



DEPARTMENT OF ECONOMICS
DISCUSSION PAPER SERIES

**RATIONAL HABITS AND UNCERTAIN PRICES: SIMULATING
GASOLINE CONSUMPTION BEHAVIOR**

K. Rebecca Scott

Number 596
March 2012

Manor Road Building, Oxford OX1 3UQ

Rational Habits and Uncertain Prices: Simulating Gasoline Consumption Behavior

K. Rebecca Scott
University of Oxford

8 July 2011

Key Words: gasoline demand, rational habits, price elasticity

JEL Classification: H30, Q40, Q41, Q50, R40

Abstract

When consumers are forward-looking with respect to their demand for a habit-forming good, traditional measures of price elasticity are misleading. In particular, such measures will underestimate sensitivity to long-run price shifts—and therefore underestimate the potential effect of policy instruments that act through price. Correcting elasticities for the behavior of the price process requires a model with forward-looking consumers, a habit-forming good, and uncertain relative prices. With appropriate restrictions on the type of price uncertainty, this paper shows that it is possible to solve for the optimal consumption path under any price process. Simulations then sketch out how habits and the price process shape demand. Gasoline demand motivates the model and illustrates its implications.

1 Introduction

Most models of gasoline demand suffer two shortcomings: they gloss over the dynamics of consumer behavior, and they ignore the dynamics of gasoline prices.

This disregard for dynamics is likely to prejudice estimates of price sensitivity toward zero. Provided gasoline is a habit-forming good and consumers are forward-looking, consumers will respond more vigorously to long-run price shifts than to short-term fluctuations. All types of price changes, however, are lumped together in traditional measures of price elasticity, and so such measures will underestimate demand's responsiveness to long-run price changes.

This systematic underestimation is of particular concern if we use price elasticity to project the effect of long-run price changes, such as the price changes we might bring about through tax increases or carbon pricing. To eliminate this prejudice, we need a model that takes into account both consumers' habits and the process that the gasoline price follows over time.

Such a model will allow us to correct price-elasticity estimates for the type of price change at hand—short-term or permanent, crude-price driven or policy-driven. It will allow us to examine how demand is shaped by the gasoline price process, and thus allow us to project how demand would change if we tweaked that process. And it will allow us, eventually, to estimate how much of the variation in price elasticity across regions is caused by variation in price processes rather than by variation in infrastructure or differences in consumers themselves.

Towering technical hurdles arise from the combination of rational habits and price uncertainty. Indeed, the ideal formulation of the model is so computationally expensive as to be nearly intractable. To lower these technical hurdles and render the model feasible, I place certain restrictions on price uncertainty. Given these restrictions, I can calculate price elasticities that take the price process into account.

This paper focuses on theory, simulating how consumers' price sensitivity varies with parameters of the price process. Later, the implications of this model can be used to estimate the effects of the price process in practice.

Naturally, although this model of demand is motivated by gasoline, it will—both in theory and in practice—apply to any habit-forming good.

2 Background

Over the years the rational habits literature has burgeoned, with papers suggesting alternative empirical structures¹; analyzing the problem in continuous time²; and applying rational habits to goods as diverse as milk³, cigarettes⁴, and cocaine⁵.

Somewhat surprisingly, only one paper within this large literature has addressed relative price uncertainty in a multigood model. Coppejans et al. (2007) examine this problem in the context of cigarette demand. They prove that when all available information about the future price distribution is contained in the current price of the habit-forming good, an increase in the variance of that distribution will reduce consumption of the habit-forming good. They go on to estimate a negative effect of price variance on the likelihood and intensity of smoking. Although this is an interesting first look at a setup with rational habits and price uncertainty, Coppejans et al. do not actually *solve* the consumer's problem. Nor do they consider the form of the price process or explore what habits imply for price responsiveness. In the case of gasoline (and other habit-forming goods whose markets invite government intervention), much of the policy interest lies where the dynamics of demand meets the dynamics of price; and it is the gap at this intersection that I will address.

3 Building a Model

3.1 Introducing of Habits by Modifying the Consumption Variable

To build a model that yields intuition about how habits influence consumers' response to gasoline price changes, we must first set up a scheme whereby past behavior affects current behavior. Either a HAD or a short-memory approach will accomplish this, and in fact both can be nested within a broader model. Roughly following Spinnewyn (1981) and Browning (1991, App. A), I will let consumers' utility for gasoline depend upon the adjusted quantity $\bar{g}_t$:

$$\bar{g}_t = g_t - \delta s_{gt} + \gamma_g \quad (1)$$

where g_t is the quantity of gasoline consumed in period t , and δ and γ_g are constant parameters. The variable s_{gt} is a gasoline habit-stock variable, which decays over time and is replenished with each period's consumption:

$$s_{gt} = \alpha g_{t-1} + (1 - \alpha)s_{g,t-1}, \quad 0 \leq \alpha \leq 1 \quad (2)$$

Using (2) to define and substitute for $s_{g,t-1}$, and iterating this process for further lags, we can write the habit stock as an infinite sum of past gasoline consumption:

$$s_{gt} = \alpha g_{t-1} + (1 - \alpha) [\alpha g_{t-2} + (1 - \alpha)s_{g,t-2}] \quad (3)$$

$$= \alpha g_{t-1} + \alpha(1 - \alpha)g_{t-2} + (1 - \alpha)^2 s_{g,t-2} \quad (4)$$

$$s_{gt} = \alpha g_{t-1} + \alpha(1 - \alpha)g_{t-2} + (1 - \alpha)^2 [\alpha g_{t-3} + (1 - \alpha)s_{g,t-3}] \quad (5)$$

$$= \alpha g_{t-1} + \alpha(1 - \alpha)g_{t-2} + \alpha(1 - \alpha)^2 g_{t-3} + (1 - \alpha)^3 s_{g,t-3} \quad (6)$$

...

$$s_{gt} = \alpha \sum_{i=1}^{\infty} g_{t-i} (1 - \alpha)^{i-1} \quad (7)$$

Thus the adjusted quantity $\bar{g}_t$ —the reference bundle—can be written as

$$\bar{g}_t = g_t - \delta \left[\alpha \sum_{i=1}^{\infty} g_{t-i} (1 - \alpha)^{i-1} \right] + \gamma_g \quad (8)$$

¹Becker, Grossman, and Murphy (1994)

²Houthakker and Taylor (1966) and Becker and Murphy (1988), for example.

³Auld and Grootendorst (2004)

⁴Baltagi and Griffin (2001), Chaloupka (1991)

⁵Grossman and Chaloupka (1998)

The parameter δ adjusts the strength of habits. If $\delta = 0$, then $\bar{g}_t = g_t + \gamma_g$, and there are no habits—only current consumption enters into the agent’s utility function in any period; preferences are time-separable. As δ increases, the effect of the habit-stock on the reference bundle $\bar{g}_t$ grows, and preferences depend more and more on past consumption. The higher the δ , the stronger an agent’s habit—that is, the stronger gasoline’s addictiveness.

To verify that this specification of the reference quantity $\bar{g}_t$ nests both the short-memory model and the HAD model, note that by the appropriate choice of α , we can recover either. First, to recover the short-memory model, let $\alpha = 1$. The habit-stock is then given by $s_{gt} = \alpha g_{t-1} + (1 - \alpha)s_{g,t-1} = g_{t-1}$, meaning the reference quantity is given by $\bar{g}_t = g_t - \delta g_{t-1} + \gamma_g$. The impact of gasoline consumption on $\bar{g}_t$ therefore lasts only one period—which is the defining characteristic of the short-memory model. To recover the HAD model, we can take a cue from Browning (1991, App. A) and let $\delta = \frac{\alpha-1}{\alpha}$. Then

$$\bar{g}_t = g_t - \frac{\alpha - 1}{\alpha} \left[\alpha \sum_{i=1}^{\infty} g_{t-i} (1 - \alpha)^{i-1} \right] + \gamma_g \quad (9)$$

$$\bar{g}_t = g_t + \sum_{i=1}^{\infty} g_{t-i} (1 - \alpha)^i + \gamma_g \quad (10)$$

$$\bar{g}_t = \sum_{i=0}^{\infty} g_{t-i} (1 - \alpha)^i + \gamma_g \quad (11)$$

The reference quantity $\bar{g}_t$ is thus a weighted sum of current and all past gasoline consumption—and this dependence on all past values of g_t is, of course, what characterizes the HAD model.

