Title
Three Essays on Strategic Considerations for Product Development

Permalink
https://escholarship.org/uc/item/44f1m78x

Author
Sawhill, James Winslow

Publication Date
2010

Peer reviewed|Thesis/dissertation

Three Essays on Strategic Considerations for Product Development

By James Winslow Sawhill

A dissertation submitted in partial satisfaction

of the requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor J. Miguel Villas-Boas, Chair
Professor James Powell
Professor James Wilcox

Fall 2010

Abstract

Three Essays on Strategic Considerations for Product Management

by

James Winslow Sawhill

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor J. Miguel Villas-Boas, Chair

This dissertation is composed of three essays focused on strategic considerations for product development. In chapter I, I address the question of whether consumers equally weigh capital and operating costs when purchasing durable goods. This trade off is important to manufacturers as they determine how much of their product design and production costs should be dedicated to keeping operating costs low. I test this empirically using data from the automobile industry. Chapter II is also an empirical study which explores whether consumers are willing to pay for socially responsible products. The answer to this question is important for firms to address in their product development process as they decide whether they will gain more market share by producing a socially responsible product with somewhat higher costs or a low cost product which does not incorporate socially responsible practices. In this case, I use data on retail sales in coffee industry using fair trade practices as an exemplar of social responsibility. Chapter III addresses the question of how durable or reliable manufacturers of durable goods should make their products. Consumers will likely want to pay more for a more durable product, yet increased durability depresses replacement demand. I attempt to gain insight into this trade off by developing an analytical model of the interplay between consumers and a monopolist manufacturer of durable goods. The remainder of the abstract provides a more detailed summary of each of these chapters in turn.

Chapter I explores whether consumers behave as if they are optimally trading off capital and operating costs when purchasing a durable good. I study this question using data on gas prices and automobile sales over the 20 year period, 1971-1990. This question is important for three reasons. First, it is interesting from a theoretical basis if consumers make this trade off optimally. Many theoretical models in marketing and economics make the fundamental assumption that consumers equally weigh current and future events when making decisions today. Yet there is some evidence, mostly from laboratory experiments, that consumers underweight future events. I attempt to explore this question in a market setting where the stakes are considerably higher. This research adds evidence to the debate about how much weight consumers place on future events when making choices today. Second, it is interesting to firms making product design decisions. If consumers underweight future events, then when making purchase decisions, consumers will view operating costs as less important than the upfront capital

cost of the product. Finally, the answer to this question informs public policy. Many have argued that there is a need for the US to reduce gasoline demand per capita. Lower gasoline consumption would reduce environmental pressures, potentially dampen inflation, and allow more foreign policy flexibility in dealing with antagonistic regimes in oil-exporting states. While these rationale for reducing gasoline consumption have a normative flavor, and reasonable people may disagree as to the validity and motivations of this goal, it will nonetheless be useful to know the relative effectiveness of different policy levers in curbing gasoline consumption. For example, if consumers underweight fuel costs during the vehicle purchase decision, then a gas tax will be relatively less effective than a tax on gas guzzling vehicles.

To study this question I develop a choice model of the automobile industry. I identify the weight the consumer places on capital v. operating costs by determining how much of the variation in automobile market share can be explained by variation in each of these two factors holding other product attributes constant. We use data from the period 1971-1990, a period over which gas prices and thereby operating costs experienced considerable variation. In order to model operating costs which are not known at the time of purchase, I account for the expectations of consumers about car usage and gas prices. I assume that consumers are aware that they will respond to changes in the gasoline prices with changes in their driving patterns. Consumers also know that gasoline prices are not stable, and need to form expectations about future gasoline prices at the time of the automobile purchase. To take account of this effect, I estimate an ARIMA model of US gasoline prices from 1960-1995, that is used as the by consumers in their expectations formation process. Taking car usage and gas price expectations together enables an estimate of future gasoline costs of operating the car in the future. I also account for consumer heterogeneity in miles driven, sensitivity to automobile price, and sensitivity to operating costs. Finally, I recognize that prices are not exogenously determined and attempt to model prices as market outcomes.

Based on the results of the model developed, I find no evidence to support behavioral theories that consumers systematically underweight the cost of future events in real market settings. However, I find significant evidence that large portions of the population are not making the trade-off optimally. Some consumers underweight future operating costs (SUV drivers) while others appear to overweight them (hybrid drivers). Conservatively, at least 30% of the population is either drastically underweighting or overweighting operating costs when purchasing a new car.

Chapter II addresses the question of whether consumers are willing to pay for corporate social responsibility(CSR). This question is important in an environment where CSR is ubiquitous, yet it is unclear that these programs actually pay off for the firms that sponsor them. For example, consider Target's program to donate 1% of all retail sales to United Way local charities. Do consumers really want their money spent this way? Are consumers happily paying 1% higher prices or are they switching to a competitor which does not donate a portion of revenues to charity?   Who is making the donation in the end, Target's shareholders or customers? From a social planner's perspective, the point is largely moot, yet to the shareholders of Target and many other firms practicing CSR, the question is crucially important.

I endeavor to study this question within the context of the coffee industry, an important and sizeable commodity market. In particular, I explore the impact of Fair-Trade (FT) certification on the retail coffee market. FT is a social and ethical movement that supports the ethical production of coffee and other products largely in third world countries. Coffee can be FT certified by adhering to FT standards. Once certified, FT coffee is distinguished from non-FT by distinctive labeling visible to the consumer who is deciding which coffee product to select from the supermarket shelf.

The analytical strategy for this paper is to first estimate the price premium commanded by FT coffee over non-FT coffees through ordinary least squares (OLS) and Fixed Effects hedonic price regressions. However, these tools do not allow us to disentangle the portion of the price premium which is due to supply considerations (i.e., FT certification costs) from the portion which is due to consumers' willingness to pay for FT coffee because they want to support socially responsible coffee production. To parse out willingness to pay from the overall price premium, I specify a brand choice model similar to the model used in chapter I.

Using hedonic price regressions I establish that FT coffee carries a price premium of $1.74 per 12-16 oz. It seems likely that at least a portion of this premium is due to increased consumer willingness to pay for FT coffee. However, I cannot rule out the possibility that the price premium is a result of the added costs associated with that FT practices. The choice model specified in this paper should enable the allocation of the cause of the price premium we have now established for FT coffee to demand v. supply considerations. I hope to estimate this model in future research.

Chapter III addresses the problem of how durable or long lasting manufacturers should they make their products. On the one hand a more durable product will be more desirable to consumers, since it will provide benefits over a longer period. Thus, a longer-lived product will command a higher price. However, it seems likely that unit production costs will increase as a product is made more durable due to the increased cost of more reliable materials and more exacting quality standards. In addition, a product which is more durable will be replaced less frequently. *Ceteris paribus,* less frequent replacement is less desirable to manufacturers, as the periodicity of the revenue stream increases. Manufacturers can trade off the benefits of durability with the costs to determine the optimal reliability or life for the goods they produce.

In some sense, this problem is a classic trade-off between quality and cost. What distinguishes the durable goods reliability problem is that increasing quality depresses replacement demand. A common anecdote is that light bulbs could easily be manufactured to last longer, but are not in order to increase replacement sales.

This question is important for manufacturers to understand from several perspectives. First, a manufacturer of a product with technology that is fairly static (e.g., light bulbs), needs to consider replacement demand in developing product designs. When technology is not static (e.g., computers), it is important to understand how the rate of technology advance will stimulate replacement demand. Should the products be designed to be more or less durable in the face of technological advance? An additional complication arises when the rate of technological advance may only be partially

observable to the consumer (e.g. golf clubs). Finally, manufacturers need to consider "buying back" used durable goods from their customers in order to stimulate replacement demand. Even if the manufacturer does not take physical possession of the old product, it may be optimal to offer price concessions on newer products in order to induce the consumer to replace her existing technology.

I study this question by developing a partial equilibrium model for durable goods where both manufacturers and consumers are forward looking over an infinite horizon. I make the simplifying assumptions that there is a single monopolist manufacturer and that consumers have homogeneous preferences.

I find that in all cases, it is optimal for manufacturers to incent consumers (buy back) to replace their existing durable goods before the end of the useful life of the product. I also find that under extremely rapid rates of technological advance, it is optimal for manufacturers to extend the life of the product and increase durability. This result runs somewhat counter to intuition in that one might think that rapid advances in technology would promote more frequent product replacement. However, if technology advances rapidly then patient consumers will want to be able to enjoy their new product (e.g., HDTV) and are willing to pay a high price for durability. This increased willingness to pay for an advanced long-lived product will more than offset the loss of replacement sales in the future. Finally, I show that when technology advances are uncertain and not directly observable by consumers, they are willing to pay more than when technology advances at a known and constant rate. The intuition for this result is that consumers are risk averse and do not want to miss out on a new breakthrough product. We believe this may be the reason that certain durable goods manufacturers (e.g. golf club manufacturers) appear to introduce new products at a far greater rate than technology is actually advancing.

Chapter I: Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market

In any durable goods market, the consumer faces two types of costs – upfront capital costs and ongoing operating costs. Therefore, when studying these markets, an important question is do consumers equally weight the capital and operating costs of the product when making the initial purchase decision. The US market for automobiles and gasoline provides a fertile research domain for addressing this important question. The automobile market covers a broad cross section of consumers, and the stakes are high with a typical low-end new car selling for around $20,000 and gasoline commonly retailing for more than $3.00 per gallon. A consumer who drives their car 15,000 miles per year at the current US Corporate Average Fuel Efficiency (CAFE) standard of 27.5 miles per gallon (mpg) for passenger cars will incur an annual fuel cost of roughly $1,650.

By developing a consumer choice model for the US automobile market we can estimate how consumers are making these tradeoffs. We don't have an *a priori* prediction as to the results of this analysis. Regardless of the outcome, the results will be informative to policy makers regarding alternative methods of modifying consumer behavior. Many have argued that there is a need for the US to reduce gasoline demand per capita. For example, President George W. Bush called for a 20% reduction in gasoline consumption over the next 10 years in his 2007 state of the union address. Lower gasoline consumption would reduce environmental pressures, potentially dampen inflation, and allow more foreign policy flexibility in dealing with antagonistic regimes in oil-exporting states. While these rationales, for reducing gasoline consumption have a normative flavor, and reasonable people may disagree as to the validity and motivations of this goal, it will nonetheless be useful to know the relative effectiveness of different policy levers in curbing gasoline consumption. If fuel costs are underweighted during the vehicle purchase decision, then a gas tax will be relatively less effective than a tax on gas guzzling vehicles. More broadly, the outcome of this research will provide insight into other policy questions such as how do we motivate individuals to increase their savings rate, or how do we encourage people to improve the energy efficiency of their homes.

In addition to the practical policy significance of this research, the analysis will provide a useful test of the hyperbolic discounting or present-biased preferences hypothesis (Laibson 1997, O'Donoghue and Rabin 2001). This theory suggests that people underweight the economic consequences of future events above and beyond a reasonable level of time discounting. Applied to the automobile market, present-biased preferences imply that gasoline costs, since they occur in the future, will be underweighted relative to capital costs in the consumer's new car purchase decision.

Support for the present-biased preferences hypothesis has been found in laboratory experiments (e.g., Lowenstein and Prelec 1994, Zauberman 2003). In addition, two empirical studies both find that consumers have a discount rate of approximately 20% when purchasing room air conditioners (Hausman 1979) and automobiles (Mannering and Winston 1985). While these studies are similar to the current research, the above

studies do not address three fundamental elements of the consumer durable markets which we explicitly model in this paper.

First, unlike the previous studies the current paper considers that there are product attributes which are unobserved (by the econometrician) attributes (*e.g.,* crash test ratings) which will likely be correlated with observed vehicle attributes (*e.g.*, price/capital costs). If these omitted determinants of vehicle utility are not modeled, estimates for the effect of observed attributes on product choice will be inconsistent, fatally limiting the ability to answer the research query at hand. Second, in this paper we account for uncertainty with respect to operating costs at the time of the original purchase. Given highly volatile energy prices, it is important to consider this uncertainty and to model consumer behavior as this uncertainty is resolved. When considering future periods, our general approach is to allow for consumers to react to the variability in the price of gasoline by dynamically adjusting their driving patterns. Finally, in this paper we consider consumer heterogeneity[1] in preferences for vehicle attributes. Without a model of heterogeneity in attribute preferences, consumers share the same expected ranking over products(cars). Consequently, any consumer facing a capital cost increase in her first choice which induces substitution will always to substitute to vehicles that are on average the most popular regardless of the characteristics of her first choice. In this framework, if the next best car is a gas guzzler then we will underestimate the elasticity of operating costs, while if the next best car has high fuel economy we will overestimate operating cost elasticity. Thus, modeling heterogeneity in attribute preferences is important for answering the research question of this paper.

To conduct this research, we employ data from the US automobile market over the period 1971-1990. These data when augmented with data on gasoline price trajectories, consumer driving patterns and a gasoline demand model will be employed to estimate a discrete choice model for automobile purchase behavior. We develop a dynamic model for calculating expected vehicle operating costs. The model will provide estimates as to the sensitivity of consumer product choice to changes in capital costs and changes in operating costs. The null hypothesis is that consumers will be equally sensitive to changes in operating and capital costs.

Section A of this chapter will outline the econometric model. In section B, we will discuss our estimation procedure. Section C, we will outline and discuss the model's data and results. Finally, we will conclude in Section D and discuss some directions for future research.

IA. The Model

IA.1) Overview

We draw upon the rich literature of discrete choice models in the automobile industry (see Berry et. al. (1995, 2004), Goldberg (1995, 1998), Petrin (2002)). The data and

---

[1] Heterogeneity in consumers is explicitly modeled by Hausman's paper on room air conditioners but not by the Mannering and Winston study of automobiles.

model being considered in this research is closest to that used by Berry et al. (1995), so we use their methodology as a useful point of departure. We specify a model for the utility, U_{ij}, that consumer i derives from product j to be linear in product and consumer attributes. Let $j = 0,...,J$ index the products competing in the market, where product $j = 0$ is an outside good (i.e., the option of not purchasing a new car). Let k index a set of observed product characteristics. A random coefficients model for utility can be written as follows

$$U_{ij} = \sum_k x_{jk} \tilde{\beta}_{ik} + \xi_j + \varepsilon_{ij}$$
(1)

with

$$\tilde{\beta}_{ik} = \bar{\beta}_k + \beta_k^u \nu_{ik}$$
(1a)

where x_{jk} and ξ_j represent the observed and unobserved product characteristics respectively, the β_{ik} represent the taste of consumer i for characteristic k, $\mathbf{v}_i$ is a vector of unobserved consumer attributes, and the ε_{ij} are idiosyncratic individual preferences assumed to be independent of the product attributes and of each other.

Two elements of the specification above are noteworthy with respect to the research question at hand. First, as discussed above it has the ability to represent consumer heterogeneity. Consider a model where consumer preferences across product characteristics are homogeneous (i.e., $\tilde{\beta}_{ik} = \bar{\beta}_k$). In this case, the only source of consumer heterogeneity comes from the independent and identically distributed ε_{ij}'s. The implication of this formulation is that all consumers share the same expected ranking over products(cars). Consequently, any consumer facing a capital cost increase in her first choice which induces substitution will always to substitute to vehicles that are on average the most popular regardless of the characteristics of her first choice. In this framework, if the next best car is a gas guzzler then we will underestimate the elasticity of operating costs, while if the next best car has high fuel economy we will overestimate operating cost elasticity. For this research, it is critically important to get an accurate read on the willingness of consumers to tradeoff capital costs with operating costs and to recognize that different consumers will make this trade off differently. Therefore, we require a model which accounts for the variability in consumer tastes across attributes.

Second, while our data set has descriptions of some vehicle attributes, there will be unobserved (by the econometrician) attributes (*e.g.,* crash test ratings) which will likely be correlated with observed vehicle attributes (*e.g.*, price/capital costs). The role of ξ in our model is to pick up these unobserved attributes. If we do not account for these omitted determinants of vehicle utility, we will obtain inconsistent estimates for the effect of observed attributes on product choice, fatally limiting our ability to answer the research query at hand. However, by instrumenting for the observed variables which are correlated with ξ we can achieve consistent estimates for the parameters of interest.

