Numerically Stable Design Optimization With Price Competition

W. Ross Morrow
Mechanical Engineering,
Iowa State University,
Ames, Iowa 50011
e-mail: wrmorrow@iastate.edu

Joshua Mineroff
Mechanical Engineering,
Iowa State University,
Ames, Iowa 50011
e-mail: mineroff@iastate.edu

Kate S. Whitefoot
National Academy of Engineering,
Washington, DC 20418
e-mail: kwhitefoot@nae.edu

Researchers in decision-based design (DBD) have suggested that business objectives, e.g., profits, should replace engineering requirements or performance metrics as the objective for engineering design. This requires modeling market performance, including consumer preferences and competition between firms. Game-theoretic "design-then-pricing" models—i.e., product design anticipating future price competition–provide an important framework for integrating consumer preferences and competition when design decisions must be made before prices are decided by a firm or by its competitors. This article concerns computational optimization in a design-then-pricing model. We argue that some approaches may be fundamentally difficult for existing solvers and propose a method that exhibits both improved efficiency and reliability relative to existing methods. Numerical results for a vehicle design example validate our theoretical arguments and examine the impact of anticipating pricing competition on design decisions. We find that anticipating pricing competition, while potentially important for accurately forecasting profits, does not necessarily have a significant effect on optimal design decisions. Most existing examples suggest otherwise, anticipating competition in prices is important to choosing optimal designs. Our example differs in the importance of design constraints, that reduce the influence the market model has on optimal designs.
[DOI: 10.1115/1.4025703]

1 Introduction

Research in DBD, Enterprise-System Design, and Value-Driven Design suggests replacing traditional engineering requirements or performance metrics with overarching business goals as the objective for engineering design [1–4]. Wassenaar and Chen [3] specifically propose using economic value to the firm—e.g., profits—as the metric to judge different designs. The recent "market-systems" theme in this literature has recognized that using profits as the objective for engineering design ultimately requires modeling the features of real markets in which products are sold. For example, product profitability is not only a function of consumer's choices, but the choices made by competing firms, retailers, and regulators. Shiau and Michalek [5] and Williams et al. [6,7] have shown that retail channel structure may influence optimal engineering designs. Several engineering studies have studied interactions between design and regulatory policy [5,8–10]. These features of real markets have been integrated into engineering design using Game Theory [11–14], following long trends in empirical economics [15–30] and marketing [31–37].

We consider design decisions when firms anticipate price competition, referred to here as the design-*then*-pricing paradigm [31,38–40]. In this model, a firm chooses designs assuming that, at some later time, prices will be determined through some competitive mechanism given these design decisions; see Fig. 1, right. We assume that the firm adopts Bertrand–Nash equilibrium pricing [31,39,41] as the representation of market competition in prices, though other equilibrium concepts have been explored [40]. Some existing models adopt a design-and-pricing paradigm that assumes firms make design-and-pricing decisions simultaneously, without accounting for "reactions" by their competitors or other market entities [8–10]; see Fig. 1, left. However, it is possible that competitors or other market entities do change some decisions in reaction to a firms design and/or pricing decisions. One prominent example occurs when modeling retail channels [5–7,34,35,42], where it is reasonable to suppose that retailers' maximize profits with their pricing decisions once they know wholesale prices.

Existing research has demonstrated the potential importance of the game-theoretic market-systems paradigm to product design, including the design-then-pricing model [5–8,39,40]. The majority of this existing work is, appropriately, based on small-scale "illustrative" examples with significant engineering detail rather than focusing on "real-scale" applications of the design-then-pricing paradigm to current differentiated product markets that may have hundreds or thousands of product variants [43]. For example, econometric studies of the new car market include hundreds of vehicles [9], and automotive market data sources routinely represent thousands of distinct vehicle models [44,45]. Unfortunately optimization methods proposed on the basis of performance for small-scale problems may be incapable of reliably or efficiently solving problems for large markets.

This article clarifies and confronts this concern. We compare three approaches for solving design-then-pricing problems: an "implicit programming" approach and two approaches based on mathematical programs with equilibrium constraints (MPECs). In implicit programming [46], equilibrium prices are treated as an intermediate quantity or internal "simulation" and computed using iterative techniques at each trial point (e.g., Refs. [31,39,47]). "MPEC" approaches treat prices as variables that must satisfy a constraint that describes equilibrium (e.g., Ref. [39]). Our primary contributions are as follows.

First, while implicit programming can reliably compute equilibrium prices, it may require prohibitively accurate computations of equilibrium prices to converge reliably when there are many product variants. Appropriately designed MPEC methods can be both more efficient and reliable for computing optimal designs when anticipating price competition. Similar phenomena have been observed in econometrics [48] and chemical engineering [49].

Contributed by the Design Automation Committee of ASME for publication in the Journal of Mechanical Design. Manuscript received September 12, 2012; final manuscript received September 9, 2013; published online June 2, 2014. Assoc. Editor: Wei Chen.

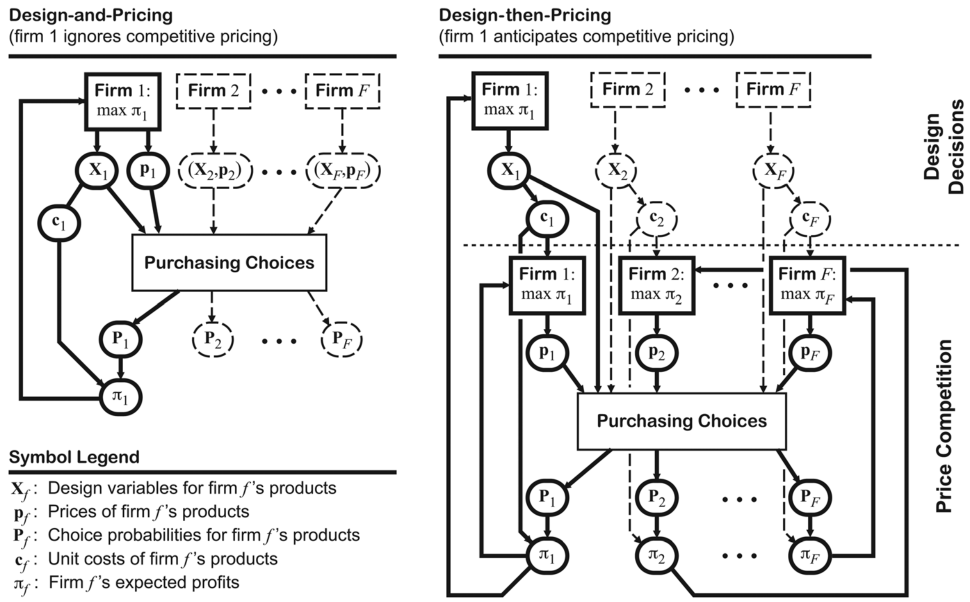

Fig. 1 Conceptual difference between a design-*and*-pricing model (left) and a design-*then*-pricing model (right) for *F* product-designing firms. Dashed lines denote fixed components in the respective model, while thick solid lines denote variable components.

Second, we show that the choice of equilibrium constraint in MPEC formulations can determine problem solvability. We demonstrate, both in theory and in a practical example, that spurious solutions to design-then-pricing problems exist and can be computed if the problem is formulated using the literal first-order condition for equilibrium as proposed in the existing literature. Applying a fixed-point representation of the first-order condition for equilibrium prices provably eliminates many of these spurious solutions [41,50,51]. We also show that spurious solutions exist and can be computed in design-and-pricing, and demonstrate a correction for this case as well.

Third, we discuss the importance of checking the second-order sufficient condition (SOSC) [52, Chap. 12] in design-then-pricing problems, and an efficient approach to undertaking such checks in practice. In the context of design-then-pricing, the use of first-order conditions to compute equilibrium prices drives the need to validate equilibrium using the SOSC [39,41]. However, the existing literature has proposed a potentially inefficient heuristic that is not based on the SOSC, requires re-optimizing firms profits with respect to prices, and may not always correctly validate equilibrium in design-then-pricing problems. Our check, based on the standard SOSC derived from local properties of the profit functions, is theoretically rigorous and can be executed using efficient linear algebra routines alone.

Fourth, new numerical results concerning a real-scale vehicle design example with 472 products are provided. This example is 2 orders of magnitude larger, in terms of market scale, than the examples currently in the literature. This example serves to demonstrate how dramatically some numerical methods can fail for large design-then-pricing problems, while simultaneously showing that such problems can be solved very efficiently when using the right methods. We also use this example to compare design-then-pricing with design-and-pricing, allowing us to assess the value of anticipating pricing competition when making design decisions. This comparison is more subtle than is currently acknowledged in the literature. Particularly, we may attach more importance to anticipating pricing competition if the accuracy of profit forecasts during design is important; however, if we only care about realizing the highest profits possible, anticipating competition may be irrelevant if physical constraints and monotonic consumer preferences drive design decisions.

The remainder of this article proceeds as follows. Section 2 introduces a framework for the design-then-pricing model considered here. Section 3 discusses numerical methods for the design-then-pricing problem. Section 4 defines a real-scale vehicle design case study and compares the three methods on this example. Section 5 provides further discussion of the model and results, and Sec. 6 concludes.

2 Framework

This section presents a mathematical framework for design-then-pricing problems. This framework has two primary components: the demand model and the firm's objective. We focus on the Mixed Logit model of demand that captures consumer heterogeneity; see Ref. [53, Chap. 6] for more information. Several recent studies have emphasized the importance of heterogeneity to engineering design [39,47,54,55]. Following the market-systems literature, we assume firms intend to maximize profits.

2.1 Mixed Logit Demand. Forecasting sales given design decisions is a pivotal step in an engineering design paradigm that includes profits as a design objective. Many researchers have integrated discrete choice analysis (DCA) [53,56,57] with engineering design optimization to serve this purpose [3,5,8,39,47,55,58–62]. DCA models can then be statistically estimated from either market observations [53] or surveys [57] and yield *choice probabilities* for each potential set of products, depending on design-and-pricing decisions. Though relatively recently introduced to engineering design, DCA has been one of the most important techniques in economics and marketing for over 30 yr [56,63,64].

We focus on the Mixed Logit model of demand [53], which can be defined as follows. Individual consumers are identified by a vector of characteristics θ from some set $\mathcal{T}$. These individual characteristics can include both observed demographic variables as well as random coefficients [16,53,65] that characterize unobserved individual-specific heterogeneity with respect to preference for product characteristics. The relative density of individual characteristic vectors in the population is described by a probability distribution μ over $\mathcal{T}$; we identify μ with the probability density function of this distribution (or probability mass function if $\mathcal{T}$ is finite).

There are J products, each of which is defined by a vector of product attributes $\mathbf{y}_j \in \mathcal{Y}$ and a price $p_j \geq 0$. An individual identified by $\theta \in \mathcal{T}$ receives the utility $U_j(\theta, \mathbf{y}_j, p_j) = u(\theta, \mathbf{y}_j, p_j) + E_j$ from purchasing product $j \in \{1, \ldots, J\}$, and $U_0(\theta) = \vartheta(\theta) + E_0$ for forgoing purchase of any of these products. Individuals choose nothing, indexed by 0, or the product $j \in \{1, \ldots, J\}$ that maximizes their utility. Here $u : \mathcal{T} \times \mathcal{Y} \times \mathcal{P} \to [-\infty, \infty)$ is a utility function, $\vartheta : \mathcal{T} \to \mathbb{R}$ is a valuation of the no-purchase option or "outside good," and $\mathbf{E} = \{E_j\}_{j=0}^{J}$ is a random vector of i.i.d. standard extreme value variables.

Demand for each product j is characterized by choice probabilities $P_j : (\mathcal{Y} \times \mathcal{P})^J \to [0,1]$ derived from utility maximization [53]. Given the distributional assumption on $\mathbf{E}$, the choice probabilities conditional on $\theta \in \mathcal{T}$ are those of the Logit model [53, Chap. 3]

$$P_j^L(\theta, \mathbf{Y}, \mathbf{p}) = \frac{\exp\{u_j(\theta, \mathbf{y}_j, p_j)\}}{\exp\{\vartheta(\theta)\} + \sum_{k=1}^{J} \exp\{u_k(\theta, \mathbf{y}_k, p_k)\}} \tag{1}$$

The Mixed Logit choice probabilities $P_j(\mathbf{p}) = \int P_j^L(\theta, \mathbf{p}) d\mu(\theta)$ follow from integrating over the distribution of individual characteristics [53, Chap. 6]. The vector of Mixed Logit choice probabilities for all products is denoted by $\mathbf{P}(\mathbf{Y}, \mathbf{p}) \in [0,1]^J$, where $\mathbf{Y}$ denotes the matrix of all product attribute vectors and $\mathbf{p}$ denotes the vector of all product prices. Differentiability of the choice probabilities is discussed in Refs. [41,50].

2.2 Utility Specification. Assumptions on the utility functions and demographic distributions are required for well-posed design-and-pricing problems. This article employs a specification closely related to that introduced by Caplin and Nalebuff [66] as refined by Morrow [67] and Morrow and Skerlos [41,68].

ASSUMPTION 1. *For all j, there exist functions $w_j : \mathcal{T} \times \mathcal{P} \to [-\infty, \infty)$ and $v_j : \mathcal{T} \times \mathcal{Y} \to (-\infty, \infty)$ such that the systematic utility function $u_j : \mathcal{T} \times \mathcal{Y} \times \mathcal{P} \to [-\infty, \infty)$ can be written $u_j(\theta, \mathbf{y}, p) = w_j(\theta, p) + v_j(\theta, \mathbf{y})$, where $v_j(\theta, \cdot) : \mathcal{Y} \to (-\infty, \infty)$ is affine over $\mathcal{Y}$. Furthermore, there exists a function $\imath : \mathcal{T} \to (0, \infty]$ such that $w_j : \mathcal{T} \times [0, \infty] \to [-\infty, \infty)$ satisfies, for all j and μ-almost every (a.e.) $\theta \in \mathcal{T}$,*

(a) *$w_j(\theta, \cdot) : (0, \imath(\theta)) \to (-\infty, \infty)$ is twice continuously differentiable, strictly decreasing, eventually decreases sufficiently quickly [67,68], and has subquadratic second derivatives [67,68],*
(b) *$w_j(\theta, p) \downarrow -\infty$ as $p \uparrow \imath(\theta)$, and $w_j(\theta, p) = -\infty$ for all $p \geq \imath(\theta)$.*
(c) *$w_j(\theta, p) + \log|(Dw_j)(\theta, p)| \to -\infty$ as $p \uparrow \imath(\theta)$*

Let $\imath_*$ denote the essential supremum of $\{\imath(\theta) : \theta \in \mathcal{T}\}$.

Several comments should clarify these technical assumptions:

Utilities: Allowing the utility functions to be product-specific allows for the inclusion of product attributes or features that are not explicitly design decisions, but are important in determining demand. Utility functions that are strictly decreasing in price are an intuitive and common economic assumption; we discuss this further below. Twice continuous differentiability of w_j and v_j will be necessary for twice continuous differentiability of the choice probabilities, and is thus, necessary to apply most techniques for continuous optimization. The "eventually decreasing sufficiently quickly," "subquadratic second derivatives," and strict monotonicity conditions on w_j are sufficient for the existence of profit-optimal prices with simple Logit models [67,68], and thus, are natural requirements to impose in the context of Mixed Logit models. These conditions are weak and satisfied for the model types used most often in practice.

Finite Purchasing Power: At first glance, including the map $\imath : \mathcal{T} \to (0, \infty]$ may seem unnecessarily complicated. Linear-in-price utilities of the form $u \sim -\alpha p$ (e.g., Refs. [9,69]), for example, have $\imath(\theta) = \infty$ for all θ. $\imath$ is, however, essential if we are to allow finite limits on individual purchasing power. For example, in the popular "Berry, Levinsohn, and Pakes" model of the new vehicle market [16,25], $\imath(\theta)$ represents individual or household income. Other metrics of purchasing power, such as household wealth or available credit, may be more appropriate than income in other contexts. Regardless of the empirical meaning of $\imath(\theta)$, Assumption 1(b) states that individuals do not purchase products that cost more than $\imath(\theta)$ *and* the probability they purchase a particular product vanishes as that product's price approaches $\imath(\theta)$. Assumption 1(c) is required to ensure that an individual's Logit choice probabilities are not only continuous in prices at $\imath(\theta)$, but also *continuously differentiable*; a requirement if *simulated* Mixed Logit choice probabilities are to be continuously differentiable; see Sec. 2.3 and Ref. [51]. Model formulations that do not need to make assumptions (b) and (c) do exist, but currently face computational challenges [70]. Addressing the case of finite purchasing power ($\imath_* < \infty$) properly can be important for robust computational methods [51].