3.2 Infeasibility of the Ideal Problem

Ideally, we would proceed by maximizing the expectation at $t = 1$ of a weighted sum of the agent’s present and future utility, subject to budget constraints and the rules governing the habit-stock, and with the distribution of gasoline prices p_t depending on information available one period previously, such as lagged realizations of the price:

$$\begin{aligned} \max_{g_1, g_2, \dots, g_T} \quad & E_1 \left[\sum_{t=1}^T \beta^{t-1} u(\bar{g}_t, c_t) \right] \\ (T \quad & \text{integer or } \infty) \\ \text{s.t.} \quad & \text{budget constraints}^6; \\ \bar{g}_t \quad & = g_t - \delta s_{gt} + \gamma_g; \quad s_{gt} = \alpha g_{t-1} + (1 - \alpha)s_{g,t-1}, \quad 0 \leq \alpha \leq 1 \\ \text{with} \quad & \\ p_t \quad & \sim f(I_{t-1}), \text{ where } I_{t-1} \text{ is information available at time } t-1 \end{aligned} \quad (12)$$

The uncertainty enters this version of the problem through the unknown future prices in the budget constraints. The variable c_t represents a general, non-habit-forming good, whose price is normalized to 1.

This seemingly-straightforward problem turns out to be less than straightforward to solve. The trouble arises from the combination of intertemporally-dependent utility and serially-dependent price distributions. No simple backward induction is possible, as optimal consumption at T will depend not only on the realization of the price at T , but on previous consumption decisions—and thus on previous expectations of the entire *path* $\{p_1, \dots, p_T\}$.

Although we could reformulate this ‘ideal’ problem as a dynamic programming problem and approach it via value function iteration or collocation, these approaches are doomed by the curse of dimensionality. To demonstrate this, I set up the problem as a Bellman equation in Appendix A and walk through the computational approaches in Appendix B.

Given the infeasibility of solving this ideal problem for a potentially-complicated price process, I turn to a version of the problem that limits agents’ learning behavior and the distribution of price uncertainty. These limits, I demonstrate in the next section, yield a model that is solvable no matter how complicated the price process.

⁶The budget constraints will depend upon the time horizon and the rules governing intertemporal borrowing.

3.3 Simplifying the Consumer's Problem Using Uniformly-Distributed Prices

The problem laid out in Section 12 begins to give way if we require that price uncertainty come from a uniform distribution. The advantage of the uniform distribution is that, when combined with conditions for the utility function that are not overly restrictive, it allows us to solve for expected utility in any time period. A uniform price distribution is not, of course, the *only* way to achieve a closed-form expression for expected utility: if we were willing to work with a discrete price distribution, for example, we could easily solve for expected utility without imposing *any* extra restrictions on the form of the utility function. At another extreme, if we were willing to impose linear utility, we could work with more complicated price distributions, including log-normal. Of the possible combinations of price distributions and utility functions that yield closed-form expected utility, however, the uniform price distribution seems a reasonable choice, allowing a sensible choice of utility function without sacrificing succinctness or *too* much plausibility in the price distribution.

To see how uniformly-distributed prices can be used to simplify the problem, let the price of gasoline in period t be distributed uniformly, with lower bound low_t and upper bound $high_t$:

$$p_t \sim U [low_t, high_t]$$

$$\text{density of } p_t : f(p_t) = \frac{1}{high_t - low_t}$$

Let the agent have a period- t utility function of the form

$$u_t = (c_t + \gamma_c)^{\sigma_c} + \theta (g_t - \delta s_{gt} + \gamma_g)^{\sigma_g}$$

which encompasses both short-memory and HAD-style habits for gasoline.

Finally, assume a budget of the form

$$p_t g_t + c_t = a_t$$

where the price of other consumption c_t is normalized to 1, and a_t represents assets available for consumption at time t . This implies a fixed budget in period t , a decision I will shortly justify in the context of a multi-period model. Substituting this budget into the utility function, we can eliminate c_t and express utility in terms of gasoline consumption and prices:

$$v_t = (a_t - p_t g_t + \gamma_c)^{\sigma_c} + \theta (g_t - \delta s_{gt} + \gamma_g)^{\sigma_g}$$

Now, assuming low_t and $high_t$ are known, we can solve for the expectation of v_t :

$$\begin{aligned} E[v_t] &= \int_{low_t}^{high_t} v_t f(p_t) dp_t \\ &= \frac{1}{high_t - low_t} \int_{low_t}^{high_t} v_t dp_t \\ &= \frac{1}{high_t - low_t} \int_{low_t}^{high_t} (a_t - p_t g_t + \gamma_c)^{\sigma_c} dp_t + \frac{1}{high_t - low_t} [\theta (g_t - \delta s_{gt} + \gamma_g)^{\sigma_g} p_t]_{p_t=low_t}^{p_t=high_t} \\ &= \frac{1}{high_t - low_t} \left[-\frac{(a_t + \gamma_c - p_t g_t)^{\sigma_c+1}}{g_t (\sigma_c + 1)} \right]_{p_t=low_t}^{p_t=high_t} + \theta (g_t - \delta s_{gt} + \gamma_g)^{\sigma_g} \\ &= \frac{1}{high_t - low_t} \left[\frac{(a_t + \gamma_c - low_t g_t)^{\sigma_c+1} - (a_t + \gamma_c - high_t g_t)^{\sigma_c+1}}{g_t (\sigma_c + 1)} \right] + \theta (g_t - \delta s_{gt} + \gamma_g)^{\sigma_g} \end{aligned}$$

The expectation of utility at time t , therefore, can be expressed as a function of known bounds of the uniform time- t price distribution, current gasoline consumption, and the habit-stock s_{gt} . In

Appendix C, I show that such a closed-form expression for expected utility exists for a variety of additively- and multiplicatively-separable utility functions.

The closed form of $E[v_t]$ makes it straightforward to maximize expected time- t utility with respect to time- t gasoline consumption. Similarly, it is now straightforward to consider a multiperiod model, in which the agent's problem is to maximize a discounted sum of present and expected future utility:

$$\begin{aligned} \text{Max}_{g_1, g_2, g_3, \dots, g_T} \quad & v_1 + \sum_{t=2}^T \beta^{t-1} E[v_t] \\ \text{s.t.} \quad & \text{non-negativity constraints (15), (16), (17), and (18)} \end{aligned} \quad (13)$$

where β is the agent's discount factor, T is his time horizon, and the current ($t = 1$) price is known.

What are the relevant non-negativity constraints? Note that the budget has been built into the expression of utility v , and that a no-borrowing rule can therefore be enforced with a simple constraint:

$$c_t = a_t - p_t g_t \geq 0 \quad \forall t \quad (14)$$

Condition (14) applies at $t = 1$:

$$c_1 = a_1 - p_1 g_1 \geq 0 \quad (15)$$

From $t = 2$ onward, however, our interest is in *expected* utility rather than utility. Since price p_t itself does not enter into $E[v_t]$, we must rethink these constraints. Mathematically speaking, we can't allow either $(a_t + \gamma_c - \text{low}_t g_t)$ or $(a_t + \gamma_c - \text{high}_t g_t)$ to be negative.

$$\begin{aligned} a_t + \gamma_c - \text{low}_t g_t &\geq 0, \quad t \geq 2 \\ g_t &\leq \frac{a_t + \gamma_c}{\text{low}_t}, \quad t \geq 2 \end{aligned}$$

$$\begin{aligned} a_t + \gamma_c - \text{high}_t g_t &\geq 0, \quad t \geq 2 \\ g_t &\leq \frac{a_t + \gamma_c}{\text{high}_t}, \quad t \geq 2 \end{aligned} \quad (16)$$

But of course $\frac{a_t + \gamma_c}{\text{high}_t} \leq \frac{a_t + \gamma_c}{\text{low}_t}$, so only $g_t \leq \frac{a_t + \gamma_c}{\text{high}_t}$ binds. This means that the agent's gasoline consumption must always be less than what he could buy if prices were at their maximum and he spent all his money on gas ($\frac{a_t}{\text{high}_t}$) plus some small constant ($\frac{\gamma_c}{\text{high}_t}$)—or, rather, *minus* some small constant if γ_c is negative and commits the agent to consuming some amount of the other good c .

