Essays in Consumption Habits and the Environment

K. Rebecca Scott

Balliol College

University of Oxford

Essays in Consumption Habits and the Environment

K. Rebecca Scott

Balliol College

University of Oxford

Submitted for the degree of
doctor of philosophy
at the
University of Oxford,
Hilary Term 2011

Abstract

The dynamics of demand for energy goods such as gasoline are complicated by investment decisions and behavioral habits. Both types of complication can be captured by a habits model, in which past consumption enters into an agent's current utility function. If the agent is forward-looking, or 'rational', then habits imply his consumption of the habit-forming good will be sensitive to his expectation of future market conditions, in particular future prices. This sensitivity implies, in turn, that demand and price elasticity will depend upon price volatility, and that the agent will respond differently to different types of price changes. Price elasticity measured over the mixture of price changes that occur in a given market will therefore underestimate the power of policy instruments that act through permanent or long-run price changes.

This thesis examines the implications of rational habits on demand behavior in a multi-good setting, drawing motivation from and considering applications to gasoline demand. Chapter 1 introduces the theory of rational habits, examines a simple perfect-foresight model, and uses microeconometric techniques to look at an application to U.S. gasoline demand. Chapter 2 builds a theory model in which agents have rational habits and future prices are uncertain. The implications of this model are then utilized in Chapter 3, which uses macro-type econometric techniques to test for rational habits in international gasoline demand. The empirical evidence suggests that rational habits may indeed shape demand for gasoline and that traditionally-measured price elasticity should not be used to project consumers' responses to policy interventions.

Word Count: Approx. 80,000

Supervisors: Martin Browning and John Muellbauer

Acknowledgements

I would like to extend my deepest thanks to my thesis supervisors, Martin Browning and John Muellbauer, whose guidance has been invaluable. For useful discussions and pointers along the way, I would also like to thank Tom Holden; Antonio Mele; and several anonymous referees, whose comments on Chapter 1 were particularly helpful.

For financial support while writing this thesis, I am grateful to the Clarendon Fund and the Linnett Scholarship in Environmental Economics.

Contents

Introduction vii
References ix

1 Rational Habits and U.S. Gasoline Demand **1**
1.1 Introduction 1
1.2 Gasoline Demand Literature 2
1.3 Habits in the Literature 4
1.4 Modeling Gasoline Demand Habits 6
1.4.1 Why Habits? 6
1.4.2 A Simplified Theoretical Model 7
1.4.3 The Primary Empirical Models 11
1.4.4 Alternative Empirical Models with Lead Price Components 12
1.4.5 BGM-Style Model 13
1.5 Empirics 13
1.5.1 Data 13
1.5.2 Estimation Issues 18
1.5.3 Implementation of GMM Estimation 27
1.5.4 GMM's Trade-offs and Least Squares 28
1.5.5 Alternative Models 28
1.6 Results 33
1.6.1 Results for Models 1 and 2 33
1.6.2 Results for Alternative Models with Lead Price Components 38
1.6.3 Results for BGM-Style Models 43
1.7 Results in Context 48
1.8 Conclusion 50
1.A Parameter Values and Assumptions Used for Simulation in Section 1.4.2 52
1.B History of Federal Tax Rates on Gasoline 52
References 53

2 Rational Habits and Uncertain Prices: Simulating Gasoline Consumption Behavior **59**
2.1 Introduction 59
2.2 Background 60
2.3 Building a Model 61
2.3.1 Introducing of Habits by Modifying the Consumption Variable 61
2.3.2 Infeasibility of the Ideal Problem 62

2.3.3 Simplifying the Consumer's Problem Using Uniformly-Distributed Prices . 63
2.3.4 Limitations of the Model 67
2.4 Applying the Model 68
2.4.1 The Price Process 68
2.4.2 Optimization Methods and Calculation of Elasticity 69
2.4.3 Beginning- and Endpoint Concerns 70
2.5 Results 74
2.5.1 Time Horizon 75
2.5.2 Durability of Price Shocks 76
2.5.3 Future Price Uncertainty 76
2.5.4 Habit Strength 79
2.5.5 Utility Weight on Habit-Forming Good 81
2.5.6 Discounting of Future Utility 81
2.6 Conclusion 83
2.A Setting up the Ideal Problem as a Dynamic Programming Problem 87
2.B Approaches to the Value Function Problem and the Curse of Dimensionality 90
2.B.1 Value Function Iteration 90
2.B.2 Collocation 92
2.C Deriving Expected Utility for Alternative Utility Functions 94
2.D Optimization Methods 96
References 97
3 Demand and Price Volatility: Rational Habits in International Gasoline Demand 99
3.1 Introduction 100
3.2 Gasoline Demand Literature Review 102
3.3 Data and Specification of Variables 106
3.3.1 Data 106
3.3.2 Measuring Price Volatility 106
3.3.3 Currency Issues 108
3.3.4 A Brief Overview of the Data 110
3.4 Models and Estimation 115
3.4.1 Least Squares 115
3.4.2 System GMM 120
3.4.3 Heterogeneous Coefficients and the Error-Correction Model 122
3.5 Results and Discussion 128
3.5.1 Least-Squares Estimates of the ADL Model 129
3.5.2 System GMM Estimates of the ADL Model 133

3.5.3 Error-Correction Model 133
3.6 Conclusion 145
3.A Further Data Exposition 147
3.A.1 Geographic Scope 147
3.A.2 Grades and Prices 147
3.B Local-currency Estimates 148
3.C Implementation of Estimation and Tests 156
3.D PMG Estimates with Individual Country-Specific Short-Run Coefficients 157
References 173

Conclusion **180**

Introduction

To keep the global average temperature increase to 2.4-2.8°C, the 2007 IPCC report estimates that we would need to reduce global CO_2-equivalent emissions 30 to 60% by 2050.[1] The corresponding per-capita reductions, particularly in the developed world, would need to be far more drastic. Achieving these reductions in emissions—or lesser, likelier reductions—will require significant cuts in our consumption of pollution-intensive goods. Key amongst these goods is fuel for road transport: as of 2004, road transport accounted for 17% of worldwide GHG emissions.[2]

Reducing transport-related emissions is synonymous with reducing consumption of petroleum products.[3] Considerable attention, therefore, is now being directed at our demand for gasoline and diesel, and at the policy options available to bridle this demand. A multitude of policies stands at the ready: more stringent fuel economy standards, for example, or new regulations on gasoline formulations, or an expansion and subsidization of public transport. The most direct approach—namely, using taxes or quantity restrictions to raise prices—has received less earnest consideration than might be expected.

In part, price instruments have been written off for their political unpalatability. But they have also been dismissed as ineffective, with this dismissal generally based on gasoline's low price elasticity. Several reports by the United States CBO, for example, base their projections of consumers' response to gas-tax increases on a long-run price elasticity of -0.4.[4] Meanwhile an FTC study (2005) argues, based on a short-run price elasticity of -0.2, that wild price spikes would be necessary to re-balance supply and demand after a local gasoline-supply shock.

Using price elasticity as the basis for projecting consumers' responses to policy-induced price changes requires the implicit assumption that all price changes are alike. It must be assumed, that is, that gasoline taxes provoke the same changes in behavior as the mixture of price changes over

[1] IPCC 2007, Working Group III, Technical Summary, Section 3: Issues related to mitigation in the long-term context - Stabilization scenarios, Table 3.5.

[2] Of energy-related CO2 emissions in 2004, 23% arose from transport; and of transport-related emissions, 74% arose from road transport. (IPCC 2007, Working Group III, Chapter 5: Transport and its infrastructure.)

[3] Petroleum provides 95% of the energy used in the transport sector. (IPCC 2007, Working Group III, Chapter 5: Transport and its infrastructure.)

[4] See, for example, CBO (2003) and CBO (2008). Note that these reports do not themselves explicitly write off price instruments.

which price elasticity is measured. This mixture includes tax-driven changes, but it is dominated by changes arising from local supply disruptions and fluctuations in the world oil price—changes that tend to be fleeting. And yet, as the very reports that use elasticity to project policy responses recognize,[5] permanent and fleeting price changes provoke different responses.

It is easy to imagine why consumers care about the expected duration of gasoline price changes: gasoline use is tangled not only with other purchase and investment decisions, but with the rotes of daily life. Consumers will only find it worthwhile to endure the costs of reducing gasoline consumption—costs of money, time, and mental exertion—if that adjustment is expected to yield benefits far into the future. The longevity of price changes matters.

Although intuition suggests that policy instruments operating through permanent price changes should reduce consumption by more than price elasticity implies, this distinction continues to be overlooked—and the potency of price instruments continues to be underestimated. This thesis begins the process of correcting those estimates. It does so by using habits to represent gasoline's enmeshment with investment decisions and many aspects of life, and by allowing consumers to recognize gasoline habits' effects on their future welfare. Such a 'rational habits' model allows the anticipated longevity of price changes to be incorporated into projections of those price changes' effects on consumption.

In this thesis, rational habits models are introduced, developed, and applied to gasoline consumption in three stand-alone chapters. Because the chapters are written to be independently intelligible, they overlap somewhat in their exposition before diverging into different material. Chapter 1 introduces the theory of rational habits, examines a very simple perfect-foresight model, and uses microeconometric techniques to look at an application to U.S. gasoline demand. Chapter 2 focuses more tightly on theory, building a rational habits model in which future prices are uncertain. The implications of this model are then utilized in Chapter 3, which uses macro-type econometric techniques to test for rational habits in international gasoline demand.

The empirical models used in Chapters 1 and 3 are quite different from each other and represent two distinct approaches to testing for the implications of rational gasoline habits. The first approach focuses on consumers' anticipatory reactions to future prices; the second, on consumers' sensitivity to price volatility. Although these approaches are technically distinct, conceptually they are closely intertwined, as price volatility may serve as a proxy for how well consumers are able to anticipate future price changes. Both approaches converge, moreover, toward the same underlying idea—namely, that the predictability of prices and the duration of price changes matter to consumers.

Taken together, these three chapters offer an overview of how rational habits shape consumers'

[5] See, for example, CBO (2005), pp. 15-16.

responsiveness to price changes of varying sources and durations. Although rational habits are discussed here solely in the context of gasoline demand, neither the rational-habits model nor its implications about the potency of price instruments is specific to gasoline. Indeed, rational habits may be appropriate for modeling demand for other energy goods, such as electricity—the use of which is, like the use of gasoline, tangled up with investment decisions and everyday behaviors. Just as habits have an important story to tell about demand for gasoline, so too may they offer an illuminating narrative about demand in the non-transport sectors—and about the power of price instruments over the remaining 77% of GHG emissions.

References

[1] Congressional Budget Office (CBO). 2002. Reducing Gasoline Consumption: Three Policy Options. 4 April 2011 <http://www.cbo.gov/ftpdocs/39xx/doc3991/11-21-GasolineStudy.pdf>.

[2] Congressional Budget Office (CBO). 2003. The Economic Costs of Fuel Economy Standards Versus a Gasoline Tax. 4 April 2011 <http://www.cbo.gov/ftpdocs/49xx/doc4917/12-24-03_CAFE.pdf>.

[3] Congressional Budget Office (CBO). 2008. Effects of Gasoline Prices on Driving Behavior and Vehicle Markets. Pub. No. 2883. 1 April 2011 <http://www.cbo.gov/ftpdocs/88xx/doc8893/01-14-GasolinePrices.pdf>.

[4] Federal Trade Commission (FTC). 2005. Gasoline Price Changes: The Dynamic of Supply, Demand, and Competition. 1 April 2011 <www.ftc.gov/reports/gasprices05/050705gaspricesrpt.pdf>.

[5] Intergovernmental Panel on Climate Change (IPCC). 2007. Climate Change 2007 (AR4): Climate Change Mitigation. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Ed. B. Metz, et al.. Cambridge UP: Cambridge, UK and New York, NY, USA. 4 April 2011 <http://www.ipcc.ch/publications_and_data/ar4/wg3/en/contents.html>.

Chapter 1

Rational Habits and U.S. Gasoline Demand

Abstract

Understanding the dynamics of demand for gasoline is essential to deciding how to rein in its consumption. Rational habits have not been broached in the gasoline demand literature, and yet they may be a key shaper of demand: if consumers are forward-looking, then habits will render their demand sensitive to price dynamics. This sensitivity implies that the usual measures of price elasticity will underproject consumers' responses to policy interventions.

This paper examines the implications of rational habits on gasoline demand. Using a simple model encompassing myopic and rational habits, I demonstrate that an agent with rational habits will respond to anticipated future price changes and react more strongly to permanent than to temporary price changes, with this distinction increasing in the strength of the habit. I then estimate several habits models using panel data on U.S. states for the years 1989-2003. Some of these model specifications yield strong evidence for rational gasoline habits.

1.1 Introduction

Gasoline demand is of interest both for its intricate structure and for its environmental implications. Although this interest has prompted research on many aspects of gasoline demand, the possibility of rational habits has not been broached. Yet rational habits may shape the dynamics of gasoline demand, rendering price elasticity a poor and biased tool for projecting the consumption response to policy interventions. In this chapter I set out a rational habits model, sketch its implications in a simplified scenario, and test whether rational habits shape demand in the U.S. gasoline market.

We may suspect that habits shape gasoline demand because many of the decisions tangled up with gasoline consumption take time and effort to adjust. These decisions include long-term commitments such as commuting distance and vehicle choice, as well as everyday behaviors such as carpooling, trip frequency, and driving style. Whether habits arise from long-term decisions or ingrained behaviors, they will link a consumer's preferences across time: at every juncture, the consumer's decisions are influenced by past gasoline consumption; and in turn, each decision commits him to some future consumption. Although at the surface the consumer may experience these commitments as habits for outputs *produced* using gasoline, in practice a habit for the final output will induce a habit for the intermediate input as well.

If a consumer is myopic, then his commitments to the future are irrelevant to his decisions, and habits will merely slow his adjustment to changing market conditions. But if the consumer is rational—if he *does* look to the future—then these commitments will shape his consumption in curious ways. When choosing a home and therefore a commuting distance, for example, he will take into account not only the current cost of fueling his commute, but the cost in years to come. The consumer's expectation of future gasoline prices will determine the merits of any decision. If today's price increase is expected to be offset by a decrease tomorrow, it may not be worthwhile to adjust consumption. If today's price increase is expected to be permanent, on the other hand, then the effort of adjusting will yield a benefit period after period—and the rational consumer may be provoked into a larger response.

Such nuances are not captured by traditional measures of price elasticity. In fact, because such measures lump permanent price changes together with temporary fluctuations, they may lead us to underestimate the potential of economic instruments to reduce gasoline consumption. If a government can raise the price of gasoline and signal credibly that the increase is permanent, then forward-looking consumers with habits should trim their gasoline use by more than the conventionally-measured price elasticity suggests.

1.2 Gasoline Demand Literature

Despite the importance that rational habits may hold in explaining gasoline demand—and in controlling it—they have received little attention up to now. The most common way of dealing with habits is to allow for "partial adjustment", i.e. to include a lag of consumption in the demand regression, but this non-structural approach merely captures some of the sluggish behavior associated with myopic habits.[1] Habits models tied more closely to theory have addressed gasoline

[1] See, for example, Hughes, Knittel, and Sperling (2006) and Small and Van Dender (2007).

consumption only tangentially, and under the assumptions of myopic consumers and/or constant prices. Pashardes (1986), for example, includes the category "fuel" when estimating a demand system that nests myopic and forward-looking habits, but he imposes constant prices and discount rates. "Gasoline and oil" is a category in Heien's (2001) demand system, but the habits he considers are strictly myopic. Breunig and Gisz (2009) explore the econometric implications of inserting an unobserved habit stock directly into a gasoline demand regression, but they do not look at this habit stock in the context of demand theory, and their regression model is only capable of capturing myopic behavior. Wirl (1991) takes a different tack and examines an adjustment cost model for liquid fuel demand, but he does not allow for imperfectly-anticipated prices or examine the theoretical implications of price dynamics. The idea of rational habits in gasoline demand remains unexplored.

Despite the uncharted gap in the area of rational habits, other areas of gasoline consumption have been mapped in detail. So much research has focused on estimating price and income elasticities of demand that Espey's (1998) meta-analysis is able to look at 41 studies, Dahl and Sterner's (1991) review at over 100 studies. The price elasticities estimated in these studies range from 0 to -2.72—from perfect inelasticity to very high elasticity (Espey 1998). Dahl and Sterner find that both the type of model and the periodicity of the data influence the resulting elasticity estimates. Static models yield an average price elasticity of -0.29 with monthly/quarterly and -0.53 with annual data, for example; and lagged endogenous models yield average short-run price elasticities of -0.24 with annual data and -0.14 with quarterly data. Elasticities are markedly greater in the long run than they are in the short run—for lagged endogenous models, on average the difference is a factor of three to five. Espey's (1998) results are similar, though she does come to a different conclusion about the speed with which consumers adjust to price changes in the short run—that is, she finds that consumers do most of their short-run adjustment in the space of a month.

Interestingly, Espey (1998) finds higher short-run price elasticities in Europe and in studies that combine U.S. with international data than in studies that address the U.S. alone. There are many possible reasons for these international differences, but rational habits offer an additional explanation: the difference may lie in the type of price changes, permanent or transient, that consumers are most likely to experience.

1.3 Habits in the Literature

Although economists have shown an interest in habits for over a century,[2] only in more recent times have they ventured formally from time-separable preferences. Time-inseparable preferences were first addressed under the assumption of myopic agents: consumers in models such as Pollak's (1970) seek to maximize utility in the present period only.

The introduction of forward-looking agents in models such as Spinnewyn's (1981) was a milestone because such an outlook yields significantly different demand behavior. Given a myopic outlook, demand will depend only on past consumption—which is the source of consumers' habits—and on current variables. Given a forward-looking outlook, under which agents choose a consumption bundle to maximize the sum of present and discounted future utility, demand will also depend upon future variables—or on expectations thereof.

Habits models can be differentiated not only by the time horizon over which agents optimize, but by the timescale over which habits linger. To introduce time-nonseparability to the traditional model in which preferences are separable across time, habits models modify the set of variables over which the agent takes his preferences. Preferences that depend only on current and once-lagged consumption yield "short-memory" models; preferences that depend on current and all previous consumption yield "habits-as-durables" (HAD) models.

In practice, both short-memory and HAD models introduce the lag(s) of consumption to the agents' utility function by modifying the consumption bundle. This modification involves subtracting from current consumption either some weighting of once-lagged consumption (in the short-memory case) or some weighted, discounted sum of all previous consumption (in the HAD case). Note that the discounted sum of previous consumption in the HAD model is analogous to the stock of a durable good: each period's consumption contributes to the habit-stock variable just as investment contributes to the stock of a durable; and each period the habit-stock decays predictably, just as a traditional durable depreciates.

Once modified, the consumption bundle is known in the literature variously as "uncommitted consumption" (Spinnewyn 1981), a set of "adjusted quantities" (Browning 1991), or a "reference bundle" (Spinnewyn 1981). As the last of these names suggests, consumers experience utility based on a "reference" level of consumption to which they have become accustomed. Consumption in any period contributes to the next period's reference level, thus breaching the time-separability of traditional models and leading to steadier consumption.

Deriving tractable demands when preferences are linked across time is far from trivial. Pol-

[2] Browning (1991) traces the primitive habits literature back to Marshall's 1890 *Principles of Economics*.

lak (1970) managed this for myopic consumers, under both short-memory and HAD habits, but forward-looking consumers complicate the problem considerably. Spinnewyn (1981) showed that the utility-maximization problem of a forward-looking agent with habits could be transposed into a more traditional utility maximization problem by adjusting prices and income to take the costs of habits into account. Browning (1991, App. A) generalized this result. Unfortunately, the necessary transposition turns the current price into a function of current and all future prices, current consumption into a function of current and all past consumption, and wealth/income into a variable defined recursively from future adjusted wealth/income (Browning 1991). The transposed problem is thus of no direct use in empirical work: in practice, variables cannot be defined recursively from a point infinitely far in the future.

Theoretically, however, the Spinnewyn/Browning transformation does yield an interesting result: the demand arising from a short-memory model will be a function of one lag of consumption, the current price, and all future prices; the demand arising from a HAD model will be a function of all past consumption, the current price, and one lead price (Browning 1991, p. 611). Although Spinnewyn's transformation does not yield a solved-out demand that can be estimated empirically, therefore, it does suggest some general forms that an empirical model might take.

Browning (1991) takes a somewhat different approach to developing a habits model that can be estimated empirically. Instead of representing preferences via a utility function, he represents them via a profit function; and instead of deriving demands directly via a constrained maximization, he works backward to find the one set of preferences that yields a particular form for demand. In this form, demand is a function of prices only as far removed as one period back and one period forward. Browning shows that the associated preferences coincide with a HAD structure.

Although this approach prevents demand from depending explicitly on variables infinitely far into the past or infinitely far into the future, it does so, unfortunately, by "roll[ing] up" the effects of prices from $t + 2$ onwards into the marginal utility of wealth—or, in the case where there is uncertainty, into the expected marginal utility of wealth. In scenarios where price dynamics and consumer adjustment are of primary concern, unfortunately, this approach will shroud the objects of interest in an unobservable variable. As an alternative to dealing with the (expected) marginal utility of wealth, Browning suggests a way to use total expenditure instead, but this method will still obscure the dynamics of interest.