Focusing on the willingness of consumers to substitute vehicle capital costs for operating costs, we rewrite the consumer utility function as:

$$U_{ij} = -\alpha_i \, (p_j^k - \gamma_i \, E(c_{ij})) + \mathbf{x_j'}\boldsymbol{\beta}_i + \xi_j + \varepsilon_{ij} \tag{2}$$

where p_j^k is the capital cost of brand j, $E(c_{ij})$ is the expected discounted operating cost of car j to consumer i over the car's lifetime, ai is a parameter representing consumer preference for more income (lower capital and operating costs), gi is a parameter which represents the relative consumer aversion to operating costs versus capital costs, and the other elements of the model are as specified above. In order to estimate the above model we need a measure of expected operating costs, which will require some effort to produce. We outline our methodology for calculating consumer expected operating costs below. Supposing for the moment that we have a model for expected operating costs in hand we can proceed to estimate the full random coefficients model above, and test the null hypothesis of consumer rationality, *i.e.*,

$$\frac{\partial U / \partial p_j^k}{\partial U / \partial E(c_{ij})} = \frac{\alpha_i \gamma_i}{\alpha_i} = \gamma_i = 1 \; \forall i \tag{3}$$

Where (3) states that the marginal utility of saving a dollar on the price of a car is equal to the marginal utility of saving a dollar on the net present value of expected operating costs over the life of the car. Since we are explicitly modeling heterogeneity through a random coefficients model, we permit some consumers to place more emphasis on capital costs than operating costs and vice versa. If all consumers are rational, then they should all tradeoff capital and operating costs dollar for dollar. However, since consumers may be heterogeneous with respect to rates of time preference then it is possible to have heterogeneity in the ratio in (3) within a population of consumers who are all trading off the price of a car with its operating costs one for one. Therefore, it is not possible to precisely disentangle heterogeneity in rates of time preferences from heterogeneity in tradeoff behavior. For purposes of this paper, we will estimate the distribution of *gi* without attempting to make precise attribution of its underlying behavior. The remainder of this section provides specifics on the model for operating costs including gas prices, and on modeling price endogeneity.

IA.2) Operating Costs – Overview

To keep the analysis manageable, the only element of vehicle operating costs considered in this analysis is the consumer's expected gasoline purchases. For most vehicles, gasoline purchases are the largest single component of operating costs. Moreover, given that gas price and EPA fuel efficiency information is readily available to consumers, they should be well-informed as to the level of expenditures required to operate a car at least in the current period. When considering future periods, our general approach is to allow for consumers to react to the variability in the price of gasoline by dynamically adjusting their driving patterns. At the time of purchase, consumers form expectations about the distribution of possible future gas price trajectories and plan to reduce their driving in

4

the event of price increases and increase their driving in the event of price decreases. The dynamic treatment of operating costs represents an extension from previous literature on automobile choice which considers the purchase decision only at a single point in time independent of decisions in other time periods. Ideally, the model could be extended to allow for dynamic vehicle replacement as well.

We develop a general approach for modeling operating costs. Model specifics are outlined in later subsections.

Denote the price per mile driven for product j in time t as:

$$p_{tj}^m = \frac{p_t^g}{f_j}$$

(4)

where p_t^g is the price of a gallon of gas at time t, f_j is the vehicle's fuel efficiency in miles per gallon which is assumed not to vary over time. Using this notation the operating cost of car j at time t is

$$c_{tj} = p_{tj}^m m(p_{tj}^m)$$

(5)

where $m(p^m)$ is the quantity of miles the vehicle is driven as a function of the price per mile. We assume that the function m(.) does not vary over time. Under the assumption that m(.) is declining linearly[2] in p^m with slope λ and intercept ψ, then

$$c_{tj} = p_{tj}^m(\psi - \lambda p_{tj}^m)$$

(6)

To model consumer expectations of future gas prices, we assume that the car is purchased at time 0. We fit an ARIMA (1,1,0) model to annual retail gasoline price data obtained from the US Department of Energy. This model, discussed in further detail below, shows that the evolution of p^g can described by the following process:

$$\Delta p_t^g = \phi \Delta p_{t-1}^g + u_t$$

(7)

Where $\Delta p_t^g \equiv p_t^g - p_{t-1}^g$, $\{u_s\}_{s=1}^t$ is a white noise process with $E(u_t) = 0$, $Var(u_t) = \sigma_u^2$ for all t , and $Cov(u_t u_s) = 0$ for $t \neq s$, and ϕ is a parameter to be estimated. Given this setup, equation (7) can be rewritten to express the gas price in any time t, p_t^g, as a function of the price in period 0, p_0^g, the change in gas prices in period 0, Δp_0^g, and the random shocks $\{u_s\}_{s=1}^t$.

---

[2] We make this assumption as it offers a good approximation for modest price changes. Moreover, the linear model offers analytical tractability, and better illustrates the effect of consumer dynamics with respect to operating costs. Nonetheless, the model is easily extended to other demand specifications.

$$p_t^g = p_0^g + \frac{\phi(1-\phi^t)}{1-\phi}\Delta p_0^g + \sum_{s=1}^t \frac{1-\phi^{t-s+1}}{1-\phi}u_s \tag{8}$$

Thus, given knowledge of f, gas prices in the current and most recent period and the variance of the random shocks to gas prices, the consumer can calculate expected gasoline costs in any future period conditional on the model of car she buys today. For clarity, we make the following simple definitions and intermediate calculations:

$$\mu_{tj} \equiv \frac{1}{f_j}p_0^g + \frac{\phi(1-\phi^t)}{1-\phi}\Delta p_0^g \text{ the deterministic component of } p_{tj}^m \tag{9}$$

$$\varepsilon_{tj} \equiv \frac{1}{f_j}\sum_{s=1}^t \frac{1-\phi^{t-s+1}}{1-\phi}u_s \text{ the random component of } p_{tj}^m \tag{10}$$

$$m_{tj}^0 \equiv m(\mu_{tj}) \text{ the miles driven under the expected gas price at time t, and} \tag{11}$$

$$\sigma_{\varepsilon tj}^2 \equiv Var(\varepsilon_{tj}) = \frac{\sigma_u^2}{f_j^2}\sum_{s=1}^t \left(\frac{1-\phi^{t-s+1}}{1-\phi}\right)^2 \text{ the variance in the price per mile of car j at time t.} \tag{12}$$

Given these definitions, we can write the expected operating costs for car j at time t as:

$$E(c_{tj}) = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} m(\mu_{tj} + \varepsilon_{tj})(\mu_{tj} + \varepsilon_{tj})dF(\varepsilon_{tj})$$

$$= m_{tj}^0\mu_{tj} - \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \lambda\varepsilon_{tj}(\mu_{tj} + \varepsilon_{tj})dF(\varepsilon_{tj})$$

$$= m_{tj}^0\mu_{tj} - \lambda\sigma_{\varepsilon tj}^2$$

$$= m_{tj}^0\mu_{tj} - \frac{\lambda\sigma_u^2}{f_j^2}\sum_{s=1}^t \left(\frac{1-\phi^{t-s+1}}{1-\phi}\right)^2 \tag{13}$$

Interpreting the result in (13), the first term is the consumer's gasoline expenditures if gas prices evolve to their expected level by time t. The second term is the option value to the consumer from being able to adjust her driving patterns if gas prices evolve to a level either above or below the expected trajectory. Note that with downward sloping demand for driving, the option value is positive (reduces expected gasoline costs), and is increasing in the sensitivity of the demand for miles driven to gas prices, λ, the time between the present and the expected date when the gas is purchased, t, and the volatility of gasoline prices, σ_u^2; while the option value is decreasing in the fuel efficiency of the vehicle, f_j.

These results are highly intuitive. Consumers derive value from being flexible in response to perturbations in mileage costs. The consumer is acquiring real options on the price of gasoline. For each future period the consumer has both a put and a call on gasoline with the strike price equal to her mean expected price of gasoline in that future

period. If the price is below the expected price, she exercises the put and uses the profits to purchase more gasoline and other products in her optimal consumption bundle. Likewise if the price of gasoline is higher than expected, she just exercises the call and again the proceeds are allocated to other goods in her optimal consumption bundle.

The literature on real options (e.g., Dixit and Pindyck 1994) tells us that the value of the option increases with the variance in the price of the underlying commodity, gasoline in this case. This result simply reflects that a greater variance in the price of gasoline increases the likelihood that the consumer will make an adjustment in her driving patterns. In addition, a person whose mileage demand is more sensitive price fluctuations is acquiring more options (or an option on more gasoline) than a person who is less price sensitive. A person also acquires more options when she purchases a less fuel efficient car, since she uses more gas to drive. Finally option value increases as the time to expiration increases, as gas prices have more time to deviate from the mean, making larger deviations more likely and the payoffs to consumers of adjusting their driving patterns increase.

If the car has life T and the consumer has discount factor δ, we get an expression for the expected gasoline costs of car j over its life, which depends only on data and parameters which are known to the consumer at the time of the new car purchase.

$$E(c_j) = \sum_{t=1}^{T} \delta^t E(c_{tj})$$
. (14)

Nonetheless, our derivation of $E(c_j)$ leaves some open questions as to how operating costs can be represented in the data. Specifically,

How do we know the life of the car?

What is the evolution of gas prices?

What is the model for mileage demand?

What is the consumer's discount factor?

We proceed to address each of these questions in turn.

IA.3) Operating Costs – Car Life

The model for operating costs developed above takes as given the number of periods that the consumer will operate the car. Casual observation of the auto market suggests that there is significant variation both in the durability of different car models and the driving intensity of consumers. Consequently, choosing a single number for car life, T, is problematic.

In addressing these problem it is helpful to think of the car j as having a residual or salvage value $r_j(t)$ which represents the secondary market price for the vehicle t periods after the original purchase. Consider two cars with identical attributes on every dimension with the exception of durability. While these cars are similar in many respects, they would have different secondary market valuations at every time period. Nonetheless if variation in durability were the only source of variation in car life, then our problem is handled relatively well by the inclusion of the unobserved component , ξ_j, in our utility specification. We could make an estimate of average car life, T, informed by available industry publications. Then for cars j and k with different durability and all other attributes the same, the difference in utility after T periods is:

$$V(\delta^T(r_j(T) - r_k(T))) = \xi_j - \xi_k , \qquad\qquad (15)$$

where V(.) is the consumer's utility function for current wealth. Since our estimation algorithm explicitly calculates ξ_j, we can be relatively comfortable that our model can account for differences in durability across cars. Moreover, since it seems reasonable that differences in durability across cars persist over time, we can estimate the model relatively well over a range of T's and still get consistent parameter estimates. For example, if a Toyota and a Honda are identical in every aspect except that the Toyota is more durable than the Honda, then it does not seem reasonable that the Toyota would have a greater market value than the Honda after 5 years but a lower market value after 6 years. If the Toyota is more durable, it will always have a higher value in the secondary market in every time period. Thus the ranking across cars of the durability component of utility is invariant to the choice of car life, T. Since utility is an ordinal measure, if the variation in durability is the only source of variation in T, then our choice of T can be somewhat arbitrary. Moreover, operating costs at the end of a car's life are discounted more heavily, so for relatively large T, the marginal impact of changing our modeling assumption regarding car life is small. Obviously, we would prefer to have an accurate a choice of T so as to apportion utility properly between unobserved vehicle heterogeneity and observed differences in fuel efficiency; however, an elaborate calculation of vehicle life does not appear to be required.

However, as mentioned above, cars also last different lengths of time due to variation in driving intensity by consumers. A car that is driven more intensively will wear out faster. Therefore, two cars, even if they are the identical model make and year, will have different secondary market values at any point in time t, as a result of heterogeneity in consumer driving intensity. Thus, the car life, which we now denote T_{ij} is specific to each consumer car combination.

To address the issue of heterogeneity in driving intensity, we assume that the residual value of the car is a function of the cumulative number of miles driven denoted as M. If two cars have both been driven M miles then we assume that the difference in residual value is solely a function of difference in the durability of the cars not a difference in

consumer driving intensity. While one could argue that a car which has been driven more intensively would have a lower residual value at a given M, since it has been put through more mechanical stress. One could also argue the reverse that the heavily utilized car is worth more at a given M, since it is by definition it is newer (closer to its original purchase date). Thus, for a given M, the effect is ambiguous. Moreover, from a practical matter, casual observation of the used car market suggests that one of the most important components of car value is odometer reading. We conclude that any effects from the simplifying assumption that a vehicle's residual value is solely a function of M are second order and are more than outweighed by the analytical tractability this assumption provides.

Given this assumption we can just replace T in equation (15) by M and get:

$$V(\delta^t (r_j(M) - r_k(M))) = \xi_j - \xi_k ,$$
(16)

Now using the same argument that we used above to conclude that the choice of T can be made in ad hoc but educated manner, we can choose a reasonably well-informed choice about the useful life of a vehicle in terms of miles, $\overline{M}$, without worrying about biasing our results. Given an estimate of $\overline{M}$, we can calculate life of car j for consumer i, T_{ij} as:

$$T_{ij} : \sum_{t}^{T_{ij}} E(m_i(p_{tj}^m)) = \overline{M}$$
(17)

where the subscript i on the mileage demand function, $m_i(p_t^m)$, represents the heterogeneity in driving patterns across consumers. Note that the calculation of T_{ij} in (17) allows for consumers' choice of car does to affect mileage demand. For example, it is possible for a consumer have a driving plan which depends on whether she buys a Cadillac or a Toyota. However, in addition to complicating the model, this added flexibility does not appear to reflect actual behavior. We believe that people plan to drive a fixed number of miles based on factors such as the current state of gas prices or their major life circumstances such as job location or family size. This fixed number of miles does not depend on the choice of car. In a specific example, Petrin(2002) finds that family size increases the likelihood that a household purchases a minivan. While this study does not show nor does it claim to show causation, upon introspection, it does seem reasonable that people have kids and then buy a minivan as opposed to the opposite sequence of buying a minivan and then deciding to have children. Thus for purposes of our analysis, we assume that consumers have an anchor for mileage

demand which is fixed for each car model for each period under the expected gasoline price trajectory, i.e.,

$$E(m_i(p_{tj}^m)) = m_i(\mu_{tj}) = m_i^0, \forall t \text{ and } j.$$

(18)

Under this assumption our expression for T_{ij} simplifies to:

$$T_i = \overline{M} / m_i^0$$

(19)

Thus to account for variation in car life as a result of heterogeneity in driving intensity, we simply need a measure for mileage life and a distribution of miles driven in the new car buying population. We use the Department of Energy Regional Transportation Energy Consumption Survey (RTECS) which records exact odometer readings and estimated gasoline expenditures for a panel of approximately 6,000 vehicles every 3 years providing the data required to calculate this distribution. For purposes of this study, we use data from the 1991 survey, which is the survey year for which has public data availability which is roughly coincident with our data from the automobile market. Note the data does not distinguish whether the car is a new car (although model year is available), nor whether the vehicle is currently being driven by its original owner.

The sample consists of 5,879 vehicles, of which 5,275 or 90% had been driven 2,000 or more miles during the year. We discarded observations where the vehicle was driven less than 2,000 miles under the assumption that 2,000 miles is the minimum driving need for someone to be in the market for a new car. This data set has mileage driven ranging from 2,003 to 38,792 miles, with a mean of 10,206 miles and a median of 9,218 miles. The empirical cumulative distribution function for this sample is shown in figure 1. However, given that we are interested in the new car market, we adjust this distribution to reflect the fact that people who drive more will be in the market for a new car more frequently. For example, someone who drives their car 30,000 miles per year will be in the market 3 times as often as someone who drives 10,000 miles per year. Thus, for each observation we compute the ratio of actual miles driven to the mean miles driven which reflects the relative frequency with which a consumer with the observed driving pattern will enter the new car market. We then multiply the observed frequency (1/5,275) in the data by the adjustment ratio calculated above to yield our estimate for the distribution of miles driven by new car buyers. The empirical cumulative distribution function for the sample after adjustment is also shown in figure I.1. Note that after adjustment, the expected value of miles driven for new car buyers is now 13,612. By comparison, the RTECS summary report finds that the mean annualized miles driven in 1991 for cars in model years 1991 and 1992 was 14,000, so we can feel comfortable that the adjustments to the empirical distribution of miles driven, m_i^0, are

reasonably accurate. We will use this distribution in our estimation algorithm discussed in section B of this chapter.

Finally, to estimate the mileage life, $\overline{M}$, the RTECS data indicate that a typical car life is 80,000-100,000 miles. In addition, in an interview with a used car sales manager in the San Francisco bay area, we found that used cars decline rapidly in value once the odometer reading reaches 85,000-90,000. Finally, as noted above, the specific choice of $\overline{M}$ is not critical to obtaining consistent parameter estimates. For this reason, we adopt a base case $\overline{M}$ of 90,000 miles and subject this assumption to a robustness check when we report results in section C of this chapter.

IA.4) Operating Costs – Gas Prices

To calculate gasoline costs over the life of a car she is considering for purchase, a consumer needs a forecast of future gas prices. We start from the assumption that consumers develop forecasts of future prices based on past prices. Given this assumption, we then adopt the Autoregressive Integrated Moving Average (ARIMA) modeling methodology made popular by Box and Jenkins (1976). For simplicity, we assume that there are no moving average components. An important issue in developing an ARIMA model is whether the process is integrated, i.e., should the data be modeled in levels or in differences. Therefore, a unit root test[3] was conducted on the retail gasoline price data obtained from the Department of Energy. The data produce a test statistic = $DF\ t$ -2.64 which does not exceed in absolute value the 5% critical value of -2.95 for the Dickey-Fuller distribution (Hamilton 1994). Based on this test, we cannot reject the null hypothesis of a unit root and that the process is integrated[4].