In addition to these budget-related constraints, we must also ensure that both gasoline consumption (17) and the "adjusted" gasoline quantity over which the agent takes his preferences (18) are non-negative:

$$g_t \geq 0 \quad \forall t \quad (17)$$

$$g_t - \delta s_{gt} + \gamma_g \geq 0 \quad \forall t \quad (18)$$

Returning to the budget constraints, why eliminate the possibility of borrowing and saving across time? This decision is crucial in allowing us to solve easily for $E[v_t]$ —without it, a_t becomes a choice variable contingent upon a maximization of expected utility across the agent's entire time horizon. The agent's allocation of spending across time will depend both on his beliefs about future prices and on the size of the gasoline habit he anticipates holding in every future period—and the latter will depend upon his beliefs about future prices. Given the possibility of borrowing and saving, therefore, the agent's problem reverts back to the computationally-infeasible dynamic programming problem discussed in Section 3.2 and Appendices A and B. From a technical viewpoint, the elimination of intertemporal allocation decisions is thus vital.

Note that although this ban on borrowing and saving imposes limits on the consumption of g and c in every period, it is not related in any practical way to the literature on rationing. In the rationing literature (see, e.g., Neary and Roberts 1980 or Ellis and Naughton 1990), limits

are imposed not on *overall* spending in a period, but on the consumption of a particular good or goods. Here, the limit is imposed on overall period expenditure: given a fixed expenditure level, spending may be allocated in any way across goods.

A convenient side-effect of the ban on borrowing and saving is that it allows us to examine the effects of habits without the confounding veil of intertemporal substitution. Imagine for a moment that the agent *could* borrow and save across time. When faced with a future increase in the gasoline price, his current reaction would be governed by two forces:

1. habits would drive him to reduce current gasoline consumption in order to reduce the cost of servicing his habit in the future, and
2. the current-future price differential would drive him to substitute cheap current gasoline for expensive future gasoline.

Our ban on borrowing prevents the second effect, allowing us to examine the habits-effect in isolation.

3.4 Limitations of the Model

Without the assumptions we've just made, we would be stuck. *With* the assumptions, we are able to examine consumer responses to a limited type of price uncertainty. It is important to bear in mind, however, what the consumer's problem outlined above actually addresses—and what it doesn't.

The first limitation of this model is, of course, the uniformity of the distribution from which prices are drawn. Ordinarily a uniform distribution would not be our first choice. Statistically, we would usually not model a process using errors drawn from a uniform rather than, say, a normal distribution; and intuitively, we'd expect gasoline prices to cluster around a mean, with larger deviations occurring less frequently.

The second limitation of the model is its lack of consumer learning. This is, in fact, a much greater concession than uniformity. By "lack of learning", I mean that the agent behaves as if *certain* of future price *distributions*. He bases his beliefs about those future distributions on information he receives in the initial period, and he does not revise those beliefs in response to information he receives in subsequent periods. This implies that the realization of tomorrow's price does not affect the distribution of prices thereafter—or, rather, that the agent does not *behave* as if he'll receive new information tomorrow and have the chance to re-optimize given his revised beliefs. In short, this is not a dynamic programming problem.

Although the agent in this model is certain about future price distributions, this certainty does not imply *constant* distributions. The sole restriction on the price process (or on the agent's beliefs thereof) is that no shock that occurs after the initial period can affect subsequent price distributions. A shock that occurs in or before the initial period can propagate along according to any model at all. The agent may believe prices to be trending upward or downward or reverting to a long-run mean; he may believe price variance to be increasing or decreasing or fluctuating; indeed, he may believe the mean and variance of future prices to be hopping along a path that is purely arbitrary, as long as he knows from the outset what this path will be. This is the redeeming beauty of this model: we can examine consumer behavior under any price process we want—and, in particular, under price processes we've estimated, and variations thereof.

In sum, this model allows us to examine a very *particular* type of price uncertainty—uncertainty arising from the variance of known, uniform future price distributions. It also allows us to examine the behavioral effects of the processes governing the mean and variance of price. We can explore, for example, what happens when we increase the duration of a shock to the mean price. We cannot, however, use this model to examine the effects of uncertainty about the future price *distribution*.

4 Applying the Model

Our modifications have given us a model that we can solve. Now we need to solve it—and put it to use exploring how consumer habits and the price process shape the elasticity of demand for gasoline.

To solve the maximization problem (13) analytically and symbolically would be unworkable, at least for reasonable values of the time horizon T . Instead, I choose parameter values and perform the consumer's optimization numerically, via search. To measure elasticity, I displace the gasoline price by some small ϵ , re-optimize, and calculate the percent change in first-period gasoline consumption. By varying the parameters of the model and price process, I can sketch out the consumer's responsiveness under a continuum of scenarios.

4.1 The Price Process

As laid out above, the consumer knows at $t = 1$ the uniform distributions from which all future prices will be drawn, and beyond this there are no restrictions on future price distributions.

For clarity of illustration, I will allow the mean of future price distributions to follow a simple mean-reverting process:

$$p_1 = p_0 + shock_1 \quad (19)$$

$$p_t = p_0 + \rho^{t-1} shock_1 + u_t, \quad t = 2, \dots, T \quad (20)$$

where

- p_1 is the known, realized price at $t = 1$;
- $shock_1$ is the portion of the period-1 price that comes as a deviation from the long-run mean p_0 ;
- $\rho \in [0, 1]$ governs the speed with which the price reverts to its long-run mean; and
- u_t is the unanticipated component of p_t , drawn from some uniform distribution centered on 0.

Note that the only "shock" that propagates through future prices is $shock_1$: each u_t is a single-period "shock" that affects the price *only* in period t . Also note that the actual *realizations* of p_2 through p_T are irrelevant to the consumer's problem: the agent cares only about the bounds of future price distributions, which are calculated as $E[p_t] \pm \frac{1}{2} spread_t$, where $spread_t$ is the length of the support of the uniform distribution at t .

This process for the mean is not intended as a realistic representation of the gasoline price process. It has been adopted, instead, because it encapsulates the persistence of price shocks in a single parameter, ρ , and thus allows us to see quite clearly how the durability of price shocks affects demand behavior.

4.2 Optimization Methods and Calculation of Elasticity

I solve for the optimal consumption path using line search methods. Details of the optimization procedure are provided in Appendix D.

To calculate elasticity, I first optimize under a base scenario in which $shock_1 = 0$. I then displace the initial price by ϵ by setting $shock_1 = \epsilon = 0.1$, propagating this shock through all subsequent time periods according to (20).⁷ If $\rho = 0$, future price distributions remain the same and only p_1 changes; if $\rho = 1$, all future distributions shift upward by ϵ ; and if $0 < \rho < 1$, future price distributions shift upward, but their means creep back toward p_0 with time.

Under this new price regime, I re-optimize. I calculate the agent's price elasticity as

$$\frac{\% \Delta g_1^*}{\% \Delta p_1} \approx \frac{\frac{g_1^* - g_1}{g_1^*}}{\frac{\epsilon}{p_0}}$$

where g_1^* is the agent's optimal period-1 gasoline consumption under the initial price scenario and g_1' is his optimal period-1 consumption after the small price displacement.

⁷Note that it is not *necessary* for the base scenario to be $shock_1 = 0$. I could just as easily base the calculation of elasticity off of a scenario in which the agent is already adjusting to another price shock.

4.3 Beginning- and Endpoint Concerns

For the time being, I wish to sketch out elasticity's theoretical relationship with various parameters. In the future it may be interesting to choose realistic parameter values and calibrate the model to real-world data, but right now it suffices to choose a convenient set of parameters and consider variations from that base.