Distribution of Individual Characteristics: Common examples of μ from the econometrics, marketing, and engineering literature include finitely supported distributions (often empirical frequency distributions for integral observed demographic variables or latent-class Conjoint-Logit models), standard continuous distributions (e.g., normal, lognormal, and χ^2), truncated standard continuous distributions, and finite mixtures or independent products of any of these types of distributions. This generality addresses a wide variety of seemingly disparate examples (e.g., conjoint models and econometric random coefficients models) within a single notation. In particular, this generality allows us to use a single framework to treat both "full" Mixed Logit models defined by some μ with unbounded, uncountable support and simulation-based approximations to such models. Conditions on the distribution μ and the associated utilities are required to ensure continuous differentiability of the choice probabilities, and thus, profits; see Ref. [50]. Most of the existing DCA literature implicitly assumes such conditions are satisfied, as similar conditions are required for model estimation and the assumption is necessary for tractable optimization methods.

Generality: Despite the detailed and technical nature of Assumption 1, it represents a very broad class of Mixed Logit models. We assume utilities are separable in prices and attributes; that is, w_j does not depend on $\mathbf{y}$. This assumption holds in most empirical models currently used but is made only for simplicity. The assumption that $v_j(\theta, \cdot)$ is affine entails no loss in generality, for the attributes can simply be redefined to ensure that this holds [66]. For example, utilities with "interaction effects" such as

$$v_j(\theta, \mathbf{y}) = \beta_1(\theta) y_1 + \beta_2(\theta) y_2 + \beta_3(\theta) y_1 y_2$$

are affine in attributes defined as $\mathbf{y}' = (y_1, y_2, y_1 y_2)$ instead of $\mathbf{y} = (y_1, y_2)$.

The key assumption that potentially limits generality is assuming that utility is a continuous function of price that decreases in price. Economists almost always use this assumption; in particular, all econometric Mixed Logit models based on revealed preferences data we are aware of satisfy Assumption 1. There are, however, Mixed Logit models that do not satisfy Assumption 1 because prices are discrete or the associated utilities are not strictly decreasing in price. For example, marketing researchers commonly build "latent-class" Logit models based on discrete attribute and price "levels"; product design-and-price optimization is different in this setting [71].

While engineering design studies have also adopted this approach, the estimated "part-worths" are commonly interpolated to obtain a continuous utility function over prices or other product attributes; see, e.g., Refs. [39,62]. Specifically, if a survey used $p^{(1)} \leq \cdots \leq p^{(L)}$ as the L levels for price and estimated the resulting part-worth coefficients $\beta^g_{p,\ell}$ for price levels $\ell = 1, \ldots, L$ and (latent) groups $g = 1, \ldots, G$, then a utility model over continuous prices in any group may be specified through any interpolator $B^g_p : [p^{(1)}, p^{(L)}] \rightarrow \mathbb{R}$ of $\{\beta^g_{p,\ell}\}^L_{\ell=1}$ ($B^g_p(p^{(\ell)}) = \beta^g_{p,\ell}$ for all ℓ). If the estimated part-worths increase in price over any successive levels, so will any interpolator of the part-worths and thus, Assumption 1(a) will not hold. For example, one model used by Shiau and Michalek [39] has part-worths that increase in price, and thus, an interpolated utility function that does as well. Note that should utility increase with price then the optimal pricing decision is trivial: make price as high as possible.

We believe that models claiming that higher prices are preferred should be approached with caution. In general prices have two functions: they are literally a cost associated with transfer of ownership, but are also a signal of shared expressions or expectations of product value. High prices can signal to a potential customer a level of product desirability and value that may be poorly captured in observable attributes. The purchase of high-priced products may also serve to signal individual characteristics (e.g., wealth) to the rest of the population, a phenomenon termed "conspicuous consumption" [72]. Increasing preferences for prices are likely to be related to the signaling function of prices, not literally the cost of ownership transfer. Empirical choice models should attempt to disentangle these two effects as much as possible. One potential empirical mechanism might be to incorporate interaction effects between price and product characteristics assuming that signaling for value is effective only conditional on characteristic values. For example, it is unlikely that the Honda Fit can benefit from luxury price signaling, while very likely the Lexus 480 h can. When preferences over prices reflect only purchasing costs it is probably safe to assume lower prices are preferred, in accordance with Assumption 1.

2.3 Simulation. Using Mixed Logit choice probabilities (and their derivatives) in practice requires *simulation* [73]: integral approximation using random draws from the demographic distribution. We use simple sample averaging by drawing $I \in \mathbb{N}$ samples of the utility coefficients θ_i from the distribution μ defined above and solving the optimal design problem with P_j replaced by the *simulator* [73]

$$\frac{1}{I}\sum_{i=1}^{I} P^L_j(\theta_i, \mathbf{X}, \mathbf{p}) \tag{2}$$

generated by these samples; more efficient techniques for simulation than sampling directly from μ exist [73]. The Law of Large numbers guarantees that as $I \uparrow \infty$, the simulated probabilities, Eq. (2), converge to the true probabilities. Unfortunately, solutions to the simulated optimal design problem using probabilities as defined in Eq. (2) may not converge to the solution of the true problem without additional conditions. Compactness of the feasible space is sufficient, as can be proven using the theory described in Ref. [74]; see also Ref. [51].

We have found computations to be more robust by neglecting the $1/I$ term in Eq. (2) as well as the corresponding derivatives of the simulated Mixed Logit choice probabilities; Sec. 4 reports results only for this case. If I is much larger than J, as could be required to ensure convergence, dividing by I could result in loss of accuracy. Changing this scaling factor does not affect solutions, but can affect solver performance; indeed, some studies also include a "market size" that scales *up* the choice probability approximations in addition to division by I (e.g., Ref. [39]).

2.4 Profits and Optimal Design. Firm f designs the products in $\mathcal{J}_f \subset \{1, \ldots, J\}$. Each product $j \in \mathcal{J}_f$ is defined by an N_j-dimensional vector of product characteristics $\mathbf{x}_j \in \mathbb{R}^{N_j}$. We assume that feasible characteristic vectors satisfy finite lower and upper bounds $\mathbf{l}_j, \mathbf{u}_j$ and some set of equality ($\mathbf{g}_j(\mathbf{x}_j) = \mathbf{0}$) and inequality ($\mathbf{h}_j(\mathbf{x}_j) \geq \mathbf{0}$) constraints. Unit costs for product j are a function of the design vector $\mathbf{x}_j : c_j : [\mathbf{l}_j, \mathbf{u}_j] \rightarrow [0, \infty)$. We allow the cost functions to be product-specific to capture situations where there are fixed product-specific features that influence costs. For each product j there is also a transformation, $\mathbf{y}_j : [\mathbf{l}_j, \mathbf{u}_j] \rightarrow \mathcal{Y}_j$, that defines product attributes given the design vector. We allow this transformation to be product-specific to capture situations where there are parameters not included in the set of design variables that influence how product performance is perceived by consumers. Both c_j and $\mathbf{y}_j$ are assumed to be twice continuously differentiable. We often simply denote $\mathbf{y}_j(\mathbf{x}_j)$ by $\mathbf{y}_j$, and $\mathbf{Y}(\mathbf{X}) = \{\mathbf{y}_j(\mathbf{x}_j)\}^J_{j=1}$ by $\mathbf{Y}$. Finally, firm f has fixed costs $c^F_f \geq 0$.

If both unit and fixed costs are independent of the quantity sold [67], expected profits are given by

$$\pi_f(\mathbf{X}, \mathbf{p}) = S\hat{\pi}_f(\mathbf{Y}(\mathbf{X}), \mathbf{c}(\mathbf{X}), \mathbf{p}) - c^F_f \tag{3}$$

for a market with S individuals, where

$$\hat{\pi}_f(\mathbf{Y}, \mathbf{c}, \mathbf{p}) = \sum_{j \in \mathcal{J}_f} P_j(\mathbf{Y}, \mathbf{p})(p_j - c_j) \tag{4}$$

for any $\mathbf{Y}$, $\mathbf{c}$, $\mathbf{p}$. Because scaling and shifting do not affect optimizers, we can neglect market size S and fixed costs c^F_f in our analysis below. Including a discount factor to harmonize revenues or costs that occur at different points in time also simply scales profits and/or costs, and is thus, left out of our formulation.

Optimal design-*and*-pricing can then be written as follows:

$$\begin{aligned} \max \quad & \hat{\pi}_f(\mathbf{X}, \mathbf{p}) \\ \text{w.r.t.} \quad & \mathbf{l}_j \leq \mathbf{x}_j \leq \mathbf{u}_j, \\ & 0 \leq p_j \leq \imath_* \text{ for all } j \in \mathcal{J}_f \\ \text{s.t.} \quad & \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0}, \ \mathbf{h}_j(\mathbf{x}_j) \geq \mathbf{0} \text{ for all } j \in \mathcal{J}_f \end{aligned} \tag{5}$$

That is, the firm chooses a design and a price for each of the products it offers, taking its competitors' designs and prices as fixed; recall Fig. 1. Note that we treat $\imath_*$ as a price bound; this can be viewed as an explicit constraint, but derives from the firm's recognition that setting p_j above $\imath_*$ would result in a zero probability of any individual choosing product j. Optimal design-*then*-pricing is written as follows:

$$\begin{aligned} \max \quad & \hat{\pi}_f(\mathbf{X}, \mathbf{p}) \\ \text{w.r.t.} \quad & \mathbf{l}_j \leq \mathbf{x}_j \leq \mathbf{u}_j \text{ for all } j \in \mathcal{J}_f, \\ & 0 \leq p_j \leq \imath_* \text{ for all } j \in \mathcal{J} \\ \text{s.t.} \quad & \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0}, \ \mathbf{h}_j(\mathbf{x}_j) \geq \mathbf{0} \text{ for all } j \in \mathcal{J}_f \\ & \mathbf{p} \in \mathcal{E}(\mathbf{Y}(\mathbf{X}), \mathbf{c}(\mathbf{X})) \end{aligned} \tag{6}$$

where $\mathbf{p} \in \mathcal{E}(\mathbf{Y}(\mathbf{X}), \mathbf{c}(\mathbf{X}))$ denotes the condition that prices, for *all* products, are in "equilibrium" given the design decisions $\mathbf{x}_j$, $j \in \mathcal{J}_f$. Section 3.1 and Appendix A define and discuss equilibrium prices in more detail. Note that firm f's design-then-pricing problem concerns its own as well as its competitors' product prices, rather than just the prices of the firm f's own products. Firm f is not literally choosing its' competitors' prices for them, but is rather "forecasting" or "anticipating" competitors' prices (as well as its own) as an equilibrium response to its design choices.

Appendix A details how to extend this framework to situations where the firm would explicitly constrain prices to lie within some bounds. This could occur if, for example, the firm was uncomfortable pricing inconsistently with historical norms, does not trust the demand model outside of some window around current prices, or has some specific "markup constraints" imposed by company policy, retailers, or even regulators [39]. Little changes in our model should explicit bounds be adopted; however, we believe that these bounds are best specified when given explicitly by the consumer choice model [70].

3 Solving Design-Then-Pricing Problems

This section discusses three methods for solving Eq. (6), the design-then-pricing problem. First is an implicit programming method that treats prices as an intermediate quantity that is computed for any set of designs. Second is a "MPEC" formulation where the combined gradient is added as a constraint on all prices. We prove that this formulation has spurious KKT points that threaten to nullify the strong convergence properties of existing solvers. Finally, we present a MPEC formulation with a coercive representation of equilibrium prices that eliminates these spurious KKT points. Before discussing these methods, we briefly review several results pertaining to equilibrium pricing that are important for the discussion that follows.

3.1 Equilibrium Prices. Methods for efficiently and reliably solving design-then-pricing problems take advantage of efficient and reliable methods for computing equilibrium prices. This section briefly reviews existing work in this area, drawn largely from Refs. [41,50,51]. See Appendix A, as well as these references, for more information.

If the design decisions $\mathbf{X}$ are fixed, then $\mathbf{Y}$ and $\mathbf{c}$ are fixed and each firm solves the pricing problem

$$\begin{aligned} &\max \quad \hat{\pi}_f(\mathbf{Y}, \mathbf{c}, \mathbf{p}) \\ &\text{w.r.t.} \quad 0 \le p_j \le \imath_* \text{ for all } j \in \mathcal{J}_f \end{aligned} \tag{7}$$

(Local) Equilibrium prices are prices $\mathbf{p}$ for which $\mathbf{p}_f = \{p_j : j \in \mathcal{J}_f\}$ (locally) solves Eq. (7) for all f. The set of all (local) equilibrium prices for given product characteristics and costs is denoted by $(\mathcal{E}^{\ell}(\mathbf{Y}, \mathbf{c}))\mathcal{E}(\mathbf{Y}, \mathbf{c})$.

If the choice probabilities are continuously differentiable in prices on $[0, \imath_*]^J$, then equilibrium prices $\mathbf{p}$ solve the combined KKT conditions for every firms' pricing problem. These combined conditions can be written as the mixed complementarity problem (MCP)[1]

$$\mathbf{0} \le \mathbf{p} \le \imath_* \mathbf{1} \perp -(\tilde{\nabla}^p \hat{\pi})(\mathbf{Y}, \mathbf{c}, \mathbf{p}) \tag{8}$$

where the "combined gradient," $(\tilde{\nabla}^p \hat{\pi})(\mathbf{Y}, \mathbf{c}, \mathbf{p})$ has components

$$((\tilde{\nabla}^p \hat{\pi})(\mathbf{Y}, \mathbf{c}, \mathbf{p}))_k = (D^p_k \hat{\pi}_{f(k)})(\mathbf{Y}, \mathbf{c}, \mathbf{p})$$

and $f(k)$ denotes the unique firm offering product k.

Unfortunately, solving Eq. (8) is an unreliable method for computing equilibrium prices [41,51]. The fundamental problem is that the profit derivatives, with respect to prices, vanish as prices approach $\imath_*$; see Lemma 1 below. This leads to the existence of spurious KKT points of Eq. (8): points that are stationary in part just because prices are large.

This problem can be corrected by solving

$$\mathbf{0} \le \mathbf{p} \le \imath_* \mathbf{1} \perp \varphi(\mathbf{Y}, \mathbf{c}, \mathbf{p}) \tag{9}$$

instead. The definition of the map $\varphi(\mathbf{Y}, \mathbf{c}, \mathbf{p})$ is somewhat technical and provided elsewhere; see Appendix A and Refs. [41,50,51]. The effect of this transformation is as follows: the Jacobian matrix of the choice probabilities with respect to prices can be "split" into a diagonal matrix and a full matrix, where the diagonal matrix contains the key components of the choice probability derivatives that vanish as prices become large. By factoring this diagonal matrix out of $(\tilde{\nabla}^p \hat{\pi})$ we obtain φ that does *not* vanish as prices become large. In this way, Eq. (9) is essentially equivalent to Eq. (8) but has solutions with components equal to $\imath_*$ only when this is profit-maximizing [51]. Moreover, for simulators φ can be continuously differentiably extended to all of $\mathbb{R}^J$ in such a way that if $\mathbf{p}$ solves $\varphi(\mathbf{Y}, \mathbf{c}, \mathbf{p}) = \mathbf{0}$ then the projection of $\mathbf{p}$ onto $[0, \imath_*]^J$ solves Eq. (9) [51]. Thus, we can presume, without loss of generality, that the *smooth nonlinear system* $\varphi(\mathbf{Y}, \mathbf{c}, \mathbf{p}) = \mathbf{0}$ is defined for all prices and provides a first-order condition for equilibrium prices that eliminates "spuriously" stationary prices that are stationary only because some prices are large. Appendix A details how this logic can be extended for problems with explicit bounds on prices, as in Ref. [39]. The reformulation in Eq. (9) also applies to problems with strategic retailers (e.g., Refs. [5–7,34,35,42]).