We also test the alternative specifications for the autoregressive component by adding additional lags in the difference terms, and do not find any significant coefficients for the difference terms beyond one lag while the unit root test still cannot be rejected. Thus, we find that when the gas price data are differenced, we can parsimoniously represent the path of the change in gas prices with an autoregressive process of order one.

$$\Delta p_t^g = \mu^g + \phi \Delta p_{t-1}^g + u_t \qquad u_t \sim WN(\sigma_u^2) \qquad \qquad (20)$$

The model in (20), also known as an ARIMA(1,1,0), is then fitted to our data using OLS to produce the estimates of μ^g, ϕ, and σ_u^2 reported in Table I.1. Since the estimate for μ^g

---

[3] Note that while the Dickey-Fuller test calculates t statistics in the standard way, the distribution of the statistic (sometimes referred to as the Dickey-Fuller distribution) is not the t distribution and does not converge to a standard normal.

[4] Note that unit root tests typically have low power and it can be difficult to reject the null hypothesis of a unit root (see Stock 1994). Nonetheless, we proceed with the integrated specification, since we cannot reject the null and because we believe it to a better model for consumer expectation formation.

is not significantly different from zero, we drop it from our model of expected operating costs. Finally, to check that the disturbance terms are truly white noise with no remaining serial correlation, we calculate sample autocorrelations for the fitted residuals of the model in (20) for up to 5 lags (see Table 1). Calculating the Box-Pierce Q statistic based on the sum of the squared sample autocorrelations yields a Q = .07, which is far below the critical value of $\chi^2_{.05,5} = 11.07$. Therefore, we are confident that the gas price model developed here is free of serial correlation.

Based on the model in (20) and the estimates produced in Table I.1, we can now forecast future gas prices as if we were consumers trying to forecast the path of gas prices over the life of their cars. Specifically, under our model the j period ahead forecast for a consumer purchasing a car at time t is given by:

$$p^g_{t+j} = p^g_t + \frac{\phi(1-\phi^j)}{1-\phi}\Delta p^g_t \tag{21}$$

While it does not seem likely that consumers actually use a sophisticated model such as the ARIMA model developed above, we believe that the model above captures the essence of the forward-looking consumer. For example, consumers might use a simple heuristic such as, "I expect current trends to continue". In addition to its desirable statistical properties the integrated model developed here has the behavioral implication (provided that ϕ is positive) that consumers expect shocks in gas prices to persist into the future. If gas prices are high today they are expected to be high in the future, and if gas prices went up last period they will go up next period as well.

IA.5) Operating Costs – Elasticity of Gasoline Demand

Preliminary investigation of existing models which estimate the price elasticity of gasoline indicates developing a careful model from scratch represents a significant research project in its own right.

Nonetheless, there is no shortage of prior research in this area. Hausman and Newey (1995) estimate a partial linear model on household level data and derive a price elasticity estimate of roughly -.8 and an income elasticity of .4. Yatchew and No (2001) conduct a similar study on a dataset of Canadian households including additional demographic variables in their model and find an income elasticity of roughly .3 and a price elasticity of approximately -.9. Finally, there are a plethora of studies of gasoline demand based on aggregate data that are reviewed by Espey(1996). Across 70 studies, price elasticity estimates range from -.02 to -1.5 with a mean of -0.53.

Note many of the studies referenced above do not hold the stock of vehicles constant in estimating gasoline elasticity. Clearly, consumers can reduce gasoline consumption in response to a price increase by a) driving an existing vehicle less or b) buying a more fuel efficient vehicle. In this study we are interested in the former of these responses or the short-run elasticity of gasoline. Espey in her analysis indicates that roughly 75% of

the response in gasoline demand to price changes takes place within one year of the change. Thus in developing a workable assumption for our model, we settle on a somewhat ad hoc base case estimate of -.6 which reflects the results from some of the more careful studies in this area (e.g., Hausman and Newey) coupled with the meta-analysis of Espey. In addition, we conduct a series of robustness checks (discussed in section C of this chapter), over the range from 0 to -1.2.

IA.6) Operating Costs – Consumer Discount Rates

For purposes of this paper, we choose a real consumer discount rate of 5% for our base case analysis. As with other assumptions made in developing model, we test the sensitivity of our findings to changes in this assumption.

IA.7) Modeling Endogeneity in Car Price

Consistent with most of the heterogeneous products demand literature, the only component of consumer utility we consider to be endogenous is price. This endogeneity arises there are unobserved (by the econometrician) attributes (*e.g.,* crash test ratings) which will likely be correlated with price. The primary reason for this correlation is that car manufacturers consider unobserved vehicle attributes when they set price. For example, if a high crash test rating is achieved by adding extra air bags then the cost of these extra air bags will likely be reflected in the price charged to the consumer.

For convenience we repeat the consumer utility specification from above

$$U_{ij} = -\alpha_i \ (p_j^k - \gamma_i \ E(c_{ij})) + \mathbf{x_j'}\boldsymbol{\beta}_i + \xi_j + \varepsilon_{ij}$$

(22)

In this model, the role of ξ is to pick up unobserved vehicle attributes. If we do not account for these omitted determinants of vehicle utility, we will obtain inconsistent estimates for the effect of price on product choice, fatally limiting our ability to determine if the effect of price is equal to the effect of operating costs. In order to achieve consistent estimates for the parameters of interest, we need to instrument for price, with a set of variables which is uncorrelated with ξ.

In determining an instrument strategy the following straightforward identity provides intuition:

$$p_j^k = mc_j + (p_j^k - mc_j)$$

(23)

where p_j^k is the price of car j and mc_j is the marginal cost of car j. Therefore, the price is just the sum of the marginal cost and the mark-up. Although this result is trivial, it

suggests two approaches for finding instruments which are correlated with price but not with ξ. The first approach commonly employed is to find data which serves as a cost shifter (e.g., steel or rubber prices) which is likely to be correlated with marginal costs (see for example Villas-Boas and Winer 1999). The second approach is to find data which is correlated with mark-ups. This second approach was developed by Berry et. al. (1995) and is the method we employ in this paper.

In developing instruments which are correlated with mark-ups, we exploit the basic principle of oligopoly pricing that goods which face close substitutes have low mark-ups, while goods without nearby competition will have higher mark-ups. Additionally, since the automobile industry is characterized by multi-product firms, the effect of substitute goods on mark-ups is a function of whether the substitute product is produced by the same firm or a rival. Consequently, the essence of our instrument strategy is to develop measures of the level of both inter- and intra-firm competition for each car.

To form these measures we take the set of exogenous attributes $\mathbf{x}_j$ for each car, the sum of exogenous attributes across all cars produced by the same firm, and the sum of exogenous attributes across all cars produced by a competitor firm. The complete vector of instruments for each car model is thus:

$$\mathbf{z}_j = \{\mathbf{x_j}, \sum_{k \in \Im} \mathbf{x_k}, \sum_{k \notin \Im} \mathbf{x_k}\} \tag{24}$$

where $\Im$ is the set of cars manufactured by the parent firm. Given this set of instruments, we make the natural assumption that the instrument vector is not correlated with the unobservable product characteristics.

$$E(\mathbf{z_j}\xi_j) = \mathbf{0} \ \forall j \tag{25}$$

The population moment assumption in (25) is used to drive a Generalized Method of Moments (GMM) procedure, which is used to estimate the model parameters and is described in more detail in the following section.

IB. Estimation

As discussed above, we specify a random coefficients model for utility and demand which we estimate based on from aggregate data on market shares, prices, and product attributes and supplement with information from the RTECS data on the distribution of

consumer driving patterns. The basic estimation procedure follows closely that used by Berry et. al. (1995). Again we start from the utility specification,

$$U_{ij} = -\alpha_i \, (p_j^k - \gamma_i \, E(c_{ij})) + \mathbf{x_j'}\boldsymbol{\beta}_i + \xi_j + \varepsilon_{ij} \tag{26}$$

In order to represent consumer heterogeneity, distributional assumptions and/or estimates are made regarding all parameters and data which vary by consumer. For the model in (28) we use,

$$\varepsilon_{ij} \sim iid \; Extreme \;\; Value$$

$$\left(\alpha_i^k, \alpha_i^c, \boldsymbol{\beta_i}\right)' \sim N\!\left(\left(\tilde{\alpha}^k, \tilde{\alpha}^c, \tilde{\boldsymbol{\beta}}\right)', \Omega\right) \quad \Omega = \begin{pmatrix} \sigma_\alpha^2 & 0 & \mathbf{0} \\ 0 & \sigma_\gamma^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \sigma_\beta^2 \mathbf{I} \end{pmatrix} \tag{27}$$

$$m_i^0 \sim RTECS \; empirical \; distribution \; \text{(see figure 1 for cdf)}$$

We also assume that ε_{ij}, $\left(\alpha_i^k, \alpha_i^c, \boldsymbol{\beta_i}\right)'$, and m_i^0 are all independent of each other. Our challenge is the estimate the parameters $\boldsymbol{\theta} = \left(\tilde{\alpha}^k, \tilde{\alpha}^c, \tilde{\boldsymbol{\beta}}, vec(\Omega)\right)'$ which would allow us to completely characterize automobile demand for the specification we are studying. Combining (26) and (27), we can decompose the utility function into a portion, denoted δ_j, where the parameters enter linearly into utility and a portion, denoted η_{ij}, where the parameters enter in a non-linear manner.

$$U_{ij} = \delta_j + \eta_{ij}$$

$$\delta_j = -\tilde{\alpha}^k p_j^k + \mathbf{x_j}'\tilde{\boldsymbol{\beta}} + \xi_j \tag{28}$$

$$\eta_{ij} = -\sigma_k v_{ki} p_j^k - (\tilde{\gamma} + \sigma_c) v_{ci} E(c_{ij}(m_i^0)) + \sum_{m=1}^{M} x_{jm} \sigma_{\beta m} v_{Bmi} + \varepsilon_{ij}$$

where $\mathbf{v} = (v_k, v_c, v_\beta)$ is a standard normal random vector. We see from (28) that the parameters $\tilde{\alpha}^k$ and $\tilde{\boldsymbol{\beta}}$ enter through the linear portion δ_j and the remaining parameters which we will denote as $\boldsymbol{\sigma}$ enter non-linearly through η_{ij}. Consumer i will purchase product j, if $U_{ij} > U_{ik} \forall k \neq j$. Conditional on $\boldsymbol{\delta} = (\delta_1, ..., \delta_J)$, we can define the set of values for η_{ij} that results in the choice of car j as $A_j(\boldsymbol{\delta}) = \{\boldsymbol{\eta_i} \mid \delta_j + v_{ij} > \delta_k + v_{ik}, \forall k \neq j\}$. Therefore, the probability that consumer i purchases product j is the probability that $\boldsymbol{\eta_i}$ falls in the set A_j, and the market share for product j can then be calculated as

$$\hat{s}_j(\boldsymbol{\delta}, \boldsymbol{\eta}) = \int_{A_j(\boldsymbol{\delta})} f(\boldsymbol{\eta}) d\boldsymbol{\eta},$$

(29)

Given the size of the market, quantities can be calculated based on the market shares calculated above. Note that we refer to the model market share as $\hat{s}_j$ to denote these values are not observed and are predicted by the model.

The integral in (29) cannot be calculated analytically and even numerical methods are problematic. The problem in this paper is compounded by the use of a non-parametric empirical distribution for consumer miles driven, m_i^0. Nonetheless, given an initial guess at the values of the parameters, we proceed to approximate the integral in (29) through a simulated set of consumers[5] drawn from the distributions specified in (27). We then apply the popular smooth simulator approach (see Nevo 1998, or Arias and Cox 1999 for a description) to calculate $\hat{\mathbf{s}} = (\hat{s}_1, ..., \hat{s}_J)$.

Note that thus far we have imposed no restrictions that the predicted market shares equal to their true values, i.e., $\hat{\mathbf{s}} = \mathbf{s}$. However, at the true values for the parameters and the distribution of $\boldsymbol{\eta}$, the correspondence between predicted and actual market shares should be exact. Berry (1994) shows that under mild regularity conditions on the distribution of $\boldsymbol{\eta}$, for any set of non-linear parameters $\boldsymbol{\sigma}$, there exists a unique vector $\boldsymbol{\delta}$, such that $\hat{\mathbf{s}} = \mathbf{s}$. Thus imposing the restriction that predicted market share equal observed market share allows us to invert out $\boldsymbol{\delta}$ for any value at the non-linear parameters $\boldsymbol{\sigma}$.

---

[5] We use a random draw of 500 consumers for all estimation routines

$$\mathbf{s} = \hat{\mathbf{s}}(\boldsymbol{\delta}, \boldsymbol{\sigma}) \Rightarrow$$

$$\boldsymbol{\delta}(\mathbf{s}, \boldsymbol{\sigma}) = \hat{\mathbf{s}}^{-1}(\mathbf{s}, \boldsymbol{\sigma}) \tag{30}$$

Once the linear component of utility is inverted out, we can then proceed to calculate values for the parameters $\tilde{\alpha}^k$ and $\tilde{\boldsymbol{\beta}}$ using standard linear regression techniques such as two-stage least squares (2SLS). More importantly, with $\tilde{\alpha}^k$ and $\tilde{\boldsymbol{\beta}}$ in hand, we can now calculate the vector of unobserved utility $\boldsymbol{\xi} = (\xi_1, ..., \xi_J)$. Interacting the unobservables with the matrix of instruments $\mathbf{Z}$, described in (24), we can then form the empirical analog to the moment conditions (25) as

$$\mathbf{g} = \mathbf{Z}' \boldsymbol{\xi} = \mathbf{0}. \tag{31}$$

Since we have more instruments than parameters, it will not be possible to satisfy (31) exactly, so we minimize a quadratic form in $\mathbf{g}$ and our parameter estimates are those which achieve this minimum.

$$\hat{\boldsymbol{\theta}} = (\tilde{\alpha}^k, \tilde{\boldsymbol{\beta}}, \boldsymbol{\sigma}) = \arg\min(\mathbf{g}'\mathbf{W}\mathbf{g}) \tag{32}$$

where $\mathbf{W}$ is a positive semi-definite weighting matrix. To solve for $\hat{\boldsymbol{\theta}}$, we can iterate over the process described above until our objective function is minimized. We summarize the process and provide an overview of the estimation technique by listing the major steps of our procedure below:

1. Start with an initial guess of the non-linear parameter vector $\boldsymbol{\sigma}$.
2. Calculate the predicted market share function $\hat{\mathbf{s}}(\boldsymbol{\delta}, \boldsymbol{\eta})$ in (31) using the smooth simulator.
3. Find the value for $\boldsymbol{\delta}$, such that $\hat{\mathbf{s}} = \mathbf{s}$.
4. Calculate estimates of the linear parameters and the unobservable $\boldsymbol{\xi}$.
5. Calculate the value for the objective function $\mathbf{g}'\mathbf{W}\mathbf{g}$.
6. Update the guess for $\boldsymbol{\sigma}$ and go back to step 2.
7. Stop when the objective function is minimized.

To choose iterations on the values **σ** and to determine when the objective function has reached a minimum we employ the Nelder-Mead (1965) simplex direct search algorithm.

Using an initial weighting matrix $\mathbf{W} = (\mathbf{Z'Z})^{-1}$ (the homoskedastic case), we run the procedure outlined above to obtain an initial set of parameter estimates. These initial estimates are used to estimate the covariance of the moment vector **g** to account for heteroskadasticity, which we denote as $\hat{\mathbf{V}}$. We then run the procedure again using a weighting matrix $\mathbf{W} = \hat{\mathbf{V}}^{-1}$ to obtain our final parameter estimates. Hansen (1982) shows that for GMM estimation procedures such as the employed here, the two-step approach to estimation we use will produce parameter estimates which are asymptotically efficient.

IC. Data and Results

IC.1) Data

The data on gas prices and driving patterns is discussed in Section A above, so we confine the discussion in this section to data on automobiles. The automobile data used for this study is identical to that used by BLP in their 1995 paper. The data on car characteristics comes from the *Automotive News Market Data Book*. Data are available for the number of doors, number of cylinders, length, width, weight, horsepower, wheelbase, and EPA miles per gallon rating. In addition dummy variables are available for whether the car has air conditioning, front wheel drive, automatic transmission and power steering as standard equipment. Price data in ($000s) is based on list prices, which is subject to some measurement error given that most transactions in the US car market are negotiated.