In particular, let us start from the following set of parameters:

Price Process Parameters	Model Parameters
$p_0 = 3$	$T = 10$
$\rho = 0.5$	$a_t = 10000 \forall t$
$spread_t = 4 \forall t \geq 2$ (i.e. $var(p_t) = \frac{4}{3}, t \geq 2$)	$\gamma_c = \gamma_g = 0.5$
	$\sigma_c = \sigma_g = 0.5$
	$\theta = 0.3$
	$\beta = 0.9$
	$\delta = 0.5$
	$\alpha = 1$

Recalling the definition of the habit stock, $s_{gt} = \alpha g_{t-1} + (1 - \alpha) s_{g,t-1}$, note that the choice $\alpha = 1$ implies short-memory habits. It also implies that $s_{g1} = g_0$, so the starting point for habits is just time-0 gasoline consumption.

The initial consumption level g_0 cannot be chosen as arbitrarily as the rest of the set of base parameters. Too high a g_0 , and the agent finds himself drastically cutting back on gasoline in period 1; too low, and the agent begins with a sudden gasoline binge. Whether a given g_0 is too high or too low depends on the other parameters in the model; and in turn, g_0 affects not only first-period elasticity, but the interactions between other parameters and first-period elasticity.

To show that g_0 affects first-period elasticity, Figure 1 plots elasticity (given the parameter values above) against a range of starting values g_0 . The agent is adjusting his initial consumption upward ($g_1^* > g_0$) through $g_0 = 55$ and downward from $g_0 = 60$, and clearly he is much more sensitive to a small price increase when in the process of *increasing* consumption than when decreasing it.

That g_0 affects other parameters' interactions with elasticity is clear from Figure 2, which plots elasticity versus ρ for three different values of g_0 . The lower the value of g_0 , the more sensitive elasticity is to ρ , which governs the durability of a price shock.

Clearly we cannot set g_0 arbitrarily. What, then, should we choose as the starting value?

Ideally we might want to set g_0 equal to the long-run equilibrium g_t^* under the base parameters and constant prices. As we will soon see, however, such an equilibrium does not exist because of the distortionary effect of the agent's finite horizon.

4.3.1 The Finite-Horizon Effect

The finiteness of the agent's horizon is necessary to a tractable model but comes at a price: it forces a gasoline binge as t approaches T .

To see why, recall the utility function, and now specify short-memory habits:

$$u_t = (c_t + \gamma_c)^{\sigma_c} + \theta (g_t - \delta g_{t-1} + \gamma_g)^{\sigma_g}$$

In most periods, gasoline's positive contribution to utility is offset by its negative habit-effect the following period. Period T gasoline consumption, however, is never penalized, as there is no period $T + 1$. This means the agent will *always* want to consume more gasoline in period T . As long as the consumer has habits ($\delta \neq 0$), moreover, he will begin increasing gasoline consumption in the periods leading up to T . In short, there is no "equilibrium" level at which the consumer would be happy to consume gasoline every period. For an illustration of the agent's consumption path, see Figure 3.




Figure 1: Elasticity versus g_0




Figure 2: Elasticity vs. ρ , for varying g_0




Figure 3: Gasoline Consumption over Time

4.3.2 Choosing g_0

Given that gasoline consumption will never settle toward a long-run equilibrium, how shall we choose g_0 ?

Observe in Figure 3 that although gasoline consumption rises sharply toward T , it does plateau around the middle of the time horizon. If we set g_0 near the level of this plateau, optimal consumption (under the constant-mean-price scenario) should be relatively stable at the beginning of the agent's horizon.

To find a value of g_0 such that g_0 is near the plateau, I iterate through the agent's problem several times. First I set $g_0 = 0$ and optimize. From the resulting consumption path, I take g_5^* and use this as my new g_0 . Re-optimizing, I take g_5^* from the new resulting consumption path and again replace g_0 . I do this a total of five times, finally setting g_0 equal to g_5^* from the fifth optimization. For the base parameter set, this process yields $g_0 = 58.9119$. Although the level of the plateau depends on the model's parameters, I will keep g_0 at this level unless otherwise noted.




Figure 4: Choosing g_0 : Consumption paths in first and fifth iterations

This process for choosing g_0 does not eliminate the apocalyptic gasoline binge, but it does ensure steady consumption over most of the horizon. With T sufficiently large, end-of-horizon behavior is hugely discounted, and its effect on first-period behavior is insignificant. (Indeed, flattening the end-of-horizon binge by introducing the restriction $g_T = g_{T-1}$ leaves first-period elasticities virtually unchanged.) In Section 5.1, I verify that $T = 10$ is sufficiently large as to render end-of-horizon behavior moot.

5 Results

Our original motivation was to discover how habits interact with the price process to shape consumption behavior. By manipulating the model parameters and the price process, we can now simulate these effects.

5.1 Time Horizon

My first concern is validating my choice of time horizon T . Although in Section 4.3.2 I demonstrated that the finite-horizon effect was not driving first-period results, I can now demonstrate that the choice of T has practically no effect beyond a certain point, and that $T = 10$ is a sufficiently long horizon.

Figure 5 shows optimal first-period consumption for models with $T = 5, 10, 25$, and 50 . Consumption drops 1.77% from $T = 5$ to $T = 10$, but after that it only falls another 0.03% from $T = 10$ to $T = 50$. Increasing the agent's horizon beyond ten periods, therefore, offers little in exchange for a high computational price.




Figure 5: Optimal first-period consumption vs. horizon T

This holds for elasticity as well as consumption levels. Figure 6 shows price elasticity at the base parameter set for $T = 5, 10, 25$, and 50 . Again, there is a slight difference between elasticity at $T = 5$ and $T = 10$, with price sensitivity decreasing somewhat with the longer horizon; but after that the differences are imperceptible. For the purposes of this model, ten and infinite are roughly equivalent.

5.2 Durability of Price Shocks

Given rational habits, consumers are sensitive not just to the current price of gasoline, but to the duration of price shocks. Figure 7 illustrates this sensitivity. At the base parameter set, the difference between the elasticity with respect to a purely-temporary shock ($\rho = 0$) and the elasticity with respect to a permanent shock ($\rho = 1$) is nearly two-fold: -0.56 vs. -0.99.

Shortly I will look at the effects on elasticity as ρ interacts with other parameters.

5.3 Future Price Uncertainty

Observing the effect of future price uncertainty is easier when θ is increased above its base value of 0.3, so I temporarily set $\theta = 0.5$. I adjust g_0 accordingly, as per the procedure outlined in Section 4.3.2.




Figure 6: Elasticity vs. time horizon, for varying ρ




Figure 7: Elasticity vs. durability of price shocks (ρ)

Figure 8 shows the effect of changing the spread of the uniform distribution from which prices are drawn for all periods $t = 2, \dots, T$. As uncertainty increases, price sensitivity slightly *decreases*. This is true whether price shocks are fleeting or permanent: no matter what process governs the future price's *mean*, increasing the future price's *variance* reduces the magnitude of elasticity. The mean path and the variance do not act entirely independently on elasticity, however. Although it is hard to observe in Figures 8 and 9, increasing ρ magnifies the effect of future uncertainty.




Figure 8: Effect of future price uncertainty, for varying price-shock durations

Just as interesting as the direction of the effect of price uncertainty is its magnitude—namely, tiny. Increasing the spread of all future price distributions from 0.1 to 6 means increasing the variance of those distributions from 0.00083 to 3—a factor of 3600—and yet only changes elasticity by 0.9 to 1.5%. Figure 9, which plots elasticity against ρ for $spread_t = 0.1$ and $spread_t = 6$, illustrates just how small the effect of price uncertainty is in comparison to the effect of price-shock duration. In the real world, of course, price uncertainty does not take the restricted form assumed here, where the means of future price distributions are known with certainty. The relatively large effect of ρ , which controls these future means, suggests that real-world uncertainty about future mean prices may give uncertainty a considerably larger effect on price responsiveness. Although price volatility and uncertainty about future price distributions are *related*, the current restrictions reveal that the latter is likely to play a larger role in shaping price elasticity.

