubé J-P, Hitsch G, Manchanda P (2004) An empirical model of advertising dynamics. *Quant. Marketing Econom.* 3(2):107–144.
- Ericson R, Pakes A (1995) Markov-perfect industry dynamics: A framework for empirical work. *Rev. Econom. Stud.* 62(1):53–82.
- Esteban S, Shum M (2007) Durable-goods oligopoly with secondary markets: The case of automobiles. *RAND J. Econom.* 38(2):1–23.
- Goettler RL, Gordon B (2011) Does AMD spur Intel to innovate more? *J. Political Econom.* 119(6):1141–1200.
- Goettler RL, Gordon B (2014) Competition and product innovation in dynamic oligopoly. *Quant. Marketing Econom.* 12(1):1–42.
- Gowrisankaran G (1999) A dynamic model of endogenous horizontal mergers. *RAND J. Econom.* 30(1):56–83.
- Gowrisankaran G, Town RJ (1997) Dynamic equilibrium in the hospital industry. *J. Econom. Management Strategy* 6(1):45–74.
- Greene WH (2012) *Econometric Analysis*, 7th ed. (Prentice Hall, Englewood Cliffs, NJ).
- Hashmi AR, Van Biesebroeck J (2016) The relationship between market structure and innovation in industry equilibrium: A case study of the global automobile industry. *Rev. Econom. Statist.* 98(1):192–208.
- Hill K, Swiecki B, Menk DM, Cregger J (2014) Just how high-tech is the automotive industry? Report, Center for Automotive Research, Ann Arbor, MI.
- J.D. Power and Associates (2008) Auto buyer clickstream study. Report, McGraw Hill Financial, Westlake Village, CA.
- Keane MP (2010) Structural vs. atheoretic approaches to econometrics. *J. Econometrics* 156(1):3–20.
- Keane MP, Wolpin KI (2007) Exploring the usefulness of a nonrandom holdout sample for model validation: Welfare effects of female behavior. *Internat. Econom. Rev.* 48(4):1531–1578.
- Liu Q, Gupta S, Venkataraman S, Liu H (2016) An empirical model of drug detailing: Dynamic competition and policy implications. *Management Sci.* 62(8):2321–2340.
- Markovitch S (2008) Snowball: A dynamic oligopoly model with indirect network effects. *J. Econom. Dynam. Control* 32(3):909–938.
- Maskin E, Tirole J (1987) A theory of dynamic oligopoly, III: Cournot competition. *Eur. Econom. Rev.* 31(4):947–968.
- Maskin E, Tirole J (1988a) A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica* 56(3):549–569.
- Maskin E, Tirole J (1988b) A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* 56(3):571–599.
- McFadden D, Talvitie AP, Cosslett S, Hasan I, Johnson M, Reid FA, Train K (1977) Validation of disaggregate travel demand models: Some tests. Urban Travel Demand Forecasting Project, Phase I Final Report, Volume 5, Institute of Transportation Studies, University of California, Berkeley, Berkeley.
- Meinero MM, Rooney B (2008) Gas price record reaches \$4 a gallon. *CNNMoney.com* (June 8), http://money.cnn.com/2008/06/08/news/economy/gas_prices/.
- Misra S, Nair HS (2011) A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quant. Marketing Econom.* 9(3):211–257.
- Mitra D, Golder PN (2006) How does objective quality affect perceived quality? Short-term effects, long-term effects, and asymmetries. *Marketing Sci.* 25(3):230–347.
- Pakes A, McGuire P (1994) Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND J. Econom.* 25(4):555–589.
- Pakes A, Ostrovsky M, Berry S (2007) Simple estimators for the parameters of discrete dynamic games (with entry-exit examples). *RAND J. Econom.* 38(2):373–399.

- Pathak PA, Shi P (2015) Demand modeling, forecasting, and counterfactuals, part I. Working paper, Massachusetts Institute of Technology, Cambridge.
- Picard RR, Cook RD (1984) Cross-validation of regression models. *J. Amer. Statist. Assoc.* 79(387):575–583.
- Qi S (2013) The impact of advertising regulation on industry: The cigarette advertising ban of 1971. *RAND J. Econom.* 44(2): 215–248.
- Raval D, Rosenbaum T, Wilson NE (2016) Industrial reorganization: Learning about patient substitution patterns from natural experiments. Working paper, Federal Trade Commission, Washington, DC.
- Rust J (1987) Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5):999–1033.
- Rust J (2014) The limits of inference with theory: A review of Wolpin (2013). *J. Econom. Literature* 52(3):820–850.
- Ryan SP (2012) The costs of environmental regulation in a concentrated industry. *Econometrica* 80(3):1019–1061.
- Shao J (1993) Linear model selection by cross-validation. *J. Amer. Statist. Assoc.* 88(422):486–494.
- Song M (2011) A dynamic analysis of cooperative research in the semiconductor industry. *Internat. Econom. Rev.* 52(4): 1157–1177.
- Su C-L, Judd K (2012) Constrained optimization approaches to estimation of structural models. *Econometrica* 80(5):2213–2230.
- Tan W (2006) The effects of taxes and advertising restrictions on the market structure of the U.S. cigarette market. *Rev. Indust. Organ.* 28(3):231–251.
- Weintraub GY, Benkard CL, Van Roy B (2008) Markov perfect industry dynamics with many firms. *Econometrica* 76(6):1375–1411.
- Xu L, Wilbur KC, Siddarth S, Silva-Risso J (2014) Price advertising by manufacturers and dealers. *Management Sci.* 60(11): 2816–2834.