Data are available for virtually every model of car marketed in the US for the twenty year period 1971-1990. In this study, a model/year is treated as a single observation, with a total of 2,217 observations. A summary table of descriptive statistics for some of the variables used in our specification is shown in Table I.2 where all price and cost data are expressed in real 1983 dollars.

Inspection of Table I.2 shows that new car prices went up significantly over the study period from an average of $7,868 in 1971 to $10,337 in 1990. Despite the 30% increase in real prices volume does not decline fluctuating between 7 and 11 million units per year. These fluctuations seem to be driven largely by business cycles. For example, the lowest volume year is 1982 with 6.8 million which coincides with a recession. The persistence of demand in the face of real price increases suggests product quality is increasing over the period. For example, the percentage of cars sold with air conditioning standard increases steadily from 0% in 1971 to 31% in 1990.

In terms of gasoline costs, we see that the industry responded to the oil price shocks of the late 70's and early 80's with a substantial improvement in fuel economy. The average fuel economy of a new car sold in 1971 was 16.6 miles per gallon (mpg) and

increased substantially to reach a peak of 26.0 mpg in 1983. As gas prices tapered off in the late in the late 80's, industry fuel economy declines reaching 22.7 mpg by 1990. However, despite the fact that by the late 80's real gas prices had returned to levels at or even below the experience of the early 70's, new cars remained significantly more efficient. This trend manifests itself in the average cost per mile which declined from 5.4 cents per mile in 1971 to 3.5 cents per mile by 1990 after reaching a high 6.8 cents per mile in 1979.

Overall the preliminary analysis of the descriptive statistics in Table I.2 provides initial evidence that consumers are very responsive to changes in operating costs. Moreover, the significant variation in both prices and operating costs over the study period should prove sufficient for us to identify the effects price v operating costs using the econometric model we have outlined above.

IC.2) Results

Before moving to the results of the full-model, we show some results from a simpler model without random coefficients for purposes of gaining initial intuition and, demonstrating the importance of a detailed model for operating costs, and testing our instrument strategy.

Without random coefficients and a detailed model of operating costs, estimation simplifies to a least squares regression of where s_o is the market share of the outside good. We estimate this model using OLS and 2SLS using the instruments as described in section IA.4. In addition to price, we include horsepower, car size, a dummy for air conditioning as standard, and the cost of driving the car 100 miles, based on current gas prices as regressors. The results from these regressions are shown in Table I.3.

Comparing the coefficient on price from the OLS to 2SLS case, we see an increase in the effect of price on demand when price is modeled as an endogenous regressor using the instrumental variables strategy outlined above. This result is consistent with the well documented finding that OLS underestimates price elasticity. To gain a sense for this effect, we compute the sales weighted average elasticity across the car models marketed in 1990. We find that under OLS the elasticity is only .88 increasing to 1.42 under 2SLS.

Clearly the OLS result is far too low given that in any differentiated products market, theory suggests an elasticity greater than 1. Moreover the 2SLS results produce an elasticity which is at the low end of a plausible range for this estimate. We hypothesize that the inability of the simple model without heterogeneity to properly account for substitution effects is biasing the 2SLS elasticity estimate downward as well, and that introducing the full random coefficients approach will increase this estimate.

The results also show that, as we expected, consumer utility is negatively affected by increases in gasoline cost. In an effort to make a unit-free comparison between the effect of price and the effect of operating costs, we compute elasticity measures for operating costs as well. These results indicate that consumers underweight operating costs relative to capital costs when they purchase a car. However, this model is highly simplified and does not account for expectations for future changes in gas prices nor does it account for the two sources of heterogeneity in consumer preferences for operating costs –

differences in tastes and differences in driving patterns. Moreover, we find it surprising that the relatively low operating cost elasticities estimated by the simple model could have generated the significant improvement in fuel efficiency across the industry. Firms were clearly responding to something. Thus we hesitate to come to any conclusion on the basis of these preliminary regressions.

Moving to the full model described in Section IA, we use the same set of explanatory variables with the exception that we now replace the cost of driving 100 miles with our detailed model of expected operating costs. Since we do not have a single explanatory variable for operating costs, but rather an empirical distribution we add miles per gallon, and its sum across the same firm and competing firms into the instrument matrix. Thus we keep the same number of instruments as in the 2SLS case, while using mpg as a proxy for operating costs.

The results for the base case are displayed in Table I.4. All of the parameters have the expected sign and with the exception of the constant, they are statistically significant. In particular, we find that horsepower, car size, and air conditioning all increase utility, while price and operating costs decrease utility. Examining further the coefficients on price and operating cost we see that the elasticity of both has increased substantially relative to the 2SLS model discussed above.

The average price coefficient estimate of 3.78 seems reasonable and is consistent with the results produced by BLP in their paper which uses the same data set. We also find that the consumers are relatively homogenous in how they respond to changes in price. The estimate for the standard deviation of the normally distributed price coefficient is .62 which implies that 95% of consumers have a price coefficient estimate across all models between 2.56 and 5.00. By comparison, we see that modal elasticity with respect to operating costs is 5.26. While this measure is slightly higher than the price effect, a Wald test of the null hypothesis that mean coefficients on price and operating costs are equal fails to reject at a 5% significance level (W=2.48< $\chi^2_{.05,1}$= 3.84).

However, unlike price effects which appear to be homogeneous across consumers, the effect of operating costs appears to vary widely in the population. A Wald test of the null hypothesis that the standard deviation parameters for price and operating costs are equal is rejected at a 5% level (W=37.84 < $\chi^2_{.05,1}$ = 3.84). Our estimate of the standard deviation parameter for operating costs of 4.51 implies a 95% of consumer average elasticities fall within the range of -3.58 to 14.1. The fact that some consumers appear to prefer higher operating costs raises a concern as to reasonableness of the results. However, only 12.2% of consumers have negative coefficients for operating costs under the parameter estimates generated by the model, which is not outside the bounds of plausibility.

Thus we find no evidence to support behavioral theories which suggest that consumers systematically underweight future events; however, we find significant evidence that some people do. In figure I.2, we plot the estimated cumulative distribution function for the difference between operating cost elasticity and price elasticity. Reading off the figure we see that 37% of the population weight prices more heavily than operating costs while 63% do not. Moreover, the degree of heterogeneity in the population is striking.

Fully 30% of new car buyers, have an absolute difference between the operating cost and price coefficient exceeding 5. Thus, a significant proportion of the population is not making a rational trade off between capital and operating costs, and we see biases in both directions.

We find the results of the model to be highly robust both to the assumptions discussed in Section IA and to the specification of other exogenous explanatory variables. In Table I.5, we show three sensitivities regarding assumptions. The first panel of the table repeats the results for the base case. The second panel shows results for the model with an assumed mileage life for a car of 110,000 miles up from 90,000 in the base case. The results are virtually unchanged from the base case, and all of the substantive conclusions we outlined above for the base case. As was hypothesized in the discussion of the model for vehicle life, the inclusion of an explicit unobservable term in our utility specification appears to picking up differences in durability across vehicles, and we are able to calculate consistent parameter estimates over a reasonable range of assumptions. Even an increase of more than 20% increase in vehicle life produces almost no change to our results.

The third panel of Table I.5 shows a set of results when the consumers short-run elasticity for gas prices is set to zero. Again all of the substantive conclusions from the base case still apply. It is also interesting to note that setting gas price elasticity to zero effectively sets the value of the consumer's option to change her driving patterns to zero. We see that this change has only a modest effect on the operating cost elasticity parameters. The mean elasticity declines to 4.95 from 5.26 in the base case. We find that the relatively low gas price elasticity results in low option values.

In addition, this result indicates that operating costs are affecting utility independently of the other variables in the model. In both sensitivities, operating costs change, and all of the parameter estimates related to attributes other than operating costs are virtually unchanged. Thus, it appears that there is a relatively fixed level of variation in vehicle choice that is explained by operating costs which is not explained by any other vehicle attributes. This effect can be seen by noting that when operating costs go up, the mean elasticity estimate goes down holding the total effect of operating costs relatively constant. We believe this is an ideal situation and provides an ideal research environment for policy experiments regarding operating costs such as determining the effect of a gasoline tax.

Finally, in panel 4 of Table I.5 we show a case where the consumer discount rate is reduced to zero from 5% in the base case. Once again all of the substantive conclusions remain the same. While we do see some shifts in the parameter values, they are in the expected direction. Reducing the discount rate increases operating costs, this in turn reduces the mean operating cost coefficient. While the mean operating cost coefficient is now slightly less than the mean price coefficient, a Wald test (W=.45< 21 , 05 .c =3.84) shows that the parameters are not significantly different from one another.

To test the robustness of the model to specification of other exogenous explanatory variables, we add variables for the number of doors, number of cylinders and a dummy variable for automatic transmission. The results are compared with the base case in

Table I.6. The coefficients on the number of doors and the presence of automatic transmission are both significantly different from zero, so there may be some merit in including them in the model. Nonetheless, the new specification does not change the main conclusion of this research that on average consumers weight operating costs and capital costs equally, but there is considerably more heterogeneity in the effect of operating costs.

ID. Summary and Directions Future Research

This paper represents a significant contribution to the understanding of whether consumers behave as if they are optimally trading off capital and operating costs when purchasing a durable good. In addition, by applying this research to the critical automotive and energy sectors, this research provides insight into one of the most pressing public policy questions facing the US today. Finally, our model of forward looking consumer behavior with respect to gas prices represents the first attempt to model consumer expectations in a choice model of the US automobile industry.

To reiterate our findings, we find no evidence to support behavioral economics theories that consumers systematically underweight the cost of future events in real market settings. However, we find significant evidence that large portions of the population are not making the trade-off optimally. Conservatively, at least 30% of the population is either drastically underweighting or overweighting operating costs when purchasing a new car.

This research has several limitations which provide a useful roadmap for future studies. First, in the model developed in this paper consumers are only forward looking as to their expectations regarding driving patterns and gasoline purchases. Ideally, the model could be extended to allow for dynamic vehicle replacement as well (see Chapter 3 of this dissertation for a start on this effort), creating a fully dynamic framework. In addition, the model developed in this paper makes strong assumptions about the functional form of the distribution of consumers in the population. By restricting consumer taste parameters for price and operating costs to be normally distributed, we are necessarily ruling out more flexible distributions of consumer behavior which may be present in the population. For example, there may be a cluster of consumers who place zero weight on operating costs and another cluster of consumers who place 2x the weight of price on operating costs. The model can be extended to allow for more flexible semi-parametric and non-parametric distributions.

Figure I.1
Comparison of the Cumulative Distribution Functions for Annual Miles Driven

All Drivers v. New Car Buyers
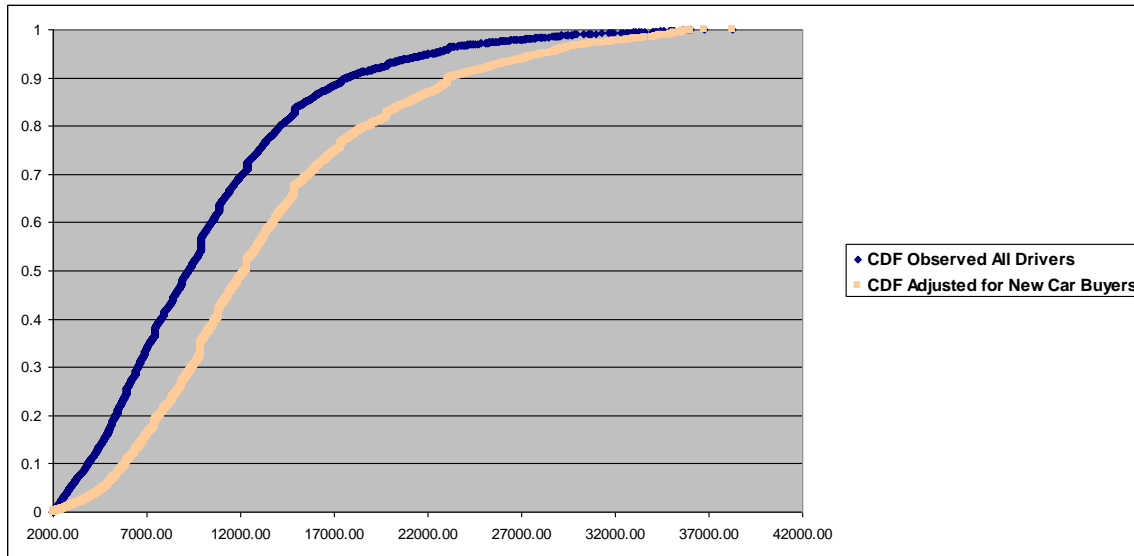
Source: Regional Transportation Energy Consumption Survey 1991.

Table I.1
Results from Fitting ARIMA (1,1,0) Model to Annual US Retail Gas Price Data 1960-1990

	Estimate	Standard Error
Constant	0.3012	2.7386
Lagged Delta Gas Price	0.3518	0.1707
Variance of Residuals	232.32	Cents sqared per gallon
Sample Autocorrelations of Residuals		
1	0.0789	
2	-0.1624	
3	-0.1574	
4	-0.0884	
5	0.0742	
Box Pierce Q Statistic	0.3434	
Critical Value (alpha=.05)	11.07	

Table I.2
Summary of US New Car Market Attributes
1971-1990

	Models	Quantity	Avg. Price	% Domestic	Avg. HP	Avg. Size	% AC	Avg. MPG	Avg. $/Mile	$/Gallon
1971	92	7,994,033	7,868	87%	165	1.50	0%	16.6	0.054	0.90
1972	89	8,777,411	7,979	89%	133	1.51	1%	16.2	0.053	0.86
1973	86	7,979,547	7,535	93%	127	1.53	2%	15.9	0.055	0.87
1974	72	7,568,571	7,506	89%	121	1.51	3%	15.7	0.069	1.08
1975	93	7,884,031	7,821	85%	117	1.48	5%	15.8	0.067	1.05
1976	99	9,244,776	7,787	88%	119	1.51	6%	17.6	0.059	1.04
1977	95	9,284,086	7,651	84%	112	1.47	3%	19.5	0.055	1.06
1978	95	9,447,225	7,645	85%	105	1.40	3%	19.8	0.052	1.03
1979	102	8,439,722	7,599	80%	99	1.34	5%	20.6	0.060	1.24
1980	103	7,371,410	7,718	77%	94	1.30	8%	22.1	0.068	1.51
1981	116	7,195,517	8,349	74%	93	1.29	9%	23.6	0.064	1.52
1982	110	6,808,210	8,831	71%	92	1.28	13%	24.4	0.055	1.34
1983	115	7,805,914	8,821	73%	94	1.28	13%	26.0	0.048	1.25
1984	113	9,710,453	8,870	78%	99	1.29	13%	24.7	0.047	1.17
1985	136	10,627,473	8,938	76%	100	1.26	14%	22.6	0.049	1.12
1986	130	10,888,255	9,382	73%	102	1.25	18%	24.2	0.035	0.85
1987	143	9,676,415	9,965	70%	108	1.25	23%	23.3	0.036	0.83
1988	150	10,061,673	10,069	72%	109	1.25	24%	23.3	0.034	0.80
1989	147	9,248,300	10,321	69%	113	1.26	29%	23.1	0.036	0.82
1990	131	8,695,428	10,337	68%	120	1.27	31%	22.7	0.035	0.80

Source: Automotive News Market Data Book; US Department of Energy
Averages and Percentages are sales weighted
Dollar figures are real 1983

Table I.3
Results from Preliminary Regressions without Random Coefficients and a Detailed Model for Operating Costs

	OLS		2SLS	
	Coefficient	Std. Error	Coefficient	Std. Error
Constant	-9.118	0.139	-8.667	0.166
HP	-0.116	0.075	0.336	0.117
Size	2.602	0.154	2.039	0.188
AC	-0.054	0.071	0.461	0.125
Price	-0.085	0.004	-0.137	0.011
$/Mile	-0.129	0.018	-0.085	0.020
1990 Avg. Price Elasticity	-0.884		-1.416	
1990 Avg $/Mile Elasticity	-0.4725		-0.3094	

Table I.4
Parameter Estimates for Base Case Model
With Random Coefficients and Detailed Model for Operating Costs

	Coefficient	Std. Error
Constant	-1.0438	0.6509
HP	0.8781	0.1358
Size	1.3365	0.2328
AC	1.163	0.1443
Price	-3.7842	0.393
Operating $	-5.2613	0.6378
Price SD	0.6202	0.3094
Operating $ SD	4.5099	0.5702