3.2 Implicit Programming. One approach to solving Eq. (6) assumes that $\mathbf{p}^* \in \mathcal{E}^{\ell}(\mathbf{Y}, \mathbf{c})$ can be computed for any $(\mathbf{Y}, \mathbf{c})$ and then treats prices as an intermediate variable

$$\begin{aligned} &\max \quad \hat{\pi}^*_f(\mathbf{X}) \\ &\text{w.r.t.} \quad \mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \text{ for all } j \in \mathcal{J}_f \\ &\text{s.t.} \quad \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0}, \ \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0} \text{ for all } j \in \mathcal{J}_f \end{aligned} \tag{10}$$

where $\hat{\pi}^*_f(\mathbf{X}) = \hat{\pi}_f(\mathbf{X}, \mathbf{p}^*(\mathbf{X}))$ and $\mathbf{p}^*(\mathbf{X}) \in \mathcal{E}^{\ell}(\mathbf{Y}(\mathbf{X}), \mathbf{c}(\mathbf{X}))$ are local equilibrium prices written as a function of $\mathbf{X}$. Local, rather than proper, equilibrium is used here because it will be impractical, in general, to *guarantee* computations of equilibria, as opposed to local equilibria. We follow Luo et al. [46] in calling solution of Eq. (10) an implicit programming method because of the theoretical reliance of this approach on some form of Implicit Function Theorem (see, e.g., Refs. [46,77]) to ensure that equilibrium prices are at least locally unique and, in some sense, differentiable in $\mathbf{X}$.

Implicit programming has been applied successfully to solve design-then-pricing problems. Shiau and Michalek [39] discuss a variant of this approach where equilibrium prices are computed via an iterated optimization strategy–"Variational Relaxation"–following Choi et al. [31] rather than solution of Eq. (8) or Eq. (9). Frischknecht et al. [47] optimize vehicle designs by computing equilibrium prices as an intermediate variable using a fixed-point method derived in Ref. [41] based on Eq. (9); we use the same technique in our numerical results. This iteration is the fastest known method for reliably computing equilibrium prices with Mixed Logit models [41,51].

Implicit programming may only be practical when stationary prices are characterized by a smooth nonlinear system. Specifically, most conventional solvers for solving Eq. (10) would formally require that $\mathbf{p}^*(\mathbf{X})$ is *twice* continuously differentiable in $\mathbf{X}_f$, to ensure that $\hat{\pi}^*_f$ is twice continuously differentiable [52]. The review in Sec. 3.1 suggests that this is not a serious obstacle for the models considered in this article; however, the differentiability requirement should formally rule out problems with prices constrained by explicit bounds. In general, the first-order conditions for equilibrium only have *piecewise*-smooth parametric solution maps and specialized methods may thus, be required to ensure global convergence to optimal designs [46]. Moreover, even if $\mathbf{p}^*(\mathbf{X})$ is defined by a smooth nonlinear system, $\mathbf{p}^*(\mathbf{X})$ may still fail to be even differentiable at some $\mathbf{X}$ if the Jacobian of the nonlinear system used to solve for $\mathbf{p}^*(\mathbf{X})$ is singular at some point. Conditions ensuring nonsingularity of the Jacobian of φ with respect to prices are not known.

[1]The real-valued MCP "$l \le x \le u \perp F(x)$" is solved by x satisfying one of the following: $x \in [l,u]$ if $F(x) = 0$, $x = l$ if $F(l) > 0$, or $x = u$ if $F(u) < 0$. Note the similarity to the KKT conditions for bound-constrained optimization. In fact, if $F(x)$ is the derivative of some function $f(x)$, then $l \le x \le u \perp F(x)$ are the KKT conditions for min $f(x)$ subject to $l \le x \le u$. See Refs. [75,76] or Appendix B for a generalization to vector-valued MCPs.

Moreover, equilibrium prices must be computed, using iterative methods, for *every* trial vector of designs chosen by an optimizer. Introducing these computations as an "inner loop" may become computationally prohibitive for large problems, especially if choice probabilities, profits, and profit gradients in equilibrium are to be evaluated accurately. In our experience, part of this expense can be mitigated by "warm starting" equilibrium computations at prices that were in equilibrium for previous design variable values. But because equilibrium prices can only be *approximated*, treating prices as an intermediate variable can also introduce noise into the optimization. This will be especially problematic if finite-differences are used, as noise in the approximate solves can amplify truncation errors associated with finite differencing. Thus, the derivatives of equilibrium prices with respect to design decisions must be computed explicitly using the standard Implicit Function Theorem if this method is to be used. Computing these derivatives, when they exist, requires defining and solving a $J \times J$ linear system. Even if this is possible it may incur a significant computational burden and introduce more numerical error when there are many products. We have found that equilibrium prices must be computed to very tight tolerances to ensure reliable solution of the design-only problem; see Sec. 4 for details.

3.3 Combined Gradient Constrained MPEC Approach. Almost two decades of research concerning MPECs suggests that the quantities in equilibrium should be considered problem variables subject to some constraints that represent equilibrium. The constraints used are typically only a first-order stationarity condition, creating a mathematical program with complementarity constraints (MPCC) [78–80]. Solving this MPCC does not guarantee that subgame variables are in equilibrium at solutions to the MPCC when the subproblems are nonconvex, as is the case with design-then-pricing. The potential advantage is that prices can be out-of-equilibrium for the majority of solver iterations. Moreover, the derivatives of the equilibrium constraints do not threaten to introduce noise into the optimization in the same manner as the approximately computed equilibrium prices and their derivatives as required in implicit programming.

In the case of design-then-pricing, substitution of Eq. (8) yields the following problem, a special case of the form used in Ref. [39]

$$\begin{aligned} \text{max} \quad & \hat{\pi}_f(\mathbf{X}, \mathbf{p}) \\ \text{w.r.t.} \quad & \mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \text{ for all } j \in \mathcal{J}_f, \ \mathbf{p} \ge \mathbf{0} \\ \text{s.t.} \quad & \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0}, \ \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0} \text{ for all } j \in \mathcal{J}_f \\ & (\tilde{\nabla}^p \hat{\pi})(\mathbf{X}, \mathbf{p}) = \mathbf{0} \end{aligned} \tag{11}$$

where $(\tilde{\nabla}^p \hat{\pi})(\mathbf{X}, \mathbf{p}) = (\tilde{\nabla}^p \hat{\pi})(\mathbf{Y}(\mathbf{X}), \mathbf{c}(\mathbf{X}), \mathbf{p})$ is a more compact notation for the combined gradient of profits (with respect to prices) as a function of design decisions. Lemma 1 proves that Eq. (8), formally a MCP, can be thought of as a smooth nonlinear system when we exclude explicit bounds on prices (and no price is profit-optimally zero [50]). Thus, while we are literally including an "equilibrium constraint", this constraint is not as difficult as the complementarity constraints typically considered in the MPCC literature. Problems with explicit bounds on prices will have proper complementarity constraints that must be handled appropriately [80]. We do not view this as a fundamental impediment to solving design-then-pricing problems. Sequential quadratic programming (SQP) solvers like SNOPT [81], applied in our numerical results, retain many of their strong convergence properties for MPCCs formulated in the correct way [78]. Interior-point (IP) [49,82] methods can also perform very well on MPCCs when the complementarity constraints are included using objective penalization [79]. Augmented Lagrangian techniques have also shown promise in their ability to handle problems with degenerate constraints, including MPCCs [83].

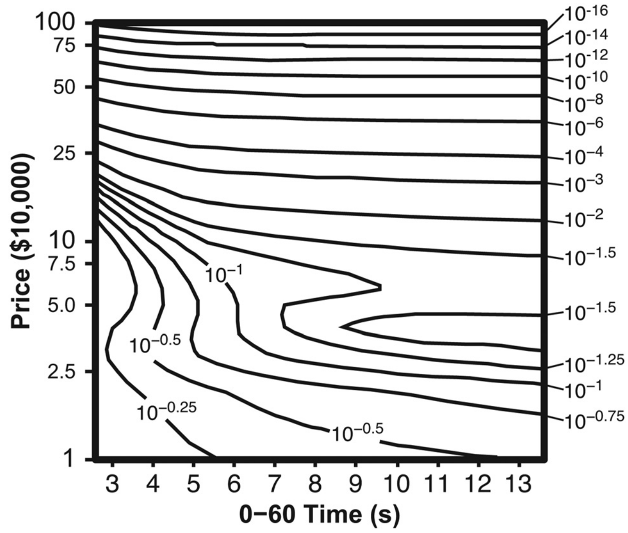

Fig. 2 Illustration of Lemma 1. Contours denote the level sets of the max norm ($\|\mathbf{x}\|_\infty = \max_n\{|x_n|\}$) of the profit gradient for a two-vehicle design-then-pricing problem as defined in Sec. 4.1. Labels denote the value of the norm over the contours drawn.

3.4 Spurious Solutions in Design-then-Pricing. Equation (11), like Eq. (8), also has spurious KKT points that can be computed by commercial solvers. We illustrate this situation with a two-vehicle instance of the model described below in Sec. 4.1. A firm offers one vehicle, and has one competitor offering one vehicle. The firm chooses the vehicle's 0-60 acceleration and price, with a constraint that uniquely defines fuel economy. Figure 2 plots the max norm of the profit gradient over feasible accelerations and prices from \$10,000 to \$1,000,000. Note that as price increases, the profit gradient vanishes for *any* value of acceleration. As we clarify below, this behavior suggests that the KKT conditions would be asymptotically satisfied for large prices and any value of acceleration. Moreover, note that both components of the profit gradient have magnitude less than 10^{-6} if price is above \$400,000, regardless of acceleration. Hence, many more points may "numerically" satisfy the KKT conditions in the sense of satisfying numerical termination criteria used in many existing solvers.

The existence of these spurious KKT points is a salient issue for design-then-pricing. Our numerical example in Sec. 4 shows that these spurious KKT points can be a significant obstacle to practical computations in large markets. However, Fig. 2 shows that spurious KKT points exist for problems of any scale, even with a single product. Moreover, the *rate* at which the magnitude of the profit gradients vanish depends sensitively on the problem. Repeating this example using a mean price coefficient, one magnitude larger than that used in Sec. 4.1, the profit gradient is "numerically" zero (i.e., $< 10^{-6}$) for prices larger than \$40,000, instead of \$400,000. Alternatively, assuming the firm faces 471 competing vehicles, instead of one, and using the original price coefficient leads to profit gradients that are numerically zero for prices larger than \$250,000. Note also that spurious KKT points could still exist when using explicit bounds on prices. Specifically, should the upper bound on prices be *larger* than the *smallest* price such that the profit gradient is numerically zero, then solvers could still terminate at spurious KKT points even with explicit bounds on prices. These issues motivate a clearer identification of the problem and its resolution.

We develop a better understanding of this issue by examining the KKT conditions for Eq. (11), which can be written as the following MCP:

$$
\left.\begin{aligned}
\mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \perp L_j^x(\mathbf{X},\mathbf{p},\mu_j,\mu_f^p) \\
\infty < \mu_j^E < \infty \perp \mathbf{g}_j(\mathbf{x}_j) \\
\mathbf{0} \le \mu_j^I \perp \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0}
\end{aligned}\right\} \quad \text{for all } j \in \mathcal{J}_f
$$
$$
\mathbf{0} = -(\nabla^p \hat{\pi}_f)(\mathbf{X},\mathbf{p}) + (D^p \tilde{\nabla}^p \hat{\pi})(\mathbf{X},\mathbf{p})^\top \mu_f^p
$$
$$
\mathbf{0} = (\tilde{\nabla}^p \hat{\pi})(\mathbf{X},\mathbf{p}) \tag{12}
$$

where $\mu_j = (\mu_j^E, \mu_j^I)$ and μ_f^p are Lagrange multipliers for the design variable equality constraints, design variable inequality constraints, and equilibrium price constraints (respectively) and

$$
\begin{aligned}
L_j^x(\mathbf{X},\mathbf{p},\mu_j,\mu_f^p) = &-(\nabla_j^x \hat{\pi}_f)(\mathbf{X},\mathbf{p}) - (D_j^x \tilde{\nabla}^p \hat{\pi})(\mathbf{X},\mathbf{p})^\top \mu_f^p \\
&- (D_j^x \mathbf{g}_j)(\mathbf{x}_j)^\top \mu_j^E - (D_j^x \mathbf{h}_j)(\mathbf{x}_j)^\top \mu_j^I
\end{aligned}
$$

where ∇_j^x and D_j^x denote the gradient and derivative with respect to the product j's design decisions. See Appendix B for more details on this MCP notation.

Taking all the multipliers to be zero, Eq. (12) reduces to

$$
\left.\begin{aligned}
\mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \perp -(\nabla_j^x \hat{\pi}_f)(\mathbf{X},\mathbf{p}) \\
\mathbf{g}_j(\mathbf{x}_j) = \mathbf{0},\ \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0}
\end{aligned}\right\} \quad \text{for all } j \in \mathcal{J}_f
$$
$$
(\nabla^p \hat{\pi}_f)(\mathbf{X},\mathbf{p}) = \mathbf{0},\ (\tilde{\nabla}^p \hat{\pi})(\mathbf{X},\mathbf{p}) = \mathbf{0}
$$

Lemma 1 below suggests that Eq. (12) can be approximately satisfied by making $\mathbf{p}_f$ sufficiently large, for *any* feasible design vectors $\mathbf{x}_j$; this shows the generality of the problem illustrated by Fig. 2.

While this will be true for most Mixed Logit models, we provide a result for *simulations* of Mixed Logit models that are applied in computations; see Sec. 2.3.

LEMMA 1. *Suppose the Mixed Logit model is a simulator of a Mixed Logit model satisfying Assumption 1. Then, as* $p_k \uparrow \iota_*$, *the following hold: (i)* $(D_k^p \hat{\pi}_f) \to 0$ *for any firm f, even if* $k \notin \mathcal{J}_f$; *(ii) if* $k \in \mathcal{J}_f, (\nabla_k^x \hat{\pi}_f) \to \mathbf{0}$; *(iii)* $(\nabla_j^x \tilde{\nabla}^p \hat{\pi})(\mathbf{X},\mathbf{p}) \to \mathbf{0}$; *(iv)* $(D_l^p D_k^p \hat{\pi}_{f(k)})(\mathbf{X},\mathbf{p}) = (D_k^p D_l^p \hat{\pi}_{f(l)})(\mathbf{X},\mathbf{p}) = 0$.

Proof. The choice probabilities are a weighted sum of Logit choice probabilities; the result thus, follows from the corresponding result for Logit models given in Appendix C. □

COROLLARY 1. *Suppose the Mixed Logit model is a simulator of a Mixed Logit model satisfying Assumption 1. As* $p_j \uparrow \iota_*$, *the KKT conditions (12) for Eq. (11) converge to those of the design-then-pricing problem without product j, regardless of whether eliminating product j from the product portfolio is actually profit-optimal. More generally, let* $\mathcal{J}_f^* \subset \mathcal{J}_f$, $\mathcal{J}_f^\circ = \mathcal{J}_f \backslash \mathcal{J}_f^*$, *and suppose* $(\mathbf{x}_j, p_j, \mu_j^E, \mu_j^I, \mu_j^p)$, *for all* $j \in \mathcal{J}_f^\circ$, *satisfy the KKT conditions for Eq. (11) assuming firm f only offered the products in* $\mathcal{J}_f^\circ$. *Then, these values appended with* $(\mathbf{x}_j, \infty, \mathbf{0}, \mathbf{0}, \mu_j^p)$ *for any feasible* $\mathbf{x}_j$ *and any* $\mu_j^p \in \mathbb{R}$, *for all* $j \in \mathcal{J}_f^*$ *forms a KKT point for Eq. (11).*

See Appendix C for the proof, that is a relatively straightforward application of Lemma 1 (despite requiring a complicated notation).