Becker, Grossman, and Murphy (1994; henceforth "BGM") take yet another approach to modeling habits empirically: they abandon the quest for traditional solved-out demands and rely instead upon something akin to an Euler equation. Their model incorporates short-memory habits by introducing one lag of consumption to the utility function. Given quadratic utility, there is then

a linear relationship between current, past, and future consumption:

$$c_t = \theta c_{t-1} + \beta\theta c_{t+1} + \theta_1 p_t + \theta_2 e_t + \theta_3 e_{t+1} \tag{1.1}$$

where c is consumption, p is the price, and e_t and e_{t+1} are "shift variables" incorporating such things as income. Like Browning (1991), BGM thus roll up any effect of future prices into one variable. For BGM, this variable is lead consumption. Of course, lead consumption is clearly not exogenous: consumers choose c_t and c_{t+1} simultaneously, and causality goes both ways. BGM address this problem by using prices as instruments for lag and lead consumption, thereby channeling the effect of future prices through c_{t+1}. Instrumenting in this way, however, imposes a linear relationship between the instruments and past and future consumption; unfortunately, BGM's model does not offer a structural relationship between prices and consumption. Moreover, as Auld and Grootendorst (2004) point out, estimating this variety of model is likely to yield spurious support for rational habits, especially when using aggregate data to examine persistent consumption.

1.4 Modeling Gasoline Demand Habits

Two motives drive the search for a rational-habits model. First, I need a way to develop intuition about how habits may affect gasoline-consumption behavior. Second, I need a framework in which to examine gasoline habits empirically.

Ideally, one model would satisfy both of these needs, and I would derive a structural model that served as both an empirical framework and a theoretical touchstone. As the preceding discussion of the literature suggests and as I will shortly demonstrate, however, there is no easy way to satisfy both these needs simultaneously: even vast theoretical simplifications fail to yield simple solved-out demand functions. Instead, I will build a simple theoretical model and use this to sketch habits' behavioral implications, then move on to a model that is tractable for estimation.

1.4.1 Why Habits?

Habits are not, of course, the only way to link decisions across time. Indeed, some of the mechanisms that link gasoline consumption across periods—e.g., vehicle choice and commuting distance—may at first glance appear more like investments than like "habits" in the traditional, behavioral sense of the word.

By allowing past behavior to influence current preferences, however, the habits framework

captures the effect of investments. Within the habits framework, that is, past consumption enters into the agent's current utility function: if an agent "invests" today in a more efficient vehicle or a shorter commute, the resulting reduction in *present* gasoline usage will reduce the amount of gasoline needed to achieve a given utility level tomorrow. Adjustments in everyday behavior—learning to consolidate trips, for instance, or developing the confidence to cycle certain routes—are captured in the same way. The habits framework is therefore more flexible than an investment framework, subsuming both investment-style and behavioral-style intertemporal links. The habits framework is also flexible in the sense that it naturally allows agents to "invest" in technology (such as SUVs) that makes them *less* fuel-efficient.

The choice to pursue a pure habits model and focus solely on *gasoline* demand, abstracting from a multitude of gasoline-related decisions, is driven by my interest primarily in the *policy* implications of these intertemporal linkages. There is plenty to be learned, of course, from taking a more descriptive approach to household transport demand and modelling gasoline as a derived demand; and indeed quite a few studies—e.g., Berkowitz et al. (1990)—take such an approach. For policy purposes, however, it is important to know how price instruments will affect gasoline consumption not only at the household level, but at an aggregate level. The pure-habits approach lends itself to looking at aggregate demands. Unlike more descriptive models, moreover, it is not restricted to modelling the obvious and measurable mechanisms through which habits may arise. The pure habits model's abstraction is in this sense an advantage, sweeping up *all* the intertemporal linkages that affect demand. By sweeping up these linkages as simply as possible, habits offer a relatively straightforward way to improve our understanding of demand's interaction with the gasoline price process and, potentially, with policy instruments.

1.4.2 A Simplified Theoretical Model

To build a model that yields intuition about how habits influence consumers' behavior, we must first devise a scheme by which past behavior influences current preferences. 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 let consumers' utility for gasoline depend upon the adjusted quantity $\overline{g}_t$:

$$\overline{g}_t = g_t - \delta_g s_{gt} - \gamma_g \tag{1.2}$$

where g_t is the quantity of gasoline consumed in period t, δ_g and γ_g are constant parameters, and s_{gt} is a gasoline habit-stock variable. This habit-stock decays over time and is replenished with

each period's consumption:

$$s_{gt} = \alpha g_{t-1} + (1-\alpha)s_{g,t-1}, \qquad 0 \leq \ \alpha \leq 1 \tag{1.3}$$

To verify that this specification of the reference quantity $\overline{g}_t$ nests both short-memory and HAD models, note that the appropriate choice of α allows us to recover either. First, to recover the short-memory model, let $\alpha = 1$. The habit-stock then simplifies to $s_{gt} = g_{t-1}$, meaning the impact of gasoline consumption on the reference quantity $\overline{g}_t$ lasts only one period. Second, to recover the HAD model, take a cue from Browning (1991, App. A) and let $\delta_g = \frac{\alpha-1}{\alpha}$. The reference quantity then becomes

$$\overline{g}_t = \sum_{i=0}^{\infty} g_{t-i}(1-\alpha)^i - \gamma_g \tag{1.4}$$

and utility depends on present consumption and a weighted sum of *all* past gasoline consumption. This dependence on all past values of g_t is, of course, precisely what characterizes an HAD model.

Whereas α controls the duration of habits, δ_g adjusts their strength. If $\delta_g = 0$, then $\overline{g}_t = g_t - \gamma_g$, and there are no habits: only contemporary consumption enters into the agent's utility function in any period. As δ_g increases, the effect of the habit-stock on the reference bundle $\overline{g}_t$ grows, and preferences depend more and more on past consumption.

The Agent's Problem

With this habit-producing reference quantity in place, we can now set up the problem. Let the agent's consumption bundle consist of two goods: gasoline, g, and all other consumption, c. For simplicity, let c be non-habit-forming, so that utility in any period t is $u_t = u\left(\overline{g}_t, c_t\right)$.

The agent's goal is to maximize, subject to some set of budget constraints, a weighted sum of present and expected future utility. In order to progress further, I now make several broad simplifying assumptions.

Uncertainty and Expectations: First, I assume that the agent has perfect knowledge of future prices and income—or at least that he behaves as if he does, forming expectations and acting as if they were certain. This assumption is drastic, but it has considerable precedent: Spinnewyn (1981), for example, assumes "deterministic expectations about future unknown parameters," and Becker and Murphy (1988) make the same assumption implicitly. More realistic treatment of uncertainty swiftly renders the optimization problem at hand unwieldy, and fortunately the model yields useful intuition even without this complication. A theoretical model that incorporates uncertainty more realistically will be developed in Chapter 2.

Agent's Horizon: It might be ideal to allow the agent's time horizon to stretch infinitely into the future: on average, the populations I will be considering empirically have long lifetimes before them. Without using Spinnewyn and Browning's transposition, however, I have not found a way to solve the infinite-horizon problem explicitly; and since that transposition is difficult to reconcile not only with data, but with intuition, I assume a finite horizon, T. In fact, three periods will be enough to allow us to compare the effects of transitory and permanent price changes.

Budget: If consumers can borrow and save across periods, then an increase in tomorrow's gasoline price will have two effects: a rational-habit effect that dampens today's gasoline consumption, and an intertemporal substitution effect as consumers substitute cheap gasoline today for expensive gasoline tomorrow. Since I am interested chiefly in the former effect, I will isolate it by banning borrowing and saving across periods. The consumer's budget in any period t is therefore given by

$$p_{gt}g_t + p_{ct}c_t \leq y_t$$

where y_t is period-t income and p_{gt} and p_{ct} are the prices of gasoline and other consumption.

Given these assumptions, the problem faced by a consumer with discount factor β is thus

$$\underset{g_1,c_1,g_2,c_2,g_3,c_3}{Max} U = \sum_{t=1}^{T} \beta^{t-1} u\left(\overline{g}_t, c_t\right)$$

$$\text{s.t. } p_{gt}g_t + p_{ct}c_t \leq y_t \ \forall \ t$$

$$g_t \geq 0, \ c_t \geq 0 \ \forall \ t \tag{1.5}$$

I will further assume Cobb-Douglas utility of the form $u_t = \theta \ln \overline{g}_t + (1-\theta)\ln\left(c_t - \gamma_c\right)$.

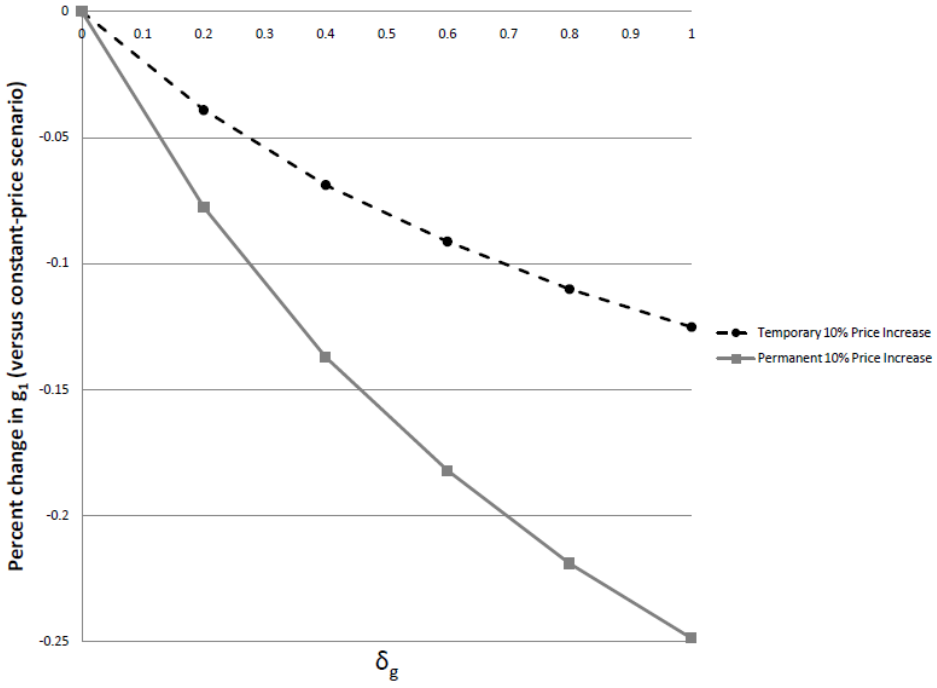
Solving the Agent's Problem and Simulating Price Changes

Even after these simplifying assumptions, closed-form, analytical solutions for the optimal consumption bundles remain out of reach. Clearly another approach is warranted, and simulation offers a simple way to sketch the model's implications. Choosing parameter values, I can solve the agent's utility-maximization problem under a series of three price regimes: one in which prices remain constant over all three periods; one in which the gasoline price rises temporarily in the second period before returning to its original level in the third; and one in which the gasoline price rises permanently from the second period onward. The price of other consumption, p_{ct}, is kept

constant across all periods and price regimes, and both the transitory and the permanent price changes are set at 10%.

The results of this simple simulation are illustrated in Figure 1.1, which plots the percent change in optimal first-period gasoline consumption (with respect to the constant-price scenario) versus the strength of habits, δ_g. Even before a price change actually occurs, habits shape the agent's demand: as the figure clearly demonstrates, habits drive the agent to reduce period-1 gasoline consumption in response to a price increase in period 2. This preemptive reduction in gasoline use increases as gasoline becomes more habit-forming, from no response when $\delta_g = 0$ to larger and larger responses as δ_g approaches 1. For all $\delta_g > 0$, moreover, a permanent increase in the future price of gasoline depresses current demand more than a temporary increase does; and this *difference*, too, increases with the strength of the gasoline habit.

Figure 1.1: First-period demand response to future price increases vs. δ_g



This model, simple as it is, thus provides several insights into how habits may shape gasoline demand. In particular, it suggests that if gasoline is habit-forming, people

- will reduce their current gasoline consumption in response to anticipated future increases in the gasoline price (and similarly, will increase their current consumption in response to future price decreases),
- will reduce their current gasoline consumption more in response to a future price increase

that is permanent than to a future price increase that is temporary, and

- will exhibit a sensitivity to the distinction between temporary and permanent price changes that rises with the strength of their habits.

1.4.3 The Primary Empirical Models

As Browning (1991) has shown, forward-looking consumers with HAD or short-memory habits will have demand functions that depend on infinite lags of consumption or infinite leads of price, respectively. In practice, the effects of lagged consumption and future prices will become negligible as their distance from the current period grows large, and so a finite approximation to these demand functions is not too great a concession.

I will look at two model variants:

$$g_t = f(g_{t-1}, y_t, p_t, p_{t+1}) \qquad \text{(Model 1)}$$

$$g_{it} = f(g_{t-2}, g_{t-1}, y_t, p_t, p_{t+1}) \qquad \text{(Model 2)}$$

where p_t is the real price of gasoline and y_t is real per capita income. Model 1, with only one lag of consumption, could be interpreted equally well as an HAD-style demand with truncated lags of consumption or as a short-memory demand with truncated leads of price. Model 2, with its two lags of consumption, would have to be interpreted as a truncated HAD-style demand function. Although these empirical models would resemble theoretical rational-habits models more closely if they included further lags of consumption or further leads of price, estimating such models would require far more data than is available.

My data consist of a fifteen-year panel of U.S. states and will be used to estimate models of the functional forms

$$g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_1 p_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it} \qquad \text{(Model 1)}$$

$$g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_1 p_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it} \qquad \text{(Model 2)}$$

where g is log per-capita gasoline consumption, p is log total real retail price of gasoline, y is log real per capita income, and the subscripts $i = 1, ...N$ and $t = 0, ..., T$ identify the state and the year, respectively. The error term ε_{it} is comprised of an unobserved state effect, η_i, and idiosyncratic error v_{it}:

$$\varepsilon_{it} = \eta_i + v_{it} \tag{1.6}$$

The idiosyncratic error v_{it} is assumed to be distributed independently across time, $E[v_{it}v_{is}] = 0$ $\forall\, t \neq s$.

1.4.4 Alternative Empirical Models with Lead Price Components

Given the role of consumers' expectations in determining gasoline demand, it seems natural to turn some attention to the expectations themselves. Although a thorough examination of consumers' price forecasting is beyond the scope of this chapter, some simple variations of the model will allow me to explore whether consumers perceive some components of the gasoline price as more important forecasting tools than others.

Suppose that the current gasoline tax provides better information than the current crude oil price about what gasoline prices will be in the future. Since gasoline taxes are fixed legislatively and not subject to the whims of the market, this does not seem an unreasonable supposition. Fixed-effects regressions of real taxes and real crude oil prices on their first lag values, indeed, yield an autocorrelation coefficient of 0.586 for real taxes and only 0.146 for the real crude oil price. If in fact current taxes are more informative about the future than the current crude oil price is, then we may expect that, when deciding current consumption, consumers weight the expected future tax level more heavily (relative to its size) than they do the expected future crude oil price.

One way to test this possibility is to replace $p_{i,t+1}$ in Models 1 and 2 with the lead crude oil price and the lead tax. By estimating these modified models using current and lagged instruments, as will be discussed Section 1.5, we can examine the relative effects of *expected* future taxes and crude oil prices. These estimates must be interpreted carefully, however, as taxes and the crude oil price may have slightly different effects on the gasoline price, and these differences may be reflected in the magnitudes of the lead-tax and lead-crude coefficients.

Another way to test whether taxes are more informative than the crude oil price is to divide the future gasoline price in Models 1 and 2 into tax and non-tax components. This approach is robust to the possibility that taxes and the crude oil price have different effects on the gasoline price; but since the non-tax component is presumably driven by oil prices, the price-component model should capture similar effects.

The specification and estimation of these models will be discussed in further detail in Section 1.5.5.

1.4.5 BGM-Style Model

For completeness, I also estimate a BGM-style model. The exact specification of this model, my methods of estimation, and my reasons for distrusting estimates based on a BGM-style specification are discussed in Section 1.5.5, and results are reported in Section 1.6.

1.5 Empirics

1.5.1 Data

My data set is a U.S. state-level panel for the years 1989 through 2003. The data set does not yet extend to more recent years because of the difficulty of obtaining state tax rates.

Sources and Definitions

Data on aggregate state gasoline consumption is available from the U.S. Energy Information Administration (EIA). I consider all grades of motor gasoline sold to end users, excluding aviation gasoline. Although this volume data includes gasoline sold to firms, the focus on *gasoline* rather than diesel effectively excludes much if not most of the fuel used by transport and distribution firms. To transform state-aggregated volumes into volumes per capita, I divide by the annual state populations estimated by the Bureau of the Census.

As an instrument for price, I use the price of crude oil, specifically the spot price of the West Texas Intermediate (WTI) stream based in Cushing, Oklahoma. The WTI is the major U.S. stream used to orient the valuation of other streams and should provide the best possible representation of the crude price faced nationally. The WTI spot price is monitored and recorded by the EIA, and for the sake of easier intuition I divide the price per barrel by 22 to give the approximate price per gallon of gasoline yielded.[3] The same crude oil price, it should be emphasized, is used for all states; and as this crude price varies only along the time dimension, it could be fitted perfectly with time dummies. Since time dummies would render the crude price instrument completely redundant, they are not included in the model—where time effects are estimated, I impose a linear trend.

The gasoline price faced by consumers is comprised of three elements: the before-tax price, a federal tax, and state taxes. Monthly average pre-tax prices are available from the EIA. Again, I consider all grades of gasoline sold to end users; the price is therefore a volume-weighted average over the various grades. To convert these before-tax prices to real terms—and later to convert all

[3] A barrel is equivalent to 42 gallons, of which about 51% is processed into gasoline (California Energy Commission 2005).

taxes, total prices, and crude oil prices—I use the U.S. CPI (base 1982-1984) for all items excluding energy.

The U.S. federal tax is levied as a flat cents-per-gallon amount. Its rate increased only five times during the sample period. Since its inception at 1 cent per gallon in 1932, in fact, the federal tax rate has seen only 14 changes. A history of these rate changes is provided in Appendix 1.B.

Whereas before-tax prices and the federal tax rate are readily available, an appropriate record of state taxes is difficult to compile. Part of the challenge arises from the complicated nature of the taxes themselves. Although most states, for example, levy a tax in cents per gallon and exempt gasoline from the usual percentage-based sales tax, several states do charge sales tax; and among these states, some apply sales tax only to the before-tax price, some apply sales tax to the before-tax price plus the federal tax, and one even applies sales tax to its own gasoline tax. For states that charge sales tax, I calculate the cents-per-gallon equivalent based upon the prevailing price and add this to the gasoline excise tax.[4]

To further complicate matters, some states levy supplementary gasoline taxes that masquerade under different names. In addition to the fee of around 16.5 cpg that it levies as a "gasoline tax," for example, New Mexico levies a 1.875 cpg "Petroleum Products Loading Fee," New Hampshire a 0.7 cpg "Oil Discharge and Disposal Cleanup Fee."[5] As far as the consumer is concerned, these taxes differ from the regular gasoline taxes in name only; and so for those states that charge them, I simply add the supplementary taxes to the regular tax level.

More problematic than the complications arising from the taxes themselves is the lack of centralized records of their historical rates. The Federal Highway Administration's annual Highway Statistics publishes the rates that prevailed on December 31 or January 1 of the corresponding year, but these rates may or may not represent the rate that prevailed throughout most of the year. To calculate annual tax rates, therefore, I have researched each state's tax history separately and compiled a list of all tax changes and the dates on which they went into effect. As the "annual" tax rate, I use a time-weighted average of the rates prevailing through the year.

In some states, it is especially difficult to compile a history of the gasoline tax rate because the state has introduced mechanisms to adjust the tax rate frequently and non-legislatively. Records of these many changes are sparse. The non-legislative mechanisms used to adjust the tax, moreover, may be affected by demand—which is a problem since tax will be used as an instrument for price. Because of these non-legislative mechanisms, I am forced to drop Connecticut, Nebraska, New

[4] Some states wave the sales tax under certain scenarios—within certain cities or counties, or to vehicles providing public transportation for the handicapped. With no data on the extent to which consumers take advantage of these exceptions, I simply ignore them; but they are unlikely to be widespread enough to have too great an effect on average final prices.

[5] These values as of 2003 (Federal Highway Administration 2006).

York, and North Carolina from the panel.

I am also forced to drop several other states. Florida is omitted because it allows individual counties to levy their own gasoline taxes—a privilege that many counties exercise. Alaska, Hawaii and California are omitted because their gasoline supplies are cut off from those of the rest of the country—Alaska and Hawaii by geography, California by state environmental policy. California law permits only a special reformulated blend of gasoline, which is not required elsewhere in the country; this isolation of the state's supply means a steep supply curve and high price variability. The Californian market is therefore likely to behave differently from the larger U.S. market.

These omissions leave the panel with 42 instead of 50 states.

A Brief Look at the Data

Before the data is examined in the framework of a demand model, it is worth examining briefly on its own.

The most striking features of the U.S. gasoline market are its exceptionally low prices and exceptionally high consumption. In 2001, indeed, America's average gasoline price was 60% lower than the U.K. price; meanwhile its per-capita consumption was almost 3.5 times as high.[6] Prices and per-capita consumption for a selection of developed countries in the year 2000 are plotted in Figure 1.2. The U.S. lies at the bottom right, its prices the lowest among these countries, its consumption surpassed only by Luxembourg's.