Figure I.2
Operating Cost Elasticity – Price Elasticity
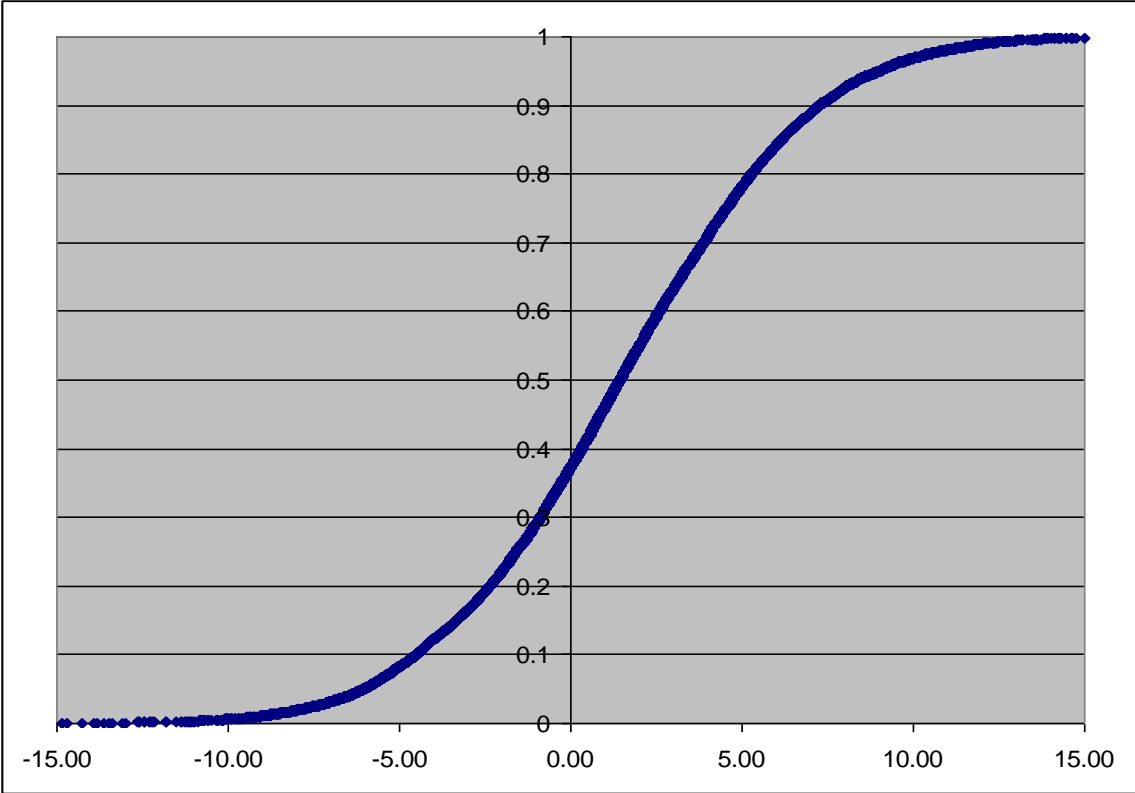**Estimated Cumulative Distribution Function**

Table I.5
Robustness Checks against Operating Cost Model Assumptions

	Base Case		Mile Life = 110,000		Gas P Elasticity = 0		Discount Rate = 0	
	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error
Constant	-1.0438	0.6509	-0.7807	0.6336	-1.0772	0.6825	-0.0484	0.7107
HP	0.8781	0.1358	0.9094	0.1353	0.8406	0.1274	0.8534	0.1271
Size	1.3365	0.2328	1.4187	0.2286	1.5359	0.259	1.9234	0.2243
AC	1.163	0.1443	1.1887	0.1434	1.1409	0.1387	1.1858	0.1366
Price	-3.7842	0.393	-3.9252	0.3843	-3.7621	0.5374	-4.8119	0.7457
Operating $	-5.2613	0.6378	-5.1323	0.595	-4.9545	0.5229	-4.2396	0.5737
Price SD	0.6202	0.3094	0.7411	0.2614	0.2762	0.5005	1.0066	0.381
Operating $ SD	4.5099	0.5702	4.3263	0.5077	3.8581	0.3522	3.1893	0.3302

Table I.6
Robustness Checks against Alternative Specification of Explanatory Variables

	Base Case		Add Doors, Cylinders, AT	
	Coefficient	Std. Error	Coefficient	Std. Error
Constant	-1.044	0.651	-1.519	0.760
HP	0.878	0.136	0.896	0.138
Size	1.337	0.233	1.136	0.281
AC	1.163	0.144	1.031	0.123
AT			0.435	0.095
Cylinders			-0.034	0.036
Doors			0.085	0.027
Price	-3.784	0.393	-3.635	0.668
Operating $	-5.261	0.638	-3.267	0.538
Price SD	0.620	0.309	0.282	0.566
Operating $ SD	4.510	0.570	2.804	0.363

Chapter I References

Arias, Carlos and Thomas L. Cox. 1999. "Maximum Simulated Likelihood: A Brief Introduction for Practitioners", University of Wisconsin, Madison Staff Paper No. 421.

Berry, Steven T. 1994. "Estimating Discrete Choice Models of Product Differentiation." Rand Journal of Economics 25 (July): 242-262.

Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica 63 (July): 841-90.

_____. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market" Journal of Political Economy 112 (February): 68-104.

Dickey, David A. and Wayne A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." Journal of the American Statistical Association 74: 427-431.

_____. 1984. "Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root." Econometrica 49 (July): 1057-1072.

Espey, Molley. 1996. "Explaining the Variation of Elasticity Estimates of Gasoline Demand in the United States: A Meta-Analysis." The Energy Journal 17 49-60.

Goldberg, Pinelopi Koujianou. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the US Automobile Industry." Econometrica 63 (July): 891-951.

_____. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the US." Journal of Industrial Economics, 46 (March): 1-33.

Hansen, Lars P. 1982."Large Sample Properties of Generalized Method of Moments Estimators." Econometrica 50(July): 1029-1054.

Hamilton, James D. 1994. *Time Series Analysis*. Princeton University Press, Princeton, NJ.

Hausman, Jerry and Whitney Newey. 1995. "Nonparametric Estimation of Exact Consumer Surplus and Deadweight Loss." Econometrica 63 (November): 1445-1476.

O'Donohue, Ted, and Matthew Rabin. 2001. "Choice and Procrastination." Quarterly Journal of Economics. 116 (February): 121-160.

Nevo, Aviv. 1998. "A Research Assistant's Guide to Random Coefficients Discrete Choice Models of Demand", NBER Technical Working Paper No. 221.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan" Journal of Political Economy 110 (August): 706-29.

Stock, James H. 1991. "Unit Roots and Trend Breaks," in Robert Engle and Daniel McFadden, eds. Handbook of Econometrics, Vol. 4. Amsterdam: North-Holland.

Yatchew, Adonis and Joungyeo Angela No. 2001. "Household Gasoline Demand in Canada." Econometrica 69 (November): 1697-1709.

Chapter II: Who Pays for Corporate Social Responsibility: Consumers or Shareholders? Evidence from the US Coffee Market[6]

Corporate social responsibility (CSR) programs are ubiquitous, yet it is unclear that these programs actually pay off for the firms that sponsor them. For example, consider Target's program to donate 1% of all retail sales to United Way local charities. Do consumers really want their money spent this way? Are consumers happily paying 1% higher prices or are they switching to a competitor which does not donate a portion of revenues to charity? If consumers are happy paying the higher prices, then they are making an indirect donation to the United Way by shopping at Target. However, if the consumers switch to a competitor then Target's shareholders are absorbing this cost. Who is making the donation in the end, Target's shareholders or customers? From a social planner's perspective, the point is largely moot[7], yet to the shareholders of Target and many other firms practicing CSR, the question is crucially important.

In this paper, we endeavor to study this question within the context of the coffee industry, an important and sizeable commodity market. In 2006, coffee was one of the most highly exported commodities in the developing world, second only to crude oil (Cailleba and Casteran 2009). In particular, we explore the impact of Fair-Trade (FT) certification on the retail coffee market. FT is a social and ethical movement that supports the ethical production, agribusiness trading, and consumption of goods in an attempt to defend producers who are pushed into poverty.   Coffee and other products can be FT certified by adhering to FT standards. Once certified, FT coffee is distinguished from non-FT by distinctive labeling visible to the consumer who is deciding which coffee product to select from the supermarket shelf.

Sales of FT certified products are growing rapidly. According to the Fair Trade Federation, sales of FT commodities have risen ~40% per year (since 2004) in North America and the Pacific Rim (Loureiro 2005). Given the conservative estimate of $500 million for the sales of FT-certified products in 2001 (Redfern and Snedker 2002), there is evidence to suggest that FT sales worldwide could reach as much as $10 billion by 2010. FT coffee provides an excellent exemplar of CSR in a market of significant and growing importance.

The analytical strategy for this paper is to first estimate the price premium commanded by FT coffee over non-FT coffees through ordinary least squares (OLS) and Fixed Effects hedonic price regressions. However, these tools do not allow us to disentangle the portion of the price premium which is due to supply considerations (i.e., FT certification costs) from the portion which is due to consumers' willingness to pay for FT coffee because they want to support socially responsible coffee production. To parse out willingness to pay from the overall price premium, we will develop a brand choice model for coffee which explicitly models brand market shares and prices as equilibrium

---

[6] This work is part of joint work with Sofia Villas-Boas and Sarah Wong.
[7] Assuming United Way charities is efficiently using the receipts from this program.

outcomes from the interaction between coffee supply and demand (see Berry, Levinsohn and Pakes 1995, Villas-Boas and Winer 1999 for methodological references). Using the results of the equilibrium brand choice model, we can answer the research questions of whether consumers are willing to pay for social welfare and in turn if CSR programs pay off for shareholders.

An existing body of related academic literature attempts to assess the effect of CSR on corporate performance. One of the earliest such studies by Moskowitz (1972) showed that the market performance of 14 "socially responsible" firms exceeded the performance of overall market indicators such as the Dow Jones Industrial Average or the NYSE composite stock index. However, the criteria for what constitutes socially responsible were not well documented and subject to considerable criticism and debate. Numerous alternative methodologies were proposed some of which supported the findings of Moskowitz (e.g., Vance 1975) and others which found no evidence of a link between CSR and firm performance(e.g., Corson and Steiner 1974).

In an attempt to create a more reliable measure for social responsibility Carroll (1979) developed a more rigorous survey instrument for corporate CEOs designed to elicit management focus in four areas of social responsibility. Aupperle et. al (1985) use Carroll's construct to measure CSR and found no relationship between social responsibility (as measured by Carroll's construct) and profitability. Continued research has continued to attempt to relate CSR programs to financial performance and the results are equivocal (see Stanwick and Stanwick 1998 for a review).

The difficulties encountered by many of the above researchers in coming up with a definitive answer to profit impact of CSR are threefold. First, as stated above there is much controversy about how to measure CSR. Second, even with standards for measuring CSR, there are many different types of CSR (e.g., diversity in hiring, fair labor practices, ecologically friendly products, community support, etc.). Finally, even if it is determined that investment in a specific CSR domain has an effect on profits, it is unclear why. Examples of how profits could be affected by CSR include an increase in consumer demand, an increase in employee morale, or preferential tax treatment due to CSR programs.

In an effort to address the difficulties above, Sen and Bhattacharya (2001) conduct a series of surveys and experiments with student subjects to assess consumer responsiveness to specific types of hypothetical CSR programs. Similarly, Mohr and Webb (2005) survey a national sample of adults to ascertain how CSR along the environmental and philanthropic domains affects company evaluation and purchase intent. Both of these studies study the effect of specific types of CSR on a specific stakeholder group (consumers), yet both studies are not based on actual marketplace behavior or programs nor do they link directly the programs to firm financial performance.

Also related to this paper is a body of work on the consumers' willingness to pay for food labeling (see McClusky and Loureiro 2003 for a review); however the labeling is not necessarily related to CSR. In addition, Cailleba, P. & Casteran (2009) conduct an extensive survey of the characteristics of French consumers of FT coffee. Perhaps the

work most closely related to our current effort is that of Maguire et al (2004), who conduct hedonic price regressions to determine the price premiums for organic baby food. Our study differs in two important ways: 1) we explore a FT not an organic food product – consumers do not necessarily purchase organic food especially for babies to promote social welfare; 2) we explicitly try to account for the portion of the price premium in FT coffee which is related to supply considerations v. demand considerations.

In summary, the present study attempts to overcome some of the difficulties encountered by researchers to date by 1) defining CSR specifically and narrowly (FT v Non-FT), 2) measuring actual consumer purchases of FT v non-FT products using scanner panel data, and 3) by relating FT to a change in market share and a change in consumer's willingness to pay. To our knowledge this paper is the first effort to use actual consumer preferences revealed in a market setting to assess the impact of a CSR program on sales and prices.

The remainder of this chapter is organized as follows. Section A provides additional background and institutional detail on the coffee industry and the FT movement. In Section B, we discuss the data used. Section C outlines the models used in this paper. Section D discusses the results and section E concludes.

IIA. Background of Coffee Industry and FT Movement

Coffee is practically a staple of the American diet. Its popularity stems from post-war America's love for the addicting and caffeine-rich black liquid that purportedly made its consumers feel more energetic and productive (Clark 2008). The United States is the largest coffee consuming country in the world with an overall market size of approximately $72 billion (Donnet 2010). Given the large size of the coffee market, it is unsurprising to find a vast variety of products. One can select from a multitude of different flavors, brands, cultivation methods, specialties, and other individual characteristics. According to Beaver, et al consumers do treat coffee as a highly differentiated product, like wine and cheese (Beaver et al 2006).

Out of this market, specialty coffee has developed into a growing and dynamic industry with a retail market of $12.26 billion in the United States alone (Donnet 2010). At first the specialty market began due to small roasters with taste and business savvy who started to differentiate their coffee blends to capture the demand of consumers who appreciated "aroma and flavour and the distinctive profiles of the various coffee origins and flavours" (Roseberry, 1996) (Donnet 2010). Now the specialty market has derivatives that extend beyond better quality and flavor. Fair Trade Certified coffee has emerged as an important market and a growing industry that strives for quality, but above all else, social change within agribusiness, and sustainable farming.

The Fair Trade movement is not the first to help marginalized producers. The idea of fair trade has been around since the early nineteenth century in niche markets. Abolitionists in the United States opened "free produce" stores that only sold items produced by non-slaves (Glickman 2004). Consumer activism was popular at the time

among abolitionists, non-abolitionists, and Christians (Cailleba and Casteran 2009). As a result a Dutch writer, E. Douwes Dekker, aka Multatuli, wrote about the activism. He published the now well-known *Max Havelaar*, which exposed the unfair labor conditions in the Dutch Indies. Though the World Wars prevented much growth of such organizations, the first "Fair Trade" shop was established in Puerto Rico in 1958 (Cailleba and Casteran 2009). Later, the *Oxfam UK* began selling crafts made by Chinese refugees. This was turned into the first Fair Trade Organization (EFTA, 2006a).

In 1988 the Max Havelaar label was introduced in the Netherlands. Not only was it the first Fair Trade label initiative, but with its targeting of large coffee companies it also marked the point in history when Fair Trade left small shops, and entered the "mainstream." The label's contract with companies said that the company could buy any percentage of their coffee on Fair Trade terms, and could in turn purchase the right to the seal to apply to their roasted coffee. Max Havelaar's strategy to target "profit-driven" roasters helped expose consumers to Fair Trade and marked the start of Fair Trade as a viable market concept. The label spread throughout Europe, and elsewhere taking on different names. For example, the organization TransFair USA pushed Fair Trade terms in coffee production and administrative practices. Importers and roasters were also certified but had to pay to be certified (the cost encompassed the certification and monitoring costs). Nonprofits who also certified Fair Trade labels altered how coffee was marketed. They made it possible for farmers to form cooperatives to market their own beans. With many small organizations there was a need for a governing body (Levi and Linton 2003).

In 1997, all the Fair Trade certification groups were united under one organization, the Fairtrade Labeling Organization (FLO) International, an organization devoted to standardizing the certification process for coffee among other Fair Trade products. FLO maintains a Fair Trade Register of producer groups approved to sell to the Fair Trade market (Levi and Linton 2003). In May 2001, it listed 363 groups from 22 countries. The majority are in the Western Hemisphere, but other countries are also represented, such as Tanzania, Uganda, Ethiopia, Indonesia, and Thailand. (Levi and Linton 2003). In addition, to indicating the growth in FT, this standardization under FLO, improves the integrity of our analysis of the impact of FT on prices and consumers. We can be assured that any coffee labeled as FT will have met the same FT standards.

IIB. DATA

The bulk of our dataset is derived from raw scanner data in 2005-2008 from a large grocery chain in California. Data on average weekly price and quantity sold for each of 499 coffee products was collected over a period of 178 weeks from 563 stores, yielding just over 50 million observations. We summarize this data by product and by market.