Corollary 1 states that the KKT conditions for Eq. (11) "nest" the KKT conditions for any subproblems formed by eliminating *any* subset of the products offered by firm *f*, *even though eliminating these products may not be profit-optimal*. Computationally the existence of these potentially spurious KKT points threaten to degrade the reliability of SQP or IP solvers, even though SQP and IP solvers do more than explicitly solve the KKT conditions. This theoretical conclusion is confirmed in our numerical results below.

3.5 A Well-Posed MPEC Approach. Using the reformulation of the stationarity condition for equilibrium prices in Eq. (9) eliminates spurious KKT points from the set of computable solutions to Eq. (11). Specifically, instead of Eq. (11), consider

$$
\begin{aligned}
\max \quad & \hat{\pi}_f(\mathbf{X},\mathbf{p}) \\
\text{w.r.t.} \quad & \mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \text{ for all } j \in \mathcal{J}_f,\ \mathbf{p} \ge \mathbf{0} \\
\text{s.t.} \quad & \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0},\ \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0} \text{ for all } j \in \mathcal{J}_f \\
& \varphi(\mathbf{X},\mathbf{p}) = \mathbf{0}
\end{aligned} \tag{13}
$$

Problems (11) and (13) are functionally equivalent because the objective has not been changed and the feasible sets are equivalent *except* for the excision of spurious KKT points: because the KKT conditions for Eq. (13) contain the equation $\varphi(\mathbf{X},\mathbf{p}) = \mathbf{0}$, that has no spurious solutions [41,50,51], there are no spurious KKT points. Formally, our assumption that the feasible space for design decisions is bounded is required to retain this property of φ without making further assumptions on the choice and cost model.

3.6 Spurious Solutions in Design-and-Pricing. Lemma 1 and its corollary regarding spurious KKT points applies equally well to design-*and*-pricing. The results in Sec. 4 show that this may be an impediment to reliable computations.

Equation (13) can be adapted to obtain a well-posed formulation of design-and-pricing

$$
\begin{aligned}
\max \quad & \hat{\pi}_f(\mathbf{X},\mathbf{p}) \\
\text{w.r.t.} \quad & \mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j,\ p_j \ge 0 \text{ for all } j \in \mathcal{J}_f \\
\text{s.t.} \quad & \mathbf{g}_j(\mathbf{x}_j) = \mathbf{0},\ \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0} \text{ for all } j \in \mathcal{J}_f \\
& \varphi_f(\mathbf{X},\mathbf{p}) = \mathbf{0}
\end{aligned} \tag{14}
$$

where φ_f ($\mathbf{X}$, $\mathbf{p}$) are the components of φ corresponding only to the products offered by firm *f*. Because $\varphi_f(\mathbf{X},\mathbf{p}) = \mathbf{0}$ implies satisfaction of the first-order condition for profit-maximizing prices for firm *f*, the additional constraint in Eq. (14) is redundant and will have zero Lagrange multipliers at any KKT point. However, because φ_f ($\mathbf{X}$, $\mathbf{p}$) $= \mathbf{0}$ rules out large prices that are not profit-maximizing, Eq. (14) does not have spurious KKT points.

3.7 Sufficient Conditions. Solving either MPEC formulation, Eqs. (11) and (13), only ensures that $\mathbf{p} \in \mathcal{S}(\mathbf{Y}(\mathbf{X}),\mathbf{c}(\mathbf{X}))$ at optimality, while the true goal is to solve for optimal designs with $\mathbf{p} \in \mathcal{E}^\ell(\mathbf{Y}(\mathbf{X}),\mathbf{c}(\mathbf{X}))$. Similarly, the implicit method solves a nonlinear system equivalent to the first-order condition for equilibrium prices at $\mathbf{X}$, that in principle only guarantees that $\mathbf{p}_*(\mathbf{X}) \in \mathcal{S}(\mathbf{Y}(\mathbf{X}),\mathbf{c}(\mathbf{X}))$. See Ref. [84] for an example design-then-pricing problem in which solvers may compute solutions that do *not* have prices in equilibrium because some firms profits are locally minimized, not maximized. Thus, verifying a solution to Eqs. (10), (11), and (13) is in fact a solution to Eq. (6) requires checking a SOSC for equilibrium prices.

Shiau and Michalek [39] also recognize this, but propose a heuristic SOSC check that does not necessarily identify local optimality. Explicitly, they propose to re-optimize firms' profits, with respect to prices alone, at a first-order solution to the design-then-pricing problem. If no higher value of profits is obtained, they assume the resulting point is an equilibrium. The primary concern with this approach is that if the first-order condition for equilibrium is solved for the prices used in the starting condition in a firm's price-only optimization then the optimizer should terminate immediately regardless of whether the firm's current prices locally maximize or minimize profits. We could perturb initial prices, but there is little to say how far we should perturb prices. Worse, if we had originally computed a spurious KKT point then both the first and second partial derivatives of profits will be very small in a neighborhood of the computed point. Thus, this re-optimization check is unlikely to resolve spurious solutions should one be computed.

A sufficient condition for local optimality is that a particular submatrix $\mathbf{H}'_f$ of every firms' Hessian matrix $\mathbf{H}_f$, with respect to changes in their own prices alone, is negative definite; see Chap. 12 in Ref. [52] for general information on the SOSC. The appropriate submatrix to use consists of those rows and columns corresponding to prices set by the firm within $(0, \imath_*)$ or, in the case of explicit bounds, within the interior of the bounded region (see Appendix A). This check is likely to be faster and more robust than re-optimizing firms' profits with respect to prices alone. Moreover, re-optimizing with SQP or IP methods uses (or approximates) the same information: local curvature information captured in firms' Hessians. Verifying the negative definiteness of $\mathbf{H}'_f$ can be accomplished with a single Cholesky factorization applied to $-\mathbf{H}'_f$. Cholesky is relatively fast, stable against numerical errors, and the standard method for testing definiteness [85–87].

Because negative definiteness of $\mathbf{H}'_f$ is only sufficient, there could be cases where $\mathbf{H}'_f$ fails to be negative definite *and* the current value of $\mathbf{p}_f$ locally maximizes firm f's profits. However, negative *semi*definiteness of $\mathbf{H}'_f$ *is* necessary [52]. Thus, if we fail to identify that $\mathbf{H}'_f$ is negative definite, showing $\mathbf{H}'_f$ to be indefinite proves that the current prices are not in equilibrium. If the original Cholesky factorization failed the factorization up to termination can often be used to determine a direction of positive curvature–a vector $\mathbf{d}$ such that $\mathbf{d}^\top \mathbf{H}_f \mathbf{d} > 0$ [88]—proving that $\mathbf{H}_f$ is indefinite rather than negative semidefinite. Alternatively, several techniques may reveal semidefiniteness: Cholesky factorization with pivoting, an inertia-revealing symmetric indefinite factorization, or identification of the null-space of $\mathbf{H}'_f$ could used [85,87]. However, identifying semidefiniteness is likely to be sensitive to two types of numerical errors: errors in the linear-algebraic methods used and in the approximation of the Hessian matrices themselves. We believe the latter category of errors are likely to dominate purely numerical errors. In fact, we may not be able to identify if $\mathbf{H}'_f$ is semidefinite but rather only "sufficiently" negative definite or "sufficiently" indefinite, relative to the level of errors involved in computing $\mathbf{H}'_f$. These issues require further research, as practical SOSC checks with "noisy" function evaluations have not yet been investigated.

Any factorization-based SOSC checks presume that the Hessian matrices for each firm are available. We have developed formulas and code for computing these matrices in a generic Mixed Logit simulator, and use this in our numerical example below. Neither formulas nor code are provided here, but both are available by request. There are "Hessian-free" methods for checking the SOSC based on the same linear-algebraic ideas, including Cholesky factorization, that are more efficient if the Hessian is not known explicitly [88].

It is not so important to verify a SOSC for the optimal design problem, Eq. (6), itself. Modern solvers cannot compute local minimizers to a maximization problem and are likely to avoid saddles. Some algorithms use of directions of negative curvature to guarantee convergence to a second-order necessary point, fully avoiding saddles [89].

4 Comparison of Methods

This section presents results for a design-then-pricing problem based on the 2006 new vehicle market, described below. A firm chooses the acceleration, fuel economy, and technology content of a subset of vehicles expecting prices for all vehicles on the market to be set in equilibrium as a result of these decisions. All computational results below were generated using the SNOPT SQP software [81]. Problem data are programmed in C and all runs were executed on a single Mac Pro tower (OS X version 10.6.8) with dual 6-core 2.66 GHz processors, 64 GB of RAM, and 1333 MHz bus. SNOPT's optimality and feasibility tolerances were both set to 10^{-6}, and 1000 major iterations were allowed. We have also found this example problem to be solvable with the KNITRO software implementing an IP method [82]; however, KNITRO can take significantly longer to solve than SNOPT. In general, SQP strategies should be expected to perform better on our example problem because the solutions are highly constrained.

4.1 Vehicle Design Model and Mixed Logit Utility Specification

This section defines our real-scale vehicle portfolio design problem. In total, 21 automakers ("firms") offer 472 vehicles (i.e., $F = 21$ and $J = 472$). We focus on firm 1, who offers 29 vehicles expecting prices for all vehicles to be set in equilibrium given their design decisions. Fixed vehicle characteristics for all other vehicles are drawn from calendar year 2006 vehicle market data; the design and choice models discussed below are adapted from Ref. [9]. This model is not intended to be a high-fidelity model of vehicle design; our intended application of this model is a comparison of numerical methods in a large market.

Each vehicle j is described by its 0-60 acceleration time (a_j, in s), fuel economy (e_j, in mpg), "technology content" (t_j, unitless) and Manufacturer's Suggested Retail Price (p, 10^4\$), referred to below simply as "price"; body style, footprint, and weight are fixed, vehicle-specific parameters. Technology content is a continuous index of the level of cost-effective technology features included in the vehicle; see Ref. [9] for more details. Vehicle j has body-style specific bounds on a_j, e_j, and t_j derived from the data discussed in Ref. [9].

Models of vehicle performance and costs are as follows: Vehicle j's fuel economy, acceleration, and technology content are related through a single equality constraint $g_j(e, a, t) = 0$ where

$$g_j(e,a,t) = \frac{1000}{e-3.46} - \beta_{b(j),1} - \beta_{b(j),2}\exp\{-a\} - \beta_{b(j),3}t - \beta_{b(j),4}a^2 t - \beta_{b(j),5}w_j - \beta_{b(j),6}w_j a \tag{15}$$

In Eq. (15), w_j is vehicle j's weight (in 1000 lb) and $b(j)$ is vehicle j's body style. The body-style specific β coefficients are given in Ref. [9]. Unit costs are also a function of acceleration/fuel economy performance, given by a function $c_j(e, a, t)$ as given in the following equation:

$$c_j(e,a,t) = \beta_{b(j),1} + \beta_{b(j),2}\exp\{-a\} + \beta_{b(j),3}t + \beta_{b(j),4}w_j + \beta_{b(j),5}w_j a \tag{16}$$

Again, coefficients are drawn from Ref. [9]. These models were estimated using detailed engineering simulations from AVL Cruise in conjunction with confidential technology production cost data provided to NHTSA by automakers in advance of the 2012–2016 fuel economy rulemaking [90].

We consider a version of the Mixed Logit model estimated in Ref. [9]. Vehicle-specific utility functions u_j are given by

$$u_j(e,a,p) = -\alpha p + \frac{\beta_e}{e} + \frac{\beta_a}{a} + \beta_f f_j + \xi_j - \vartheta \tag{17}$$

where f_j is vehicle j's footprint (in 1000 in^2), ξ_j is a product-specific utility determined in estimation (see, e.g., Ref. [16]), $\vartheta = -8.0903$ represents a constant utility of the outside good, and the coefficients $\alpha, \beta_e, \beta_a, \beta_f$ are random. The price coefficient is $\alpha = |0.4591 + 0.1756/\imath - 0.0377V|$, where $\imath$ denotes household income and $V \sim N(0,1)$, where $N(\mu, \sigma)$ denotes a normal distribution with mean μ and variance σ^2. We assume income is greater than \$10,000 and is drawn from an empirical frequency distribution based on data from 2006's Current Population Survey [91]. The attribute coefficients have the following distributions: $\beta_e \sim N(-36.77, 0.022)$, $\beta_a \sim N(11.262, 0.0321)$, and $\beta_f \sim N(2.4541, 0.0393)$. See Ref. [9] for more information. The results for all methods discussed below are obtained using a single simulator derived from 1000 samples of α, β_e, β_a, β_f.

4.2 Computational Statistics

Figure 3 illustrate statistics concerning 1000 attempted solves of the vehicle design-then-

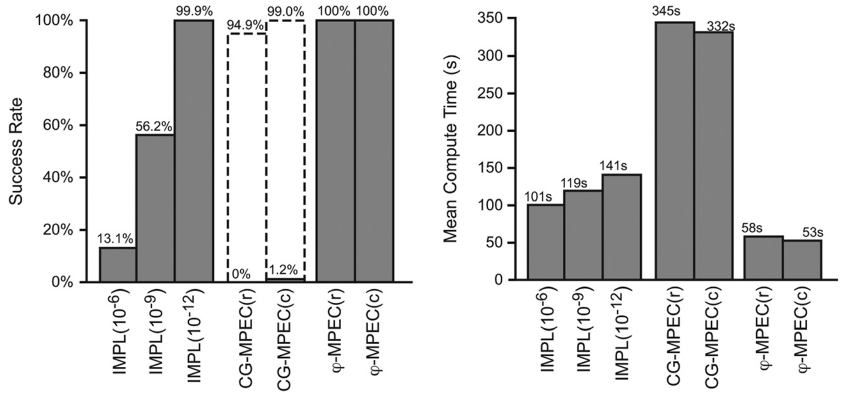

Fig. 3 Computational comparison of performance of the implicit programming and MPEC methods on 1000 trials started at different initial conditions for a single set of 1000 samples. Implicit programming was implemented computing equilibrium to tolerances of 10^{-6}, 10^{-9}, and 10^{-12}, abbreviated IMPL(10^{-6}), IMPL(10^{-9}), and IMPL(10^{-12}), respectively. Optimal design problems solved to tolerances of 10^{-6}. CG-MPEC was started at both random and the "smarter" initial conditions, abbreviated CG-MPEC(r) and CG-MPEC(c), respectively; similarly with φ-MPEC. (Left) Success rate captures both SNOPT successes and computation of an equilibrium, that are the same for all methods except the CG-MPEC approach. Dashed boxes represent the SNOPT success rate for CG-MPEC including spurious solutions. (Right) Mean compute times include only successful runs, with the exception of CG-MPEC(r) for which there were no successful runs.

Table 1 Computational comparison of performance of the implicit and MPEC methods from 1000 trials starting at two types of random initial conditions. See Sec. 4 for more details.

		IMPL(10^{-6})	IMPL(10^{-10})	IMPL(10^{-12})	CG-MPEC	φ-MPEC
Problem size statistics						
Number of variables	(#)	116	116	116	588	588
Number of constraints	(#)	58	58	58	530	530
Number of nonzero derivatives	(#)	261	261	261	265,422	265,422
Random initial conditions						
SNOPT successes	(%)	13.1	56.2	99.9	94.4	100.0
True successes	(%)	13.1	56.2	99.9	0.0	100.0
Computed a spurious KKT Point	(%)	0.0	0.0	99.9	94.4	0.0
SNOPT failures	(%)	86.9	43.8	0.1	5.6	0.0
Mean (median) time, all runs	(s)	103.3 (92.4)	122.7 (104.2)	145.4 (126.0)	344.9 (162.8)	58.1 (53.0)
Mean (median) time, successes	(s)	100.7 (82.9)	119.3 (103.2)	140.8 (125.9)	– (–)	58.1 (53.0)
"Smarter" initial conditions						
SNOPT successes	(%)	—	—	—	98.9	100.0
True successes	(%)	—	—	—	1.2	100.0
Computed a trivial KKT point	(%)	—	—	—	97.7	0.0
SNOPT failures	(%)	—	—	—	1.1	0.0
Mean (median) time, all runs	(s)	—	—	—	189.2 (104.3)	52.5 (47.4)
Mean (median) time, successes	(s)	—	—	—	331.5 (139.9)	52.5 (47.4)

pricing example presented above with different initial conditions; see also Table 1 in Appendix D. Each method discussed in Sec. 3 is employed. Implicit programming (Sec. 3.2) is abbreviated as "IMPL," an MPEC formulation with the combined gradient (Sec. 3.3) as "CG-MPEC," and an MPEC formulation with the φ map (Sec. 3.5) as "φ-MPEC." Spurious KKT points were identified with explicit verification of the negative definiteness of firms' profit Hessians with respect to prices. Two types of initial conditions were employed in our trials. In random initial conditions, we drew designs and prices as independent, uniformly distributed draws. Design variables were drawn from within their body-style specific bounds, and prices were drawn between \$0 and \$100,000. In "smarter" initial conditions we also drew designs uniformly randomly from their bounds, but drew vehicle j's price from \$0 and c_j where c_j is the cost for vehicle j corresponding to the randomly drawn design variables. In this way, the "smarter" initial conditions have initial prices that are no greater than costs, and thus, strictly less than equilibrium prices for the initial designs [41,50]. Because the spurious solutions identified in Sec. 3.4 concern prices that are too high, setting initial prices less than costs guarantees that any spurious solutions obtained are not a consequence of poor initial conditions alone.