The U.S. can attribute its low prices largely to low taxes. In January 2000, for example, Australia's gasoline tax was more than twice as high as the average U.S. tax; Germany's, more than five times higher; and the U.K.'s, almost eight times higher.[7] Taxes and total prices for the U.S. and several European countries are graphed in Figure 1.3. From this picture it is clear that the before-tax price is roughly the same across countries; at any given time, taxes are the main source of cross-country price variation.

Taxes are a major source of price variation *within* the U.S. as well—both across states and across time. In nominal terms, the state excise taxes in my sample range from 7.5 cpg (Georgia for all years) to 30 cpg (Rhode Island from July 2002). The average nominal state excise tax for all years is 19.2 cpg; if sales taxes are included, the average creeps up to 19.9 cpg.

Over time, nominal tax rates follow a "step" pattern, with increases in the tax rates interspersed between plateaus. As a consequence of this pattern in the nominal tax rates, *real* tax rates tend to jump up discretely, then trail back down as they're eroded by inflation. For an illustration of

[6] World Resources Institute, "Transportation: Pump prices for super gasoline" and "Transportation: Motor gasoline consumption per capita."

[7] OECD - European Environmental Agency

Figure 1.2: Price versus Consumption for Selected Countries, 2000

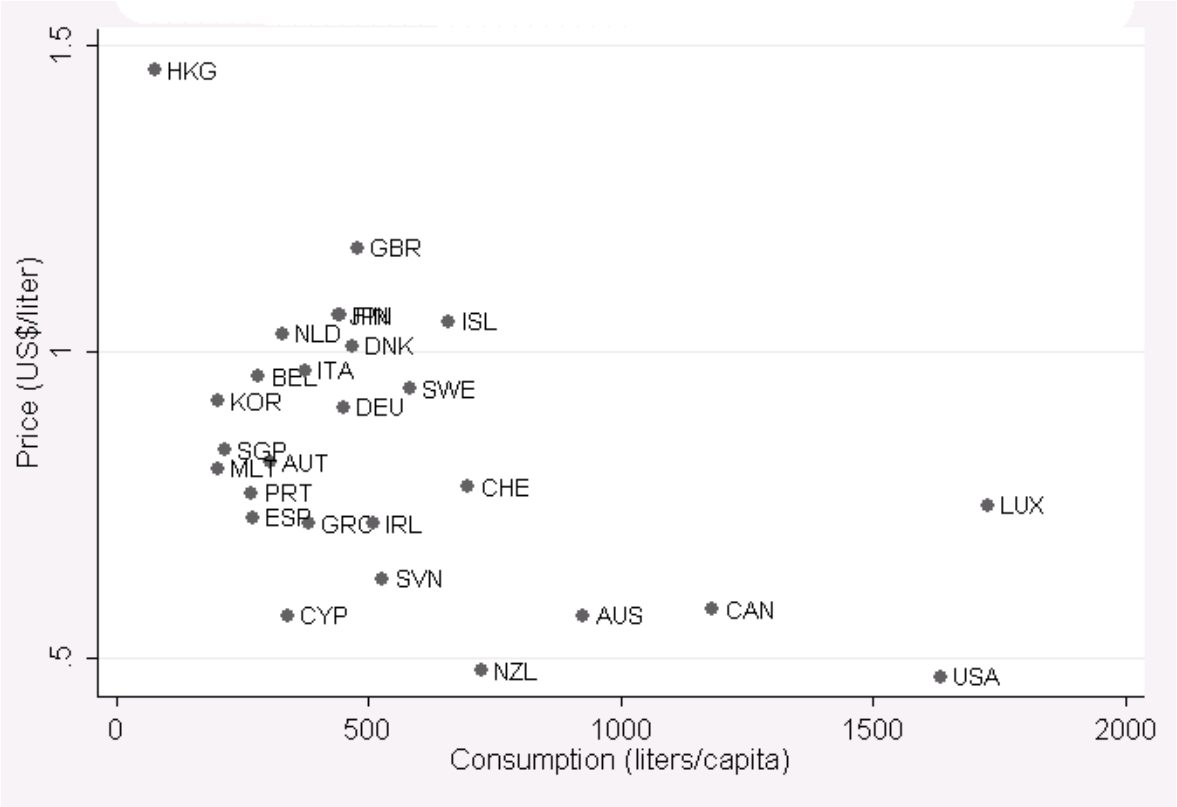
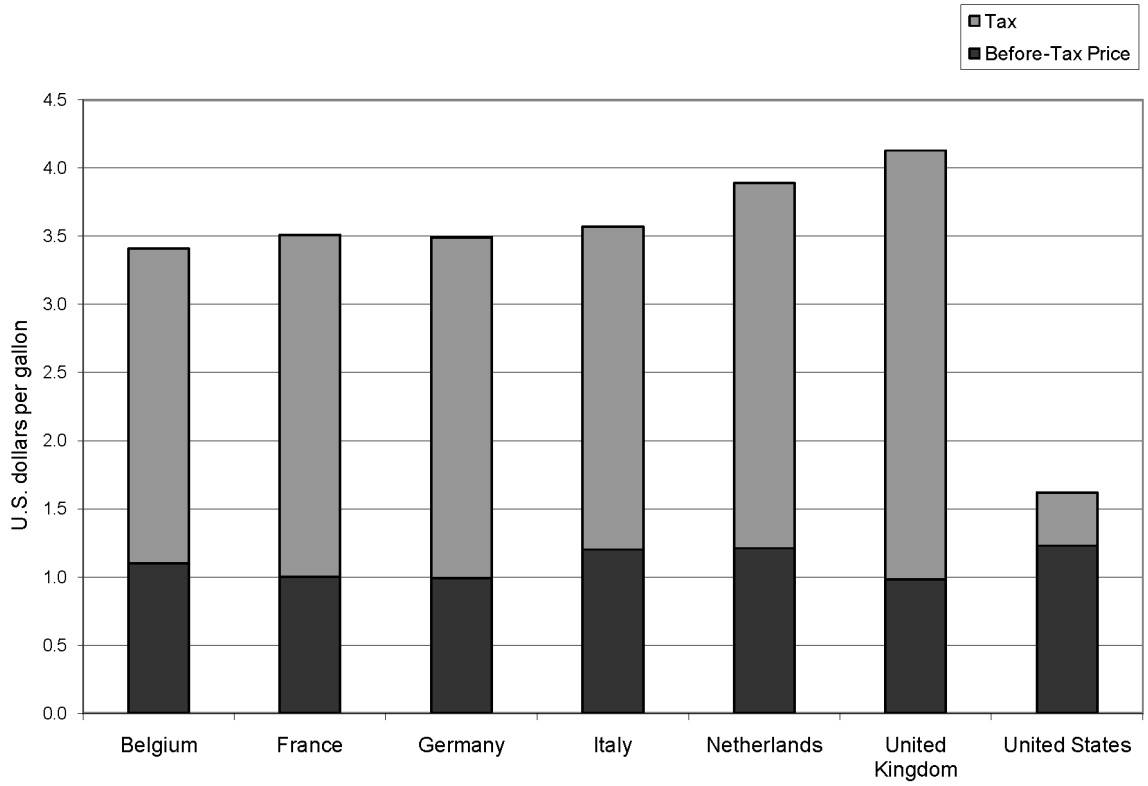


Figure 1.3: Gasoline Prices and Taxes in Selected Countries, 2001



this pattern, see Figures 1.4 and 1.5, which plot the nominal and real tax rates of Pennsylvania and North Dakota over time. To get a general view of the variation in real taxes across states and over time, see Figure 1.6.

Taxes are not, of course, the only source of price variation over time; indeed, input prices explain a good deal of this variation. As Figure 1.8 makes clear, the price of crude oil moves quite closely with the price of gasoline. This is significant because it will allow us to use the price of crude oil as an instrument for the gasoline price: there is little question of the instrument's relevance.

On the consumption side, the average quantity observed in the sample is 497.4 gallons of gasoline per person per year. In West Virginia in 1997, however, per-capita consumption dipped as low as 274.1 gallons per person per year; and in Oklahoma in 1992, it reached as high as 740.8 gallons per person per year. The time-path of each state's consumption level is plotted in Figure 1.7.

1.5.2 Estimation Issues

To estimate Models 1 and 2 empirically, it will be necessary to address several issues, including the endogeneity of price, the inclusion of the *future* price, and the dynamic-panel concerns arising from the unobserved state effects and the inclusion of a lagged dependent variable as a regressor.

Instruments for Price

Because price and consumption are determined simultaneously, we should expect the price of gasoline to be endogenous—that is, $E[p_{it}v_{it}] \neq 0$. This endogeneity might be inconsequential if our data were on individual consumers, each of whose demand had only a negligible effect on the market price. With data on *aggregate* consumption, however, the endogeneity cannot safely be ignored, and it is likely to bias the coefficient on price toward 0. The natural recourse is to find instruments for price that are independent of the idiosyncratic error v_{it}. Two obvious instrument candidates are the crude oil price and the tax level.

Crude oil is, of course, the main input into the production of gasoline. As Figure 1.8 illustrates, the retail price of gasoline rises and falls with the price of crude oil, and therefore the price of crude should be relevant as an instrument for the price of gasoline. As long as the effect of one U.S. state's gasoline demand on the world crude-oil market is negligible, moreover, the price of crude should be uncorrelated with the error term v_{it}. If this is the case, then price of crude oil will be a valid instrument for the price of gasoline.

The exogeneity of the crude oil price should not be taken simply as a given, however: as of

Figure 1.4: Nominal and Real Levels of the Pennsylvania Gasoline Excise Tax

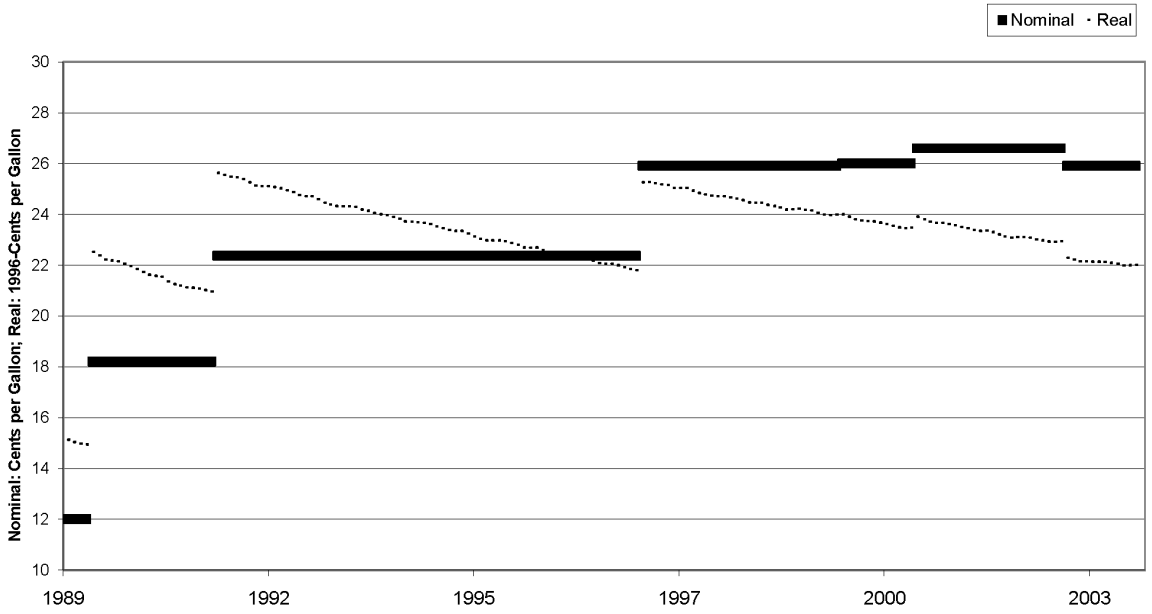


Figure 1.5: Nominal and Real Levels of the North Dakota Gasoline Excise Tax

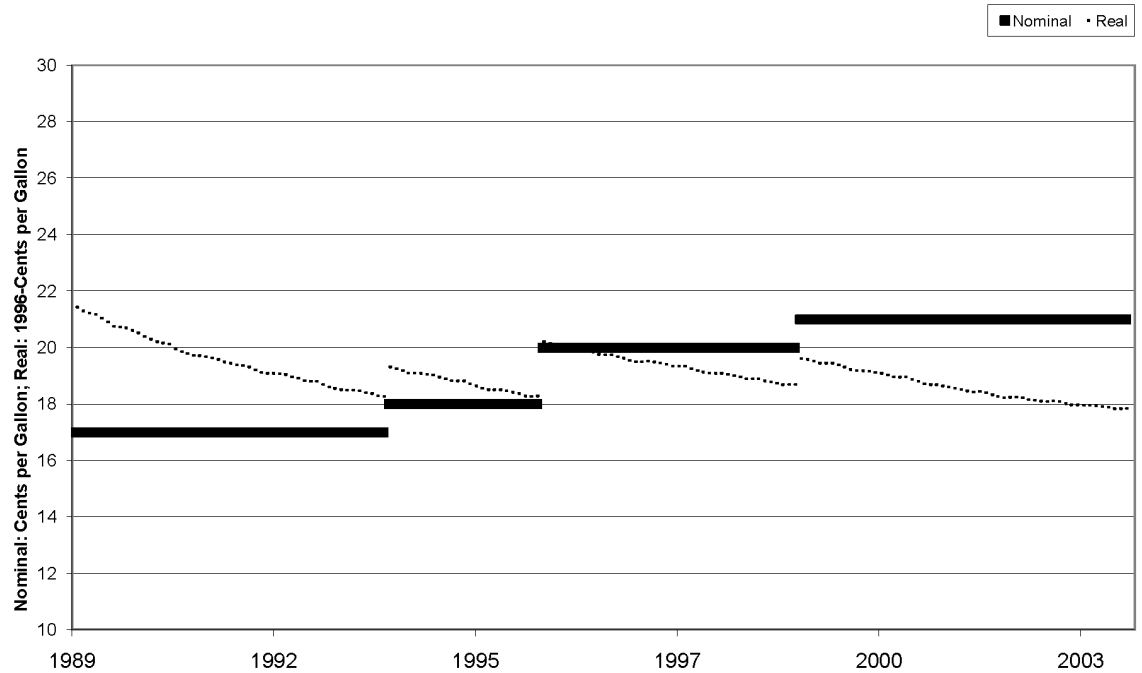
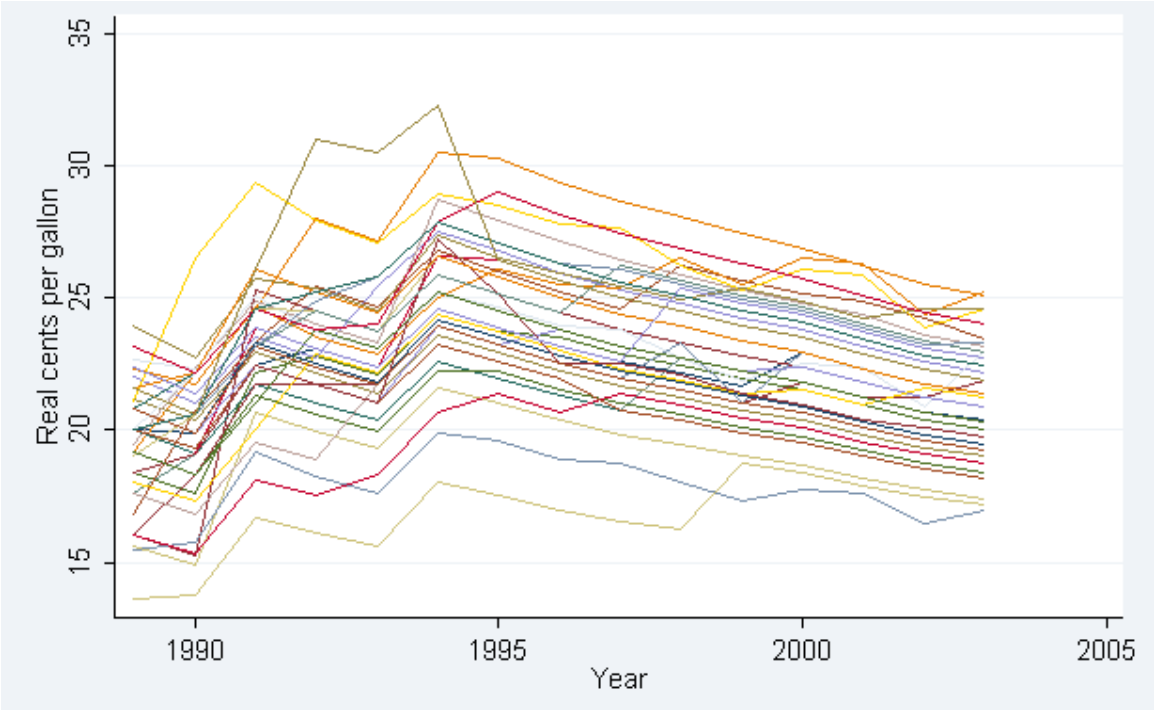


Figure 1.6: Total Real Gasoline Tax by State



2006, the United States constituted almost 25% of the world market for crude oil (EIA c). If demand shocks are correlated across states, then the crude oil price may be correlated with v_{it}, after all. The exogeneity of the crude oil price might also be compromised if an individual state's demand is not, in fact, negligible on the world market. Given that Texas burned through more than 6.5 million gallons of gasoline per day in 2003 (EIA b)—the production of which required approximately 300 thousand barrels of crude oil,[8] or 0.35% of the world's daily production—we cannot dismiss this concern out of hand. Fortunately, specification tests later on confirm that in empirical reality, the crude oil price is a good instrument.

The other instrument, the tax level, is a major source of price variation within the U.S., both across states and across time. Figure 1.9 plots the real level of the gasoline tax over time, by state. Since federal and state taxes gasoline taxes are generally levied in nominal cents per gallon and not as a percentage of the price, taxes should—barring demand-driven political pressures on the tax level—be a valid instrument. Note that political pressure on tax levels does *not* threaten taxes' validity as an instrument as long as that pressure is not driven by demand. Although studies such as Goel and Nelson's (1999) have uncovered evidence that gasoline prices influence gasoline taxes (with higher prices putting downward pressure on tax levels), these higher prices are likely to have arisen from changes on the gasoline-supply side. Unless *demand* shocks induce price changes that

[8] At 22 gallons of gasoline per barrel of crude oil (California Energy Commission, 2005).