In addition to price uncertainty's effect on the price elasticity of demand, it has a small effect on the *level* of demand. Figure 10 illustrates this effect: increasing the price spread from 0.1 to 6 for all $t \geq 2$ decreases first-period consumption by 0.37%. This effect is consistent with Coppejans et al.'s (2007) finding that increasing the variance of future prices depresses demand.

5.4 Habit Strength

All other things equal, stronger habits imply lower price sensitivity. This relationship is clear in Figure 11, which plots elasticity against the habit-strength parameter δ . When $\rho = 0.5$, increasing




Figure 9: Elasticity vs. price-shock duration, for varying uncertainty




Figure 10: First-period consumption vs. price uncertainty




Figure 11: Elasticity vs. habit-strength, for varying ρ

δ from 0 (no habits) to 1 (strong habits) decreases the magnitude of elasticity from 1.95 to 0.020—from very elastic to almost inelastic. As we would expect, the process of the mean price has no

effect on elasticity when there are no habits. As Figure 11 suggests, ρ also becomes unimportant as δ approaches 1, rendering last period's consumption as relevant to current utility as this period's. At intermediate levels of δ , the duration of price shocks has an appreciable effect on elasticity, with higher ρ implying higher price sensitivity. At $\delta = 0.5$, the difference between a single-period shock ($\rho = 0$) and a permanent shock ($\rho = 1$) leads, as previously noted, to a near-twofold difference in elasticity.

Like price uncertainty, habit strength δ affects optimal consumption levels as well as price sensitivity. Figure 12 plots optimal first-period consumption against habit strength: first-period consumption decreases with δ until about $\delta = 0.7$, whereupon the trend reverses.

5.5 Utility Weight on Habit-Forming Good

The weight on gasoline in the agent's utility function, θ , increases price sensitivity. This relationship is sketched out in Figure 13, which also makes clear θ 's decreasing marginal effect: as θ approaches 1 and the weights on gasoline and all other consumption approach equality, the elasticity curve becomes flatter and flatter. The weight θ could, in theory, exceed 1; but in practice this would imply that gasoline were more important to utility than all other consumption combined—an outlandish supposition.

Predictably, utility's weight on gasoline also affects optimal consumption levels, with first-period consumption g_1^* increasing with θ : see Figure 14.

Also predictably, a higher weight on gasoline in the utility function increases elasticity's sensitivity to the duration of price shocks. Figure 15 illustrates the cross-effects between θ and ρ . Note how the slope of the elasticity-vs.- ρ curve deepens for higher θ . This deepening is greatest for small θ , decreasing on the margin as θ steps down to 1.




Figure 12: Optimal first-period consumption vs. habit strength




Figure 13: Elasticity vs. utility's weight on gasoline




Figure 14: Optimal first-period consumption vs. utility's weight on gasoline




Figure 15: Elasticity vs. price-shock durability, ρ , for varying θ

5.6 Discounting of Future Utility

The final parameters to consider, β and T , are those that control the agent's view of the future.

As the agent's discount factor β increases—and with it, his concern for his future utility—both first-period gasoline consumption and price sensitivity decrease. Figures 16 and 17 illustrate these effects.




Figure 16: Optimal first-period consumption vs. discount factor β

As β increases, the duration of price shocks becomes more important to elasticity. This fits with intuition, and Figure 18 illustrates the effect.

6 Conclusion

By assuming a very particular type of price uncertainty, we have been able to solve and apply a model of demand with rational habits, multiple goods, and uncertain relative prices. This approach has allowed us to examine, in particular, the combined effect of rational habits and the price process on demand's *responsiveness* to prices—an area that has been completely unexplored until now.

Through simulations, we've sketched how the parameters of the model and the price process shape gasoline demand. Knowing the effects of these parameters has the potential for many practical uses. If policymakers want to make consumers more price-responsive, for example, they might try to reduce consumers' habit strength, δ , through strategies that lower the barriers to behavioral change. Schemes to make public transport schedules more navigable or provide instruction on fuel-efficient driving might reduce habit strength; so too would lowering the frictional costs associated with changing the distance of a commute. (Stamp duty in the U.K., for example, discourages consumers from relocating closer to a job; policies that reduce job mobility, such as employer-tied health insurance in the U.S., discourage consumers from switching to jobs closer to home.)

A more likely use of knowing how the model's parameters shape gasoline demand comes in predicting consumption responses to price instruments. In a world conforming to the model's




Figure 17: Elasticity vs. discount factor β , for varying ρ




Figure 18: Elasticity vs. price-shock duration, for varying β

assumptions, for example, policymakers contemplating a permanent change in the gasoline tax should forecast the response using $\rho = 1$ rather than, say, $\rho = 0.5$: under the base parameter set, this would imply an elasticity of -0.99 instead of -0.71. More generally, policymakers should not use elasticities based on consumers' responses to all manner of price changes to project their response to a permanent, policy-induced price change.

One of the more interesting implications of this model is that the effect of expected price volatility around a known mean is tiny: elasticity responds *far* less to a reduction in future price variance than to an increase in the duration of a change in the mean price. Given the restricted nature of uncertainty in this model, this does not necessarily imply that price uncertainty is unimportant in practice. In the current model, all future price distributions are *known*. In the real world, where uncertainty exists about future price *distributions*, that uncertainty may play a greater role in shaping demand. Indeed, the importance of price shocks' persistence (ρ)—and by extension, the importance of future mean prices—suggests that uncertainty about price *distributions* is likely to be important. The current restrictions allow us to see that uncertainty about the mean of future prices is likely to be more important to price responsiveness than volatility around those future means. If policymakers wish to maximize the effectiveness of price instruments, they should concentrate on making changes in the mean price credibly permanent rather than on reducing price volatility—particularly as reducing price volatility has the knock-on effect of increasing the *level* of demand.

Of course, the model's restrictions on price uncertainty are not its only differences from the real world: also significant is its imposition of a ban on borrowing and saving. In reality, consumers may make tradeoffs not only between gasoline consumption and other contemporaneous consumption, but between gasoline consumption and savings. How this ability to make intertemporal tradeoffs will interact with the price process to shape demand is not clear-cut. On the one hand, if the gasoline price increases today and $\rho < 1$ —so that the shock is expected to decay—then consumers may save relatively more today in order to take advantage of the future's lower prices. If that extra savings comes out of gasoline consumption, then the ability to borrow and save may *magnify* the immediate response to a gasoline-price increase beyond the response predicted by the model. On the other hand, intertemporal allocation opens the possibility that consumers may draw from savings in order to cushion themselves and service their habits during temporary price shocks; and this would render consumers' immediate price-responsiveness *lower* than that predicted by the model.

Although intertemporal substitution may somewhat *affect* the effects of habits, however, it will not eliminate them. Even in a world with flexible intertemporal budgets, the model supplies useful insights, giving us an idea of what to expect relative to a no-habits scenario. The model's implications about the importance of price uncertainty and the duration of price shocks, in particular, remain relevant—and testable—in the real world.

To realize the model's practical uses will of course require reconciling it with real-world data. An obvious approach, if we were unconcerned by the model's simplifications of price uncertainty and intertemporal allocation, would be to model the gasoline price process and then choose parameters such that the model-predicted behavior mimicked observed behavior. Once calibrated or estimated, the model would allow us to better predict consumers' responsiveness to price changes that differed from the typical price fluctuations over which elasticity is measured. It would allow us to see how much of the variation in price elasticity across regions was merely the product of regional differences in price processes. Continuing in this vein, it would allow us to see how much of the variation in elasticity across regions remained to be explained by other factors—and whether attacking these other sources of heterogeneity would offer greater scope for manipulating consumer responsiveness than merely adjusting the gasoline price process.

For now, the model identifies biases that we must beware when using traditional measures of price elasticity with habit-forming goods.