The reliability with which design-then-pricing problems can be solved depends greatly on the method chosen; see Fig. 3, left. With a tolerance of 10^{-6} on computations of equilibrium prices, implicit programming is unreliable solving the problem in just over 13% of trials; achieving a nearly 100% success rate (failing in only one trial out of 1000) requires computing equilibrium prices to a tolerance of 10^{-12}, with correspondingly increased compute times. The feasibility of using such tight tolerances on equilibrium pricing computations will depend greatly on problem size, scaling, and the associated limit on numerical accuracy that can be expected in computing φ as a function of prices. CG-MPEC performs very poorly, never computing a nonspurious solution with random initial conditions and computing a true solution in only 1% of trials with the "smarter" initial conditions. It is important to emphasize that when CG-MPEC computes a spurious solution, SNOPT still terminates successfully. That is, if we relied on solver termination conditions alone, we would be misled into believing we had solved the design-then-pricing problem in 94% and 99% of runs with the random and "smarter" initial conditions (respectively) when, in fact, we had computed a spurious solution. φ-MPEC succeeds in computing optimal designs with prices in equilibrium in all trials for both types of initial conditions.

Mean run times also vary significantly across methods. Fig. 3, right, plots mean compute times for all methods; Table 1 includes range statistics. φ-MPEC is the fastest method, comparing compute times either across all runs (i.e., including failures) or comparing across successful runs only. When starting from a random initial point φ-MPEC takes under 1 min while the implicit (10^{-12}) method takes over 2 min, on average; CG-MPEC takes 5 3/4 min on average, but again only computes spurious solutions. Note that the speed-up observed for φ-MPEC occurs despite the fact that it literally solves a much larger problem than the implicit programming approach, with 5 times the variables, 9 times the constraints, and 1000 times the number of nonzero objective and constraint derivatives (Table 1). In reality the problem sizes are comparable; the implicit method "hides" the large problem size from the solver by computing equilibrium prices and their derivatives within more expensive function evaluations. From the "smarter" initial points φ-MPEC takes 47 s on average, while CG-MPEC takes 330 s (when successful).

4.3 Solution Details. In all, we computed roughly 307 distinct solutions to Eq. (6). Without product-specific vehicle parameters left fixed in the design problem, permuting the positions of any two-vehicle designs (from within sets of vehicles with a specific body style) in the list of vehicles leads to an equivalent solution. Thus, permutations in product order might create a large number of equivalent solutions up to permutation. Because we have product-specific parameters, however, we see a large number of distinct solutions.

Figure 4 shows the likelihood of computing profits within a given percentage of the apparent globally maximal profits using φ-MPEC. Each step up denotes a local optimum. Over 30% of all successful runs converged to the global maximum, while roughly 80% converge to a solution that has a value of profits within 90% of the global maximum. All local optima in this problem are

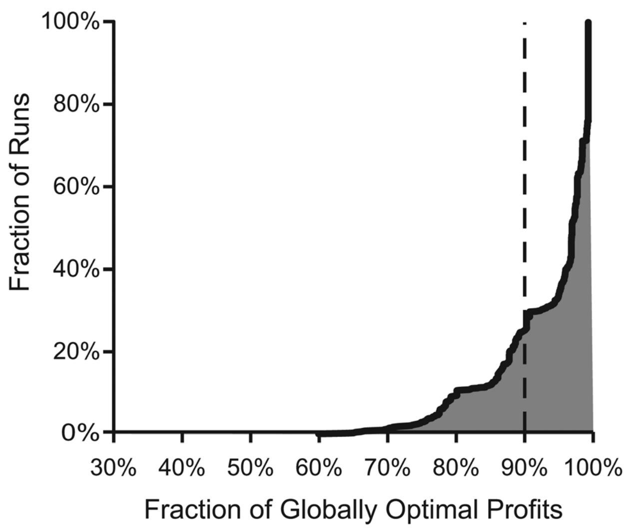

Fig. 4 Cumulative distribution function for computing profits in the design-then-pricing model within a given percentage of the apparent globally maximal profits over 1000 trials. The gray shading represents the area under this curve for comparison with Fig. 7.

characterized by 17 of firm 1's 29 vehicles with particular body styles (minicompact, subcompact, and compact cars) with technology content at various combinations of the upper and lower bounds: not adopting any technology ($t_j = 0$) or complete adoption of technology ($t_j = 42$). In principle, there might be $2^{17} = 131{,}072$ such local solutions (compared to the 307 we believe we compute). Vehicles offered by firm 1 with body styles other than these three had unique designs over all local solutions, solutions that were again bound-constrained; some of these vehicles had no technology adopted, and some had all technology adopted. Whenever a local solution has technology content at its lower bound (0) on some minicompact, subcompact, or compact car, fuel economy for that vehicle is high; when a local solution has technology content at its upper bound (42), fuel economy is low. This represents a "configuration" trade-off: if technology is adopted, the cost and demand models suggest that this investment is best utilized to maximize acceleration performance, rather than fuel economy. This tradeoff is generally reflective of trends that appear in real vehicle markets, but is certainly a feature determined by the simplified model we adopt to enable an example with a large-scale market.

4.4 Comparison to Design-and-Pricing. Design-and-pricing outcomes depend on what the firm assumes regarding fixed competitors' prices. We investigate four "scenarios" regarding competitors' product prices labeled as follows: "∞": competitors price their products at ∞ or, equivalently, the firm ignores competitors; "$\mathbf{p}_*(\mathbf{X}_0)$": competitors price their products in equilibrium given designs as in the data; "$\mathbf{p}_0$": competitors price their products as given in the data; "$\mathbf{0}$": competitors give away their products for free. These scenarios reflect the extremes of the various assumptions that could be made regarding competitor pricing behavior without anticipating reactions to design changes. For each scenario, we compute optimal designs and prices for firm f's products holding competitors' designs and prices fixed; i.e., we solve Eq. (5). Any solution defines an "expected" profits from optimal design "π_f^e", ignoring potential impacts of pricing competition at a later date. Note that we do not use the term "expected" here to mean an average over possible outcomes weighted by their likelihood, as is meant in probability theory. We then also compute equilibrium prices for all vehicles using the computed optimal designs for firm f's vehicles, obtaining an associated profits "π_f^*" in equilibrium after pricing competition occurs.

In comparing design-and-pricing to design-then-pricing, there are three important perspectives to keep in mind: First, Sec. 3 suggests that even design-and-pricing problems may not be solved reliably without regularization. Second, "expected" profits from design-and-pricing may change when pricing competition actually occurs. Third, profits in equilibrium from optimal design decisions made *without* anticipating pricing competition may, or may not, differ from profits obtainable by anticipating pricing competition even if expected profits differ. Our results are focused on these three perspectives.

Figure 5 plots success rates for 1000 attempted solves of Eq. (5) with random initial conditions and a tolerance of 10^{-6} for both optimality and feasibility. As with design-then-pricing, SNOPT successfully solves Eq. (5) in every run *but does not always compute prices that are profit-optimal*: when the firm ignores its competitors (scenario "∞"), 82.7% of these solves resulted in designs and prices that were *not* profit-optimal because some price was too large; when the firm assumes that its competitors give away their products (scenario "$\mathbf{0}$"), 21.7% of these solves resulted in designs and prices that were *not* profit-optimal because some price was too large. Regularized computations using Eq. (14) always succeeded and always computed prices that were locally profit-optimal.

Figure 6 plots the mean *anticipated* gain (or loss) in profits over the 1000 attempted solutions with random initial conditions, taking results from the regularized solves; the successful solves

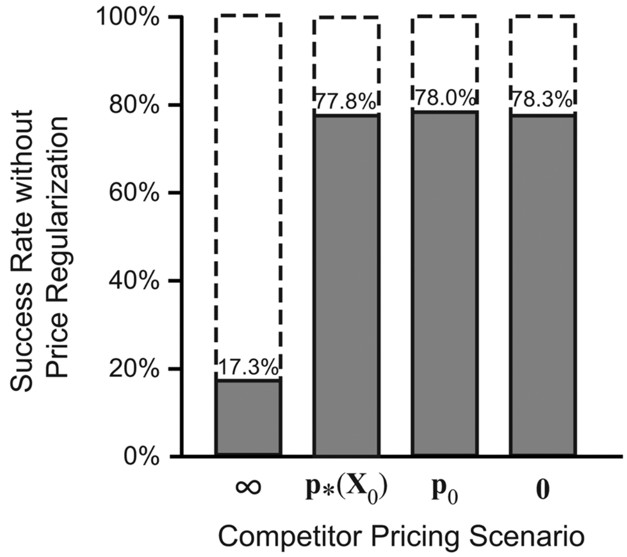

Fig. 5 Success rates for the unregularized design-and-pricing problem out of 1000 trials with random initial conditions. Scenarios are labeled as described in the text. As in Fig. 3, dashed boxes represent SNOPT success rates (all 100%) and gray boxes represent successful computations of a nonspurious solution.

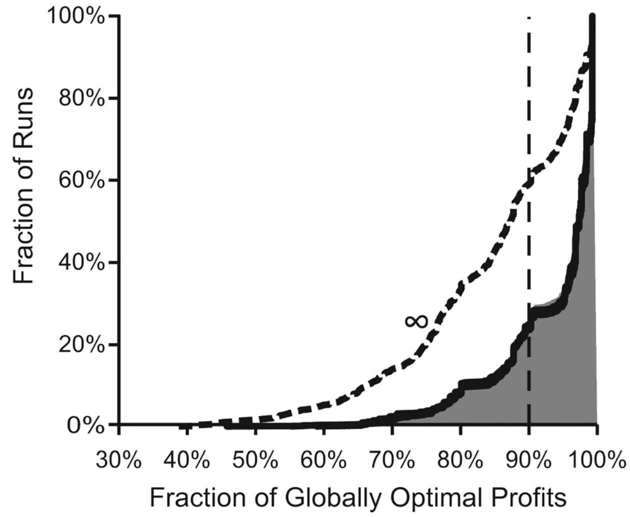

Fig. 7 Cumulative distribution function for computing profits in the design-and-pricing model within a given percentage of the apparent globally maximal profits over 1000 trials. The dark, solid black curve represents overlapped CDFs of all design-and-pricing scenarios except for "∞"; the CDF for this scenario is represented by a dashed curve. The gray area is carried over from Fig. 4.

without regularization have similar statistics. These results illustrate that the "accuracy" of expected profits when neglecting eventual pricing competition depends sensitively on the assumption held about competitors' current prices. If competitors' products are ignored (scenario "∞"), almost all profits expected from design-and-pricing decisions are lost. Failure to anticipate pricing reactions would appear, after-the-fact, to be a serious mistake in this extreme case. If competitors' are assumed to give away their products (scenario "**0**"), four times the expected profits may be earned. Here the profits obtained by a failure to anticipate equilibrium pricing could be interpreted after-the-fact as a "windfall." However, if competitors' are assumed to price according to market data (scenario "$\mathbf{p}_0$") or in equilibrium (scenario "$\mathbf{p}_*(\mathbf{X}_0)$"), expected profits from design-and-pricing are likely to appear accurate after pricing competition occurs.

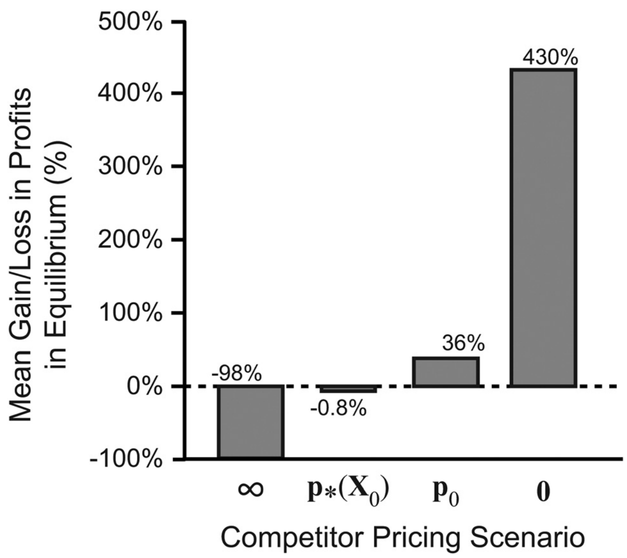

Fig. 6 Mean perceived gain or loss in profits when choosing designs and prices without anticipating pricing competition that ultimately occurs, with regularized computations. Scenarios are labeled as described in the text.

Figure 7 addresses the final question: is anticipating pricing competition important to finding the optimal designs? Unless competitors' products are ignored (scenario "∞"), 1000 attempted solves of the design-and-pricing problem result in an almost identical distribution of local profit maximizers to those obtained using the design-then-pricing model. Considering profits alone there appears to be little benefit to adopting the design-then-pricing paradigm in this vehicle design example. This is a feature clearly driven by aspects of the vehicle design and consumer decision models: the design solutions are primarily determined by the equality constraints and bounds. In this case, anticipating competition has a significant influence on vehicle prices and expected profits when the firm ignores competition in prices, but no influence on which designs are ultimately profit-optimal.

5 Discussion

This section discusses two important aspects of the article. We first examine our assumption of Bertrand–Nash equilibrium as a representation of price competition. Then, we discuss what our comparison of design-and-pricing with design-then-pricing says about the value of anticipating price competition during design.

5.1 On Bertrand–Nash Equilibrium in Design-then-Pricing. Some researchers may take issue with the design-then-pricing model applied here because it is based on classical game-theoretic notions of "Nash" equilibrium. For example, Wang et al. [40] raise three concerns about Nash equilibrium concepts: Because design-then-pricing (as presented here) contains a Nash equilibrium model of pricing, it would seem to lead to a "design solution that is only guaranteed to be optimal when market players take actions simultaneously" and "ignore[s] how the equilibrium is [attained]." [40, p. 2]. These are certainly relevant critiques of Nash equilibrium pricing and the associated design-then-pricing paradigm as a high-fidelity model of real markets. Depending on the time it takes to "reach" equilibrium in a price competition game, and the length of time the game is subsequently played in or "near" equilibrium, this full trajectory of prices could be relevant to optimal profits. Worse, there is little to say that repeated games would converge to an equilibria at all. Equilibria need not be (locally) "stable" in the sense of attractive under best-response dynamics [13]. Dynamic pricing models [92,93] might better capture the repeated nature of competition in prices in markets where designs stay fixed for a long time, and a more recent topic of research in operations research and management science.

The perspective represented by Wang et al. [40] expresses a very literal interpretation of equilibrium. Generally speaking, equilibria are sets of decisions that are fixed under certain behavioral rules guiding those decisions. Bertrand–Nash equilibria in prices represent a state of simultaneous optimality with respect to myopic pricing decisions. Wang et al. formulate an agent-based model with learning, simulating that model until convergence to an equilibrium representing the consequences of design; the "no-regret" equilibrium thus, obtained is also, when properly formulated, a fixed-point under the behavioral rules chosen for the agent-based simulation. Designing firms could indeed choose from a myriad of equilibrium concepts to represent the outcome of competitive interactions that determine the profitability of their design decisions. The practical question firms face is which equilibrium concept, if any, is *both* a reasonable representation of reality *and* yields a tractable paradigm for analysis of real problems.