Figure 1.7: Gasoline Consumption over Time, by State

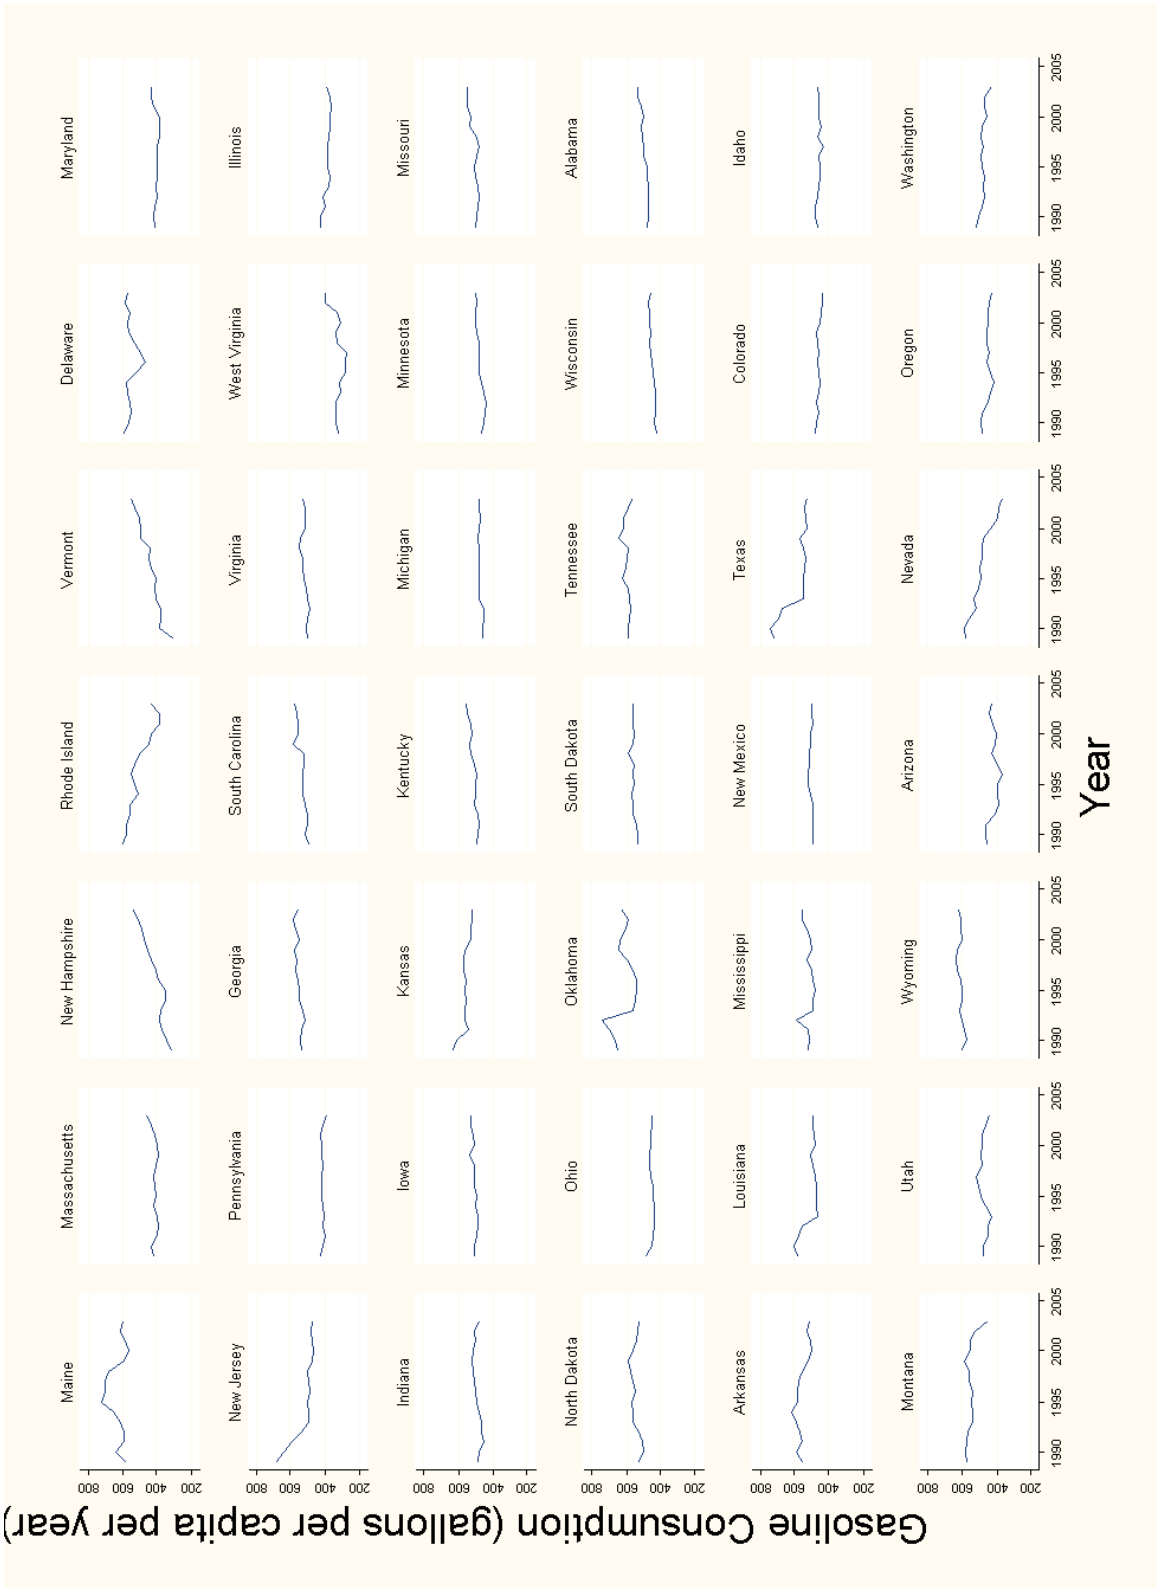
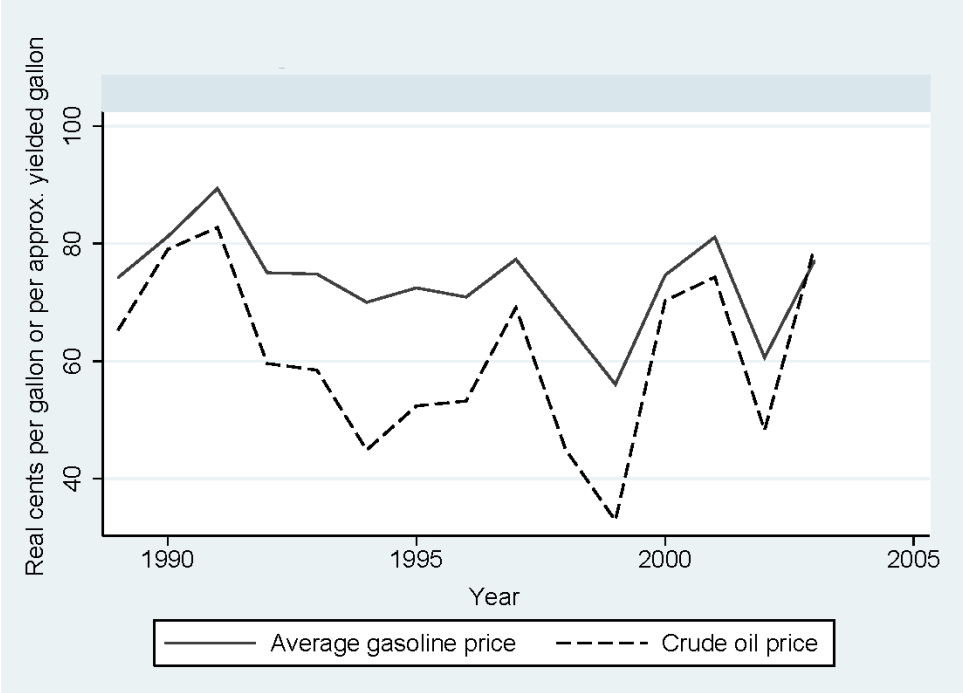


Figure 1.8: Average gasoline price and crude oil price over time



then create (or relieve) political pressure on tax levels, the use of taxes as instruments for prices remains valid.

In the models I consider, prices enter in log form, and therefore the crude oil price *level* and the tax *level* will yield a less-than-ideal fit. In order to achieve a better fit for log prices, I try two approaches. The first approach is ad hoc, using both the level and the log of the crude oil price and the gasoline tax as instruments. The second approach relies on the decomposition of the price into tax- and before-tax components:

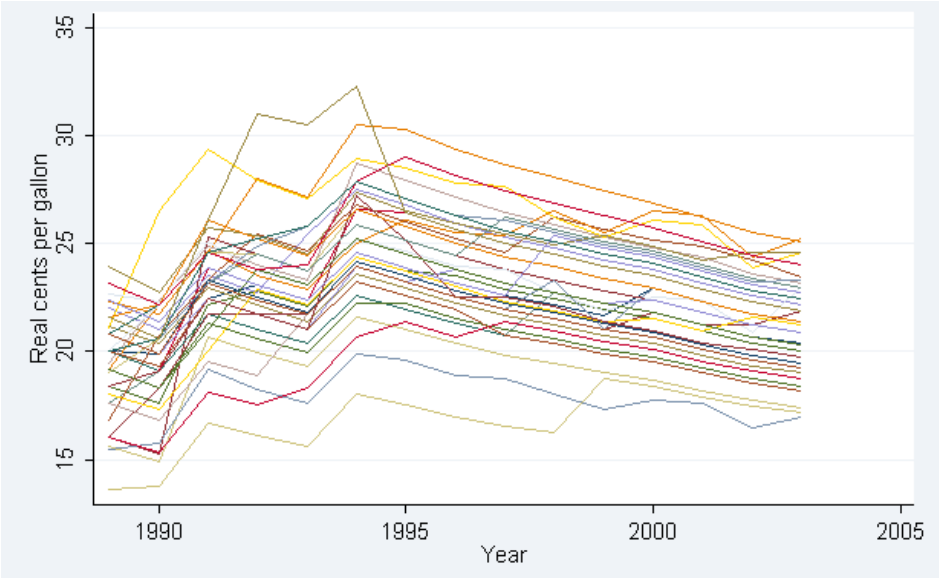
$$\begin{aligned} P_{it} &= Tax_{it} + BeforeTax_{it} \\ &= \left(1 + \frac{Tax_{it}}{BeforeTax_{it}}\right) BeforeTax_{it} \end{aligned} \tag{1.7}$$

$$p_{it} = \ln\left(1 + \frac{Tax_{it}}{BeforeTax_{it}}\right) + \ln BeforeTax_{it} \tag{1.8}$$

$$= \ln(1 + \tau_{it}) + \ln BeforeTax_{it} \tag{1.9}$$

where $\tau_{it} \equiv \frac{Tax_{it}}{BeforeTax_{it}}$. For the second term in (1.8), an obvious instrument is the log of the crude oil price, $\ln p_{crude,t}$. To create an instrument that fits the first term in (1.8), I begin by predicting before-tax prices ($\widehat{BeforeTax_{it}}$) based on a fixed-effects regression of state-specific before-tax prices

Figure 1.9: Total real gasoline tax by state



on the national crude oil price. This should remove the endogenous portion of the before-tax price. I then generate the tax instrument as

$$TaxIV_{it} = \ln\left(1 + \frac{Tax_{it}}{\widehat{BeforeTax}_{it}}\right)$$

In addition to the usual assumption that the price instruments are independent of the idiosyncratic error, my eventual estimator will rely on the assumption that the price instruments are independent of the state fixed effects η_i.[9] For the crude-based instruments, which are the same across states, this assumption is trivial. For the tax-based instruments, this assumption implies a belief that gasoline tax levels are not systematically related to the state fixed effects. This belief is quite plausible, as tax rates vary widely across states with similar infrastructure and population densities: Rhode Island, with the highest mean real tax level in the sample, is not unlike New Jersey, with one of the lowest; Wyoming, with the lowest mean real tax level in the sample, is quite similar to Montana, with one of the highest. This extra assumption does not, therefore, appear to be problematic.

Future Price

Including the lead price, $p_{i,t+1}$, as a regressor introduces the potential for several econometric complications. Unless otherwise specified, I avoid these complications by instrumenting for $p_{i,t+\ 1}$

[9] That is, for the model in levels (See Table 1.2) I rely on the moment conditions $E[w_{it}\varepsilon_{it}] = E[w_{it}(\eta_i + v_{it})] = E[w_{it}\eta_i] + E[w_{it}v_{it}] = 0$, where w denotes each price instrument in turn. On top of the standard assumption that $E[w_{it}v_{it}] = 0$, therefore, I must assume that $E[w_{it}\eta_i] = 0$.

using the same instruments I use for the current price.

This manner of instrumenting essentially transforms $p_{i,t+1}$ into an *expected* lead price. Since this expectation is based only upon information available at time t, econometrically it can be treated as a current rather than a lead variable. In terms of the model, this allows consumers to form an expectation of next year's price based upon this year's tax level and crude price, then choose their consumption level given this expectation. At the same time as this instrumenting takes care of econometric problems, therefore, it also provides a model of expectations that is more realistic than assuming consumers have perfect knowledge of future prices.

Dynamic Panel Concerns

The inclusion of lagged consumption as a regressor introduces an additional endogeneity problem. By definition, $g_{i,t-1}$ is positively correlated with the unobserved state fixed effect η_i—and η_i is part of the error term in all periods, so $g_{i,t-1}$ (and, similarly, $g_{i,t-2}$) will be correlated with the current error term ε_{it}. Ignoring this endogeneity will lead to biased estimators, and in particular

to a downward bias on the within-groups (e.g. fixed-effects) estimator of the coefficient on $g_{i,t-1}$.

For within-group estimators, the size of the bias in α_1 (the coefficient on $g_{i,t-1}$) decreases with T (that is, the estimator is consistent as $T \to \infty$) and increases with α_1. Nickell's (1981) derivation of this bias in an AR(1) with no outside regressors suggests that, if $\alpha_1 = 0.8$ and $T = 15$, $\widehat{\alpha}_1$ is biased downward by around 0.14. The inclusion of outside regressors worsens this bias (Nickell 1981), so 0.14 might be considered the minimum bias size. Since the bias depends upon the true value of α_1, however, and the true value of α_1 is unknown, it is perhaps more useful to note that a pooled-OLS estimate of α_1 will be biased upward (Bond 2002). Pooled estimators can therefore be used to establish the maximum value of α_1 and rule out the possibility that $\alpha_1 = 1$, which would have troubling implications for the dynamics of consumers' price-response behavior. Table 1.1 reports pooled estimates of Models 1 and 2, and fortunately even Model 1's upward-biased estimates of α_1 are statistically-significantly different from 1. Similarly, pooled estimates of Model 2's lagged-consumption coefficients imply stable behavior. Later estimates in which α_1 is near or greater than 1 should not, therefore, cause alarm that in reality α_1 is truly greater than 1 and dynamics are explosive.

To directly address the endogeneity of lagged consumption, we could turn to one of several dynamic panel techniques. Following Anderson and Hsiao (1981), we could take first differences, then use the twice-lagged dependent variable as a 2SLS-style instrument for the lagged differenced

Table 1.1: Pooled 2SLS and OLS estimates of Models 1 and 2

		(1) Model 1 Pooled 2SLS	(2) Model 2 Pooled 2SLS	(3) Model 1 Pooled OLS	(4) Model 2 Pooled OLS
g_{t-1}	α_1	0.867**	1.022**	0.918**	0.971**
		(0.0490)	(0.0651)	(0.0194)	(0.0512)
		[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2		-0.122+		-0.0513
			(0.0682)		(0.0476)
			[0.073]		[0.281]
y_t	β_2	0.0959	0.0364	-0.0180	-0.00400
		(0.0724)	(0.0257)	(0.0143)	(0.0149)
		[0.185]	[0.156]	[0.207]	[0.788]
p_t	β_0	0.554	0.0699	-0.0906**	-0.119**
		(0.375)	(0.0989)	(0.0206)	(0.0209)
		[0.140]	[0.480]	[0.000]	[0.000]
p_{t+1}	β_1	-1.319+	-0.446*	-0.00562	-0.0249
		(0.755)	(0.204)	(0.0205)	(0.0200)
		[0.080]	[0.028]	[0.784]	[0.213]
R^2		0.416	0.871	0.925	0.927
Endogeneity test, price and lead price		15.95	8.055	.	.
P-val		0.000345	0.00103	.	.
Underidentification test		3.902	10.77	.	.
P-val		0.0482	0.0178	.	.

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

dependent variable. For Model 1, this would imply estimating

$$\Delta g_{it} = \alpha_1 \Delta g_{i,t-1} + \beta_0 \Delta p_{it} + \beta_1 \Delta p_{i,t+1} + \beta_2 \Delta y_{it} + \Delta \varepsilon_{it} \tag{1.10}$$

using $g_{i,t-2}$ as an instrument for $\Delta g_{i,t-1}$. Note that $g_{i,t-2}$ is uncorrelated with $\Delta\varepsilon_{it} = (\eta_i + v_{it}) - (\eta_i + v_{i,t-1}) = \Delta v_{it}$ but correlated with $\Delta g_{i,t-1} = g_{i,t-1} - g_{i,t-2}$, so it should be a valid instrument for $\Delta g_{i,t-1}$.

Since my panel covers 15 years, however, I can exploit the data more fully using GMM methods that take advantage of additional lags of the dependent variable. Difference GMM, suggested by Holtz-Eakin, Newey, and Rosen (1988) and developed by Arellano and Bond (1991), allows all lags of the dependent variable beyond some starting point to be used as instruments. In particular, for a model such as Model 1 or 2 in which the v_{it} are assumed to be serially uncorrelated, difference GMM relies on the moment conditions

$$E\left[g_{i,t-k}\Delta\varepsilon_{it}\right] = 0,\ k = 2, ..., T \tag{1.11}$$

If some amount of serial correlation in the v_{it} were suspected, we would simply increase the lower bound for k to the point at which $g_{i,t-k}$ were expected to be uncorrelated with the differenced error term. The sample analogues of these moment conditions can then be used to construct a GMM estimator that is consistent in small-T, large-N situations.[10] (Blundell and Bond 1998)

Although the difference GMM estimator is consistent, however, its small-sample behavior may be less than desirable—especially when applied to a process that is highly persistent. A persistent process renders lagged consumption a weak instrument for differenced consumption. To see this, subtract $g_{i,t-1}$ from each side of Model 1

$$\Delta g_{it} = (\alpha_1 - 1)\, g_{i,t-1} + \beta_0 p_{it} + \beta_1 p_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it} \tag{1.12}$$

and consider what happens when α_1 is near 1: the correlation between $g_{i,t-1}$ and Δg_{it} is near zero.[11] Since difference GMM relies entirely on this correlation to identify Δg_{it}, lagged consumption is a very weak instrument for the evolution in consumption. In practice, weak instruments lead to biased estimates. Indeed, Monte Carlo studies by Alonso-Borrego and Arellano (1999) and Blundell and Bond (1998), among others, confirm that difference GMM yields a downward bias in

[10] If we are unwilling to make further assumptions, then (at least in the case of an AR(1) model) the two-step GMM estimator based upon these conditions is not only consistent, but asymptotically efficient. (Blundell and Bond 1998)

[11] This illustration of the potential weak-instruments problem follows Blundell and Bond's (1998) example of weak instruments in the AR(1) case. For further exposition, see also Blundell, Bond, & Windmeijer (2000).

estimates of the autoregressive coefficient in an AR(1) process. This bias increases in severity as the true value of the autoregressive parameter approaches 1.

The weak-instrument problem when α_1 is near 1 can be mitigated—at the price of some additional assumptions—by the introduction of further moment conditions. This approach, developed by Arellano and Bover (1995) and Blundell and Bond (1998) is known as system GMM. With system GMM, we would use the lagged difference of consumption as an instrument for the current level of consumption, introducing the extra set of moment conditions

$$E\left[\varepsilon_{it}\Delta g_{i,t-1}\right] = 0 \tag{1.13}$$

In a highly-persistent process, past evolutions in the dependent variable are useful in predicting the current level of the dependent variable; and so these new moment conditions will be especially helpful when α_1 is near 1.

Using these additional moment conditions, of course, requires an additional assumption. Note that

$$E\left[\varepsilon_{it}\Delta g_{i,t-1}\right] = E\left[\eta_i\Delta g_{i,t-1}\right] + E\left[v_{it}\Delta g_{i,t-1}\right] \tag{1.14}$$

and that since we've already assumed the v_{it} are serially uncorrelated, the second term in (1.14) is 0. In order to reach the moment condition in (1.13), we must assume that $E\left[\eta_i\Delta g_{i,t-1}\right] = 0$, or $E\left[\eta_i g_{i,t-1}\right] = E\left[\eta_i g_{i,t-2}\right]$. In other words, we must be willing to assume that the correlation between the unobserved state fixed effect η_i and the level of consumption g_{it} is constant over time.[12] Equivalently, we might think of the additional necessary assumption as a stationarity restriction on the initial condition for g: that is, given our prior assumptions, it is enough to assume that first-period consumption (g_{i1})'s deviation from its long-run equilibrium is independent of the state fixed effects η_i.[13] This will be satisfied as long as those initial deviations from equilibrium are random across states—and there seems to be little reason to suspect otherwise, as 1989 did not herald major state-specific disruptions in the long history of state gasoline demands. Given the high persistence in demand, therefore, I will employ the extra system GMM moment conditions.

1.5.3 Implementation of GMM Estimation

Table 1.2 spells out the moment conditions used in estimating Models 1 and 2, along with the assumptions required for their validity. "IV-style" moments impose the outside instruments for

[12] This follows Roodman (2006, pp. 28-29).

[13] See Blundell and Bond (1998), who provide an explanation of how the assumptions necessary for System GMM relate to the initial conditions and stationarity of the data.

current and future gasoline prices; "differences," the difference GMM moment conditions; and "levels," the additional system GMM moment conditions. As Table 1.2 notes, I use two different sets of IV-style instruments (IV-1 and IV-2). Results are insensitive to the choice between these instrument sets and also appear robust to further alternatives.

In estimating the models I rely on the two-step version of the GMM estimator, which is asymptotically efficient, and apply the Windmeijer (2005) correction to counteract the downward bias[14] of the two-step standard errors. One further estimation concern arises from the large number of instruments available in difference and system GMM: the abundance of instruments may lead to overfitting, as well as weaken certain specification tests (Roodman 2006, 2007). To reduce overfitting, I could simply limit the depth of lags used as instruments, but Roodman (2007) also outlines a way to reduce the instrument count by "collapsing" blocks within the instrument matrix. This has the advantage of allowing us to continue exploiting the information in all the available lags. To address potential concerns about overfitting, therefore, I estimate my models using both the standard "uncollapsed" instrument matrix and the "collapsed" version.

1.5.4 GMM's Trade-offs and Least Squares

Although system GMM has a distinct advantage over other methods—namely, that it addresses the endogeneity of lagged consumption—it comes with two trade-offs. First, GMM is not as precise as least-squares methods. (Of course, this is not a *true* trade-off given that the assumptions necessary for least squares are not entirely valid, and therefore the precision is at least partially spurious.) Second, system GMM does not have the same intuitive clarity that 2SLS has in removing the endogenous and unanticipated portions of price. For these reasons, I report fixed-effects 2SLS estimates of Models 1 and 2 in addition to system GMM estimates. These estimates use $TaxIV_{it}$ and $\ln p_{crude,t}$ as instruments for current and lead prices, with the "perfect foresight" variation also using lead instruments, $TaxIV_{i,t+1}$ and $\ln p_{crude,t+1}$. Results are reported in Table 1.10.

1.5.5 Alternative Models

Alternative Model with Lead Price Components

Section 1.4.4 suggested exploring how consumers form price expectations by splitting the lead price in Models 1 and 2 into tax and before-tax components. If taxes, which form a relatively-stable component of the gasoline price, are a more useful predictor of future gasoline prices than the crude oil price is, then we may expect consumers to respond more strongly to the tax component

[14] This problem is noted in Arellano & Bond (1991) and further examined in Blundell and Bond (1998, App. A).