At the product level we observe, Brand name, Specialty indicator, Fair Trade indicator, Rainforest indicator, and mean price across all stores over the 178 week period. Prices are expressed in US dollars. The sizes of coffee products tracked range from 12 oz to 16

oz.[8] Summary statistics are provided in Tables II.1 and II.2.

Table II.1. FT v Non-FT Mean Price Comparison

	Obs	Mean Price	Std. Err.	95% Confidence Interval
Non-FT	463	6.30	.1178	[6.07, 6.53]
FT	35	9.64	.1937	[9.25, 10.04]
				t= 7.75

The mean price of non-FT coffee is $6.30 with a standard error of .1178. The mean price of FT coffee is higher at $9.64 with a standard error of .1937. The 95% confidence interval for the mean price of non-FT is [$6.07, $6.53], while the 95% confidence interval for FT is [$9.25, $10.04]. These values indicate that the raw mean price of FT coffee is higher than non-FT. This difference is statistically significant at the <1% level with a t-statistic of 7.75. We can see this difference as well as the spread in a histogram provided in Figure II.1.

Figure II.1.

Graphs by FT

---
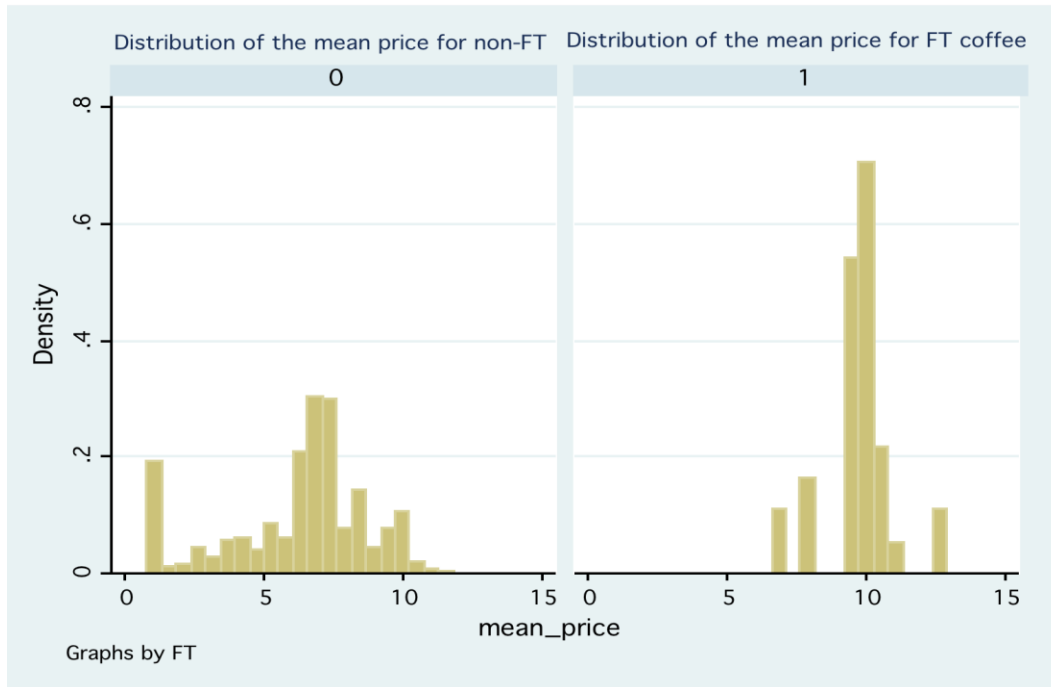
[8] We were unable to accurately find the average size of coffee across all products listed in the dataset.

Notice that both histograms give the appearance of following the normal distribution. Moreover, notice that there is less spread in the mean price for FT coffee (right graph) opposed to non-FT coffee (left graph). This may indicate that there is similar pricing within FT coffee, but may also just show the heterogeneity within the more numerous non-FT coffee observations.

Table II.2 shows the summary statistics for Specialty v. non-Specialty coffee. The mean price of Specialty coffee was $7.46 with a standard error of .1333. The mean price of non-Specialty coffee was lower at $5.07 with a standard error of .1689. The price difference between specialty and non-specialty is statistically significant at a <1% significance level with a t-statistic of 11.11. The data suggest that specialty coffee is more expensive than non-specialty coffee. Since both Specialty and FT coffee are, on average, more expensive than non-Specialty and non-FT coffee, it is likely that Specialty coffee is correlated to FT coffee and may need to be controlled for using multivariate analysis. In fact, upon inspection of the data, all FT coffees are also indicated as Specialty.

Table II.2 Specialty v Non Specialty Mean Price Comparison

Specialty	Obs	Mean Price	Std. Err.	95% Confidence Interval
Specialty	193	5.07	.1689	[4.74, 5.40]
Non-Specialty	305	7.46	.1333	[7.19, 7.72]
				t= 11.11

The market level data are aggregated into 41 (monthly) time periods across 616 clusters of 5 digit zip codes. A market is defined as monthly sales within all stores for a given zip code cluster. Within each market we observe the mean price and units sold for each product, along with the product level data described above. By defining markets as above we are able to observe equilibrium outcomes across 25,256 markets. Observing how these outcomes vary in response to changes in the product set available will permit identification of consumer's willingness to pay for FT labeling.

In addition to scanner panel data we collected data on monthly gasoline prices from 2005-2008 and for Brazilian wage rates in US dollars. Collecting these data enables us to identify supply shocks over time (across markets as defined above).

IIC. Modeling Approaches

As a first step in this analysis we want to determine if there is a price premium for FT coffee and if so how big is it? The data above strongly suggest that there is a premium, but we need to control for other factors such as brand differences and product quality

which could also be causing FT coffee to be priced higher than non-FT. To conduct this analysis we conduct OLS and a brand Fixed Effects hedonic price regressions.

The OLS model is shown below:

$$mean\_price_i = x_i'\beta + \epsilon_i \qquad\qquad (1)$$

Mean_price is the average price for a coffee product i measured in US\$ for sizes ranging from 12~16oz. The vector $x_i = (1\ FT_i\ Specialty_i)'$ represents a constant, a dummy variable FT_i denoting if product i is fair trade, and a dummy variable $Specialty_i$ denoting if product i is a specialty coffee. The error term ϵ_i is assumed to have mean zero and variance σ^2. A final crucial assumption of OLS is that the conditional mean of the error vector conditional on the X matrix is zero (i.e., $E(\epsilon|X) = 0$).

This assumption will be violated if there is any component of the error term which is correlated with any of the components of **X**. For example, suppose that consumers generally prefer a coffee with a high aroma index and are willing to pay more for it. However, aroma index is not observed by the researcher, but is observed by the consumer. Thus in our OLS specification above aroma index gets loaded into the error term. If the aroma index tends to be higher for specialty coffee products, violating our key OLS assumption (i.e., E(aroma index|specialty)>0, then some component of the consumers' willingness to pay for aroma will be misattributed to a price premium for specialty coffees. The example above illustrates the well-known omitted variable bias.

In an attempt to address this bias, we introduce brand fixed effects. We assume that unobservable product attributes (e.g. aroma) do not vary within a brand. If this holds then, we can eliminate any omitted variable bias by including a dummy variable for each brand in our dataset. The brand fixed effects specification is:

$$mean\_price_i = x_i'\beta + f_i'\gamma + \epsilon_i \qquad\qquad (2)$$

This model is almost identical to the OLS model in (1). There are J brands in our data set, hence f_i is a (J-1) x 1 vector of dummy variables. If product i is of brand j then f_i will include a 1 in row j and a 0 in all other rows.

As mentioned earlier, the hedonic price regressions above allow us to identify the price differential between FT and non-FT coffee. However, we are unsure if the price differential is driven by increased consumer demand or the increased costs of producing coffee to meet FT standards. In order to separate these two effects, we need to explicitly model the price and quantity observed in each market as equilibrium outcomes.

Starting with a model of demand we assume that a consumer n in market m derives

indirect utility U_{nim} from purchasing product i. We assume that the utility function is linear in the product attributes as shown below:

$$U_{nim} = x'_{im}\beta_n + f'_{im}\gamma_n + \alpha p_{im} + \xi_{im} + \epsilon_{nim} \qquad (3)$$

As in the hedonic regressions, x_{im} and f_{im} are vectors of product attributes and brand dummies respectively. The mean price of product i in market m is now denoted by p_{im}, ξ_{im} represents the utility derived from unobserved attributes of product i, ϵ_{nim} is the idiosyncratic utility of person n for product i and is assumed to follow the extreme value distribution, and α denotes the marginal utility of income which is assumed not to vary across consumers. The vectors β_n and γ_n are random coefficients and indicate attribute and brand preferences of person n respectively. Both these vectors are assumed to be multivariate normal with means $\bar{\beta}$ and $\bar{\gamma}$. The assumed covariance matrices for the random coefficients are shown below:

$$\Sigma_\beta = \begin{pmatrix} \sigma_{\beta 1}^2 & 0 & 0 \\ 0 & \sigma_{\beta 2}^2 & \sigma_{\beta 23} \\ 0 & \sigma_{\beta 23} & \sigma_{\beta 3}^2 \end{pmatrix} \qquad \Sigma_\gamma = \begin{pmatrix} \sigma_{\gamma 1}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{\gamma 2}^2 & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_{\gamma J-1}^2 \end{pmatrix}$$

We allow a consumer's preference for FT and Specialty coffee to be correlated, which was suggested by our hedonic price regressions. However we assume a consumer's preferences are not correlated across brands.

The average consumer's willingness to pay for FT coffee is $\frac{dp}{dFT} = -\frac{\partial U/\partial FT}{\partial U/\partial p} = -\frac{\bar{\beta}_3}{\alpha}$. The willingness to pay for FT coffee has variance $\frac{\sigma_{\beta 3}^2}{\alpha^2}$. Thus, by obtaining consistent estimates of the utility function above, we will be able to answer our research question of how much consumers are willing to pay for FT coffee, as well as get an estimate of the distribution of the willingness to pay within the population.

The difficulty in obtaining consistent estimates stems from the fact that price is not exogenous and is the endogenous outcome of the interaction between consumer preferences and supplier costs. In order to control for supply considerations, we assume the retail price for product i observed in market m is:

$$p_{im} = x'_{im}\lambda + f'_{im}\eta + z'_m\varpi + \varepsilon_{im} \qquad (4)$$

Again $\mathbf{x}_i$ and $\mathbf{f}_i$ are vectors of product attributes and brand dummies respectively, and we introduce $\mathbf{z}_m$ which is a vector of cost shifters which includes Brazilian wage rates[9] and U.S. gasoline prices. We assume that $\mathbf{x}_{i,}, \mathbf{f}_{i,}$ and $\mathbf{z}_m$ are all uncorrelated with the utility derived from product i's unobserved product attributes ξ_{im}.

This assumption implies population moments $E(\xi|X) = \mathbf{0}, E(\xi|F) = \mathbf{0},$ and $E(\xi|Z) = \mathbf{0}.$ We exploit these population moments to estimate the parameters of the model using the Generalized Method of Moments (GMM). The specifics of the estimation procedure used in this paper can be found in Berry, Levinsohn, and Pakes (1995) and are quite similar to those used in Chapter I.

IID. RESULTS

Table II.3. OLS Regression Results

	Coef.	Std. error	T	P (coef=0)	95% Confidence Interval
Fair Trade	2.47	.4010	6.10	0.000	[1.68, 3.27]
Specialty	2.10	.2125	9.89	0.000	[1.68, 2.52]
Constant	5.07	.1623	31.26	0.000	[4.75, 5.39]

R-Squared= .2552   Adjusted R-Squared= .2522

The OLS regression results are shown in Table II.3. According to the model there is a significant price premium for Specialty coffee ($2.10) and an additional ($2.47) premium for FT coffee. Thus the price of a coffee which is both specialty and FT is $4.57 per 12-16oz higher than non-Specialty and non-FT coffee. The coefficients on FT and Specialty are both statistically different from zero at the <1% significance level.

As stated above, we are concerned that the OLS model includes omitted variable bias. We suspect that there are some attributes which are positively correlated with mean price and either FT or specialty coffees. If this is true then the coefficient estimates in Table 3 are likely too high. Our suspicions are further aroused by a previous study by Loureiro (2005), who found there was a $1.45/lb price premium for FT coffee versus non-FT coffee, roughly $1/lb less than our OLS estimate.

---

[9] Brazil is far and away the largest producer of coffee in the world accounting for roughly 1/3 of total coffee production.

Table II.4. Fixed Effects Regression Results

Mean Price	Coef.	Std. error	t	P(coef=0)	95% Confidence Interval
Fair Trade	1.74	.4955	3.51	0.000	[.77, 2.71]
Specialty	2.10	.2566	8.17	0.000	[1.59, 2.6]
constant	5.42	.2501	21.68	0.000	[4.93, 5.91]

The fixed effects regression results are show in Table II.4. According to the model there is a significant price premium for Specialty coffee ($2.10) and an additional ($1.74) premium for FT coffee. Thus the price of a coffee which is both specialty and FT is $3.84 per 12-16oz higher than non-Specialty and non-FT coffee. The coefficients on FT and Specialty are both statistically different from zero at the <1% significance level.

Comparing these results to the OLS results, the most striking difference is the downward change in the FT coefficient. As we suspected, there is likely an unobserved attribute which is correlated with FT coffee which was not included in the OLS specification. We suspect that within brand these unobserved attributes do not vary so much which significantly reduces the omitted variable bias. We see no other evidence of bias in the estimates when comparing the results of the two models[10].

IIE. CONCLUSION

The goal of this research is to determine if consumers are willing to pay for corporate social responsibility programs. We selected the FT coffee market as a suitable testing ground for this research. We collected data on the prices and quantity sold for 499 coffee products over the period 2005-2008 from a large supermarket chain in California over the period 2005-2008. We developed an analytical framework which should enable the identification of the consumer's willingness to pay for FT coffee, a socially responsible practice.

Thus far, we have established that FT coffee carries a price premium of $1.74 per 12-16 oz. We strongly suspect that at least a portion of this premium is due to increased consumer willingness to pay for FT coffee. However, we cannot rule out the possibility

---

[10] One may notice that the constant coefficient is now $5.42 in the fixed effects model v $5.07 in the OLS specification. Comparing these two numbers is a bit like comparing apples and oranges. In the OLS case, $5.07 is the base price for non-FT non-specialty coffee across all brands. However, in the fixed effects model we effectively include dummies for each brand except one (the base brand). Thus, in the fixed effects case $5.42 is the base price for non-FT non-specialty coffee of the base brand.

that the price premium is a result of the added costs associated with that FT practices. In a subsequent paper, we hope to allocate the cause of the price premium we have now established for FT coffee to demand v. supply considerations by estimating the equilibrium model outlined in section IIC.

Chapter II References

Aupperle, Kenneth E., Archie Carroll and John Hatfield. 1985. "An Empirical Examination of the Relationship Between Corporate Social Responsibility and Profitability." Academy of Management Journal 28(2):446-63.

Beaver, J. (2006) "Charting the Coffee Course" *The Specialty* Coffee Chronicle (pp. 14-15).

Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica 63 (July): 841-90.

Cailleba, P. & Casteran, H. (2009) "A Quantitative Study on the Fair Trade Coffee Consumer" Journal of Applied Business Research 25(6):31-46.

Carroll, A. B. 1979. "A Three Dimensional Conceptual Model of Corporate Social Performance." Academy of Management Review 4: 497-505.

Clark, T. (2008). A Double Tall Tale of Caffeine, Commerce, and Culture. New York, NY: Back Bay Books.

Corson, J. J. and G. A. Steiner. 1974. "Measuring Business' Social Performance: The Corporate Social Audit." New York: Committee for Economic Development.

Donnet, M.L., Weatherspoon D.D., & Moss, C.B. (2010) "Measuring Food Product Differentiation by Quality Ratings: A Cross-Entropy Analysis of Specialty Coffee e-Auctions" *Journal of Agricultural Economics* 61(1):122-137.

European Fair Trade Association. (2006) 60 Years of Fair Trade. http://www.eftafairtrade.org (accessed March 4, 2010).

European Fair Trade Association. (2001) Fair Trade in Europe 2001. EFTA, Maastricht.

Fair Trade Federation. (2002) 2002 Report on Fair Trade Trends in the U.S. and Canada. Fair Trade Federation, Washington.

Glickman, L.B. (2004) "'Buy for the Sake of the Slave': Abolitionism and the Origins of American Consumer Activism" American Quarterly. 56(4):889-912.

Levi, M. & Linton, A. (2003) "Fair Trade: A Cup at a Time?" Politics Society. 31(3):407-432.

Loureiro, M.L. & Lotade, J. (2005) "Do Fair Trade and Eco-Labels in Coffee Wake Up the Consumer Conscience?" Ecological Economics. 53:129-138.