Bertrand–Nash equilibrium might be a "good" choice for design-then-pricing models for two reasons.

First, equilibrium pricing as considered here is relatively parsimonious in a space where it is difficult to forecast market outcomes precisely. More complex models and simulations of pricing competition could be very sensitive to assumptions on behavioral rules and parameters governing them. Furthermore, introducing more complex rules and parameters in a model does not always yield better predictions or judgements [94], in part because more complex models can easily capture more "noise" than "signal" [95]. Moreover, as suggested by our comparison to design-and-pricing, if design constraints determine optimal designs then anticipating pricing competition may be irrelevant to the profits ultimately obtained (though not necessarily to after-the-fact perceptions of the worth or accuracy of the decisions thus made).

Second, as we emphasize in this article, Bertrand–Nash equilibria in prices can be represented with systems of equations or complementarity problems. The resulting model of competition, when treated correctly, appears to be particularly tractable for use in making strategic design decisions in large problems. There will no doubt be a trade-off between the detail required to "accurately" resolve pricing competition and the scale at which strategic decisions whose profitability depends on pricing competition can be considered. The ultimate value of decisions made with the design-then-pricing model considered here certainly lies in the degree to which Bertrand–Nash equilibrium prices reflect fundamental aspects of pricing competition with fixed designs. By providing techniques that make this design-then-pricing model tractable for realistic examples we make an important technical contribution to addressing this important open question.

5.2 Should Firms Anticipate Pricing Competition? Does anticipating pricing competition influence the optimality of chosen designs? The result in Sec. 4.4 speak to this question and present a more subtle relationship between design-and-pricing and design-then-pricing than has been discussed thus far in the literature. Within a firm, there may be connotations associated with perceived after-the-fact gains or losses from design-and-pricing strategy regardless of whether the decisions are, in fact, better or worse than if price competition were anticipated when making design decisions. For example, suppose an engineer proposed a set of designs along with profit projections for those designs that turned out to be low because of unanticipated competition. When the "losses" appear, the engineer's career might be damaged or at risk. However, this engineer may not have been capable of making a better decision by anticipating competition; all that was ultimately "wrong" could be the profit forecast, and associated pricing decisions. This possibility is suggested by our result that anticipating equilibrium prices can appreciably change the profit forecast and pricing, but not optimal designs. One of the three case studies in Ref. [39] has the same result. On the other hand, the engineer could make the same decisions assuming *overly* competitive behavior (e.g., scenario "**0**") and appear to have made a miraculous decision when, again, they simply had made the wrong profit forecast. These interpretations are certainly specific to the model and problem explored here, but suggest an interesting relationship between the *perceived* value of specific decisions and the *objective* value of the same decisions when made in a dynamic, competitive environment. We believe this is worth exploring further by market-systems researchers.

What, specifically, about our example implies that there is no value to anticipating pricing competition? Existing results certainly suggest that there is a positive value to anticipating pricing competition when it occurs. Shiau and Michalek [39], for example, discuss three case studies: design of a painkiller (from Ref. [31]), a weighing scale (from Refs. [96,97]), and an angle grinder (from Ref. [42]). In the first and second case studies, ignoring price competition when choosing designs leads to 1.5% and 1.4% decreases in profitability, respectively. While small, these decreases in profitability suggest that anticipating price competition can be valuable. However, their third case is similar to our results: failing to anticipate pricing competition for the angle grinder results in overestimating profitability from optimal design decisions by 21%, but entails no "real" loss in profits after pricing competition occurs. Our hypothesis is that the presence of design constraints can reduce the potential value of anticipating pricing competition. In our case study, most local solutions are "fully constrained" in the sense that acceleration, fuel economy, and technology content for each vehicle are defined by a combination of bounds and the equality constraint in Eq. (15). Thus, optimal designs will be insensitive to changes in the market model that do not enter these constraints, including both the form of demand model and competitive interactions; prices and anticipated profitability may differ. A formal proof can be based on the KKT conditions. The angle grinder case examined by Shiau and Michalek, by including only discrete decisions about which product features to adopt, is similar: small changes in the market model are unlikely to change which features to include, though they may change optimal prices and/or the profits realized. Thus, the value of anticipating competition in prices when making design decisions depends on how deeply design constraints influence optimality.

6 Conclusions

This article describes a theoretical and computational study of models of optimal design-then-pricing; that is, design with price competition modeled using Bertrand–Nash equilibrium prices. We have drawn a distinction between an implicit programming approach that treats prices as intermediate quantities computed in equilibrium, and "MPEC" formulations that treat prices as variables constrained to satisfy a first-order condition for equilibrium. Implicit programming can be effective, but may require prohibitively accurate computations of equilibrium prices to reliably compute optimal designs. An MPEC formulation that uses the literal first-order condition for equilibrium prices is shown to possess spurious KKT points that are not solutions to the problem posed and can be computed by a state-of-the-art solver. Using a reformulated representation of equilibrium prices in an MPEC provably eliminates these spurious KKT points and leads to the most efficient and reliable computations on a real-scale vehicle portfolio design-then-pricing example. A comparison to design problems that ignore future pricing competition shows that the value of including pricing competition depends on (i) the importance of accuracy in profit forecasts, as opposed to realized profits regardless of forecast, and (ii) the importance of constraints in determining optimal designs.

Acknowledgment

This research was supported by Iowa State University and an EAGER award from NSF's EDI Program (#1135329). An earlier version of this article appears in the Proceedings of ASME IDETC 2012 (Paper No. DETC2012-71201). The authors wish to thank reviewers at ASME's IDETC conference and reviewers for JMD for their invaluable assistance in improving the research discussed here. Jeremy Michalek for suggesting the "smarter" initial conditions, and Michael Saunders provided insightful discussions regarding SNOPT. The findings presented in this article are the authors' own and do not represent the views of the National Academy of Engineering.

Nomenclature

$\mathbf{c}_f$ = vector of unit costs for firm f's products
c_f^F = fixed costs for firm f
c_j = unit cost of product j
D_j^p = differentiation with respect to price of product J
f = index; firm f
F = number of firms
$\mathbf{g}_j$ = equality constraints for product j
$\mathbf{h}_j$ = inequality constraints for product j
I = number of "individuals" in simulated Mixed Logit model
$\mathbf{l}_j$ = lower bound on design variables for product j
j = index; product j
J = number of products
J_f = indices of products offered by firm f
$\mathbf{p}$ = vector of all product prices
$\mathbf{p}_f$ = vector of prices for firm f's products
$\mathbf{P}_f$ = vector of choice probabilities for firm f's products
p_j = price of product j
P_j = Mixed Logit choice probability for product j
P_j^L = logit choice probability for product j
$\mathbf{p}^*$ = local equilibrium prices
$\mathcal{P}$ = set of possible product prices; $\mathcal{P} = [0, \infty)$
$\mathcal{T}$ = set of possible individual characteristics
u_j = systematic component of utility for product j
$\mathbf{u}_j$ = upper bound on design variables for product j
U_j = random utility for product j
U_0 = utility of an outside good
$\mathbf{x}_j$ = vector of design variables for product j
$\mathbf{X}_f$ = "matrix" of design vectors for firm f's products
$\mathbf{Y}$ = "matrix" of all product attribute vectors
$\mathbf{y}_j$ = vector of attributes of product j
$\mathbf{y}_j$ = transform characteristics of product j into attributes
$\mathcal{Y}$ = set of possible product attributes
$(\tilde{\nabla}^p\hat{\pi})$ = combined gradient of firm profits
$\mathcal{E}_j$ = random component of utility for product j
$\mathcal{E}_0$ = random component of utility for outside good
θ = vector of individual's characteristics
θ_i = realization of θ for the ith sample from μ
ϑ = utility of the outside good
μ = probability distribution of θ over $\mathcal{T}$
$\hat{\pi}_f$ = firm f's expected profits
$\hat{\pi}_f$ = firm f's profits at a given set of designs and prices
$\hat{\pi}_f^*$ = firm f's profits at local equilibrium prices
φ = map in the reformulated MCP for equilibrium prices; see Appendix A

Symbols used in our example

a_j = 0-60 time of product j (s)
$\mathbf{a}$ = vector of acceleration performance for all products
$b(j)$ = product j's body style
$\mathbf{e}$ = vector of fuel efficiency for all products
e_j = fuel economy of product j (mpg)
f_j = footprint of product j
$\mathbf{t}$ = vector of technology content for all products
t_j = technology content of product j
w_j = weight of product j
α = price coefficient in the utility function
β = body-style specific utility coefficients
ξ = product-specific utility

Appendix A: Computing Equilibrium Prices

Suppose the choice probabilities are continuously differentiable in prices; see Refs. [50,41] for conditions on the model to ensure this holds. Then, at any local equilibrium prices, each firm's prices satisfy the stationarity condition $(D_k^p\hat{\pi}_f)(\mathbf{Y},\mathbf{c},\mathbf{p}) = 0$ for all $k \in \mathcal{J}_f$, where D_k^p denotes differentiation with respect to the kth product's price. Combining the stationarity condition for each firm we obtain the *simultaneous stationarity condition* (SSC), a first-order necessary condition for local equilibrium prices:

Theorem 1 (SSC [41]). *Suppose* $\mathbf{P} : (0,\iota_*)^J \to [0,1]^J$ *is continuously differentiable. Let* $(\tilde{\nabla}^p\hat{\pi})(\mathbf{Y},\mathbf{c},\mathbf{p})$ *denote the "combined gradient" with components* $((\tilde{\nabla}^p\hat{\pi})(\mathbf{Y},\mathbf{c},\mathbf{p}))_j = (D_j^p\hat{\pi}_{f(j)})(\mathbf{Y},\mathbf{c},\mathbf{p})$ *where f(j) denotes the (unique) index of the firm offering product j. If* $\mathbf{p} \in (0,\iota_*)^J$ *is a local equilibrium, then*

$$(\tilde{\nabla}^p\hat{\pi})(\mathbf{Y},\mathbf{c},\mathbf{p}) = (\tilde{D}^p\mathbf{P})(\mathbf{Y},\mathbf{p})^\top(\mathbf{p}-\mathbf{c}) + \mathbf{P}(\mathbf{Y},\mathbf{p}) = \mathbf{0}. \quad \text{(B1)}$$

where $(\tilde{D}^p\mathbf{P})(\mathbf{Y},\mathbf{p})$ *is the* $J \times J$ *"intrafirm" Jacobian matrix of price derivatives of the choice probabilities defined, component-wise, by* $((\tilde{D}^p\mathbf{P})(\mathbf{Y},\mathbf{p}))_{j,k} = (D_kP_j)(\mathbf{p})$ *if products j and k are offered by the same firm and* $((\tilde{D}^p\mathbf{P})(\mathbf{Y},\mathbf{p}))_{j,k} = 0$ *otherwise.*

This follows from the standard necessary optimality conditions for unconstrained optimization; see Ref. [41]. Prices $\mathbf{p}$ satisfying Eq. (B1) are called *simultaneously stationary*, and the set of all such prices is denoted $\mathcal{S}(\mathbf{Y},\mathbf{c})$. Note that $\mathcal{S}(\mathbf{Y},\mathbf{c}) \supset \mathcal{E}^\ell(\mathbf{Y},\mathbf{c}) \supset \mathcal{E}(\mathbf{Y},\mathbf{c})$. Morrow and Skerlos [41] provide an example that has simultaneously stationary prices that are *not* a local equilibrium; i.e., $\mathcal{S}(\mathbf{Y},\mathbf{c}) \neq \mathcal{E}^\ell(\mathbf{Y},\mathbf{c})$.

When $\iota_* = \infty$ (and there are no formal problems evaluating utilities for negative prices) Theorem 1 above gives a first-order condition for equilibrium prices in terms of a nonlinear system. In principle, this system can be solved with globally convergent Newton methods (e.g., Refs. [98–100]) to compute candidates for equilibrium. However, Morrow and Skerlos [41] demonstrate that Newton methods applied to solve Eq. (B1) provide unreliable computations of equilibrium prices. The reason is that the combined gradient $(\tilde{\nabla}^p\hat{\pi})(\mathbf{Y},\mathbf{c},\cdot)$ is not *coercive* as a function of prices. A coercive function has a norm that increases without bound as the norm of its argument increases without bound [101, Chap. 6]. While coercivity is rarely discussed, it underlies most of the current convergence theory for optimization problems [98], nonlinear systems [98], and complementarity problems [102]. Specifically, without coercivity, there is no guarantee that descent-type global convergence strategies will generate bounded sequences and thus, may not contain convergent subsequences. For the equilibrium pricing problem, application of globally convergent methods may appear successful but may, in fact, compute points at which some prices are essentially infinite even though this is not profit-optimal.

Morrow and Skerlos resolve this problem by showing that two fixed-point equations equivalent to Eq. (B1) *are* coercive and thus, provide well-posed formulations for the application of Newton-type methods. Here we extend and apply one of these equations:

Theorem 2. *If* $\mathbf{p} \in (0,\iota_*)^J$ *is a local equilibrium, then* $\varphi(\mathbf{Y},\mathbf{c},\mathbf{p}) = \mathbf{0}$ *where*

$$\varphi(\mathbf{Y},\mathbf{c},\mathbf{p}) = \mathbf{p} - \mathbf{c} - \zeta(\mathbf{Y},\mathbf{c},\mathbf{p}) \quad \text{(B2)}$$

$$\zeta(\mathbf{Y},\mathbf{c},\mathbf{p}) = \Lambda(\mathbf{Y},\mathbf{p})^{-1}(\widetilde{\Gamma}(\mathbf{Y},\mathbf{p})^\top(\mathbf{p}-\mathbf{c}) - \mathbf{P}(\mathbf{Y},\mathbf{p})) \tag{B3}$$

The $J \times J$ matrices Λ and $\widetilde{\Gamma}$ are derived from the derivatives of the Mixed Logit choice probabilities; specifically,

$$(\tilde{D}^p\mathbf{P})(\mathbf{Y},\mathbf{p}) = \Lambda(\mathbf{Y},\mathbf{p}) - \widetilde{\Gamma}(\mathbf{Y},\mathbf{p})$$

See Refs. [41,50] for explicit definitions of these matrices. See Refs. [41,50,51] for a proof of this result. When $\imath_* = \infty$ and there are no problems evaluating utilities for negative prices, Newton methods can be applied to solve $\varphi(\mathbf{Y},\mathbf{c},\mathbf{p}) = \mathbf{0}$.

Theorem 2 really only characterizes equilibrium prices that are all less than the population's limit on purchasing power, $\imath_*$. When purchasing power is finite ($\imath_* < \infty$), it may be profit-optimal to "price some products out of the market" by setting $p_j = \imath_*$; for examples, see Refs. [51,68]. In general, we might hope to solve the KKT conditions for optimal pricing restricted to $[0, \imath_*]$, which can be written as in Eq. (8). Unfortunately, Eq. (8) is poorly posed because $(D^p_k\hat{\pi}_{f(k)})(\mathbf{p}) \to 0$ as $p_k \uparrow \imath_*$ (Lemma 1). This is discussed further in Ref. [51].

Equation (B2) remains useful when $\imath_* < \infty$, at for the simulated Mixed Logit models that would be applied in practice.