Table 1.2: Moment conditions and assumptions for the estimation of Models 1 and 2

Model 1: $g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_1 p_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it}$
Model 2: $g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_1 p_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it}$

IV-Style Instruments

- For the model in differences: $E\left[\Delta w_{it} \underbrace{\Delta\varepsilon_{it}}_{=\Delta v_{it}}\right] = 0$
- For the model in levels: $E\left[w_{it}\varepsilon_{it}\right] = 0$

where w_{it} denotes, in turn, each of the outside instruments for the current and future prices: either

- tax_{it}, $p_{crude,t}$, $\ln tax_{it}$ and $\ln p_{crude,t}$, or (IV-1)
- $TaxIV_{it}$ and $\ln p_{crude,t}$ (IV-2)

Assumptions:

- Each instrument correlated with the price and the lead price:
$$E\left[w_{it}p_{it}\right] \neq 0,\ E\left[w_{it}p_{it+1}\right] \neq 0$$
- Each instrument uncorrelated with the idiosyncratic error:
$$E\left[w_{it}v_{it}\right] = 0$$

Additional assumption for the model in levels:
Each instrument w uncorrelated with the state fixed effects η_i.

GMM-Style Instruments

- For the model in differences:
$$E\left[g_{i,t-k} \underbrace{\Delta\varepsilon_{it}}_{=\Delta v_{it}}\right] = 0,\ k \geq 2$$

Assumptions:
$E\left[v_{it}\right] = E\left[v_{it}v_{is}\right] = 0\ \forall\ t \neq s$
(*No serial correlation in the idiosyncratic error term* v_{it}.)

- For the model in levels:
$$E\left[\underbrace{\varepsilon_{it}}_{=\eta_i+v_{it}} \Delta g_{i,t-1}\right] = 0$$

Additional assumptions:
$E\left[\Delta g_{it}\eta_i\right] = 0\ \forall\ i,t$
(*That is,* $E\left[g_{it}\eta_i\right]$ *constant over time.)*

of price than to the before-tax component.

The same decomposition of prices used to design price instruments (see equation 1.8) can now be used to divide the lead price:

$$p_{i,t+1} = \ln(1 + \tau_{i,t+1}) + \ln BeforeTax_{i,t+1} \tag{1.15}$$

Substituting (1.15) into Models 1 and 2 gives us two alternative models with lead price components:

$$g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_\tau \ln(1+\tau_{i,t+1}) + \beta_{bt} \ln BeforeTax_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it} \quad \text{(Alternative Model 1)}$$

$$g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_\tau \ln(1+\tau_{i,t+1}) + \beta_{bt} \ln BeforeTax_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it}$$

(Alternative Model 2)

I estimate these models with the same system GMM approach used to estimate Models 1 and 2. Moment conditions are identical to those listed in Table 1.2 (IV-2), and results are reported in Table 1.11. Table 1.12, additionally, reports system GMM estimates with the lead tax instrument and the lead log crude oil price included as instruments. As I did for the standard Models 1 and 2, I also report least-squares estimates of the alternative models. The instruments used in each least-squares variation are described in Table 1.3, and results are reported in Table 1.13.

In addition to these alternative models, I consider secondary alternative models that were Section 1.4.4's initial suggestion for dealing with price components. That is, I substitute lead taxes and the lead crude oil price directly into Models 1 and 2:

$$g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_T Tax_{i,t+1} + \beta_c p_{crude,t+1} + \beta_2 y_{it} + \varepsilon_{it} \quad \text{(Secondary Alternative Model 1)}$$

$$g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_T Tax_{i,t+1} + \beta_c p_{crude,t+1} + \beta_2 y_{it} + \varepsilon_{it}$$

(Secondary Alternative Model 2)

Note that both taxes and the crude oil price appear here as *levels*, not logs. Since prices do not break down in an intuitive way into log-tax-level and log-crude-price components, this avoids some fuzziness in the meaning of β_T and β_c. The previously-mentioned caveat still applies, however, that taxes and the crude oil price may have slightly different-magnitude effects on the gasoline price, and β_T and β_c may reflect this. These secondary alternative models therefore serve just as a check on the intuition driving the alternative models. I estimate these models with system GMM, using the same moment conditions described in Table 1.2 (IV-2) but with the addition of the tax and crude oil price *levels* as IV-style instruments. Results are reported in Table 1.14. I also estimate the secondary alternative models using fixed-effects 2SLS; details of this approach

Table 1.3: Notes on instrumenting in least squares estimation of the alternate models with lead price components

Alternative Model 1:
$g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_\tau \ln(1+\tau_{i,t+1}) + \beta_{bt} \ln BeforeTax_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it}$

Alternative Model 2:
$g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_\tau \ln(1+\tau_{i,t+1}) + \beta_{bt} \ln BeforeTax_{i,t+1} + \beta_2 y_{it} + \varepsilon_{it}$

Variables treated as endogenous: $p_{it}, \ln(1+\tau_{i,t+1}), \ln BeforeTax_{i,t+1}$

Instruments
"Standard" 2SLS (columns 1 and 2): current and lagged $\ln p_{crude}$ and $TaxIV$
2SLS with lead instruments (columns 3 and 4): current and lead $\ln p_{crude}$ and $TaxIV$

Table 1.4: Notes on instrumenting in least squares estimation of reduced-form alternate models

Secondary Alternative Model 1:
$g_{it} = \alpha_1 g_{i,t-1} + \beta_0 p_{it} + \beta_c Tax_{i,t+1} + \beta_c p_{crude,t+1} + \beta_2 y_{it} + \varepsilon_{it}$

Secondary Alternative Model 2:
$g_{it} = \alpha_1 g_{i,t-1} + \alpha_2 g_{i,t-2} + \beta_0 p_{it} + \beta_T Tax_{i,t+1} + \beta_c p_{crude,t+1} + \beta_2 y_{it} + \varepsilon_{it}$

Variables treated as endogenous: $p_{it}, Tax_{i,t+1}, p_{crude,t+1}$

Instruments
$TaxIV_{it}, \ln p_{crude,t}, Tax_{it}, p_{crude,t}$

are provided in Table 1.4 and results are reported in Table 1.15.

BGM-Style Model

For thoroughness, I consider a BGM-style model of the form

$$g_{it} = \alpha g_{i,t-1} + \alpha_f g_{i,t+1} + \theta_1 p_{it} + \theta_2 y_{it} + \eta_i + v_{it} \tag{1.16}$$

where the variables and error components are defined as they were for Models 1 and 2. As before, we can difference this model to eliminate the unobserved state fixed effects η_i:

$$\Delta g_{it} = \alpha \Delta g_{i,t-1} + \alpha_f \Delta g_{i,t+1} + \theta_1 \Delta p_{it} + \theta_2 \Delta y_{it} + \Delta v_{it} \tag{1.17}$$

Unlike before, however, we cannot use lagged values of g as instruments, because (1.16) implies that every g_{it} is correlated not only with the current and lagged idiosyncratic errors (v_{it} and $v_{i,t-1}$), but with all future idiosyncratic errors ($v_{i,t+j}$, $j = 1, ...T$). Any lagged value of g, therefore, will

be correlated with the differenced error term Δv_{it} in (1.17).

Instead, BGM (1994) and Baltagi and Griffin (2001) turn to *prices* as instruments for consumption. BGM use a 2SLS estimator in which the lag and lead of the cigarette price serve as instruments for the lag and lead of consumption in levels. Notably, although BGM use aggregate state consumption data—in which the simultaneous determination of price and quantity is likely to lead to a correlation between the price and the error term—they make no attempt to address the potential endogeneity of p_{it} or of the lagged and lead prices they use as instruments. Baltagi and Griffin (2001) ignore this problem as well. In fact, they employ a difference GMM estimator that relies on the assumption that prices are uncorrelated with the error terms in all periods. This reliance on an assumption of price exogeneity is potentially devastating, as Auld and Grootendorst (2004) find that even small amounts of price endogeneity will lead to a huge bias in the coefficient on lead consumption.

Instead of using prices themselves as instruments, therefore, I use the tax and crude-price instruments for price discussed in Section 1.5.2. If these instruments are truly exogenous, then the bias arising from price endogeneity is not a problem. By excluding or including lead instruments alongside lagged and current instruments, moreover, I can vary my assumptions about consumers' foresight about prices—foresight that BGM assume to be perfect.

Even if these instruments are perfectly exogenous, however, the resulting estimates may be quite biased. Auld and Grootendorst (2004) conduct a Monte Carlo experiment in which consumption and prices are serially correlated. The series are constructed such that prices are perfectly exogenous but have a negative effect on consumption. Even though these prices are technically-valid instruments for lag and lead consumption, however, overidentified 2SLS estimators suffer considerable small-sample bias in favor of rational habits. This bias is increasing in the degree of serial correlation in consumption and in the weakness of the price instruments, and it does not disappear quickly as the sample size grows. Note that "weakness" of the price instruments refers here to the ability of prices (or, in my case, instruments for prices) to explain variation in *consumption*. My tax and crude-price instruments may be weak instruments for *consumption* even though they are strong instruments for prices, particularly if consumers are not very price-responsive.

Estimating the model in differences, Auld and Grootendorst suggest, may mitigate the small-sample bias generated by the autocorrelation of consumption. The use of within-transformed (fixed-effects) 2SLS or difference GMM accomplishes this. The potential weak instruments problem, however, remains.

Table 1.16 reports GMM estimates of (1.16) based on the moment conditions detailed in Table 1.5, and Table 1.17 reports 2SLS estimates based on the instrument sets detailed in Table 1.6.

Table 1.5: Moment conditions for difference and system GMM estimation of the BGM-style model

BGM-style model: $g_{it} = \alpha g_{i,t-1} + \alpha_f g_{i,t+1} + \theta_1 p_{it} + \theta_2 y_{it} + \underbrace{\eta_i + v_{it}}_{\varepsilon_{it}}$

Moment Conditions

For the model in differences: (A)

"Standard" conditions:

$E\left[\Delta \ln p_{crude,t} \Delta v_{it}\right] = 0,\ E\left[\Delta TaxIV_{it} \Delta v_{it}\right] = 0$

$E\left[\Delta \ln p_{crude,t-1} \Delta v_{it}\right] = 0,\ E\left[\Delta TaxIV_{i,t-1} \Delta v_{it}\right] = 0$

Additional lead price IVs:

$E\left[\Delta \ln p_{crude,t+1} \Delta v_{it}\right] = 0,\ E\left[\Delta TaxIV_{i,t+1} \Delta v_{it}\right] = 0$

For the model in levels: (B)

"Standard" conditions:

$E\left[\ln p_{crude,t} \varepsilon_{it}\right] = 0,\ E\left[TaxIV_{it} \varepsilon_{it}\right] = 0$

$E\left[\ln p_{crude,t-1} \varepsilon_{it}\right] = 0,\ E\left[TaxIV_{i,t-1} \varepsilon_{it}\right] = 0$

Additional lead price IVs:

$E\left[\ln p_{crude,t+1} \varepsilon_{it}\right] = 0,\ E\left[TaxIV_{i,t+1} \varepsilon_{it}\right] = 0$

Difference GMM uses moment conditions (A);
system GMM uses moment conditions (A) and (B).

Table 1.6: Notes on instrumenting in least squares estimation of the BGM-style model

BGM-style model: $g_{it} = \alpha g_{i,t-1} + \alpha_f g_{i,t+1} + \theta_1 p_{it} + \theta_2 y_{it} + \underbrace{\eta_i + v_{it}}_{\varepsilon_{it}}$

Variables treated as endogenous: $g_{i,t-1}, g_{i,t+1}, p_{it}$

Instruments

"Standard" 2SLS (col. 1): current and lagged $\ln p_{crude}$ and $TaxIV$

2SLS with lead instruments (col. 2): current, lagged, and lead $\ln p_{crude}$ and $TaxIV$

For reasons just discussed, however, I do not put great trust in these estimates; and my preferred specifications remain those that augment a traditional demand model with lead price or price components rather than lead consumption.

1.6 Results

1.6.1 Results for Models 1 and 2

All estimation has been performed in Stata, with Roodman's (2011) *xtabond2* used for GMM estimates and Schaffer's (2010) *xtivreg2* used for least-squares estimates.

Table 1.7: System GMM estimates of Models 1 and 2, instrument set IV-1

		(1) Model 1 (uncollapsed)	(2) Model 1 (collapsed)	(3) Model 2 (uncollapsed)	(4) Model 2 (collapsed)
g_{t-1}	α_1	0.952**	0.990**	0.974**	0.997**
		(0.0242)	(0.0405)	(0.0474)	(0.0570)
		[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2			-0.0498	-0.0272
				(0.0428)	(0.0685)
				[0.251]	[0.694]
y_t	β_2	-0.0105	0.00629	0.00767	0.0180
		(0.0290)	(0.0371)	(0.0258)	(0.0341)
		[0.719]	[0.866]	[0.768]	[0.601]
p_t	β_0	-0.0944**	-0.0976**	-0.128**	-0.131**
		(0.0226)	(0.0229)	(0.0240)	(0.0232)
		[0.000]	[0.000]	[0.000]	[0.000]
p_{t+1}	β_1	-0.0176	-0.0198	-0.0331+	-0.0254
		(0.0180)	(0.0193)	(0.0174)	(0.0188)
		[0.334]	[0.311]	[0.065]	[0.183]
Instrument count		95	18	94	18
Arellano-Bond test, AR(2)		0.696	0.696	1.112	0.970
P-val		0.486	0.487	0.266	0.332
Sargan/Hansen test (all instruments; χ^2_{df})		38.64	20.95	40.20	19.64
df		90	13	88	12
P-val		1.000	0.0739	1.000	0.074
Difference Sargan test (IV-style instruments; χ^2_4)		0.12	2.30	2.38	4.83
P-val		0.998	0.682	0.667	0.305

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

The system GMM estimates of Models 1 and 2 are presented in Tables 1.7 and 1.8. These estimates offer overwhelming evidence in favor of at least myopic habits in gasoline demand—or, rather, of gradual adjustment. In Model 1, the estimated coefficient on g_{t-1}, $\widehat{\alpha}_1$, is clearly positive and statistically significant. In Model 2, the same is true of the sum of the coefficients on the lags of consumption, $\widehat{\alpha}_1 + \widehat{\alpha}_2$. All other things equal, higher gasoline consumption last year increases demand for gasoline this year.

The evidence for *rational* habits in the standard versions of Models 1 and 2 is not as strong. The results are, to be sure, certainly not *inconsistent* with forward-looking agents. Indeed, in Model 1 the coefficient on the lead price, $\widehat{\beta}_1$, is negative whether or not the instrument matrix

Table 1.8: System GMM estimates of Models 1 and 2, instrument set IV-2

		(1) Model 1 (uncollapsed)	(2) Model 1 (collapsed)	(3) Model 2 (uncollapsed)	(4) Model 2 (collapsed)
g_{t-1}	α_1	0.947**	1.008**	0.981**	1.007**
		(0.0288)	(0.0485)	(0.0497)	(0.0592)
		[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2			-0.0480	-0.0437
				(0.0420)	(0.0700)
				[0.260]	[0.536]
y_t	β_2	-0.00704	0.0277	0.00582	0.0516
		(0.0252)	(0.0391)	(0.0252)	(0.0399)
		[0.782]	[0.482]	[0.818]	[0.203]
p_t	β_0	-0.0935**	-0.0962**	-0.128**	-0.132**
		(0.0213)	(0.0232)	(0.0249)	(0.0239)
		[0.000]	[0.000]	[0.000]	[0.000]
p_{t+1}	β_1	-0.0118	-0.0147	-0.0352*	-0.0227
		(0.0183)	(0.0190)	(0.0164)	(0.0183)
		[0.521]	[0.443]	[0.038]	[0.221]
Instrument count		93	16	92	16
Arellano-Bond test, AR(2)		0.676	0.666	1.10	1.14
P-val		0.499	0.505	0.274	0.256
Sargan/Hansen test (all instruments; χ^2_{df})		39.26	19.20	37.80	16.88
df		88	11	86	10
P-val		1.000	0.0576	1.000	0.077
Difference Sargan test (IV-style instruments; χ^2_2)		-0.41	0.65	-0.08	2.27
P-val		1.000	0.721	1.000	0.322

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

is collapsed—but the statistical significance of this negativity is not very great. In Model 2, the coefficient on the future price is negative with relatively good statistical significance ($p = 0.065$ or 0.038, depending on the price instruments used)—but only if the "traditional" form of the instrument matrix is used. Collapsing the instrument set leads to an estimate of β_1 that is indistinguishable from 0.

Surprisingly, income does not have a statistically-significant effect on consumption. This statistical insignificance is consistent across models and estimation methods. One candidate explanation is that rising incomes are also capturing a downward trend in per-capita consumption, driven perhaps by increasing fuel efficiency. Indeed, adding a trend to the models and estimating by 2SLS results in a positive (but not statistically significant) coefficient on income (see Table 1.9). When the same models are estimated using system GMM, however, the inclusion of the trend has practically no effect on income's coefficient. In fact, the trend itself is not statistically significant when added to Models 1 or 2 or the alternative versions thereof.

Overall, the models appear reasonably well-specified. Tables 1.7 and 1.8 include results of a Sargan/Hansen test for the joint validity of all the moment conditions, as well as a difference-in-Sargan test for the validity of the IV-style instruments. Given that the p-values of the Sargan/Hansen tests for the uncollapsed versions are a suspiciously-high 1.000, we might worry that many instruments have rendered the tests weak. Somewhat reassuringly, however, after a fivefold reduction in the instrument count, the Sargan/Hansen tests for the collapsed versions look much more reasonable but still fail to reject the instruments' validity at a 5% significance threshold. That said, the P-values of 0.0739 and 0.074 (IV-1) and 0.0576 and 0.77 (IV-2) for the collapsed versions of Models 1 and 2, respectively, may cause some concern. If there were an endogeneity problem with the lagged values of g and Δg used as instruments, the standard recourse would be to restrict these instruments by a further lag; but this does not make a difference in either the Sargan/Hansen tests or the coefficient estimates. There does not appear to be an endogeneity problem with the price instruments, moreover, as the difference-in-Sargan test fails to reject the validity of the IV-style instruments in any of the specifications. Despite my ex-ante concern, the exogeneity of the crude oil price is not, in practice, compromised by U.S. states' large market share.

Tables 1.7 and 1.8 present the results of one further specification test, the Arellano-Bond test for serial correlation in v_{it}. Serial correlation in v_{it} would render one or more of the recent lags of consumption currently used as difference GMM instruments invalid.[15] The Arellano-Bond tests[16] do not yield any cause for concern about serial correlation in v_{it}.

[15] See Roodman (2006, p. 34) for an intuitive explanation.
[16] For technical details of the test, see Arellano and Bond (1991, p. 282 and Appendix).

Table 1.9: Fixed-effects 2SLS estimates of Models 1 and 2 with trend

	(1) Model 1	(2) Model 2	(3) Model 1 Perfect Foresight	(4) Model 2 Perfect Foresight
g_{t-1}	0.721**	0.782**	0.759**	0.803**
	(0.0508)	(0.129)	(0.0446)	(0.0662)
	[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}		-0.0456		-0.0586
		(0.0821)		(0.0533)
		[0.579]		[0.272]
y_t	0.119	0.0920	0.231**	0.0915
	(0.107)	(0.0790)	(0.0777)	(0.0778)
	[0.265]	[0.244]	[0.003]	[0.239]
p_t	-0.264**	-0.124	-0.0867**	-0.0922**
	(0.0810)	(0.163)	(0.0222)	(0.0224)
	[0.001]	[0.447]	[0.000]	[0.000]
p_{t+1}	0.426*	0.0626	0.00509	-0.0202
	(0.187)	(0.432)	(0.0202)	(0.0196)
	[0.023]	[0.885]	[0.801]	[0.301]
t	-0.00290	-0.000851	-0.00375**	-0.000513
	(0.00177)	(0.00225)	(0.00140)	(0.00145)
	[0.101]	[0.706]	[0.007]	[0.723]
R^2	0.364	0.591	0.628	0.601
Endogeneity test, price and lead price[a]	6.997	1.725	0.250	1.917
P-val	0.0302	0.422	0.883	0
Underidentification test[b]	12.40	1.647	146.5	140.7
P-val	0.000430	0.199	0	0.383

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a]Difference-in-Sargan/Hansen; $\chi^2(2)$ under H_0.

[b]Kleibergen-Paap (2006) rk statistic; $\chi^2(3)$ under H_0 of underidentification.

Overall, therefore, Models 1 and 2 appear to be reasonably well-specified. The major remaining concern is that $\widehat{\alpha}_1$ in Model 1 and $\widehat{\alpha}_1 + \widehat{\alpha}_2$ in Model 2 are very close to 1, and preferably I would test explicitly for a unit root. Unit root tests are set out in Bond, Nauges, and Windmeijer (2002) as well as in Arellano, Hansen, and Sentana (1999), but they are presented only for the case of an AR(1) model without additional regressors. Propitiously, however, it appears that the small-sample bias in the system GMM estimator of an AR(1)-type model is *upward* when the autoregressive parameter is near 1 (Blundell and Bond 1998, Table 2), so such bias is probably not *hiding* a unit root. The previously-discussed pooled OLS estimates of α_1 and α_2 bolster the case that the model is, in fact, stable.

Finally, fixed-effects 2SLS estimates are presented in Table 1.10. As expected, these models yield higher estimates of the speed of adjustment (that is, lower coefficients on lagged consumption) and higher precision. Unlike the GMM estimates, however, these estimates provide no evidence at all of rational habits; indeed, in Model 1 the coefficient on lead price is *positive* and statistically significant. This finding is nonsensical but anomalous, and I put greater trust in the system GMM estimates than the least-squares.