Maguire, K.B., Owens, N. & Simon, N.B. (2004) "The Price Premium for Organic Babyfood: A Hedonic Analysis" Western Agricultural Economics Association. 29(1):132-149.

McCluskey, J.J. & Loureiro, M.L. (2003) "Consumer Preferences and Willingness to Pay for Food Labeling: A Discussion of Empirical Studies" *Journal of Food Distribution Research* 34(3):95-102.

Mohr, Lois A. and Deborah J. Webb (2005), "The Effects of Corporate Social Responsibility and Price on Consumer Responses," Journal of Consumer Affairs, 39, (1), 121-47.

Moore, G. (2004) "The Fair Trade Movement: Parameters, Issues, and Future Research" *Journal of Business Ethics* 53:73-86.

Moskowitz, M. 1972. "Choosing Socially Responsible Stocks." Business and Soiciety Review. [1]: 71-75.

Raynolds, L.T., Murray, D., & Taylor, P.L. (2004) "Fair Trade Coffee: Building Producer Capacity Via Global Networks" Journal of International Developme*nt* 16:1109-1121.

Redfern, A. & Snedker, P. (2002) Creating Market Opportunities for Small Enterprises: Experiences of the Fair Trade Movement. ILO, Geneva.

Roseberry, W. (1996) "The Rise of Yuppie Coffees and the Re-Imagination of Class in the United States" American Anthropologist. 98:762–775.

Sen, Sankar and C. B. Bhattacharya. 2001. "Does Doing Good Always Lead to Doing Better?" Journal of Marketing Research. Vol. 38(May): 225-43.

Stanwick, Peter A. and Sarah D. Stanwick. 1998. "The Relationship Between Corporate Social Performance and Organizational Size, Financial Performance, and Environmental Performance: An Empirical Examination." Journal of Business Ethics. 17(2): 195-204.

TransFair USA. (2008) *Fair Trade Overview*. http://www.transfairusa.org/content/about/overview.php (accessed April 3, 2010).

Vance, S. C. 1975. "Are Socially Responsible Firms Good Investment Risks?" Management Review. 64(8): 19-24.

Villas-Boas, J. Miguel and Russell Winer. 1999. "Endogeneity in Brand Choice Models." *Management Science*. 45: 1324-38.

Chapter III: How Durable Should Manufacturers Make Their Durable Products

An interesting problem for manufacturers is how durable or long lasting should they make their products. On the one hand a more durable product will be more desirable to consumers, since it will provide benefits over a longer period. Thus, a longer-lived product will command a higher price. However, it seems likely that unit production costs will increase as a product is made more durable due to the increased cost of more reliable materials and more exacting quality standards. In addition, a product which is more durable will be replaced less frequently. *Ceteris paribus,* less frequent replacement is less desirable to manufacturers, as the periodicity of the revenue stream increases. Manufacturers can trade off the benefits of durability with the costs to determine the optimal reliability or life for the goods they produce.

In some sense, this problem is a classic trade-off between quality and cost. Standard economic theory informs us that if consumers have diminishing marginal utility for quality and manufacturers have increasing marginal costs for quality, then there exists an equilibrium level of quality which maximizes producer profits. What distinguishes the durable goods reliability problem is that increasing quality depresses replacement demand. A common anecdote is that light bulbs could easily be manufactured to last longer, but are not in order to increase replacement sales.

This question is important for manufacturers to understand from several perspectives. First, a manufacturer of a product with technology that is fairly static (e.g., light bulbs), needs to consider replacement demand in developing product designs. When technology is not static (e.g., computers), it is important to understand how the rate of technology advance will stimulate replacement demand. Should the products be designed to be more or less durable in the face of technological advance? In many situations the rate of technology advance may only be partially observable to the consumer (e.g. golf clubs). For example, when a new model of golf club is introduced, consumers line up to buy it in hopes of improving their score. Yet, in spite of all the purported advances in golfing technology, data from the United States Golf Association indicate that the average US golfer has not improved over the past 10 years. Of course scores could be constant, because the advancement in technology has been accompanied by an offsetting decrease in practice and effort. However, another hypothesis is that consumers only receive a noisy signal for technology advance and are risk adverse in not wanting to miss out on the "next big thing". Finally, manufacturers need to consider "buying back" used durable goods from their customers in order to stimulate replacement demand. Even if the manufacturer does not take physical possession of the old product, it may be optimal to offer price concessions on newer products in order to induce the consumer to replace her existing technology.

The existing literature on durable goods in Marketing and Economics has addressed many other interesting problems. Swan(1970) shows under certain conditions a monopolist producer will produce the optimal amount of durability. This result is logical as long as the monopolist decision for durability and price are disjoint. Stokey (1981) and Bulow (1982) show that the "Coase" conjecture (Coase 1972), the idea under

rational expectations, a durable goods monopolist will price its product at marginal cost due the intertemporal competition between goods sold today and goods sold in the future. Rust (1985) studies the equilibrium conditions in secondary markets for durable goods. Purohit has addressed many strategic considerations in durable goods markets including optimal distribution strategies (Purohit and Staelin 1994), competition between dealers and rental agencies, (Purohit 1998) and lease v. sell decision (Desai and Purohit 1998, 1999).

Much of the existing empirical work on product adoption addresses the diffusion of new technology (durable goods) following in the tradition of the Bass model (Bass 1969). Some relevant extensions consider the diffusion of successive generations of a technology product (Norton and Bass 1987) and the optimal introduction timing of new product generations (Wilson and Norton 1989). Some incorporate product replacement into forecasting models of durable goods sales (Bayus 1988; Bayus, Hong, and Labe 1989; Steffens 2003).

The more recent empirical literature on durable goods has used structural models to study the demand for high tech good adoption including Chintagunta and Song (digital cameras 2003), Erdem et. al. (personal computers 2005), and Nair (video game consoles 2007). However all of the papers above fail to consider product replacement as a driver for product sales. Rust (bus engines 1987) and Gordon (personal computers 2008) explicitly consider product replacement in an empirical setting.

The work in this paper is closest to Levinthal and Purohit (1989) which studies durable goods and product obsolescence in a two-period setting. It also differentiates itself from Swan(1970) in that we do not assume that durability and pricing decisions are independent. The ultimate goal of this work is to distinguish itself from previous research along two important dimensions: a) the present work examines an infinite rather than two-period time horizon and b) this paper explores the information asymmetries that exist between manufacturers and consumers about product life and technological advance. Finally this paper contributes to the management literature at the interface of marketing and operations management by demonstrating how the manufacturing and selling of durable goods needs to be considered in an integrated framework.

The remainder of this chapter is outlined as follows. In section A, we develop a base case model for durable good replacement under the assumption of static technology. Section B relaxes the static technology assumption and considers deterministic rates of technological advance. In Section C, technology is assumed to advance in intermittent jumps which are known to the manufacturer but not the consumer. Section D concludes.

IIIA. Base Case Model

To proceed with this research, we will develop insights with a simple model which makes the following economic assumptions:

1. The decay in a product's performance over time is deterministic

2. Technology is static. The rate of technology advancement is zero.

3. Consumers are homogeneous

4. Information is symmetric

5. There is only one monopoly manufacturer.

After analyzing the base case, we will relax assumptions 2-4 and analyze their impact on the problem. Clearly some settings will lend themselves more to the simpler model (e.g., the market for ironing boards), while others will be better suited to a more complex setting (e.g., the market for personal computers).

The timing of the problem is as follows. Consumers begin at time, t = 0, with an endowment of one new product (i.e., the product has just been replaced). In each subsequent period, the following decisions are made:

1. The manufacturer decides the number of units of the product to be manufactured.

2. The manufacturer sets the price, p, which will be offered to the market. There are no transactions costs.

3. Consumers decide whether to accept the manufacturer's offer.

This problem is regenerative in that once a consumer decides to accept the manufacturer's offer (i.e., replace the product), the problem reverts to conditions which existed at t=0. In steady state the consumer will replace the product every T periods. Consumers can have at most one product. The product provides consumers with indirect net flow utility of $U(t,\theta,P)$ according to:

$$U(t,\theta,P) = \lambda - P \text{ if t=0 (the product is replaced)}$$

(1)

$$U(t,\theta,P) = \lambda - \theta t \text{ otherwise}$$

where $t \in [0, \lambda/\theta]$ represents time and $\theta \in [\theta_l, \theta_h]$ represents the rate at which the product depreciates. λ is assumed to be exogenous and reflects the existing level of technology for the product. Consumers have a discount rate of r and discount factor $\delta_c = \frac{1}{1+r}$. The consumer's problem is to maximize her infinite horizon net utility of:

$$V = \frac{1}{1-\delta_c^T} \int_0^T U(t,\theta,P)e^{-rt}dt$$

(2)

48

The monopolist manufacturer has a discount factor of δ_m. The unit cost function for the monopolist is given by $C(\theta) = c - \alpha ln\theta$. The manufacturer's objective is to maximize infinite horizon profits of:

$$\Pi = \frac{1}{1-\delta_m^T}(P - C(\theta)) \tag{3}$$

Lemma 1: The highest price P, that the consumer will be willing to pay to replace a product with depreciation rate θ that is T periods old is:

$$P^*(T,\theta) = \int_0^T (\lambda - \theta t)e^{-rt}dt - \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt \tag{4}$$

Proof:

a) The consumer is not willing to pay more than P*. Suppose not. If P>P*then

$$\int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt > \int_0^T (\lambda - \theta t)e^{-rt}dt - P$$

The utility of keeping the existing product until the end of its life is greater than the utility of purchasing a replacement product over its life. So the consumer will never pay more than P* to replace her product.

b) The consumer is willing to pay P*. If the consumer receives an offer of P* at time T she has the option of accepting or waiting. If the accepts, she pays P*, she realizes surplus of $\int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt$ every T periods. If she waits to time U>T the manufacturer can offer $P^*(U,\theta) = \int_0^U (\lambda - \theta t)e^{-rt}dt - \int_U^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt$ and the consumer will realize surplus $\int_U^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt < \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt$. The manufacturer can continue to offer this price, because if the consumer continues to wait the product will fully depreciate and she will receive surplus of zero and be forced to buy the new product. By waiting, the consumer will always receive less surplus than by accepting. Therefore, the consumer should accept an offer of P*. Proof Complete.

We can interpret the consumer's willingness to pay as the net utility from owning the good for T time periods less the residual utility still remaining in the currently owned unit. We can think of the first term in (4) as the gross purchase price and the second term as the trade-in or buyback price. Clearly, the manufacturer would prefer to charge

a higher price which just includes the utility realized by the consumer of T periods, but this price would not be accepted since the consumer would be better off continuing to use her existing product. For this reason, the manufacturer is unable to capture the consumer's entire surplus.

For any given combination of T and θ, a profit maximizing monopolist will charge the highest price possible. Thus, we can rewrite the manufacturer's profit function in terms of T and θ as:

$$\Pi(T,\theta) = \frac{1}{1-\delta_m^T}\left(\int_0^T (\lambda - \theta t)e^{-rt}dt - \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt - c + \alpha ln\theta\right) \tag{5}$$

The manufacturer problem is to choose T and θ to maximize the infinite horizon profits in (5). While it is the consumer who decides the replacement timing the manufacturer effectively determines T by setting the price exactly equal to the consumer's willingness to pay for a good that operates for T periods and then is bought back[11] as shown in equation (4).

Taking first order conditions of (5) yields:

$$\frac{\partial \Pi}{\partial T} = \frac{\delta_m^T ln\delta_m}{(1-\delta_m^T)^2}\left(\int_0^T (\lambda - \theta t)e^{-rt}dt - \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt - c + \alpha ln\theta\right)$$
$$+\frac{2}{1-\delta_m^T}(\lambda - \theta T)e^{-rT} = 0 \tag{6}$$

$$\frac{\partial \Pi}{\partial \theta} = \frac{1}{1-\delta_m^T}\left(\int_T^{\frac{\lambda}{\theta}} te^{-rt}dt + \frac{\alpha}{\theta} - \int_0^T te^{-rt}dt\right) = 0 \tag{7}$$

Assumption[12] 1: There exist conditions under which the manufacturer's profit function in (5) is globally concave.

---

[11] Note the manufacturer does not need to take physical possession of the bought back item. The firm effectively bribes the consumer to stop using her existing product. Throughout, this paper we use the term buy back to refer to this practice.

[12] Deriving the conditions for a unique equilibrium will be pursued in subsequent research. We also can conjecture based on the analysis below that these conditions are extremely likely to exist.

Proposition 2: Under a scenario with no technological advances, it is always optimal for the monopolist to buyback the existing stock of durable goods before it is fully depreciated.

Proof: $\delta_m \in (0,1) \Rightarrow \frac{\delta_m^T \ln \delta_m}{(1-\delta_m^T)^2} < 0$. Any solution to the manufacturer's problem must produce positive profits each cycle, so the first term in first order condition (6) is negative. Consequently, the second term must be positive. Thus, $\lambda - \theta T > 0 \Rightarrow T < \frac{\lambda}{\theta}$. Proof complete.

From an economic perspective, the first term of (6) represents the marginal profit lost on forgone replacement sales by extending the life of the product by one period. The second term represents the marginal price increase which can be realized by extending the life of the product. Note this term is twice the discounted flow utility at time T reflecting that extending the life of the product both increases the gross purchase price and reduces the buyback price by equal amounts. As T increases, the price increase that can be realized diminishes as the extra time the consumer can use the product gets pushed further out into the future and is discounted more heavily. Thus, the optimal T will occur prior to the end of the useful life of the product, and buybacks will always occur.

Examining the first order condition in (7), we see that $\int_T^{\frac{\lambda}{\theta}} te^{-rt} dt + \frac{\alpha}{\theta} = \int_0^T te^{-rt} dt$ in equilibrium. The economic intuition for this result, is that for any given product life T, increasing θ (reducing durability) both reduces the buyback price by shortening the remaining life of the product and reduces the unit product cost of the product. However reducing durability decreases the value of the product while it is in use, since it depreciates more quickly. Hence reducing durability reduces the gross purchase price.

IIIB. Exogenous technological advance

We now proceed to relax the assumption that technology is static. In the base model the level of technology was represented by λ. We now allow λ to grow at a known and constant rate γ. We assume that technology advances are exogenous and occur at no cost to the firm[13]. Based on this new assumption, the price the consumer is willing to pay and the manufacturer's objective function can be reformulated as:

$$P^*(T, \theta; \gamma) = \int_0^T \lambda e^{(\gamma - r)t} dt - \int_0^T \theta t e^{-rt} dt - \int_T^{\frac{\lambda}{\theta}} (\lambda - \theta t) e^{-rt} dt \qquad (8)$$

$$\Pi(T, \theta; \gamma) = \frac{1}{1-\delta_m^T} \left(\int_0^T \lambda e^{(\gamma - rt)} dt - \int_0^T \theta t e^{-rt} dt - \int_T^{\frac{\lambda}{\theta}} (\lambda - \theta t) e^{-rt} dt - c + \alpha \ln \theta \right) \qquad (9)$$

---

[13] This happens in practice due to spillover from basic research.

With technology advancing, the gross price for the new product has increased; however, since technology advances do not affect the value of products already in the hands of consumers, the buyback price has remained the same. Also notice that the first order condition with respect to θ is identical to the static technology case. For example, the advent of anti-lock brakes made cars safer and more valuable to consumers; however it had no effect on the rate at which the engine performance degrades. Thus the marginal impact of durability on profits is independent of the rate of technology advance. Note that the base case is just a special case of the technology advance case with $\gamma = 0$. As in the base case, extending the interval between product replacements will both increase the price consumers are willing to pay, as well as reduce the profits derived from future replacement sales. This tradeoff can be seen in the first order condition with respect to T below:

$$\frac{\partial \Pi}{\partial T} = \frac{\delta_m^T \ln \delta_m}{(1-\delta_m^T)^2}\left(\int_0^T \lambda e^{(\gamma-r)t}dt - \int_0^T \theta t e^{-rt}dt - \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt - c + \alpha \ln \theta\right) + \frac{1}{1-\delta_m^T}\left((\lambda - 2\theta T)e^{-rT} + \lambda e^{(\gamma-r)T}\right) = 0 \qquad (10)$$

Prposition 3: An increase in the rate of technological advance, γ, has an ambiguous impact on the optimal depreciation rate, θ, and the optimal interval between product replacements, T, chosen by the manufacturer.

Specifically, if the solution to the manufacturer's problem in (9) is denoted by (T',θ'), then it is optimal to extend the product life and increase durability in the face of an increase in the rate of technological advance if:

$$\frac{\delta_m^{T'} \ln \delta_m}{(1-\delta_m^{T'})^2}\int_0^{T'} \lambda t e^{(\gamma-r)t}dt + \frac{1}{1-\delta_m^{T'}}\lambda T' e^{(\gamma-r)T'} > 0 \qquad (11)$$

and optimal to decrease durability and shorten product life otherwise.