Theorem 3 [51]. *Suppose the Mixed Logit model is a simulator with $\imath_s = \imath(\theta_s)$, $\imath_S = \max_s\{\imath_s\} = \imath_*$. If $\mathbf{p}_* \in [0, \imath_*]^J$ is a local equilibrium, then $\boldsymbol{p}_*$ solves Eq. (9). Conversely, suppose $\boldsymbol{p}_*$ is a strictly complementary solution to Eq. (9). If each firm's profit Hessians are negative definite with respect to changes in prices less than $\imath_*$ only, then $\boldsymbol{p}_*$ is a local equilibrium. Moreover, $\zeta(\mathbf{Y},\mathbf{c},\cdot) : [0,\imath_*]^J \to \mathbb{R}^J$ can be extended to a continuously differentiable function $\mathbf{z}(\mathbf{Y},\mathbf{c},\cdot) : \mathbb{R}^J \to \mathbb{R}^J$ such that $\mathbf{p} - \mathbf{c} - \mathbf{z}(\mathbf{Y},\mathbf{c},\mathbf{p}) = \mathbf{0}$ if, and only if, proj{**p**} solves Eq. (9), where "proj" denotes the Euclidean projection onto $[0, \imath_*]^J$.*

These ideas extend rather readily to explicit bounds on prices; that is, bounds on prices that are less than purchasing power ($\imath_*$). Shiau and Michalek [39] consider such a model, claiming such bounds may be "imposed by manufacturer, retailer, consumer, or government policies, and they may also be used to indicate model domain bounds." Anecdotally, we do know that such bounds are included in some price optimization software currently used in industry. We can write firm f's pricing-stage profit maximization problem with explicit price bounds as

$$\begin{array}{ll} \max & \hat{\pi}_f(\mathbf{Y},\mathbf{c},\mathbf{p}) \\ \text{w.r.t.} & \ell_j \le p_j \le \nu_j \quad \text{for all } j \in \mathcal{J}_f \end{array} \tag{B4}$$

for some ℓ_j, ν_j such that $0 \le \ell_j < \nu_j \le \imath_*$. The combined KKT conditions can be written as the MCP

$$\ell \le \mathbf{p} \le \nu \perp -(\tilde{\nabla}\hat{\pi})(\mathbf{Y},\mathbf{c},\mathbf{p})$$

Formally, this MCP no longer has the spurious solution problem so long as $\nu_j < \imath_*$. However, when solving this equation *numerically*, difficulties may arise when the profit derivatives become *small* as well as strictly zero, as seen in Fig. 2. This fact makes it more difficult to know whether bounds chosen *a priori* will generate spurious solutions.

Because the sign of $-(D^p_k\hat{\pi}_{f(k)})(\mathbf{Y},\mathbf{c},\mathbf{p})$ is the sign of $\varphi_k(\mathbf{Y},\mathbf{c},\mathbf{p})$ so long as $p_k < \imath_*$ [41,50], and MCPs are invariant over changes in the map on the right-hand side of the "$\perp$" symbol that do not change the map's sign, we can equivalently solve the MCP

$$\ell \le \mathbf{p} \le \nu \perp \varphi(\mathbf{Y},\mathbf{c},\mathbf{p})$$

We believe MCP is less likely to have numerical difficulties for "large" prices, though either map could be employed.

Finally, we address second-order conditions. If no prices are bound constrained, then the SOSC is simply to verify the negative definiteness of Hessian of profits with respect to prices. Because optimal pricing problems are simply bound constrained, any active constraints are simply bounds. Active constraint gradients are then sets of standard basis vectors (with either positive or negative sign) corresponding to those prices that are bound constrained. The subspace tangent to those constraint gradients is thus spanned by the standard basis vectors that are *not* bound constrained. Verifying the SOSC–negative definiteness of the Hessian matrix over the subspace tangent to the active constraint gradients–thus reduces to verifying negative definiteness of the submatrix of the Hessian with those rows and columns corresponding to products with prices that are not bound constrained.

Appendix B: MCP Form of the KKT Conditions

This appendix proves the following result:

Lemma 2. *Consider a generic minimization problem in positive-null form [52]*

$$\begin{array}{ll} \min & f(\mathbf{x}) \\ \text{w.r.t.} & \mathbf{l} \le \mathbf{x} \le \mathbf{u} \\ \text{s.t.} & \mathbf{c}^E(\mathbf{x}) = \mathbf{0},\ \mathbf{c}^I(\mathbf{x}) \ge \mathbf{0} \end{array}$$

The KKT conditions for this problem can be written as the MCP

$$\begin{array}{l} \mathbf{l} \le \mathbf{x} \le \mathbf{u} \perp (\nabla f)(\mathbf{x}) - (D\mathbf{c}^E)(\mathbf{x})^\top \mu^E - (D\mathbf{c}^I)(\mathbf{x})^\top \mu^I \\ \infty < \mu^E < \infty \perp \mathbf{c}^E(\mathbf{x}) \\ \mathbf{0} \le \mu^I \perp \mathbf{c}^I(\mathbf{x}) \ge \mathbf{0} \end{array}$$

The "$\perp$" symbol means the following: $\mathbf{l} \le \mathbf{x} \le \mathbf{u} \perp \mathbf{F}(\mathbf{x})$ if $F_i(\mathbf{x}) > 0$ implies $x_i = l_i$, $F_i(\mathbf{x}) < 0$ implies $x_i = u_i$, and $F_i(\mathbf{x}) = 0$ implies $x_i \in [l_i, u_i]$ for all i. Note that this definition implies that $\infty < \mu^E < \infty \perp \mathbf{c}^E(\mathbf{x})$ is the nonlinear system $\mathbf{c}^E(\mathbf{x}) = \mathbf{0}$. See Refs. [75,76,102–104] for definitions and extensive discussion related to MCPs.

Proof. The KKT conditions for this problem are as follows [52]: $\mathbf{x} \in [\mathbf{l},\mathbf{u}]$, $\mathbf{c}^E(\mathbf{x}) = \mathbf{0}$, $\mathbf{c}^I(\mathbf{x}) \ge \mathbf{0}$, and there exist multipliers $\mu^E \in \mathbb{R}^{M^E}$, $\mu^I \in \mathbb{R}^{M^I}$, $\mu^I \ge \mathbf{0}$, satisfying $\mu^I \perp \mathbf{c}^I(\mathbf{x})$, and $\lambda^L, \lambda^U \in \mathbb{R}^N$, $\lambda^L, \lambda^U \ge \mathbf{0}$ satisfying $\lambda^L \perp \mathbf{x} - \mathbf{l}$ and $\lambda^U \perp \mathbf{u} - \mathbf{x}$ such that

$$\nabla f(\mathbf{x}) - (D\mathbf{c}^E)(\mathbf{x})^\top \mu^E - (D\mathbf{c}^I)(\mathbf{x})^\top \mu^I - \lambda^L + \lambda^U = \mathbf{0}$$

See Nocedal and Wright [52], Chap. 12. Thus

$$\nabla f(\mathbf{x}) - (D\mathbf{c}^E)(\mathbf{x})^\top \mu^E - (D\mathbf{c}^I)(\mathbf{x})^\top \mu^I = \lambda^L - \lambda^U = \Lambda \tag{C1}$$

Feasibility with respect to the constraints and the μ^I, $\mathbf{c}^I(\mathbf{x})$ complementarity is equivalent to

$$\begin{array}{c} \infty < \mu^E < \infty \perp \mathbf{c}^E(\mathbf{x}) \\ \mathbf{0} \le \mu^I \perp \mathbf{c}^I(\mathbf{x}) \ge \mathbf{0} \end{array}$$

Moreover, note that only one of λ^L_n, λ^U_n is nonzero at any x_n feasible with respect to the bounds $[l_n, u_n]$, implying that

$$\Lambda_n \begin{cases} \ge 0 & \text{if } x_n = l_n \\ = 0 & \text{if } x_n \in (l_n, u_n) \\ \le 0 & \text{if } x_n = u_n \end{cases}$$

or, written in MCP form, $\mathbf{l} \le \mathbf{x} \le \mathbf{u} \perp \Lambda$ as desired. □

Appendix C: Derivatives of Profits

Lemma 3. *Assume the model is a simple Logit model; i.e., $\mathcal{T}$ is a singleton. Then, as $p_k \uparrow \imath_*$, $(D^p_k\hat{\pi}_g) \to 0$ for any firm g and if $k \in \mathcal{J}_f$, $(\nabla^x_k\hat{\pi}_f) \to \mathbf{0}$.*

Proof. The first claim concerning price derivatives is proved by the formulae in Refs. [50,68].

Note that

$$(D^x_{k,n}\hat{\pi}_f)(\mathbf{X},\mathbf{p}) = \sum_{j\in\mathcal{J}_f}(D^x_{k,n}P^L_j)(\mathbf{X},\mathbf{p})(p_j - c_j(\mathbf{x}_j)) - P^L_k(\mathbf{X},\mathbf{p})(D^x_n c_k)(\mathbf{x}_k)$$

where we consider the choice probabilities as a function of $\mathbf{X}$, instead of $\mathbf{Y}$. This entails no loss of generality, and the chain rule can be invoked to compute the literal derivative of choice probabilities as defined as a function of $\mathbf{Y}$. Moreover

$$(D^x_{k,n}\hat{\pi}_f) = (D^x_n v_k)P^L_k\left(p_j - c_j - \sum_{j\in\mathcal{J}_f}P^L_j(p_j - c_j)\right) - P^L_k(D^x_n c_k)$$

where we neglect the arguments for simplicity. Now both $P^L_k(\mathbf{X},\mathbf{p})$ and $P^L_k(\mathbf{X},\mathbf{p}) \downarrow 0$ as $p_k \uparrow \imath_*$ [68], and thus, $(D^x_{k,n}\hat{\pi}_f)(\mathbf{X},\mathbf{p})$ to 0 as $p_k \uparrow \imath_*$. $\square$

Proofs for second derivatives of profits are similar.

We now prove Corollary 1.

Proof. First note that as $p_j \uparrow \imath_*$,

$$L^x_j(\mathbf{X},\mathbf{p},\mu_j,\mu^p_f) \to -(D^x_j\mathbf{g}_j)(\mathbf{x}_j)^\top\mu^E_j - (D^x_j\mathbf{h}_j)(\mathbf{x}_j)^\top\mu^I_j$$

because $(\nabla^x_j\hat{\pi}_f)(\mathbf{X},\mathbf{p}) \to \mathbf{0}$ and $(D^x_j\tilde{\nabla}^p\hat{\pi})(\mathbf{X},\mathbf{p})^\top\mu^p_f \to \mathbf{0}$ (Lemma 1). Thus, the systems

$$\begin{aligned}\mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \perp L^x_j(\mathbf{X},\mathbf{p},\mu_j,\mu^p_f)\\ \infty < \mu^E_j < \infty \perp \mathbf{g}_j(\mathbf{x}_j)\\ \mathbf{0} \le \mu^I_j \perp \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0}\end{aligned}$$

become

$$\begin{aligned}\mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \perp -(D^x_j\mathbf{g}_j)(\mathbf{x}_j)^\top\mu^E_j - (D^x_j\mathbf{h}_j)(\mathbf{x}_j)^\top\mu^I_j\\ \infty < \mu^E_j < \infty \perp \mathbf{g}_j(\mathbf{x}_j)\\ \mathbf{0} \le \mu^I_j \perp \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0}\end{aligned}$$

which is solved by any feasible $\mathbf{x}_j$ given $\mu^E_j = \mu^I_j = \mathbf{0}$. Similarly, the rows of

$$\begin{aligned}\mathbf{0} &= -(\nabla^p\hat{\pi}_f)(\mathbf{X},\mathbf{p}) + (D^p\tilde{\nabla}^p\hat{\pi})(\mathbf{X},\mathbf{p})^\top\mu^p_f\\ \mathbf{0} &= (\tilde{\nabla}^p\hat{\pi})(\mathbf{X},\mathbf{p})\end{aligned}$$

can be written componentwise as

$$\begin{aligned}0 &= -(D^p_j\hat{\pi}_f)(\mathbf{X},\mathbf{p}) + \sum_{k=1}^{J}(D^p_jD^p_k\hat{\pi}_{f(k)})(\mathbf{X},\mathbf{p})(\mu^p_f)_k\\ 0 &= (D^p_j\hat{\pi}_{f(j)})(\mathbf{X},\mathbf{p})\end{aligned}$$

However, $(D^p_j\hat{\pi}_f), (D^p_jD^p_k\hat{\pi}_{f(k)})$, and $(D^p_j\hat{\pi}_{f(j)})$ each vanishes as $p_j \uparrow \imath_*$ and thus, these systems are trivially solved by $p_j = \imath_*$.

Thus, given $p_j = \imath_*$ for all $j \in \mathcal{J}_f$, the KKT conditions reduce to

$$\left.\begin{aligned}\mathbf{l}_j \le \mathbf{x}_j \le \mathbf{u}_j \perp L^x_j(\mathbf{X},\mathbf{p},\mu_j,\mu^p_f)\\ \infty < \mu^E_j < \infty \perp \mathbf{g}_j(\mathbf{x}_j)\\ \mathbf{0} \le \mu^I_j \perp \mathbf{h}_j(\mathbf{x}_j) \ge \mathbf{0}\end{aligned}\right\}\quad \text{for all } j \in \mathring{\mathcal{J}}_f$$

and the rows of

$$\begin{aligned}\mathbf{0} &= -(\nabla^p\hat{\pi}_f)(\mathbf{X},\mathbf{p}) + (D^p\tilde{\nabla}^p\hat{\pi})(\mathbf{X},\mathbf{p})^\top\mu^p_f\\ \mathbf{0} &= (\tilde{\nabla}^p\hat{\pi})(\mathbf{X},\mathbf{p})\end{aligned}$$

corresponding to indices $j \in \mathring{\mathcal{J}}_f$. $\square$

Appendix D: Complete Results Table

Table 1 lists more details concerning the results in Sec. 4.

References

[1] Hazelrigg, G. A., 1998, "A Framework for Decision-Based Engineering Design," ASME J. Mech. Des., **120**(4), pp. 653–658.
[2] Collopy, P. D., 2001, "Economic-Based Distributed Optimal Design," Technical Report No. 2001-4675, American Institute of Aeronautics and Astronautics.
[3] Wassenaar, H. J., and Chen, W., 2003, "An Approach to Decision-Based Design With Discrete Choice Analysis for Demand Modeling," ASME J. Mech. Des., **125**(3), pp. 490–497.
[4] Cooper, A. B., Georgiopolous, P., Kim, H. M., and Papalambros, P. Y., 2006, "Analytical Target Cascading: An Enterprise Context in Optimal Product Design," ASME J. Mech. Des., **128**(1), pp. 4–13.
[5] Shiau, C.-S. N., and Michalek, J. J., 2009, "Should Designers Worry About Market Systems?," ASME J. Mech. Des., **131**(1), p. 011011.
[6] Williams, N., Azarm, S., and Kannan, P. K., 2010, "Multicategory Design of Bundled Products for Retail Channels Under Uncertainty and Competition," ASME J. Mech. Des., **132**(3), pp. 1–15.
[7] Williams, N., Kannan, P. K., and Azarm, S., 2011, "Retail Channel Structure Impact on Strategic Engineering Product Design," Manage. Sci., **57**(5), pp. 897–914.
[8] Michalek, J. J., Papalambros, P. Y., and Skerlos, S. J., 2004, "A Study of Fuel Efficiency and Emission Policy Impact on Optimal Vehicle Design Decisions," ASME J. Mech. Des., **126**(6), pp. 1062–1070.
[9] Whitefoot, K., Fowlie, M., and Skerlos, S. J., 2011, "Product Design Repsonse to Industrial Policy: Evaluating Fuel Economy Standards Using an Engineering Model of Endogenous Product Design," Technical Report No. EI @ Haas WP 214, Haas School of Business, University of California Berkeley, Berkley, CA.
[10] Whitefoot, K. S., and Skerlos, S. J., 2012, "Design Incentives to Increase Vehicle Size Created From the U.S. Footprint-Based Fuel Economy Standards," Energy Policy, **41**, pp. 402–411.
[11] von Neumann, J., and Morgenstern, O., 1947, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, NJ.
[12] Luce, R. D., and Raiffa, H., 1957, *Games and Decisions: A Critical Survey*, John Wiley and Sons, Inc., New York.
[13] Fudenberg, D., and Tirole, J., 1991, *Game Theory*, MIT Press, Cambridge, MA.
[14] Myerson, R. B., 1997, *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA.
[15] Bresnahan, T., 1987, "Competition and Collusion in the American Automobile Industry: The 1955 Price War," J. Ind. Econ., **35**(4), pp. 457–482.
[16] Berry, S., Levinsohn, J., and Pakes, A., 1995, "Automobile Prices in Market Equilibrium," Econometrica, **63**(4), pp. 841–890.
[17] Goldberg, P. K., 1995, "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," Econometrica, **63**(4), pp. 891–951.
[18] Verboven, F., 1996, "International Price Discrimination in the European Car Market," RAND J. Econ., **27**(2), pp. 240–268.
[19] Goldberg, P. K., 1998, "The Effects of the Corporate Average Fuel efficiency Standards in the US," J. Ind. Econ., **46**(1), pp. 1–33.
[20] Verboven, F., 1999, "Product Line Rivalry and Market Segmentation—With an Application to Automobile Optional Engine Pricing," J. Ind. Econ., **47**(4), pp. 399–425.
[21] Fershtman, C., Gandal, N., and Markovich, S., 1999, "Estimating the Effect of Tax Reform in Differentiated Product Oligopolistic Markets," J. Public Econ., **74**, pp. 151–170.
[22] Nevo, A., 2000, "Mergers With Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," RAND J. Econ., **31**(3), pp. 395–421.
[23] Nevo, A., 2001, "Measuring Market Power in the Ready-to-Eat Cereal Industry," Econometrica, **69**(2), pp. 307–342.
[24] Hausman, J. A., and Leonard, G. K., 2002, "The Competitive Effects of a New Product Introduction," J. Ind. Econ., **50**(3), pp. 237–263.
[25] Petrin, A., 2002, "Quantifying the Benefits of New Products: The Case of the Minivan," J. Polit. Econ., **110**(4), pp. 705–729.
[26] Verboven, F., 2002, "Quality-Based Price Discrimination and Tax Incidence: Evidence From Gasoline and Diesel Cars," RAND J. Econ., **33**(2), pp. 275–297.
[27] Goolsbee, A., and Petrin, A., 2004, "The Consumer Gains From Direct Broadcast Satellites and the Competition With Cable TV," Econometrica, **72**(2), pp. 351–381.
[28] Kleit, A., 2004, "Impacts of Long-Range Increases in the Corporate Average Fuel Economy (Cafe) Standard," Econ. Inquiry, **42**(2), pp. 279–294.