1.6.2 Results for Alternative Models with Lead Price Components

Estimates of the alternative models with price components, presented in Tables 1.11, 1.12, and 1.13, provide much more compelling evidence of rational habits than did estimates of the standard Models 1 and 2. In fact, all estimation methods—system GMM and fixed-effects 2SLS, with and without lead price instruments—yield negative coefficients on both the lead tax-component of price and the lead before-tax component. These coefficients are, moreover, statistically significant for most models and most estimation methods.

As hypothesized, consumers appear more responsive to the lead tax-component of price than to the lead non-tax component. The ratio of their coefficients tends to be about 3:1, with F-tests consistently rejecting that the coefficients are equal. This result holds, interestingly, whether or not consumers are assumed to have information about future taxes and crude oil prices, though it is very slightly stronger when only current information is assumed. Also interestingly, consumers appear more responsive to anticipated future price components than they are to the present price: the collapsed-matrix estimator of Model 2 (Table 1.11), for instance, yields an elasticity of -0.13 with respect to the current price, -0.76 with respect to the lead tax component, and -0.24 with respect to the lead before-tax component. The weight of all the future, this suggests, may be greater than the weight of the single present period.

Table 1.10: Fixed-effects 2SLS estimates of Models 1 and 2

		(1)	(2)	(3)	(4)
				Model 1	Model 2
		Model 1	Model 2	Perfect Foresight	Perfect Foresight
g_{t-1}	α_1	0.721**	0.780**	0.773**	0.804**
		(0.0548)	(0.132)	(0.0448)	(0.0663)
		[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2		-0.0416		-0.0569
			(0.0883)		(0.0529)
			[0.637]		[0.282]
y_t	β_2	-0.0478	0.0478	0.0386	0.0658*
		(0.0431)	(0.0890)	(0.0275)	(0.0284)
		[0.267]	[0.591]	[0.159]	[0.021]
p_t	β_0	-0.297**	-0.126	-0.0656**	-0.0917**
		(0.0894)	(0.166)	(0.0217)	(0.0220)
		[0.001]	[0.450]	[0.003]	[0.000]
p_{t+1}	β_1	0.540**	0.0720	0.0134	-0.0198
		(0.199)	(0.452)	(0.0203)	(0.0196)
		[0.007]	[0.873]	[0.510]	[0.312]
R^2		0.201	0.588	0.623	0.601
Endogeneity test, price and lead price[a]		9.898	1.427	0.00287	1.511
P-val		0.00709	0.490	0.999	0
Underidentification test[b]		14.80	1.517	139.8	151.4
P-val		0.000120	0.218	0	0.470

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a]Difference-in-Sargan/Hansen; $\chi^2(2)$ under H_0.

[b]Kleibergen-Paap (2006) rk statistic; $\chi^2(3)$ under H_0 of underidentification.

Table 1.11: System GMM estimates of alternative models

		(1) Alt. Model 1 (uncollapsed)	(2) Alt. Model 1 (collapsed)	(3) Alt. Model 2 (uncollapsed)	(4) Alt. Model 2 (collapsed)
g_{t-1}	α_1	0.871**	0.608**	0.886**	0.574**
		(0.0412)	(0.179)	(0.0628)	(0.137)
		[0.000]	[0.002]	[0.000]	[0.000]
g_{t-2}	α_2			-0.0311	-0.0260
				(0.0447)	(0.0501)
				[0.490]	[0.606]
y_t	β_2	-0.0201	-0.0400	0.00929	-0.103
		(0.0214)	(0.0947)	(0.0212)	(0.0991)
		[0.353]	[0.675]	[0.664]	[0.306]
p_t	β_0	-0.0832**	-0.0842**	-0.119**	-0.129**
		(0.0242)	(0.0285)	(0.0231)	(0.0242)
		[0.001]	[0.005]	[0.000]	[0.000]
$\ln(1+\tau_{t+1})$	β_τ	-0.300**	-0.634+	-0.300**	-0.764**
		(0.0963)	(0.338)	(0.0940)	(0.262)
		[0.003]	[0.068]	[0.003]	[0.006]
$\ln BeforeTax_{t+1}$	β_{bt}	-0.0937**	-0.194+	-0.108**	-0.241**
		(0.0319)	(0.105)	(0.0313)	(0.0825)
		[0.005]	[0.073]	[0.001]	[0.006]
Instrument count		93	16	92	16
Arellano-Bond test, AR(2)		0.458	0.192	0.89	0.70
P-val		0.647	0.848	0.371	0.484
Sargan/Hansen test (all instruments; χ^2_{df})		37.85	25.06	36.80	10.74
df		87	10	85	9
P-val		1.000	0.00523	1.000	0.294
Difference Sargan test (IV-style instruments; χ^2_2)		-0.97	1.10	-2.62	0.38
P-val		1.000	0.576	1.000	0.828

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

Table 1.12: System GMM estimates of alternative models, with lead price instruments

		(1) Alt. Model 1 (uncollapsed)	(2) Alt. Model 1 (collapsed)	(3) Alt. Model 2 (uncollapsed)	(4) Alt. Model 2 (collapsed)
g_{t-1}	α_1	0.888**	0.757**	0.892**	0.706**
		(0.0305)	(0.139)	(0.0669)	(0.132)
		[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2			-0.0290	-0.0463
				(0.0507)	(0.0484)
				[0.570]	[0.345]
y_t	β_2	-0.0174	-0.0574	-0.0102	-0.0822
		(0.0231)	(0.0605)	(0.0277)	(0.0726)
		[0.454]	[0.348]	[0.714]	[0.264]
p_t	β_0	-0.0793**	-0.0938**	-0.119**	-0.134**
		(0.0227)	(0.0258)	(0.0234)	(0.0254)
		[0.001]	[0.001]	[0.000]	[0.000]
$\ln(1+\tau_{t+1})$	β_τ	-0.253**	-0.425	-0.303**	-0.617*
		(0.0803)	(0.288)	(0.0947)	(0.275)
		[0.003]	[0.148]	[0.003]	[0.030]
$\ln BeforeTax_{t+1}$	β_{bt}	-0.0792**	-0.133	-0.112**	-0.199*
		(0.0282)	(0.0876)	(0.0327)	(0.0842)
		[0.008]	[0.135]	[0.001]	[0.023]
Instrument count		95	18	94	18
Arellano-Bond test, AR(2)		0.469	0.427	0.87	1.05
P-val		0.639	0.669	0.386	0.294
Sargan/Hansen test (all instruments; χ^2_{df})		40.01	24.52	35.96	14.35
df		89	12	87	11
P-val		1.000	0.0173	1.000	0.214
Difference Sargan test (IV-style instruments; χ^2_4)		1.28	2.68	-4.01	3.96
P-val		0.864	0.613	1.000	0.411

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

Table 1.13: Least-squares estimates of alternative models

		(1) Model 1 FE 2SLS	(2) Model 2 FE 2SLS	(3) Model 1 FE 2SLS lead instruments	(4) Model 2 FE 2SLS lead instruments	(5) Model 1 FE OLS	(6) Model 2 FE OLS
g_{t-1}	α_1	0.786**	0.820**	0.776**	0.804**	0.776**	0.802**
		(0.0453)	(0.0697)	(0.0442)	(0.0657)	(0.0443)	(0.0656)
		[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
g_{t-2}	α_2		-0.0657		-0.0515		-0.0473
			(0.0547)		(0.0528)		(0.0526)
			[0.230]		[0.329]		[0.369]
y_t	β_2	-0.0944*	0.0671+	-0.00789	0.0317	-0.00623	0.0194
		(0.0481)	(0.0386)	(0.0314)	(0.0309)	(0.0308)	(0.0308)
		[0.050]	[0.082]	[0.802]	[0.305]	[0.840]	[0.529]
p_t	β_0	-0.0816**	-0.0701**	-0.0731**	-0.0966**	-0.0747**	-0.109**
		(0.0268)	(0.0264)	(0.0216)	(0.0219)	(0.0218)	(0.0221)
		[0.002]	[0.008]	[0.001]	[0.000]	[0.001]	[0.000]
$\ln(1+\tau_{t+1})$	β_τ	-0.758**	-0.143	-0.244+	-0.205	-0.228+	-0.248+
		(0.229)	(0.179)	(0.132)	(0.131)	(0.132)	(0.132)
		[0.001]	[0.425]	[0.064]	[0.119]	[0.085]	[0.061]
$\ln BeforeTax_{t+1}$	β_{bt}	-0.249**	-0.102	-0.0684	-0.0786+	-0.0616	-0.0879*
		(0.0814)	(0.0721)	(0.0436)	(0.0432)	(0.0451)	(0.0442)
		[0.002]	[0.158]	[0.117]	[0.069]	[0.172]	[0.047]
R^2		0.613	0.593	0.625	0.603	0.625	0.603
Endogeneity test, price and lead price[a]		6.511	8.357	0.0731	5.947	.	.
P-val		0.0892	0.0392	0.995	0.114	.	.
Underidentification test[b]		100.5	74.23	139.9	153.6	.	.
P-val		0.000	0.000	0.000	0.000	.	.

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a]Difference-in-Sargan/Hansen; $\chi^2(3)$ under H_0.

[b]Kleibergen-Paap (2006) rk statistic; $\chi^2(2)$ under H_0 of underidentification.

As was the case for the standard models, the system GMM estimators for the alternative models appear well-specified by some metrics, with Arellano-Bond tests uncovering no evidence of serial correlation in the idiosyncratic errors and difference-Sargan tests failing to reject the validity of the IV-style instruments for price, whether or not lead price instruments are used. As happened for the standard models, however, in Alternative Model 1 the Sargan/Hansen test for all instruments has a concerningly-low p-value when the instrument matrix is collapsed. Once more, restricting the lagged values of g and Δg used as instruments by a further lag does not overwhelmingly improve this Sargan/Hansen test result. Somewhat reassuringly, however, the same concern does not extend to the collapsed estimator of Alternative Model 2, for which the p-value of the Sargan/Hansen test is a comfortable 0.29. Overall, it appears that the alternative models are at least as well-specified as the standard models.

Given the alternative model's finding that consumers respond differently to the future tax and non-tax components of price, the secondary alternative model bolsters the case that this difference in fact arises because the before-tax component is driven by fluctuating crude oil prices. System GMM estimates of the secondary alternative models are presented in Table 1.14.[17] Although the magnitudes of the coefficients on the tax level and the crude-oil price level cannot be compared, the signs and statistical significance of the coefficients are revealing: the coefficient on lead tax is negative and statistically significant (with a maximum p-value of 0.048); the coefficient on the lead crude price varies in sign depending on the estimator and is only remotely statistically significant in the collapsed version of Secondary Alternative Model 2. Consumers are responsive to their expectations of taxes but not so responsive to their expectations of the crude oil price, and this explains why they respond more strongly to the lead tax component of price than to the lead before-tax component.

1.6.3 Results for BGM-Style Models

Estimates of the BGM-style models provide limited support for the rational habits hypothesis. In GMM estimates (Table 1.16), the coefficient on lead consumption is positive and statistically significant only if moment conditions in both differences and levels are utilized. If the estimator is restricted to moment conditions for the first-differenced model, the coefficient on lead consumption is statistically indistinguishable from 0. This is consistent with Auld and Grootendorst's (2004) observation that differencing the model should reduce small-sample bias toward finding evidence of rational habits. It appears, moreover, that in practice the estimator using only the first-differenced moment conditions is more appropriate, as the Sargan/Hansen test rejects the joint validity of the

[17] Table 1.15 presents the corresponding fixed-effects 2SLS estimates, which are not very informative.

Table 1.14: System GMM estimates of secondary alternative models

		(1) Secondary Alt. Model 1 (uncollapsed)	(2) Secondary Alt. Model 1 (collapsed)	(3) Secondary Alt. Model 2 (uncollapsed)	(4) Secondary Alt. Model 2 (collapsed)
g_{t-1}	α_1	0.849**	0.590**	0.840**	0.431**
		(0.0433)	(0.170)	(0.0749)	(0.155)
		[0.000]	[0.001]	[0.000]	[0.008]
g_{t-2}	α_2			-0.0410	-0.0350
				(0.0416)	(0.0363)
				[0.330]	[0.341]
y_t	β_2	-0.0381	-0.0427	-0.0107	-0.137
		(0.0299)	(0.0803)	(0.0334)	(0.0817)
		[0.210]	[0.598]	[0.751]	[0.101]
p_t	β_0	-0.0819**	-0.0868**	-0.116**	-0.117**
		(0.0203)	(0.0259)	(0.0194)	(0.0196)
		[0.000]	[0.002]	[0.000]	[0.000]
$Tax_{i,t+1}$	β_T	-0.00452**	-0.00923*	-0.00584**	-0.0131**
		(0.00130)	(0.00453)	(0.00188)	(0.00380)
		[0.001]	[0.048]	[0.004]	[0.001]
$p_{crude,t+1}$	β_c	-1.54e-05	0.000106	-0.000247+	-4.24e-05
		(0.000163)	(0.000221)	(0.000133)	(0.000167)
		[0.925]	[0.633]	[0.071]	[0.801]
Instrument count		95	18	94	18
Arellano-Bond test, AR(2)		0.340	0.0443	0.88	0.60
P-val		0.734	0.00887	0.380	0.549
Sargan/Hansen test (all instruments; χ^2_{df})		37.36	26.58	34.48	8.57
df		89	12	87	11
P-val		1.000	0.965	1.000	0.662
Difference Sargan test (IV-style instruments; χ^2_4)		-0.09	2.41	-1.32	4.18
P-val		1.000	0.661	1.000	0.382

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

Table 1.15: Least-squares estimates of reduced-form alternate models

		(1) Model 1 FE 2SLS	(2) Model 2 FE 2SLS
g_{t-1}	α_1	0.706**	0.758**
		(0.0871)	(0.0691)
		[0.000]	[0.000]
g_{t-2}	α_2		-0.0427
			(0.0563)
			[0.448]
y_t	β_2	-0.0790	-0.00591
		(0.0523)	(0.0592)
		[0.131]	[0.920]
p_t	β_0	-0.238	-0.159**
		(0.160)	(0.0397)
		[0.136]	[0.000]
$Tax_{i,t+1}$	β_T	0.0116	0.00480
		(0.0196)	(0.00344)
		[0.555]	[0.163]
$p_{crude:i,t+1}$	β_c	0.00458	0.00200+
		(0.00454)	(0.00115)
		[0.314]	[0.083]
R^2		0.199	0.487
Underidentification test[a]		1.875	16.82
P-val		0.392	0.000223

$N = 42$, $T = 15$ (14 years available for Model 1, 13 for Model 2)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a]Kleibergen-Paap (2006) rk statistic; $\chi^2(2)$ under H_0 of underidentification.

Table 1.16: GMM estimates of BGM-style model

		(1) Difference GMM Standard IVs	(2) System GMM Standard IVs	(3) Difference GMM With Lead IVs	(4) System GMM With Lead IVs
g_{t-1}	α	0.719*	0.582**	0.733*	0.829**
		(0.274)	(0.161)	(0.279)	(0.153)
		[0.012]	[0.001]	[0.012]	[0.000]
g_{t+1}	α_f	-0.0970	0.508**	-0.132	0.222
		(0.239)	(0.163)	(0.230)	(0.143)
		[0.687]	[0.003]	[0.569]	[0.128]
p_t	θ_1	-0.124**	-0.0352	-0.121**	-0.0636*
		(0.0294)	(0.0227)	(0.0285)	(0.0263)
		[0.000]	[0.129]	[0.000]	[0.020]
y_t	θ_2	1.383**	0.356**	1.297**	0.253**
		(0.350)	(0.108)	(0.259)	(0.0861)
		[0.000]	[0.002]	[0.000]	[0.005]
Instrument count		4	5	6	7
Arellano-Bond test, AR(2)		1.578	2.336	1.415	2.208
P-val		0.115	0.0195	0.875	0.0272
Sargan/Hansen test (all instruments; χ^2_{df})				0.27	6.55
df		0	0	2	2
P-val				0.875	0.038

$N = 42$, $T = 15$ (13 years available for estimation)

Standard errors with Windmeijer correction in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

instruments in the system GMM version of the estimator incorporating lead-price instruments. (This test cannot be applied to the just-identified estimators that do not use lead-price instruments.) Meanwhile, when the BGM-style model is estimated using fixed-effects 2SLS (Table 1.17), the coefficient on lead consumption is positive and statistically significant—but for reasons previously discussed, this cannot be treated as strong evidence for the existence of rational habits.

Although I do not favor the BGM-style model for testing whether rational habits drive gasoline demand, estimating it has served a useful purpose—namely, demonstrating anew that estimates of such a model are likely to produce overwhelming "evidence" of rational habits when alternative models and more careful estimates are not so unambiguous.

Table 1.17: Least-squares estimates of BGM-style model

		(1) FE 2SLS	(2) 2SLS with lead instruments	(3) FE no instrumenting
g_{t-1}	α	0.599**	0.684**	0.499**
		(0.122)	(0.106)	(0.0502)
		[0.000]	[0.000]	[0.000]
g_{t+1}	α_f	0.450*	0.218	0.491**
		(0.196)	(0.179)	(0.0445)
		[0.022]	[0.224]	[0.000]
p_t	θ_1	-0.0420*	-0.0559**	-0.0364*
		(0.0211)	(0.0209)	(0.0166)
		[0.046]	[0.008]	[0.028]
y_t	θ_2	-0.00874	0.0163	-0.00893
		(0.0327)	(0.0328)	(0.0183)
		[0.789]	[0.620]	[0.626]
R^2		0.768	0.725	0.775
Underidentification test[a]		5.610	10.89	.
P-val		0.0605	0.0279	.

$N = 42$, $T = 15$ (13 years available for estimation)

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a]Kleibergen-Paap (2006) rk statistic; $\chi^2(2)$ in (1) and $\chi^2(4)$ in (2)

1.7 Results in Context

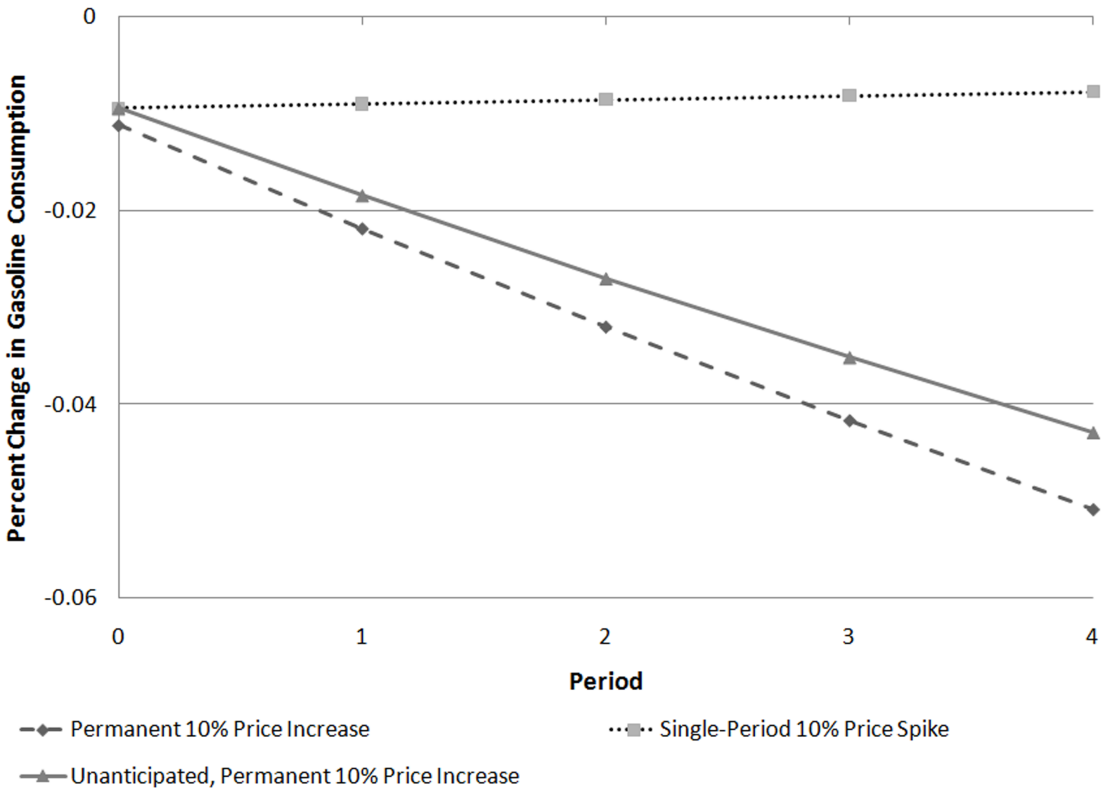
The short-run current-price elasticities produced by Models 1 and 2 fall on the lower end of short-term price elasticity estimates in the literature, which tend to run from -0.2 to -0.4[18] and average around -0.3 in Goodwin's (1992) review. Some of the studies upon which these ranges and averages are based, however, include cross-sectional and international data, which tend to yield higher price elasticities; and thus it seems very reasonable that our U.S panel data yields values of $\widehat{\beta}_0$ around -0.10 for Model 1 and -0.13 for Model 2. Unlike the elasticities from other studies, moreover, the current-price elasticities of Models 1 and 2 explicitly exclude the effect of future prices. Taking $\widehat{\beta}_1$ into account, consumers look a bit more price-responsive.