A Setting up the Ideal Problem as a Dynamic Programming Problem

We can reformulate the problem given in (12) as a Bellman equation:

$$V_t(A_t, s_t, p_t, x_{1t}, \dots, x_{nt}) = \max_{g_t, c_t} \{u(\bar{g}_t, c_t) + \beta E[V_{t+1}(A_{t+1}, s_{t+1}, p_{t+1}, x_{1,t+1}, \dots, x_{n,t+1})]\} \quad (21)$$

where

- $\bar{g}_t = g_t - \delta s_t$;
- A_i indicates assets in period i ;
- V_i is the value in period i ;
- $x_{1t}, \dots, x_{nt}$ are some variables that are relevant to predicting the future price; and
- the equations of motion for assets and the habit-stock are given by:

$$\begin{aligned} A_{t+1} &= (1+r)(A_t - c_t - p_t g_t) \\ s_{t+1} &= \alpha g_t + (1-\alpha)s_t \end{aligned}$$

To get equations of motion for the remaining state variables, let us suppose that the price process is given by

$$p_{t+1} = f(x_{1t}, \dots, x_{it}, \dots, x_{nt}) e^{\varepsilon_{t+1}}$$

where ε_{t+1} is random and serially uncorrelated. This particular form for the price process is not *necessary*, but it is flexible—and it is also convenient, as the random variation enters in a way that cannot lead to negative prices. The x variables used to predict future prices may be or may depend upon lagged prices or lagged x :

$$x_{i,t+1} = h_i(I_t) e^{\varepsilon_{xi,t+1}}$$

where the function h_i maps the information I available at time t to the next period's expected value of x_i , and $\varepsilon_{xi,t+1}$ can be either deterministic or a random component specific to x_i . Note, in particular, that if $\varepsilon_{xi,t+1} = 0$, $x_{i,t+1}$ could be the price p_t or a lag thereof.

We now have equations of motion for the gasoline price,

$$p_{t+1} = f(p_t, x_{1t}, \dots, x_{nt}) (e^{\varepsilon_{t+1}})$$

and the price-predicting state variables,

$$x_{i,t+1} = h_i(I_t) e^{\varepsilon_{xi,t+1}}$$

Substituting the equations of motion into (21), we have

$$V_t(A_t, s_t, p_t, x_{1t}, \dots, x_{nt}) = \max_{g_t, c_t} \left\{ u(g_t - \delta s_t, c_t) + \beta E \left[V_{t+1} \left[\begin{array}{c} \underbrace{(1+r)(A_t - c_t - p_t g_t)}_{A_{t+1}}, \\ \underbrace{\alpha g_t + (1-\alpha)s_t}_{s_{t+1}}, \\ \underbrace{f(x_{1t}, \dots, x_{it}, \dots, x_{nt}) (e^{\varepsilon_{t+1}})}_{p_{t+1}}, \\ \underbrace{h_1(I_t) e^{\varepsilon_{1,t+1}}, h_2(I_t) e^{\varepsilon_{2,t+1}}}_{x_{1,t+1}, x_{2,t+1}}, \\ \dots, \underbrace{h_i(I_t) e^{\varepsilon_{i,t+1}}}_{x_{i,t+1}}, \dots, \underbrace{h_n(I_t) e^{\varepsilon_{n,t+1}}}_{x_{n,t+1}} \end{array} \right] \right] \right\}$$

Given this value function equation, we can derive first-order conditions for the maximum of the right-hand side as well as envelope conditions that indicate how the maximized value V varies with each of the state variables. These are standard conditions to consider and follow, for example, Deaton (1992, pp. 21-37), who derives them for a simpler problem.

First-Order Conditions

There will be two first-order conditions, as the agent derives utility from two goods:

1. FOC g :

$$\begin{aligned} u_{\bar{g}} + \beta E[V_{1,t+1}(1+r)(-p_t) + V_{2,t+1}\alpha] &= 0 \\ \text{or} \\ u_{\bar{g}} &= \beta(1+r)(p_t)E[V_{1,t+1}] - \beta\alpha E[V_{2,t+1}] \end{aligned} \quad (22)$$

2. FOC c :

$$\begin{aligned} u_c + \beta E[V_{1,t+1}(1+r)(-1)] &= 0 \\ \text{or} \\ E[V_{1,t+1}] &= \frac{u_c}{\beta(1+r)} \end{aligned} \quad (23)$$

Combining these first-order conditions allows us to write $E[V_{2,t+1}]$ as a function of current prices and marginal utilities:

$$E[V_{2,t+1}] = \frac{1}{\beta\alpha}(u_c p_t - u_{\bar{g}}) \quad (24)$$

Envelope Conditions

There are $n+3$ envelope conditions, one for each state variable:

1. Marginal utility of wealth: $V_{A,t} = V_{1,t} = \beta(1+r)E[V_{1,t+1}]$

Combining this with (23), we see that the marginal utility of wealth is equal to the marginal utility of c :

$$V_{1,t} = u_c$$

2. Marginal "utility" of habit-stock: $V_{s,t} = V_{2,t} = -\delta u_{\bar{g}} + \beta(1-\alpha)E[V_{2,t+1}]$

Combining this with (24), we can write

$$V_{2,t} = \frac{1-\alpha}{\alpha}u_c p_t - \left(\delta + \frac{1-\alpha}{\alpha}\right)u_{\bar{g}}$$

3. Condition with respect to current price p_t :

$$V_{p,t} = V_{3,t} = \beta E \left[V_{1,t+1}(1+r)(-g_t) + V_{3,t+1} \left[\sum_{j=1}^n f_j(x_{1t}, \dots, x_{it}, \dots, x_{nt}) \frac{\partial x_j}{\partial p_t} \right] (e^{\varepsilon_{t+1}}) + \sum_{j=1}^n V_{3+j,t+1} \frac{\partial h_j}{\partial p_t} e^{\varepsilon_{j,t+1}} \right]$$

Combining this with (23) yields

$$\begin{aligned} V_{p,t} &= V_{3,t} = -g_t u_c + \beta E \left[V_{3,t+1} \left[\sum_{j=1}^n f_j(x_{1t}, \dots, x_{it}, \dots, x_{nt}) \frac{\partial x_j}{\partial p_t} \right] (e^{\varepsilon_{t+1}}) \right] \\ &\quad + \beta E \left[\sum_{j=1}^n V_{3+j,t+1} \frac{\partial h_j}{\partial p_t} e^{\varepsilon_{j,t+1}} \right] \end{aligned}$$

4. n conditions with respect to price-prediction variables

$$x_{kt}, k = 1, \dots, n : V_{3+k,t} = \beta E \left[V_{3,t+1} f_k(x_{1t}, \dots, x_{it}, \dots, x_{nt}) (e^{\varepsilon_{t+1}}) + \sum_{j=1}^n V_{3+j,t+1} \frac{\partial h_j}{\partial x_{kt}} e^{\varepsilon_{j,t+1}} \right]$$

B Approaches to the Value Function Problem and the Curse of Dimensionality

The problem set up in Appendix A now looks like a traditional (albeit complicated) stochastic programming problem, to which one could apply standard techniques.

B.1 Value Function Iteration

The first technique to which we might generally turn is value function iteration.⁸ To apply this method to our problem with continuous state and continuous control variables, we would first need to discretize the state space. Ignoring for the moment any difficulties with this discretization, let us imagine that we discretize each of the $n + 3$ state variables (indexed by s) into some number of points, d_s . This gives us an $(n + 3)$ -by- $(n + 3)$ matrix containing all the possible combinations of these points, with a total of $D = \prod_{s=1}^{n+3} d_s$ non-blank gridpoints. Let $q = 1, \dots, D$ be an index of these gridpoints.