[29] Thomadsen, R., 2005, "The Effect of Ownership Structure on Prices in Geographically Differentiated Markets," RAND J. Econ., **36**(4), pp. 908–929.
[30] Beresteanu, A., and Li, S., 2011, "Gasoline Prices, Government Support, and the Demand for Hybrid Vehicles in the U.S.," Int. Econ. Rev., **52**(1), pp. 161–182.
[31] Choi, S. C., DeSarbo, W. S., and Harker, P. T., 1990, "Product Positioning Under Price Competition," Manage. Sci., **36**(2), pp. 175–199.
[32] Choi, S. C., DeSarbo, W. S., and Harker, P. T., 1992, "A Numerical Approach to Deriving Long-Run Equilibrium Solutions in Spatial Positioning Models," Manage. Sci., **38**(1), pp. 75–86.
[33] Choi, S. C., and DeSarbo, W. S., 1993, "Game Theoretic Derivations of Competitive Strategies in Conjoint Analysis," Mark. Lett., **4**(4), pp. 337–348.
[34] Besanko, D., Gupta, S., and Jain, D., 1998, "Logit Demand Estimation Under Competitive Pricing Behavior: An Equilibrium Framework," Manage. Sci., **44**(11), pp. 1533–1547.
[35] Sudhir, K., 2001, "Competitive Pricing Behavior in the Auto Market: A Structural Analysis," Market. Sci., **20**(1), pp. 42–60.
[36] Sudhir, K., 2001, "Structural Analysis of Manufacturer Pricing in the Presence of a Strategic Retailer," Mark. Sci., **20**(3), pp. 244–264.
[37] Dube, J.-P., Chintagunta, P., Petrin, A., Bronnenberg, B., Goettler, R., Seetharaman, P. B., Sudhir, K., Thomadsen, R., and Zhao, Y., 2002, "Structural Applications of the Discrete Choice Model," Mark. Lett., **13**(3), pp. 207–220.
[38] Rhim, H., and Cooper, L. G., 2005, "Assessing Potential Threats to Incumbent Brands: New Product Positioning Under Price Competition in a Multi-Segmented Market," Int. J. Res. Mark., **22**(2), pp. 159–182.
[39] Shiau, C.-S. N., and Michalek, J. J., 2009, "Optimal Product Design Under Price Competition," ASME J. Mech. Des., **131**(7), p. 071003.
[40] Wang, Z., Azarm, S., and Kannan, P. K., 2011, "Strategic Design Decisions for Uncertain Market Systems Using an Agent Based Approach," ASME J. Mech. Des., **133**(4), p. 041003.
[41] Morrow, W. R., and Skerlos, S. J., 2011, "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand," Oper. Res., **59**(2), pp. 328–345.
[42] Luo, L., Kannan, P. K., and Ratchford, B. T., 2007, "New Product Development Under Channel Acceptance," Mark. Sci., **26**(2), pp. 149–163.
[43] Bajari, P., Benkard, C. L., and Levin, J., 2007, "Estimating Dynamic Models of Imperfect Competition," Econometrica, **75**(5), pp. 1331–1370.
[44] Wards, 2006, *Wards Automotive Yearbooks*, Wards Reports, Inc.
[45] EPA, 2012, www.fueleconomy.gov
[46] Luo, Z.-Q., Pang, J. S., and Ralph, D., 1996, *Mathematical Programs With Equilibrium Constraints*, Cambridge University Press, Cambridge, MA.
[47] Frischknecht, B. D., Whitefoot, K., and Papalambros, P. Y., 2010, "On the Suitability of Econometric Demand Models in Design for Market Systems," ASME J. Mech. Des., **132**(12), pp. 1–11.
[48] Dube, J.-P., Fox, J. T., and Su, C.-L., 2011. Improving the Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation, Working Paper, University of Chicago.
[49] Biegler, L. T., and Zavala, V. M., 2009, "Large-Scale Nonlinear Programming Using IPOPT: An Integrating Framework for Enterprise-Wide Dynamic Optimization," Comput. Chem. Eng., **33**(3), pp. 575–582.
[50] Morrow, W. R., and Skerlos, S. J., 2010, "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand: A Technical Framework for Analysis and Efficient Computational Methods," Technical Report, Iowa State University.
[51] Morrow, W. R., 2012, "Computing Bertrand-Nash Equilibrium Prices With Finite Purchasing Power," Technical Report, Iowa State University.
[52] Nocedal, J., and Wright, S. J., 2006, *Numerical Optimization*, Springer-Verlag, New York.
[53] Train, K., 2009, *Discrete Choice Methods With Simulation*, 2nd ed., Cambridge University Press, Cambridge.
[54] Donndelinger, J. A., Robinson, J. A., and Wissmann, L. A., 2008, "Choice Model Specification in Market-Based Engineering Design," Proceedings of the 34th ASME Design Automation Conference, Brooklyn, NY.
[55] Hoyle, C., Chen, W., Wang, N., and Gomez-Levi, G., 2001, "Understanding and Modeling Heterogeneity of Human Preferences for Engineering Design," Taylor & Francis, **22**(8), pp. 583–601.
[56] McFadden, D. L., 1981, "Econometric Models of Probabilistic Choice," *Structural Analysis of Discrete Data and Econometric Applications*, C. F. Manski and D. L. McFadden, eds., MIT Press, Cambridge, MA, pp. 199–272.
[57] Louvierre, J. J., Hensher, D. A., and Swait, J. D., 2000, *Stated Choice Methods: Analysis and Applications*, Cambridge University Press, Cambridge, MA.
[58] Wassenaar, H. J., Chen, W., Cheng, J., and Sudjianto, A., 2005, "Enhancing Discrete Choice Modeling for Decision-Based Design," ASME J. Mech. Des., **127**(4), pp. 514–523.
[59] Wassenaar, H. J., Deepak, K., and Chen, W., 2006, "Discrete Choice Demand Modeling for Decision-Based Design," *Decision Making in Engineering Design*, K. Lewis, W. Chen, and L. Schmidt, eds., ASME Press, New York, pp. 86–108.
[60] Hoyle, C., Chen, W., Ankenman, B., and Wang, N., 2009, "Optimal Experimental Design of Human Appraisals for Modeling Consumer Preferences in Engineering Design," ASME J. Mech. Des., **131**(8), pp. 1–9.
[61] Kumar, D., Hoyle, C., Chen, W., Wang, N., Gomex-Levi, G., and Koppelman, F., 2009, "A Hierarchical Choice Modeling Approach for Incorporating Customer Preferences in Vehicle Package Design," Int. J. Prod. Dev., **8**(3), pp. 228–251.
[62] MacDonald, E., Whitefoot, K., Allison, J., Papalambros, P. Y., and Gonzalez, R., 2010, "An Investigation of Sustainability, Preference, and Profitability in Design Optimization," Proceedings of the ASME 2010 International Design Engineering Technical Conferences.
[63] Ben-Akiva, M., 1985, *Discrete Choice Analysis: Theory and Application to Travel Demand*, MIT Press, Cambridge, MA.
[64] Green, P. E., Krieger, A. M., and Wind, Y., 2001, "Thirty Years of Conjoint Analysis: Reflections and Prospects," Interfaces, **31**(3), pp. S56–S73.
[65] Nevo, A., 2000, "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," J. Econ. Manage. Strategy, **9**(4), pp. 513–548.
[66] Caplin, A., and Nalebuff, B., 1991, "Aggregation and Imperfect Competition: On the Existence of Equilibrium," Econometrica, **59**(1), pp. 25–59.
[67] Morrow, W. R., 2008, "A Fixed-Point Approach to Equilibrium Pricing in Differentiated Product Markets," Ph.D. thesis, Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI.
[68] Morrow, W. R., 2012, "On the Existence of Bertrand-Nash Equilibrium Prices Under Logit Demand," Technical Report No. arXiv:1012.5832, Iowa State University, Ames, IA.
[69] Boyd, J. H., and Mellman, R. E., 1980, "The Effect of Fuel Economy Standards on the U.S. Automotive Market: An Hedonic Demand Analysis," Transp. Res. A, **14A**, pp. 367–378.
[70] Morrow, W. R., Long, M., and MacDonald, E. F., 2012, "Consider-Then-Choose models in Decision-Based Design optimization," Proceedings of ASME Paper No. IDETC 2012.
[71] Belloni, A., Freund, R., Selove, M., and Simester, D., 2008, "Optimal Product Line Design: Efficient Methods and Comparisons," Manage. Sci., **54**(9), pp. 1544–1552.
[72] Trigg, A. B., 2001, "Veblen, Bourdieu, and Conspicuous Consumption," J. Econ. Issues, **XXXV**(1), pp. 99–115.
[73] Train, K., 2009, *Discrete Choice Methods With Simulation*, 2nd ed., Cambridge University Press, Cambridge, MA.
[74] Shapiro, A., 2003, "Monte Carlo Sampling Methods," *Handbooks in Operations Research and Management Science*, A. Ruszczynski and A. Shapiro, eds., Vol. 10, Elsevier Science, New York/Amsterdam, pp. 353–425.
[75] Dirkse, S. P., and Ferris, M. C., 1995, "The PATH Solver: A Non-Monotone Stabilization Scheme for Mixed Complementarity Problems," Optim. Methods Software, **5**, pp. 123–156.
[76] Munson, T., 2000, "Algorithms and Environments for Complementarity," Ph.D. thesis, University of Wisconsin-Madison, Madison, WI.
[77] Munkres, J. R., 1991, *Analysis on Manifolds*, Westview Press, Boulder, CO.
[78] Fletcher, R., Leyffer, S., Ralph, D., and Scholtes, S., 2006, "Local Convergence of SQP Methods for Mathematical Programs With Equilibrium Constraints," SIAM J. Optim., **17**(1), pp. 259–286.
[79] Leyffer, S., Lopez-Calva, G., and Nocedal, J., 2006, "Interior Methods for Mathematical Programs With Complementarity Constraints," SIAM J. Optim., **17**(1), pp. 52–77.
[80] Ralph, D., 2008, "Mathematical Programs With Complementarity Constraints in Traffic and Telecommunications Networks," Philos. Trans. R. Soc. London, **366**(1872), pp. 1973–1987.
[81] Izmailov, A., Solodov, M., and Uskov, E., 2012, "Augmented Lagrangian Methods Applied to Optimization Problems With Degenerate Constraints, Including Problems With Complementarity Constraints," SIAM J. Optim., **22**, pp. 1579–1606.
[82] Gill, P. E., Murray, W., and Saunders, M. A., 2005, "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," SIAM Rev., **47**(1), pp. 99–131.
[83] Waltz, R. A., Morales, J. L., Nocedal, J., and Orban, D., 2006, "An Interior Algorithm for Nonlinear Optimization That Combines Line Search and Trust Region Steps," Math. Program., Ser. A, **107**(3), pp. 391–408.
[84] Morrow, W. R., Mineroff, J., and Whitefoot, K. S., 2012, "Numerically Stable Design Optimization With Price Competition," Paper No. ASME J. of Mech Des (in press).
[85] Golub, G. H., and Loan, C. F. V., 1996, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD.
[86] Trefethen, L. N., and Bau, D., 1997, *Numerical Linear Algebra*, SIAM, Philadelphia, PA.
[87] Higham, N. J., 2002, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, PA.
[88] Morrow, W. R., 2012, "Hessian-Free Methods for Checking the Second-Order Sufficient Conditions in Equality-Constrained Optimization and Equilibrium Problems, Technical Report arXiv:1106.0898, Iowa State University, Ames, IA.
[89] Conn, A. R., Gould, N. I. M., and Toint, P. L., 2000, *Trust Region Methods*, SIAM, Philadelphia, PA.
[90] NHTSA, 2008, "Prelimnary Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2011-2015 Passenger Cars and Light Trucks," Technical Report, U.S. Department of Transportation.
[91] U.S. Department of Labor, Bureau of Labor Statistics, 2006, Current Popoulation Survey, http://www.bls.gov/cps/
[92] Gallego, G., and van Ryzin, G., 1994, "Optimal Dynamic Pricing of Inventories With Stochastic Demand Over Finite Horizon," Manage. Sci., **40**(8), pp. 999–1020.
[93] Gallego, G., and van Ryzin, G., 1997, "A Multiproduct Dynamic Pricing Problem and Its Applications to Network Yield Management," Oper. Res., **45**(1), pp. 24–41.
[94] Gigerenzer, G., and Gaissmaier, W., 2011, "Heuristic Decision Making," Annu. Rev. Psychol., **62**, pp. 451–82.

[95] Silver, N., 2012, *The Signal and the Noise: Why So Many Predictions Fail - but Some Don't*, Penguin Press, New York.
[96] Michalek, J. J., Feinberg, F. M., and Papalambros, P. Y., 2005, "Linking Marketing and Engineering Product Design Decisions via Analytical Target Cascading," J. Prod. Innovation Manage., **22**(1), pp. 42–62.
[97] Michalek, J. J., Ceryan, O., Papalambros, P. Y., and Koren, Y., 2006, "Balancing Marketing and Manufacturing Objectives in Product Line Design," ASME J. Mech. Des., **128**(6), pp. 1196–1204.
[98] Dirkse, S. P., and Ferris, M. C., 1996, "A Pathsearch Damped Newton Method for Computing General Equilibria," Ann. Oper. Res., **68**, pp. 211–232.
[99] Ferris, M. C., and Pang, J. S., 1997, "Engineering and Economic Applications of Complementarity Problems," SIAM Rev., **39**, pp. 669–713.
[100] Ferris, M. C., Kanzow, C., and Munson, T. S., 1999, "Feasible Descent Algorithms for Mized Complementarity Problems," Math. Program., **86**, pp. 475–497.
[101] Dennis, J. E., and Schnabel, R. B., 1996, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, SIAM, Philadelphia, PA.
[102] Brown, P. N., and Saad, Y., 1990, "Hybrid Krylov Methods for Nonlinear Systems of Equations," SIAM J. Sci. Stat. Comput., **11**(3), pp. 450–481.
[103] Pernice, M., and Walker, H. F., 1998, "NITSOL: A Newton Iterative Solver for Nonlinear Systems," SIAM J. Sci. Comput., **19**(1), pp. 302–318.
[104] Ortega, J. M., and Rheinboldt, W. C., 1970, *Iterative Solution of Nonlinear Equations in Several Variables*, Society for Industrial and Applied Mathematics, Philadelphia, PA.