To sketch what these results imply for the dynamics of gasoline consumption, I work out the consequences of several hypothetical price changes. Totally differentiating Model 1 yields an expression for the percent change in current gasoline consumption. This can be plugged into next period's totally-differentiated Model 1 to get an expression for the percent change in next period's gasoline consumption—and so on, iteratively, allowing us to estimate how various price changes should alter demand over time. Specifically, let us consider a single-period 10% price spike; a permanent, anticipated 10% price increase; and a permanent 10% price increase that consumers continually believe will vanish in the next period. Figure 1.10 plots the path that gasoline consumption would take under each of these scenarios (given constant income), assuming the uncollapsed coefficient estimates for Model 1. In period 0, when the price change begins, we would predict a slightly greater reaction to the permanent, anticipated price change than to the other price scenarios. In the following periods, consumption under the temporary-price-increase scenario would rise back toward its initial level. Under the permanent price changes, on the other hand, consumption in succeeding periods would fall further as consumers acclimated to lower and lower levels of consumption. Consumption would decrease more under anticipated than under the perpetually-unanticipated permanent price change, with the difference in consumption levels growing larger with time. This difference—represented in Figure 1.10 by the vertical distance between the solid and dashed lines—holds a special meaning for a regulator aiming to reduce gasoline consumption: it is the extra reduction he can bring about simply by assuring consumers that a given price increase is permanent.

In the long run, consumption will, of course, approach a new steady state. In the case of the perpetually-unanticipated price change, in which the expected lead price is always the initial price,

[18] See Graham and Glaister (2002), who summarize the survey results of Dahl (1995), Sterner et al. (1992), and Goodwin (1992).

Figure 1.10: Consumer responses to hypothetical price changes: an extrapolation



this new steady state is represented (for Model 1) by the long-run elasticity

$$\frac{\beta_0}{1-\alpha_1}$$

or, plugging in the uncollapsed coefficient estimates of Model 1, -1.76 (SE 1.04). In the case of the anticipated, permanent price change, this new steady state is represented by the long-run elasticity

$$\frac{\beta_0+\beta_1}{1-\alpha_1}$$

or, plugging in coefficients from the same model, -1.98 (se 1.16). In the long run, that is, credible permanence gives a price change a "bonus" effect of $\frac{\beta_1}{1-\alpha_1}$.

This 'bonus' to signaling that a price change is permanent is more dramatic according to the alternative models with lead-price components. According to the alternative Model 1 (Table 1.11, Column 1), for instance, a 10% price increase brought about by a tax increase would immediately reduce gasoline consumption by 3.8%, whereas the same price increase expected to last only one period would reduce gasoline consumption by only 0.8%. In an intermediate case, with the 10%

price change expected to have the same duration as a typical crude-driven price change, the immediate reduction in consumption would still be only 1.8%. In the long run, the elasticity with respect to a perpetually-unanticipated price change would be -0.645 (se 0.682), the elasticity with respect to a typical crude-driven price change would be -1.37 (se 0.399), and the elasticity with respect to a credibly-permanent tax-induced price change would be -2.97 (se 0.326). This is a tremendous difference. Note that although the elasticity of -3 with respect to a credibly-permanent, tax-driven price change is two to three times the long-run price elasticities generally observed in the literature, the long-run elasticity of -1.37 with respect to "typical" price changes is not far out of line with the literature. Although 1.37 falls on the higher-magnitude end of long-run elasticity estimates (which averaged to 0.58 in Espey (1998), 0.80 in Dahl and Sterner (1991), and 0.84 in Brons et al. (2008)), it is well within the range observed by Espey (1998). The model's plausible prediction of the responsiveness to a "typical" price change makes it easier to trust its prediction of the responsiveness to a credibly-permanent, tax-induced price change. And this responsiveness is very high.

We should not, therefore, expect policy-driven price changes to induce the same consumption responses as natural, market-driven price changes—and nor should we assume policy interventions will be as ineffectual as typical price-elasticity measures sometimes suggest.

1.8 Conclusion

Commuting distance, vehicle stock, the rotes of daily life—there are many reasons to suspect that habits play a role in gasoline demand, and U.S. data provides compelling evidence that habits for gasoline—at least myopic ones—are strong. Evidence on whether consumers are myopic or forward-looking with respect to their gasoline habits is not quite as clear-cut: when the future price is divided into tax and before-tax components, each of these has a statistically-significant effect on consumption; but in models that include the future price as a whole, the coefficient on lead price is only significant under certain specifications. To reach a firmer conclusion about consumers' outlook, it might be useful in the future to consider household-level data.

Though the evidence is not entirely unambiguous, it is certainly possible—and, it seems, likely—that gasoline demand is shaped by rational habits. Rational habits mean that consumers will reduce gasoline consumption more in response to a price increase they perceive as permanent than to a price increase they perceive as temporary. This increases the scope of standard economic instruments to reduce gasoline consumption, implying that credible permanence gives any price increase an extra kick. This extra kick is overlooked by traditional measures of price elasticity, which do

not differentiate between permanent and fleeting price changes, and which may thus underestimate the potential of price-boosting measures to temper gasoline consumption. Conventional economic instruments may be a more potent method than previously anticipated.

1.A Parameter Values and Assumptions Used for Simulation in Section 1.4.2

Parameters: $\theta = 0.4$; $\beta = 0.99$; $\alpha = 0.01$; $\gamma_g = -100$; $\gamma_c = -100$; $y_1 = y_2 = y_3 = 100$

Past Values: $s_{g0} = 30$; $g_o = 3$

Prices

Other consumption good, c: $p_{c1} = p_{c2} = p_{c3} = 1$ in all cases

Gasoline:

Constant prices: $p_{g1} = p_{g2} = p_{g3} = 1$

Temporary 10% price increase: $p_{g1} = 1$, $p_{g2} = 1.1$, $p_{g3} = 1$

Permanent 10% price increase: $p_{g1} = 1$, $p_{g2} = 1.1$, $p_{g3} = 1.1$

1.B History of Federal Tax Rates on Gasoline

Effective from	Tax (nominal cpg)
June 21, 1932	1
June 17, 1933	1.5
January 1, 1934	1
July 1, 1940	1.5
November 1, 1951	2
July 1, 1956	3
October 1, 1959	4
April 1, 1983	9
January 1, 1987	9.1
December 1, 1990	14.1
October 1, 1993	18.4
January 1, 1996	18.3
October 1, 1997	18.4

References

[1] Alonso-Borrego, Cesar, and Manuel Arellano. 1999. Symmetrically normalized instrumental-variable estimation using panel data. Journal of Business & Economic Statistics 17(1): 36-49.

[2] Anderson, T. W., and Cheng Hsiao. 1981. Estimation of dynamic models with error components. Journal of the American Statistical Association 76(375): 598-606.

[3] Arellano, Manuel, and Steven Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. The Review of Economic Studies 58(2): 277-297.

[4] Arellano, Manuel, and Olympia Bover. 1995. Another look at the instrumental variable estimation of error-components models. Journal of Econometrics 68 : 29-51.

[5] Arellano, Manuel, Lars Peter Hansen, and Enrique Sentana. 1999. Underidentification? Econometric Society: Econometric Society World Congress 2000 Contributed Papers, no. 1824 (July 12, 1999).

[6] Auld, M. Christopher, and Paul Grootendorst. 2004. An empirical analysis of milk addiction. Journal of Health Economics 23: 1117–1133.

[7] Becker, Gary S., and Kevin M. Murphy. 1988. A theory of rational addiction. The Journal of Political Economy 96(4): 675-700.

[8] Becker, Gary S., Michael Grossman, and Kevin M. Murphy. 1994. An empirical analysis of cigarette addiction. The American Economic Review 84(3): 396-418.

[9] Berkowitz, Michael K., Nancy T. Gallini, Eric J. Miller, and Robert A. Wolfe. 1990. Disaggregate analysis of the demand for gasoline. Canadian Journal of Economics 23(2): 253-275.

[10] Blundell, Richard, Steve Bond, and Frank Windmeijer. 2000. Estimation in dynamic panel data models: Improving on the performance of the standard GMM estimator. IFS Working Papers W00/12, Institute for Fiscal Studies.

[11] Blundell, Richard, and Steven Bond. 1998. Initial conditions and moment restrictions in dynamic panel data models. Journal of Econometrics 87: 115-143.

[12] Bond, Stephen R. 2002. Dynamic panel data models: a guide to micro data methods and practice. Portuguese Economic Journal 1: 141-162.

[13] Bond, Steven, Céline Nauges, and Frank Windmeijer. 2002. Unit roots and identification in autoregressive panel data models: A comparison of alternative tests. Paper presented at International Conferences on Panel Data, 10th International Conference on Panel Data, Berlin.

[14] Breunig, Robert V., and Carol Gisz. 2009. An Exploration of Australian Petrol Demand: Unobservable Habits, Irreversibility and Some Updated Estimates. Economic Record 85(268): 73-91.

[15] Browning, Martin. 1991. A simple nonadditive preference structure for models of household behavior over time. The Journal of Political Economy 99(3): 607.

[16] California Energy Commission. 2005. Petroleum products yielded from one barrel of crude oil in California. Available from http://www.energy.ca.gov/gasoline/whats_in_barrel_oil.html (accessed 2006, 2008).

[17] Dahl, Carol. 1995. Demand for transportation fuels: A survey of demand elasticities and their components. The Journal of Energy Literature.

[18] Dahl, Carol, and Thomas Sterner. 1991. Analysing gasoline demand elasticities: A survey. Energy Economics 13(3): 203-210.

[19] Dahl, Carol A. 1979. Consumer adjustment to a gasoline tax. The Review of Economics and Statistics 61(3): 427-32.

[20] Energy Information Administration (EIA) a. Petroleum prices: Spot prices: Definitions, sources and explanatory notes. Available from http://tonto.eia.doe.gov/dnav/pet/tbldefs/pet_pri_spt_tbldef2.asp (accessed 5 June 2010).

[21] ——— b. Refiner motor gasoline sales volumes. Available from http://tonto.eia.doe.gov/dnav/pet/pet_cons_refmg_d_STX_VTR_mgalpd_a.htm (accessed 2 Apr 2008).

[22] ——— c. United States energy profile. Available from http://tonto.eia.doe.gov/country/country_energy_data.cfm?fips=US (accessed 2 Apr 2008).

[23] Espey, Molly. 1998. Gasoline demand revisited: An international meta-analysis of elasticities. Energy Economics 20(3): 273-95.

[24] Federal Highway Administration. 2006. Tax rates on motor fuel - 2003. Highway Statistics 2003. Available from http://www.fhwa.dot.gov/policy/ohim/hs03/htm/mf121t.htm (accessed 2 Apr 2008).

[25] Goel, Rajeev K. 1994. Quasi-experimental taxation elasticities of US gasoline demand. Energy Economics 16(2): 133-137.

[26] Goel, Rajeev K., and Michael A. Nelson. 1999. The political economy of motor-fuel taxation. Energy Journal 20(1): 43-59.

[27] Goodwin, P. B. 1992. A review of new demand elasticities with special reference to short and long run effects of price changes. Journal of Transport Economics and Policy (May 1992): 155-169.

[28] Graham, Daniel J., and Stephen Glaister. 2002. The demand for automobile fuel: A survey of elasticities. Journal of Transport Economics and Policy 36(1): 1-26.

[29] Heien, Dale. 2001. Habit, seasonality, and time aggregation in consumer behaviour. Applied Economics 33: 1649-1653.

[30] Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. 1988. Estimating vector autoregressions with panel data. Econometrica 56(6): 1371-95.

[31] 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.

[32] Nickell, Stephen. 1981. Biases in Dynamic Models with Fixed Effects. Econometrica 49(6): 1417-1426.

[33] Pashardes, Panos. 1986. Myopic and forward looking behavior in a dynamic demand system. International Economic Review 27(2): 387-97.

[34] Pollak, Robert A. 1970. Habit formation and dynamic demand functions. The Journal of Political Economy 78(4, Part 1): 745-763.

[35] Puller, Steven L., and Lorna A. Greening. 1999. Household adjustment to gasoline price change: An analysis using 9 years of US survey data. Energy Economics 21(1): 37-52.

[36] Roodman, David. 2006. How to do xtabond2: An introduction to "Difference" and "System" GMM in Stata. Center for Global Development, Working Paper No. 103. Available from http://www.cgdev.org/files/11619_file_HowtoDoxtabond6_12_1_06.pdf.

[37] Roodman, David. 2007. A short note on the theme of too many instruments. Center for Global Development, Working Paper No. 125. Available from http://www.cgdev.org/content/publications/detail/14256.

[38] Roodman, David. 2011 (first version 2003). XTABOND2: Stata module to extend xtabond dynamic panel data estimator. Center for Global Development, Washington. http://econpapers.repec.org/software/bocbocode/s435901.htm.

[39] Schaffer, Mark E. 2010. xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. Available from http://ideas.repec.org/c/boc/bocode/s456501.html.

[40] 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.

[41] Spinnewyn, Frans. 1981. Rational habit formation. European Economic Review 15: 91-109.

[42] StataCorp. 2009. Stata Statistical Software: Release 11. College Station, TX: StataCorp LP.

[43] Sterner, Thomas, Carol Dahl, and M. Franzen. 1992. Gasoline tax policy: Carbon emissions and the global environment. Journal of Transport Economics and Policy 26: 109-19.

[44] Windmeijer, Frank. 2005. A finite sample correction for the variance of linear efficient two-step GMM estimators. Journal of Econometrics 126(1): 25-51.

[45] Wirl, Franz. 1991. Energy demand and consumer price expectations: An empirical investigation of the consequences from the recent oil price collapse. Resources and Energy 13: 241-262.

[46] World Resources Institute. EarthTrends: Environmental Information. Available at http://earthtrends.wri.org. Washington, DC: World Resources Institute. Series: "Transportation: Motor gasoline consumption per capita" and "Transportation: Pump prices for super gasoline."

Data Sources

Consumer Price Index

[47] Bureau of Labor Statistics. Labstat. “Consumer Price Index, All Urban Consumers, U.S. City Average, All Items Less Energy, 1982-84 = 100.” <http://www.bls.gov/data/> (9 Feb 2006).

Crude Oil Price

[48] Energy Information Administration. “Spot Prices.” Series: Crude Oil: WTI – Cushing, Oklahoma. <http://tonto.eia.doe.gov/dnav/pet/xls/pet_pri_spt_s1_m.xls> (7 Feb 2006).

[49] Federal Gasoline Excise Tax Rates: Louis Alan Talley, “RL30304: The Federal Excise Tax on Gasoline and the Highway Trust Fund: A Short History.” <http://ncseonline.org/NLE/CRSreports/Transportation/trans-24.cfm> (Nov 2005).

[50] Federal Highway Administration, U.S. Department of Transportation. "Federal Tax Rates on Motor Fuel and Lubricating Oils." Highway Statistics 2000. Oct 2001. <http://www.fhwa.dot.gov/ohim/hs00/fe101a.htm> (Nov 2005).

Gasoline Price

[51] Energy Information Administration. "Gasoline Prices by Formulation, Grade, Sales Type." Series: Sales to End Users, Average, Through Retail Outlets. <http://tonto.eia.doe.gov/dnav/pet /pet_pri_allmg_a_EPM0_PTA_cpgal_m.htm> (27 Nov 2006).

Gasoline Sales Volumes

[52] Energy Information Administration. "Prime Supplier Sales Volumes." Series: Motor Gasoline. <http://tonto.eia.doe.gov/dnav/pet/ pet_cons_prim_a_EPM0_P00_Mgalpd_m.htm> (27 Nov 2006).

International Prices and Taxes

[53] Energy Information Administration. "Weekly Retail Premium Motor Gasoline Prices Including Taxes. Data for January 1, 1996 to Present." <http://www.eia.doe.gov/emeu/international/Gas1.xls> (26 Mar 2006).

[54] OECD - European Environmental Agency. Unleaded Petrol - Euro. <http://www2.oecd.org/ecoinst/queries/UnleadedPetrolEuro.pdf> (25 March 2008).

State Populations

[55] U.S. Bureau of the Census, Population Estimates Branch. "Intercensal Estimates of the Total Resident Population of States: 1980 to 1990." Aug 1996. <http://www.census.gov/popest/archives/1980s/st8090ts.txt> (22 Feb 2006).

[56] ———. "Table 1: Annual Estimates of the Population for the United States and States, and for Puerto Rico: April 1, 2000 to July 1, 2004." NST-EST2004-01. 22 Dec 2004. <http://www.census.gov/popest/states/tables/NST-EST2004-01.xls> (22 Feb 2006).

[57] ———. "Time Series of Intercensal State Population Estimates: April 1, 1990 to April 1, 2000." Table CO-EST2001-12-00. 11 Apr 2002. <http://www.census.gov/popest/archives/2000s/vintage_2001/CO-EST2001-12/CO-EST2001-12-00.xls> (22 Feb 2006).

State Gasoline Excise and Sales Tax Rates

Rates for all states, Jan. 1 or Dec. 31:

[58] Federal Highway Administration, U.S. Department of Transportation. "State Motor-Fuel Tax Rates, 1981-1995." Highway Statistics Summary to 1995. "Section I – Motor Fuel." 19 Feb 2003. <http://www.fhwa.dot.gov/ohim/summary95/mf205.pdf> (5 Jan 2006).

[59] ———. "State Motor-Fuel Tax Rates, 1988-2003." Highway Statistics 2003. Table MF-205. <http://www.fhwa.dot.gov/policy/ohim/hs03/htm/mf205.htm> (8 Nov 2005).

Dates of tax changes and intra-year tax changes: 39 additional sources; full list available from the author.

Chapter 2

Rational Habits and Uncertain Prices: Simulating Gasoline Consumption Behavior

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 chapter 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.

2.1 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 chapter 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.

2.2 Background

Over the years the rational habits literature has burgeoned, with papers suggesting alternative empirical structures[1]; analyzing the problem in continuous time[2]; and applying rational habits to goods as diverse as milk[3], cigarettes[4], and cocaine[5]. An overview of this literature is provided in Section 1.3.

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

[1] Becker, Grossman, and Murphy (1994)
[2] Houthakker and Taylor (1966) and Becker and Murphy (1988), for example.
[3] Auld and Grootendorst (2004)
[4] Baltagi and Griffin (2001), Chaloupka (1991)
[5] Grossman and Chaloupka (1998)

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.

2.3 Building a Model

2.3.1 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 $\overline{g}_t$:

$$\overline{g}_t = g_t - \delta s_{gt} + \gamma_g \tag{2.1}$$

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:

$$s_{gt} = \alpha g_{t-1} + (1-\alpha)s_{g,t-1}, \qquad 0 \leq \ \alpha \leq 1 \tag{2.2}$$

Using (2.2) 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:

$$s_{gt} = \alpha g_{t-1} + (1-\alpha)\left[\alpha g_{t-2} + (1-\alpha)s_{g,t-2}\right] \tag{2.3}$$

$$= \alpha g_{t-1} + \alpha(1-\alpha)g_{t-2} + (1-\alpha)^2 s_{g,t-2} \tag{2.4}$$

$$s_{gt} = \alpha g_{t-1} + \alpha(1-\alpha)g_{t-2} + (1-\alpha)^2\left[\alpha g_{t-3} + (1-\alpha)s_{g,t-3}\right] \tag{2.5}$$

$$= \alpha g_{t-1} + \alpha(1-\alpha)g_{t-2} + \alpha(1-\alpha)^2 g_{t-3} + (1-\alpha)^3 s_{g,t-3} \tag{2.6}$$

...

$$s_{gt} = \alpha \sum_{i=1}^{\infty} g_{t-i}(1-\alpha)^{i-1} \tag{2.7}$$

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

$$\overline{g}_t = g_t - \delta \left[\alpha \sum_{i=1}^{\infty} g_{t-i}(1-\alpha)^{i-1} \right] + \gamma_g \tag{2.8}$$

The parameter δ adjusts the strength of habits. If $\delta = 0$, then $\overline{g}_t = g_t + \gamma_g$, 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 $\overline{g}_t$ 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 $\overline{g}_t$ 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 $\overline{g}_t = g_t - \delta g_{t-1} + \gamma_g$. The impact of gasoline consumption on $\overline{g}_t$ 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

$$\overline{g}_t = g_t - \frac{\alpha - 1}{\alpha} \left[\alpha \sum_{i=1}^{\infty} g_{t-i}(1-\alpha)^{i-1} \right] + \gamma_g \tag{2.9}$$

$$\overline{g}_t = g_t + \sum_{i=1}^{\infty} g_{t-i}(1-\alpha)^{i} + \gamma_g \tag{2.10}$$

$$\overline{g}_t = \sum_{i=0}^{\infty} g_{t-i}(1-\alpha)^{i} + \gamma_g \tag{2.11}$$

The reference quantity $\overline{g}_t$ 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.

2.3.2 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:

$$\max_{g_1,g_2,\ldots,g_T} E_1\left[\sum_{t=1}^{T}\beta^{t-1}u\left(\overline{g}_t,c_t\right)\right] \tag{2.12}$$

$$(T = \text{integer or } \infty)$$

s.t. budget constraints[6];

$$\overline{g}_t = g_t - \delta s_{gt} + \gamma_g;\ s_{gt} = \alpha g_{t-1} + (1-\alpha)s_{g,t-1}, \qquad 0 \leq \alpha \leq 1$$

with

$$p_t \sim f\left(I_{t-1}\right), \text{ where } I_{t-1} \text{ is information available at time } t-1$$

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* $\{p_1,...,p_T\}$.