The method of value function iteration hinges on the stationarity of the value function V : for any set of state variables $\mathbf{x}$, $V_t(\mathbf{x}) = V_{t+1}(\mathbf{x}) = V(\mathbf{x})$. In other words, stationarity implies that our set of state variables captures *any and all* influences on value that can vary over time. We can rewrite the value function without the time subscripts on V :

$$V(A_t, s_t, p_t, x_{1t}, \dots, x_{nt}) = \max_{g_t, c_t} \left\{ u(g_t - \delta s_t, c_t) + \beta E \left[V \left[\begin{array}{c} \underbrace{(1+r)(A_t - c_t - p_t g_t)}_{A_{t+1}}, \underbrace{\alpha g_t + (1-\alpha)s_t}_{s_{t+1}}, \\ \underbrace{f(x_{1t}, \dots, x_{it}, \dots, x_{nt})}_{p_{t+1}} (e^{\varepsilon_{t+1}}), \\ \underbrace{h_1(I_t) e^{\varepsilon_{1,t+1}}}_{x_{1,t+1}}, \underbrace{h_2(I_t) e^{\varepsilon_{2,t+1}}}_{x_{2,t+1}}, \\ \dots, \underbrace{h_i(I_t) e^{\varepsilon_{i,t+1}}}_{x_{i,t+1}}, \dots, \underbrace{h_n(I_t) e^{\varepsilon_{n,t+1}}}_{x_{n,t+1}} \end{array} \right] \right] \right\}$$

or

$$V(\mathbf{x}_t) = \max_{g_t, c_t} \{ u(g_t - \delta s_t, c_t) + \beta E[V(\mathbf{x}_{t+1})] \}$$

where the transition from $\mathbf{x}_t$ to $\mathbf{x}_{t+1}$ is governed by the previously-discussed equations of motion. For convenience, write the probability of transitioning from the state given by gridpoint q to the state given by gridpoint r as π_{qr} :

$$\pi_{qr} = \Pr[\mathbf{x}_{t+1} = \mathbf{x}_r | \mathbf{x}_t = \mathbf{x}_q]$$

To apply value function iteration, we would first make some initial guess, V^0 , about the form of the value function. For now, it will be helpful to think of value as a function of the *controls* c and g , given a particular state q : that is, imagine that we guess the functions $V_q^0(c_t, g_t)$, $q = 1, \dots, D$. Using these guesses together with the probabilities π_{qr} , we can write a guess of expected future value, given starting state-gridpoint q , as the weighted average $\sum_{r=1}^D \pi_{qr} V_r^0$. We can then *improve* our guess at the value function by solving the maximization problem

$$V_q^1 = \max_{g_t, c_t} \left[u(g_t - \delta s_t, c_t) + \beta \sum_{r=1}^D \pi_{qr} V_r^0 \right], \quad q = 1, \dots, D \quad (25)$$

⁸For a more detailed and technical exposition of value function iteration, see Judd (1998), Chapter 12; Miranda and Fackler (2002), Chapters 8 and 9; and/or Ljungqvist and Sargent (2004), Chapters 3 and 4. This walk-through of how value function iteration might be applied to the problem at hand relies most heavily on Judd (1998), with additional background garnered from Mele (2009).

The resulting improved guesses, V_q^1 , $q = 1, \dots, D$, can then be substituted into the right-hand side of (25), and we can again solve the maximization problem to improve our guess at the value function:

$$V_q^2 = \max_{g_t, c_t} \left[u(g_t - \delta s_t, c_t) + \beta \sum_{r=1}^D \pi_{qr} V_r^1 \right], \quad q = 1, \dots, D$$

This process of substituting and re-maximizing can be iterated indefinitely:

$$V_q^{a+1} = \max_{g_t, c_t} \left[u(g_t - \delta s_t, c_t) + \beta \sum_{r=1}^D \pi_{qr} V_r^a \right], \quad q = 1, \dots, D \quad (26)$$

If our value function and probability distributions meet a variety of conditions, then these iterations should converge toward the true value function: we can just continue iterating until $|V^{a+1} - V^a|$ is sufficiently small, then calculate the policy associated with this value function.

I will not venture into *proving* that this value function iteration procedure would indeed converge upon the true value function, nor do I wish to explore in any detail the conditions that value function iteration might necessitate imposing. My primary concern is, instead, the workability of value function iteration *in practice*. No matter how sublimely the approach of value function iteration might solve our problem *in theory*, in practice the approach will be cursed to uselessness by the many dimensions of the state space.

The curse of dimensionality is a common barrier in value function iteration. The problem arises from the discretization of our continuous state variables. Recall that we have $n + 3$ state variables, each discretized into d_s points, for a total of $D = \prod_{s=1}^{n+3} d_s$ discrete states. If we chose d_s to be the same for all s , $d_s = d \forall s$, we would have $D = d^{n+3}$. Also recall that we must perform D maximizations for every iteration of the value function (see (26)). Even in the simplest scenario, with a rough discretization of the state variables ($d = 100$, say) and a Markov process for prices ($n = 0$), we would have to perform a million maximizations for every iteration. To consider a slightly more realistic price process—one that included, for example, two lags of the price and a lag each of the local tax level and the world crude oil price—we would have to perform a billion maximizations for every iteration. The sheer computational intensity of value function iteration renders it prohibitive for the realistic price processes I wish to consider.

Of course, the commonness of the curse of dimensionality means it has received considerable attention. There are tricks and variations that reduce the amount of computation necessary: if we turned to policy function iteration, for example, we might be able to reduce the number of iterations required before the value function converged (Mele 2009, Lecture 2). Policy function iteration would not, however, do anything to reduce the amount of computation required for each iteration.

B.2 Collocation

A more promising way to reduce computing time is collocation, a type of projection method.⁹ Collocation approaches the problem in a fundamentally different way: instead of concentrating on the value function itself, it exploits the first-order conditions and envelope conditions derived in Section A.

To use collocation, first imagine that the *solution* we're looking for—the policy function, giving optimal consumption decisions as a function of the state variables—can be written as $g(\mathbf{x})$, where $\mathbf{x}$ is, once again, a set of state variables. Note that if we knew the optimal policy, then by the very definition of its optimality we could substitute it back into the Bellman equation to find the value function. (That is, $\max_{g_t, c_t} \Psi = \Psi(g_t^*, c_t^*)$.)

Next, rewrite the first-order and envelope conditions as functions of the optimal policy $g(\mathbf{x})$. For simplicity, organize these conditions such that the right-hand side of each equation is 0. Denote

⁹For a more detailed treatment of collocation, see Judd (1998), Chapter 11, upon which the following exposition is chiefly based. Additional background from Mele (2009) has also been useful. Miranda and Fackler (2002, pp. 141-144, 227-237, 291-295) touch on collocation methods, as well.

the left-hand sides of the conditions as the operator $\mathbf{T}$. The set of first-order and envelope conditions can thus be written succinctly as

$$\mathbf{T}(g(\mathbf{x})) = \mathbf{0} \quad (27)$$

We'd like to know the policy function $g(\mathbf{x})$ that satisfies (27). Although we don't know the *form* that the function $g(\mathbf{x})$ takes, we can approximate it as a weighted average of some number of functions with *known* forms. Let us use k of these known functions, called basis functions. Denote each of these basis functions φ_i , and call the corresponding weight w_i . Our approximation of $g(\mathbf{x})$, $\hat{g}(\mathbf{x})$, can therefore be written

$$\hat{g}(\mathbf{x}) = \sum_{i=1}^k w_i \varphi_i$$

By an appropriate choice of weights w_i , we can construct an approximation $\hat{g}(\mathbf{x})$ that closely follows the first-order and envelope conditions:

$$\mathbf{T}(\hat{g}(\mathbf{x})) \approx \mathbf{0} \quad (28)$$

To find weights such that $\mathbf{T}(\hat{g}(\mathbf{x}))$ is close to $\mathbf{0}$, we can minimize with respect to these weights the norm of $\mathbf{T}(\hat{g}(\mathbf{x}))$:

$$\min_{w_i, i=1, \dots, k} \|\mathbf{T}(\hat{g}(\mathbf{x}))\|$$

If we wished to place more importance on satisfying some of the conditions than on others, we could, of course, minimize a projection of $\mathbf{T}(\hat{g}(\mathbf{x}))$ other than the norm.

Actually putting the collocation method into practice involves a mire of considerations, including—for a start—how many and what form of basis functions to use, which projection of $\mathbf{T}(\hat{g}(\mathbf{x}))$ to minimize, and whether and how to approximate the functional $\mathbf{T}$. Broadly speaking, however, the method approaches our problem by finding an approximation to a consumption-decision policy that satisfies the first-order and envelope conditions of the optimal policy.