Proof: Taking the derivative of the first order condition in (10) with respect to γ, yields the expression in (11). If the expression in (11) is greater than zero then the optimal T needs to increase in response to an increase in γ. Since $\frac{\partial^2 \Pi}{\partial T \partial \theta}(T', \theta') = \frac{-2T'}{1-\delta^{T'}}e^{-rT'} < 0$, any increase in T' must be accompanied by a decrease in θ'. Therefore if (11) is positive then T' increases and θ' decreases and vice versa. Proof Complete.

Initially, intuition might suggest that an increase in the rate of technological advance would result in a shortening of the interval between product replacements, thus the result in proposition 3 is somewhat surprising. However, the benefit to the firm of lengthening the product replacement interval is an increase in the price the consumer is willing to pay. Even for a large advance in technology, the consumer will not be willing to pay a high price if she thinks the product will need to be replaced shortly. This is especially true for patient consumers with low discount rates[14]. Thus, if the rate of

---

[14] It can be shown that (11) is increasing in γ, decreasing in r and decreasing in δ.

technological advance is high relative to the consumer's discount rate, then the marginal benefit of extending the product's life by making it more durable is significant to the consumer. This increased willingness to pay for a product with high durability will more than offset the loss of replacement sales in the future.

The manufacturer also benefits from extending the product replacement cycle if the manufacturer is impatient relative to the consumer. Since lost replacement sales occur in the future, an impatient manufacturer will not care too much if they decrease.

Figure 1 illustrates various plausible scenarios where an increase in γ would result in both longer lived products and shorter lived products. The figure demonstrates that as γ rises above a critical level, it becomes beneficial for the manufacturer to extend the replacement cycle.

The result in proposition 3 may help explain why home electronics products such as televisions are surprisingly durable and long-lived in spite of a rapid rate of technological advance. According to an informal interview with the video entertainment manager of a New Jersey electronics store, consumers rarely purchase high definition televisions (HDTV) to replace other HDTVs. Rather consumers will rotate TVs in their homes. For example, in a home with 2 TVs, the newer TV might be in the family room where most viewing takes place and the older TV would be located in the bedroom and gets less frequent use. When a new HDTV is purchased, the bedroom TV is discarded (bought back), the existing family room TV moves to the bedroom, and the new HDTV takes center stage in the family room. In effect the new TV is replacing the oldest TV. This result is also consistent with the observation that TV prices have been increasing over time. Consumers are willing to pay a higher price for the new technology, since they anticipate that they will be able to use the new unit for a considerable length of time.

Now consider a case where the replacement cycle appears to be getting shorter and prices appear to be falling – the personal computer (PC) industry. One might question the validity of the result above, given that PC technology is advancing very rapidly. Yet, upon closer examination, it appears that relative to advances in computer software, the rate of technological advance in hardware is quite modest. While a new computer may run older software versions much faster than its predecessor, the actual performance gain experienced by the consumer is small since the newest version of the operating system and standard applications such as spreadsheets take up so much more memory and CPU. Consequently, the net gain in performance perceived by the consumer is quite small. Indeed, casual interviews with PC shoppers at a New Jersey electronics store suggested that most consumers plan to recycle the old computer the new computer is replacing. Witness, the fact that PC industry leader Dell promotes free recycling kits with the purchase of new computers from the Company website. Thus, we believe that anecdotal evidence from the PC industry is also consistent with the theoretical result obtained above.

IIIC. Asymmetric Information on technological advance

Now consider the situation where technology evolves stochastically. To make things simple, assume that technology either advances at rate γ/β with probability β or remains static with probability (1-β). As in section IIIB, the expected rate of technology advance is γ. This scenario seems more realistic for industries where much of the time technology is stable and advances only intermittently in breakthrough fashion.

In addition, we now assume that the consumer only knows the evolutionary process of technology, but cannot observe when actual advances are made. There are many cases where manufacturers introduce new and allegedly improved products where the improvement is very difficult for the consumer to actually observe. For example, consider the following new product introductions: Tire manufacturer Goodyear announces its new "Aquatred" tire; a car manufacturer announces its new rack and pinion steering; a golf club manufacturer introduces its new titanium alloy golf clubs. In each of these instances the consumer has a difficult time observing the actual performance improvement until after she has already purchased the product. Even after purchase, it may be difficult to measure objectively how much better a car is handling or how much straighter and further a golf ball is flying.

Assuming that consumers are expected utility maximizers, the price the consumer is willing to pay and the manufacturer's profit function can be written as:

$$P^*(T,\theta) = \frac{\int_0^T \beta(T)\lambda e^{\left(\frac{\gamma}{\beta(T)}-r\right)t}dt + \int_0^T (1-\beta(T))\lambda e^{-rt}dt - \int_0^T \theta t e^{-rt}dt -}{\int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt} \tag{12}$$

$$\Pi(T,\theta) = \frac{1}{1-\delta_m^T}\left(\begin{array}{c}\int_0^T \beta(T)\lambda e^{\left(\frac{\gamma}{\beta(T)}-r\right)t}dt + \int_0^T (1-\beta(T))\lambda e^{-rt}dt - \int_0^T \theta t e^{-rt}dt - \\ \int_T^{\frac{\lambda}{\theta}}(\lambda - \theta t)e^{-rt}dt - c + \alpha \ln\theta\end{array}\right) \tag{13}$$

Note the probability of a technological breakthrough, β(T), is a function of the interval between product replacements. β(T) is assumed to be concave in T. The longer the interval the more likely it is that a breakthrough has occurred, and there are diminishing returns to time spent waiting for a breakthrough. In this paper, we do not endeavor to characterize fully the solution to this problem[15]. Nonetheless, there are some observations worth making.

Lemma 2: For a given rate of expected technological advance, γ, uncertainty increases the price the consumer is willing to pay.

---

[15] An excellent topic for further study.

Proof: Recall that under a certain rate of technological advance the consumer was willing to pay $P^*(T,\theta;\gamma) = \int_0^T \lambda e^{(\gamma-r)t} dt - \int_0^T \theta t e^{-rt} dt - \int_T^{\frac{\lambda}{\theta}} (\lambda - \theta t) e^{-rt} dt$. Taking the first two derivatives of this expression with respect to γ, we see that $\frac{\partial P}{\partial \gamma} = \int_0^T \lambda t e^{(\gamma-r)t} dt > 0$ and $\partial 2 P \partial \gamma 2 = 0T\lambda t2e(\gamma-r)tdt>0$, hence the price the consumer is willing to pay is convex in γ. Consequently, $E\big(P^*(T,\theta;\gamma)\big) \geq P^*(T,\theta;E(\gamma))$, with a strong inequality if γ is stochastic. Proof complete.

The intuition behind Lemma 2 is simple. Consumers are risk averse. If they are uncertain about whether a technological advance has occurred they want to make sure they don't miss out on the next big thing. In a random survey of golf club equipment in 2008, the market research organization, the Darrel Survey, reported that 41% of golfers are playing with drivers[16] that are less than two years old.

The implication for manufacturers is that more frequent product introductions, even if no technological breakthroughs have occurred can induce more rapid product replacement by consumers. Thus, we conjecture that in equilibrium, uncertainty with respect to technological advance leads to shorter replacement cycles and less durability. One of the best examples of this type of firm behavior occurs in the golf club industry. Manufacturer Taylor Made has introduced at least two new models of clubs each year for the past seven. It is hard to believe that there have been more than 14 significant breakthroughs in golf club technology in the last 7 years, especially in light of the fact that the average player's golf scores have not improved significantly over the past decade according to the United States Golf Association. Even in the face of the engineering and marketing costs associated with new product development and launches, this anecdotal evidence from the golf club industry appears to support the conjecture above.

IIID. Conclusion

In this paper we have developed a partial equilibrium model for durable goods where both manufacturers and consumers are forward looking over an infinite horizon. We see that in all cases, it is optimal for manufacturers to incent consumers (buy back) to replace their existing durable goods before the end of the useful life of the product. Perhaps this is why we see so many old refrigerators in garages and basements keeping extra drinks cool.

We also find that under extremely rapid rates of technological advance, it is optimal for manufacturers to extend the life of the product and increase durability. This result which runs somewhat counter to intuition in that one might think that rapid advances in technology would promote more frequent product replacement. However, if technology advances rapidly then patient consumers will want to be able to enjoy their new product (e.g., HDTV) and are willing to pay a high price for durability. This increased

---

[16] A driver is the golf club used by golfers to tee off a hole and typically propels the ball further than any other club due to the shallow loft on its club face. A new driver can retail for as much as $600.

willingness to pay for an advanced long-lived product will more than offset the loss of replacement sales in the future.

Finally, we show that when technology advances are uncertain and not directly observable by consumers, they are willing to pay more than when technology advances at a known and constant rate. The intuition for this result is that consumers are risk averse and do not want to miss out on a new breakthrough product. We believe this may be the reason that certain durable goods manufacturers (e.g. golf club manufacturer) appear to introduce new products at a far greater rate than technology is actually advancing.

While this paper answers a number of interesting questions, it also leaves open areas for further study. We believe more analysis needs to be conducted on the implications of stochastic rates of technological advance; moreover, it may be fruitful to model product depreciation rates as stochastic as well. In addition, this paper could be extended by introducing a model of heterogeneous consumers and competition among firms.

FIGURE III.1
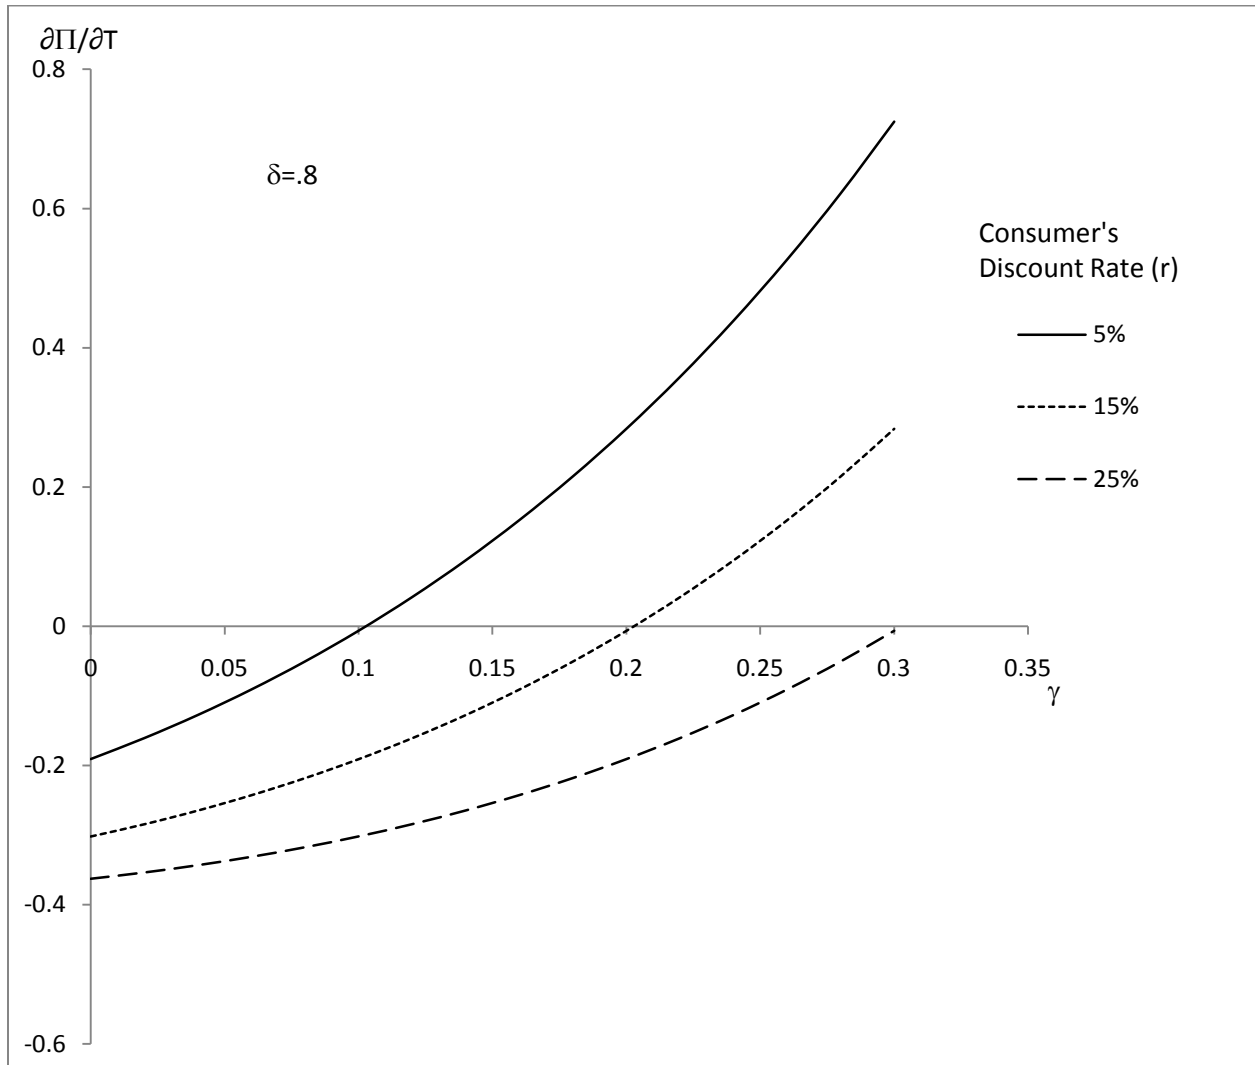
Chapter III References

Bass, F. M. (1969), "A New Product Growth Model for Consumer Durables," Management Science, 15, 215-227.

Bayus, B. L. (1988), "Accelerating the Durable Replacement Cycle with Marketing Mix Variables," Journal of Product Innovation Management, 5, 216-226.

Bayus, B. L. (1991), "The Consumer Durable Replacement Buyer," Journal of Marketing, 55(1), 42-51.

Bayus, B. L., S. Hong, and R. P. Labe Jr. (1989), "Developing and Using Forecasting Models of Consumer Durables," Journal of Product Innovation Management, 6, 5-19.

Bulow, J. (1982), "Durable Goods Monopolists," Journal of Political Economy, 90, no. 2, 314-32.

Coase, Ronald. (1972), "Durability and Monopoly," Journal of Law and Economics, vol. 15(1), 143–49.

Desai, Preyas and D. Purohit.(1999), "Competition in Durable Goods Markets: The Strategic Consequences of Leasing and Selling," Marketing Science, Vol. 18, No. 1, 42-58.

Desai, Preyas and D. Purohit. (1998), "Leasing and Selling: Optimal Marketing Strategies for a Durable Goods Firm," Management Science, Vol. 44, No. 11, S19-S34.

Erdem, T., S. Imai, and M. P. Keane (2003), "Brand and Quality Choice Dynamics Under Price Uncertainty," Quantitative Marketing and Economics, 1, 5-64.

Gordon, B. (2009), "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," Marketing Science, 2009, 28(5), 846-867.

Levinthal, D. and D. Purohit (1989), "Durable Goods and Product Obsolescence," Marketing Science, 8(1), 35-56.

Nair, H. (2007), "Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video Games," Quantitative Marketing and Economics, 5(3), 239-292.

Norton, J. A. and F. M. Bass (1987), "A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High-Technology Products," Management Science, 33(9), 1069-1086.

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," Econometrica, 55(5), 999-1033.

Rust, J. (1985), "Stationary Equilibrium in a Market for Durable Assets," Econometrica, 53(4) 783-805.

Song, I. and P. Chintagunta. (2003), "A Micromodel of New Production Adoption with Heterogeneous and Forward-Looking Consumers: Application to the Digital Camera Category," Quantitative Marketing and Economics, 1, 371-407.

Stokey, N. (1981), "Rational Expectations and Durable Goods Pricing," *Bell* Journal of Economics, 12, 112-128.

Swan, P. (1970), "Market Structure and Technological Progress: The Influence of Monopoly on Product Innovation," Quarterly Journal of Economics, 84: 627-638.

Purohit, D. and R. Staelin. (1994). "Rentals, Sales and Buybacks: Managing Secondary Distribution Channels," Journal of Marketing Research, 31, 3 (August), 325-338

Purohit, D. (1997), "Dual Distributions Channels: The Competition Between Rental Agencies and Dealers," *Marketing Science*, 16, 3, 228-245.

Steffens, P. R. (2003), "A Model of Multiple Ownership as a Diffusion Process," Technological Forecasting and Social Change, 70(9), 901-917.

Wilson, L. O. and J. A. Norton, (1989), "Optimal Entry Time for a Product Line Extension," Marketing Science", 8, 1-17.