The computational advantage of collocation arises from its use of the basis functions to interpolate between points in our grid of state variables—thereby allowing us to get by with a sparser grid. But even with a sparser grid, the number of discrete states will explode exponentially with the number of state variables. Recall from before that with $n + 3$ state variables, each discretized into d points, we have $D = d^{n+3}$ discrete states. Even if collocation allowed us to reduce the density of the discretization to, say, $d = 50$, we'd still need a third of a billion gridpoints for a problem with five state variables. Indeed, Malin, Kubler and Krueger (2010) note that standard collocation is "infeasible for three or more dimensions."

It might be possible, using more advanced collocation methods such as Malin, Kubler and Krueger's (2010) Smolyak-collocation method, to solve our problem when there are as many as twenty state variables. This might allow for a reasonable representation of the price process, especially if prices were modelled using annual rather than more frequent data. But such methods would *still* put limits on the model for the petrol price process, whose influence on demand is my chief interest. Moreover, many issues that I have simply assumed away in this brief treatment—such as the existence and differentiability of various functions, along with the enforcement of my model's many constraints—would have to be successfully addressed. I therefore leave attempts to solve Model 21 via value function iteration or collocation to the future, and focus in this thesis on the version of the model that trades limits on the *type* of price uncertainty for computational ease and total flexibility in the price process.

C Deriving Expected Utility for Alternative Utility Functions

If the agent's utility function is additively or multiplicatively separable, of the form

$$\begin{aligned} u_t &= h(c_t) + m(\bar{g}_t) \text{ [Additive Separability]} \\ \text{or} \\ u_t &= h(c_t)m(\bar{g}_t) \text{ [Multiplicative Separability]} \end{aligned}$$

and if prices are distributed uniformly,

$$p_t \sim U [low_t, high_t]$$

$$\text{density of } p_t : f(p_t) = \frac{1}{high_t - low_t}$$

and if the period- t budget is fixed,

$$p_t g_t + c_t = a_t$$

then as long as $\int_{low_t}^{high_t} h(a_t - p_t g_t) dp_t$ exists, expected utility can be expressed in closed form.

To see this, substitute the budget into these utility functions and integrate over the price distribution:

$$\begin{aligned} \text{Additive Separability} \quad & : \quad v_t = h(a_t - p_t g_t) + m(\bar{g}_t) \\ E[v_t] \quad & = \int_{low_t}^{high_t} v_t f(p_t) dp_t \\ & = \int_{low_t}^{high_t} \frac{1}{high_t - low_t} [h(a_t - p_t g_t) + m(\bar{g}_t)] dp_t \\ & = \frac{1}{high_t - low_t} \int_{low_t}^{high_t} h(a_t - p_t g_t) dp_t + \frac{1}{high_t - low_t} \int_{low_t}^{high_t} m(\bar{g}_t) dp_t \\ & = \frac{1}{high_t - low_t} \int_{low_t}^{high_t} h(a_t - p_t g_t) dp_t + m(\bar{g}_t) \end{aligned}$$

Multiplicative Separability: $v_t = h(a_t - p_t g_t) m(\bar{g}_t)$

$$\begin{aligned} E[v_t] \quad & = \int_{low_t}^{high_t} v_t f(p_t) dp_t \\ & = \int_{low_t}^{high_t} \frac{1}{high_t - low_t} h(a_t - p_t g_t) m(\bar{g}_t) dp_t \\ & = \frac{m(\bar{g}_t)}{high_t - low_t} \int_{low_t}^{high_t} h(a_t - p_t g_t) dp_t \end{aligned}$$

In either case, as long as we can solve for $\int_{low_t}^{high_t} h(a_t - p_t g_t) dp_t$, we can find $E[v_t]$.

Amongst the common utility functions for which we can express $E[v_t]$ in closed form is Cobb-Douglas, $u_t = c_t^\alpha \bar{g}_t^\beta$:

$$E[v_t] = \frac{1}{high_t - low_t} \frac{\bar{g}_t}{(1 + \alpha) g_t} \left[(a_t - low_t g_t)^{1+\alpha} - (a_t - high_t g_t)^{1+\alpha} \right]$$

D Optimization Methods

To perform the constrained optimization needed to find the optimal consumption path $g_1, \dots, g_T$, I use the command "fmincon" from Matlab's Optimization Toolbox. My algorithm is interior-point, with the following tolerances and evaluation limits:

- Termination tolerance on the function value (TolFun): 1e-12
- Termination tolerance on [gasoline consumption,] x (TolX): 1e-12
- Termination tolerance on the constraint violation (TolCon): 0
- Maximum number of function evaluations allowed (MaxFunEvals): 100000

References

- [1] Auld, M. Christopher, and Paul Grootendorst. 2004. An empirical analysis of milk addiction. *Journal of Health Economics* 23: 1117–1133.
- [2] Baltagi, Badi H., and James M. Griffin. 2001. The econometrics of rational addiction: The case of cigarettes. *Journal of Business & Economic Statistics*, 19(4): 449-454.
- [3] Becker, Gary S., and Kevin M. Murphy. 1988. A theory of rational addiction. *The Journal of Political Economy* 96(4): 675-700.
- [4] Becker, Gary S., Michael Grossman, and Kevin M. Murphy. 1994. An empirical analysis of cigarette addiction. *The American Economic Review* 84(3): 396-418.
- [5] Browning, Martin. 1991. A simple nonadditive preference structure for models of household behavior over time. *The Journal of Political Economy* 99(3): 607.
- [6] Chaloupka, Frank. 1991. Rational addictive behavior and cigarette smoking. *The Journal of Political Economy* 99(4): 722-742.
- [7] Coppejans, Mark, Donna Gilleskie, Holger Sieg, and K Coleman Strumpf. 2007. Consumer demand under price uncertainty: Empirical evidence from the market for cigarettes. *The Review of Economics and Statistics* 89(3): 510-521.
- [8] Deaton, Angus. 1992. *Understanding Consumption*. Oxford: Oxford UP.
- [9] Ellis, Christopher J., and Barry J. Naughton. 1990. On the theory of household saving in the presence of rationing. *Journal of Comparative Economics* 14: 269-285.
- [10] Grossman, Michael, and Frank Chaloupka. 1998. The demand for cocaine by young adults: A rational addiction approach. *Journal of Health Economics* 17: 427–474.
- [11] Houthakker, Hendrik S., and Lester D. Taylor. 1966. *Consumer demand in the United States, 1929-1970: Analyses and projections*. Cambridge, MA: Harvard UP.
- [12] Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2006. Evidence of a shift in the short-run price elasticity of gasoline demand. CSEM Working Paper 159. September 2006.
- [13] Judd, Kenneth L. *Numerical Methods in Economics*. 1998. Cambridge, MA: MIT Press.
- [14] Ljungqvist, Lars, and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory*. 2nd Ed. Cambridge, MA: MIT Press.

- [15] Malin, Benjamin A., Dirk Krueger and Felix Kubler. 2010. Computing Stochastic Dynamic Economic Models with a Large Number of State Variables: A Description and Application of a Smolyak-Collocation Method and Krueger. Working paper prepared for special issue of the Journal of Economic Dynamics and Control. 6 Jan 2010. 1 Apr 2010 <<http://www.econ.upenn.edu/~dkrueger/research/MKK.pdf>>.
- [16] Mele, Antonio. Lecture Notes: Computational Macroeconomics. 2009. Oxford University. "Lecture 1: Introduction and Dynamic Programming" and "Lecture 2: Numerical Methods for Dynamic Programming". 3 Apr 2010 <http://computationalmacro.files.wordpress.com/2009/10/lecture_1.pdf> and <http://computationalmacro.files.wordpress.com/2009/10/lecture_2.pdf>.
- [17] Miranda, Mario J., and Paul L. Fackler. 2002. Applied Computational Economics and Finance. Cambridge, MA: MIT Press.
- [18] Neary, Peter J., and Kevin W.S. Roberts. 1980. The theory of household behaviour under rationing. *European Economic Review* 13: 25-42.
- [19] Small, Kenneth A., and Kurt Van Dender. 2007. Fuel efficiency and motor vehicle travel: The declining rebound effect. *The Energy Journal* 28(1): 25-51.
- [20] Spinnewyn, Frans. 1981. Rational habit formation. *European Economic Review* 15: 91-109.