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 2.A and walk through the computational approaches in Appendix 2.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.

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

The problem laid out in Section 2.12 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

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

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 $high_t$:

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

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

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

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

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:

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

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

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

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 2.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:

$$\begin{aligned}
&\underset{g_1, g_2, g_3, \ldots, g_T}{\text{Max}} \; v_1 + \sum_{t=2}^{T} \beta^{t-1} E[v_t] \qquad (2.13) \\
&s.t. \text{ non-negativity constraints (2.15), (2.16), (2.17), and (2.18)}
\end{aligned}$$

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:

$$c_t = a_t - p_t g_t \geq 0 \;\forall\; t \qquad (2.14)$$

Condition (2.14) applies at $t = 1$:

$$c_1 = a_1 - p_1 g_1 \geq 0 \qquad (2.15)$$

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 $(a_t + \gamma_c - low_t g_t)$ or $(a_t + \gamma_c - high_t g_t)$ to be negative.

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

$$\begin{aligned} a_t + \gamma_c - high_t g_t &\geq 0,\ t \geq 2 \\ g_t &\leq \frac{a_t + \gamma_c}{high_t},\ t \geq 2 \end{aligned} \tag{2.16}$$

But of course $\frac{a_t+\gamma_c}{high_t} \leq \frac{a_t+\gamma_c}{low_t}$, so only $g_t \leq \frac{a_t+\gamma_c}{high_t}$ 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 ($\frac{a_t}{high_t}$) plus some small constant ($\frac{\gamma_c}{high_t}$)—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 (2.17) and the "adjusted" gasoline quantity over which the agent takes his preferences (2.18) are non-negative:

$$g_t \geq 0 \ \forall\ t \tag{2.17}$$

$$g_t - \delta s_{gt} + \gamma_g \geq 0 \ \forall\ t \tag{2.18}$$

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 2.3.2 and Appendices 2.A and 2.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.

2.3.4 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*.

2.4 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 (2.13) 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.

2.4.1 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:

$$p_1 = p_0 + shock_1 \tag{2.19}$$

$$p_t = p_0 + \rho^{t-1} shock_1 + u_t, \;\; t = 2, ..., T \tag{2.20}$$

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;
- $\rho \, \epsilon \, [0, 1]$ 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.

2.4.2 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 2.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 (2.20).[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

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

[7] Note 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.

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

2.4.3 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 2.1 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 2.2, which plots elasticity versus ρ for three different values of g_0. The lower the value of g_0, the more

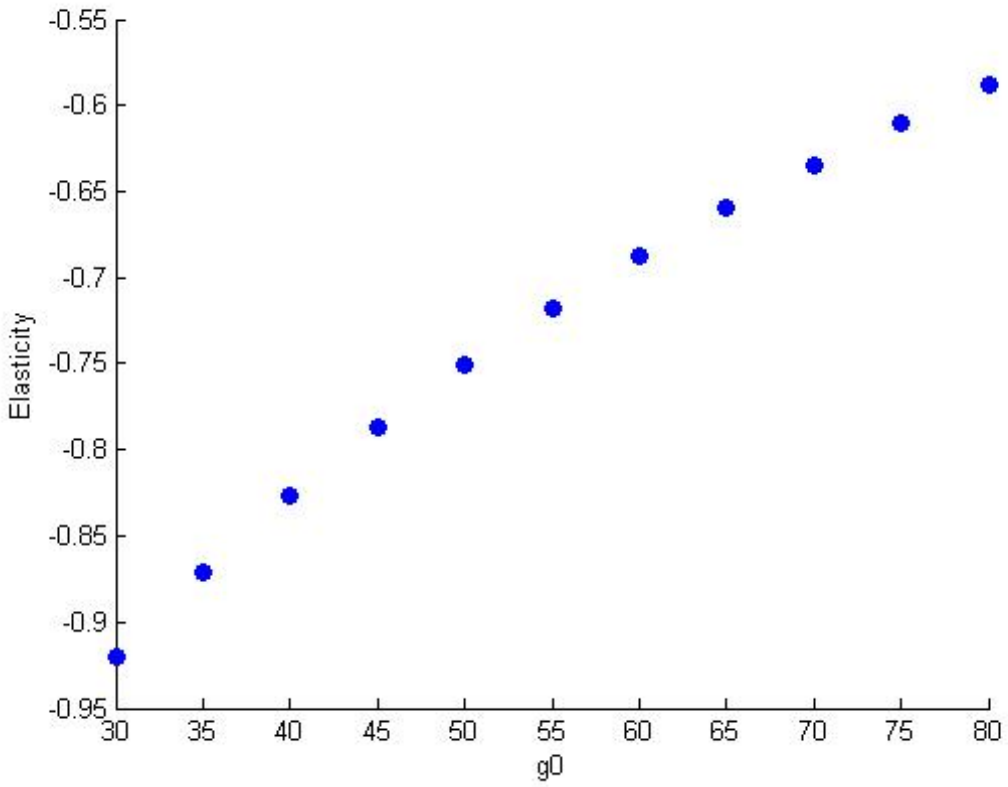


Figure 2.1: Elasticity versus g_0

sensitive elasticity is to ρ, which governs the durability of a price shock.

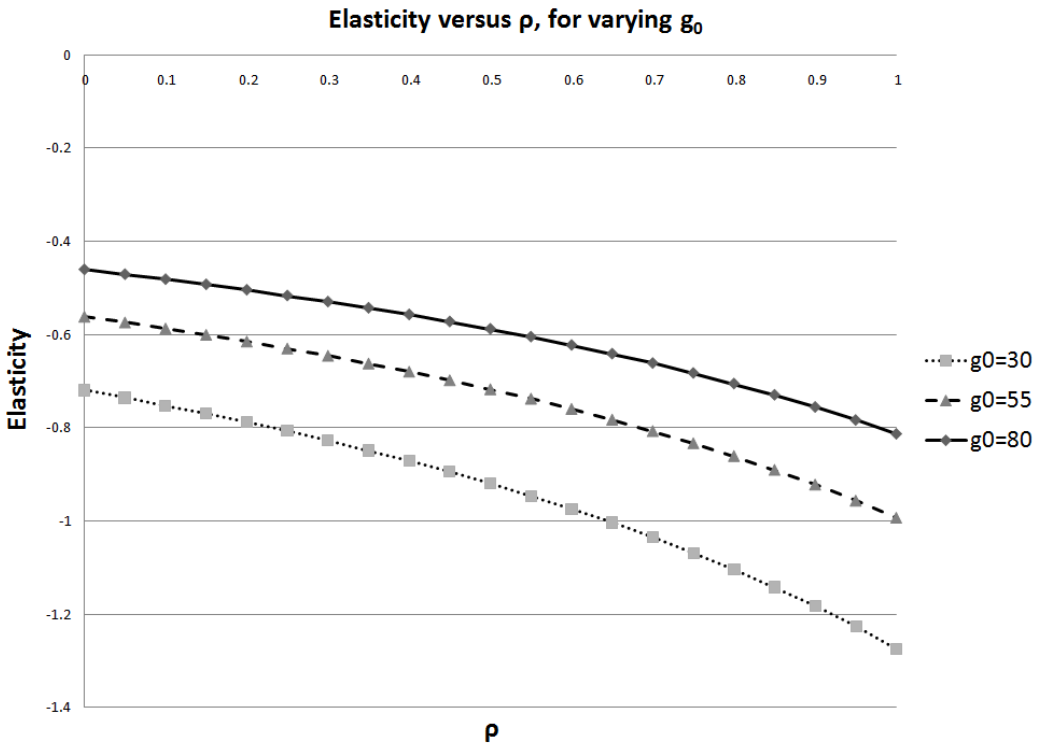


Figure 2.2: Elasticity vs. ρ, for varying g_0

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.

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:

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

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 2.3.

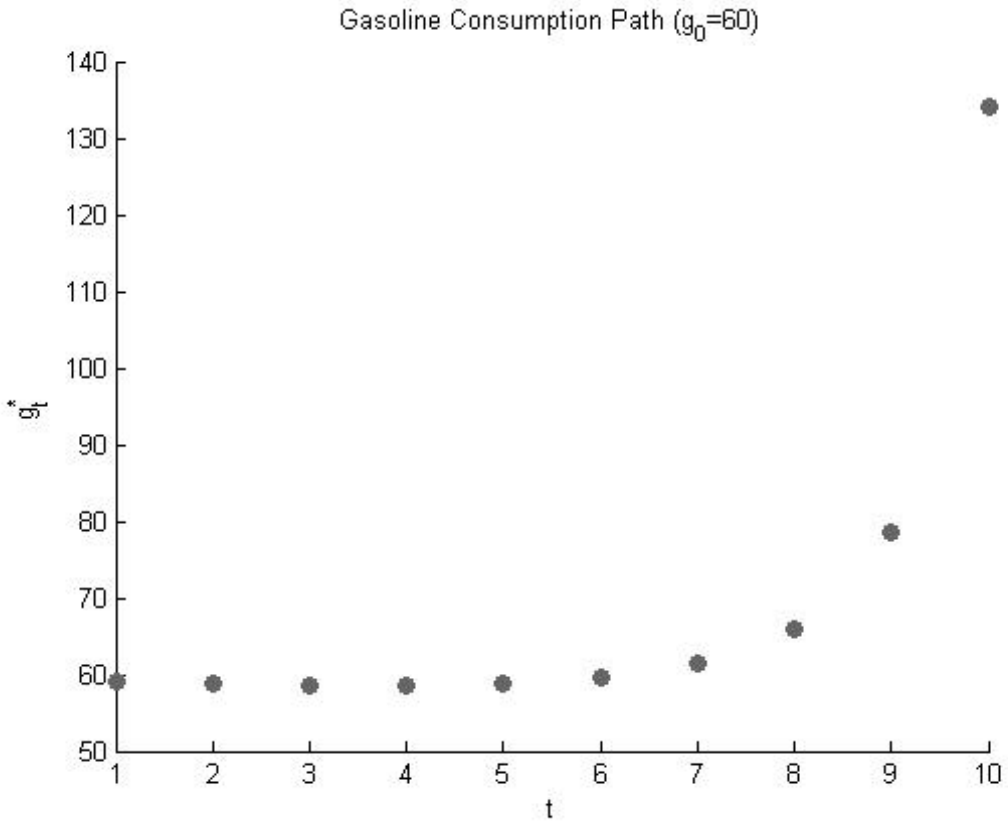


Figure 2.3: Gasoline Consumption over Time

Choosing g_o

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

Observe in Figure 2.3 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.

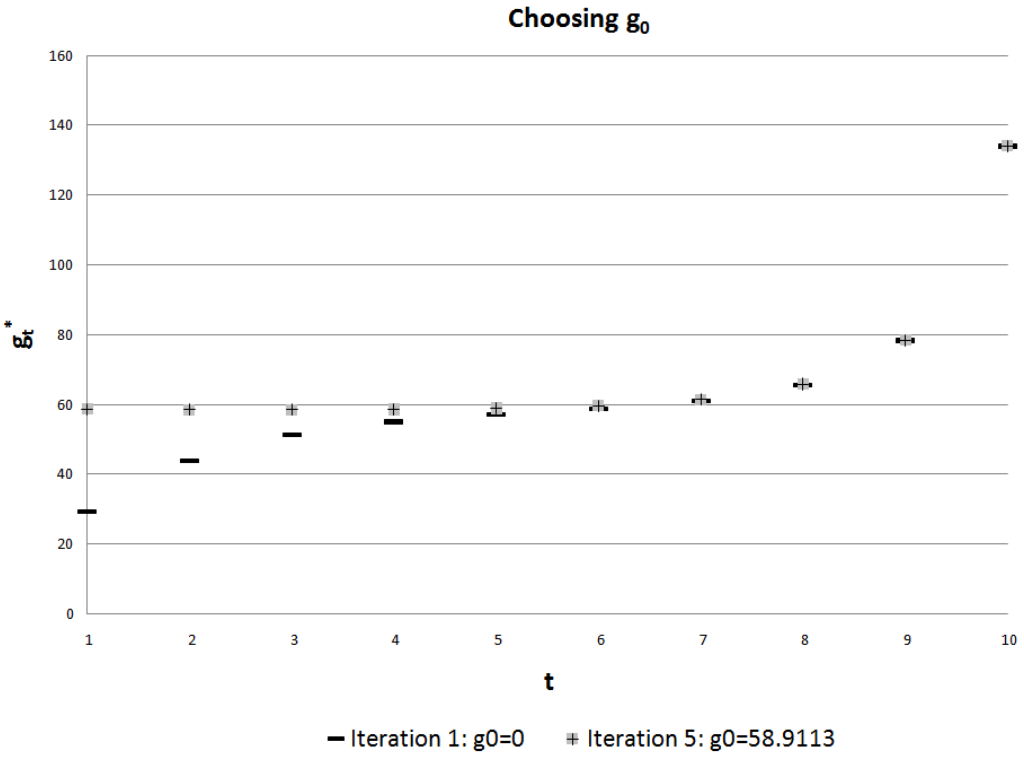


Figure 2.4: Choosing g_0: Consumption paths in first and fifth iterations

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 2.5.1, I verify that $T = 10$ is sufficiently large as to render end-of-horizon behavior moot.

2.5 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.

2.5.1 Time Horizon

My first concern is validating my choice of time horizon T. Although in Section 2.4.3 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 2.5 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.

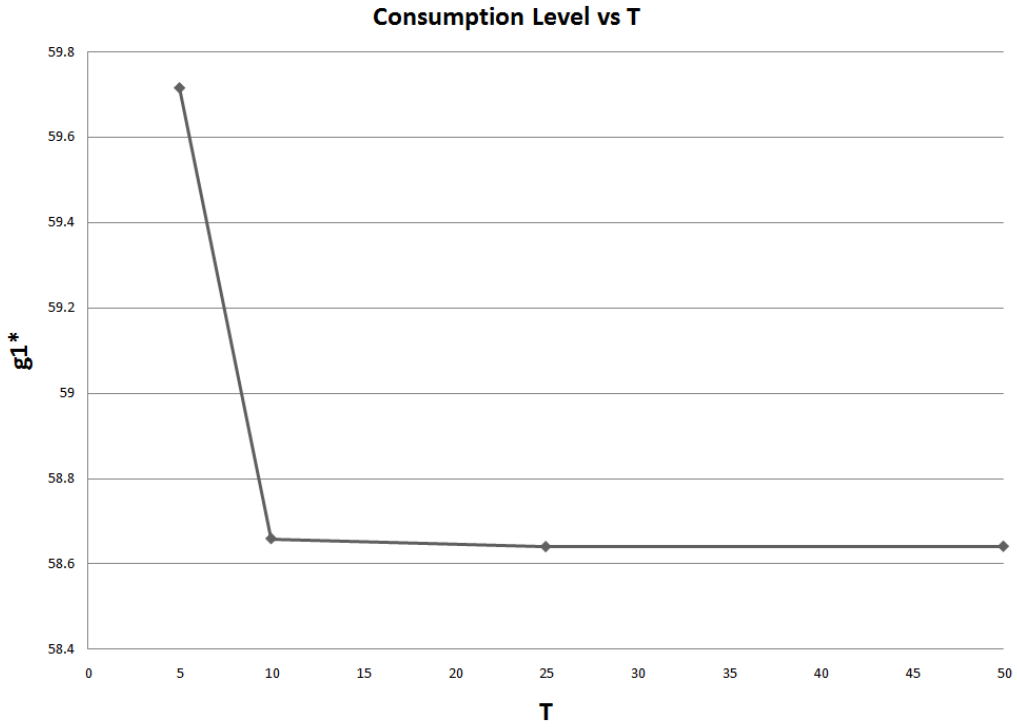


Figure 2.5: Optimal first-period consumption vs. horizon T

This holds for elasticity as well as consumption levels. Figure 2.6 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.

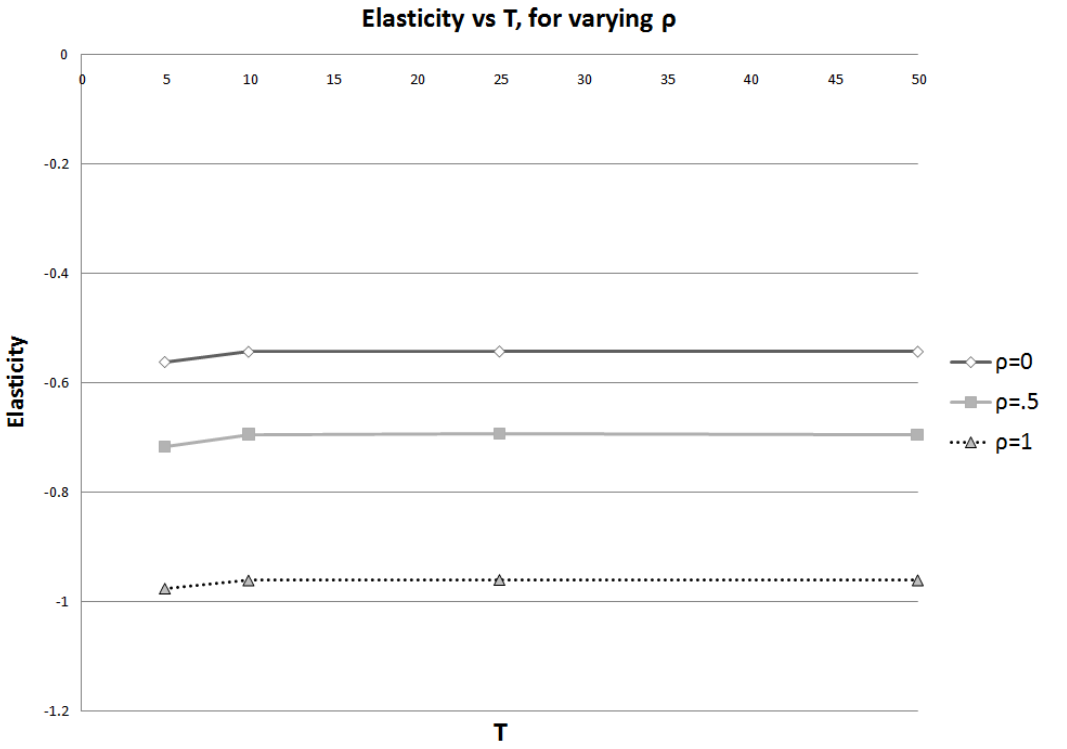


Figure 2.6: Elasticity vs. time horizon, for varying ρ

2.5.2 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 2.7 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.

2.5.3 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 2.4.3.

Figure 2.8 shows the effect of changing the spread of the uniform distribution from which prices are drawn for all periods $t = 2, ..., T$. 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

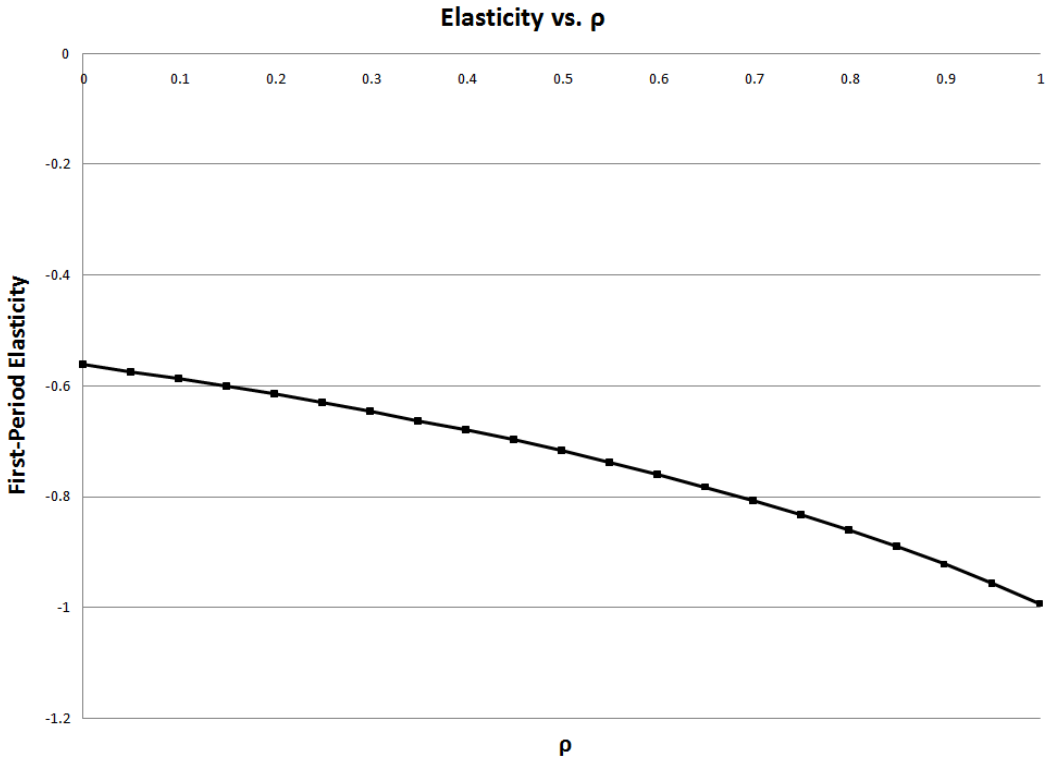


Figure 2.7: Elasticity vs. durability of price shocks (ρ)

is hard to observe in Figures 2.8 and 2.9, increasing ρ magnifies the effect of future uncertainty.

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 2.9, 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 2.10 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.

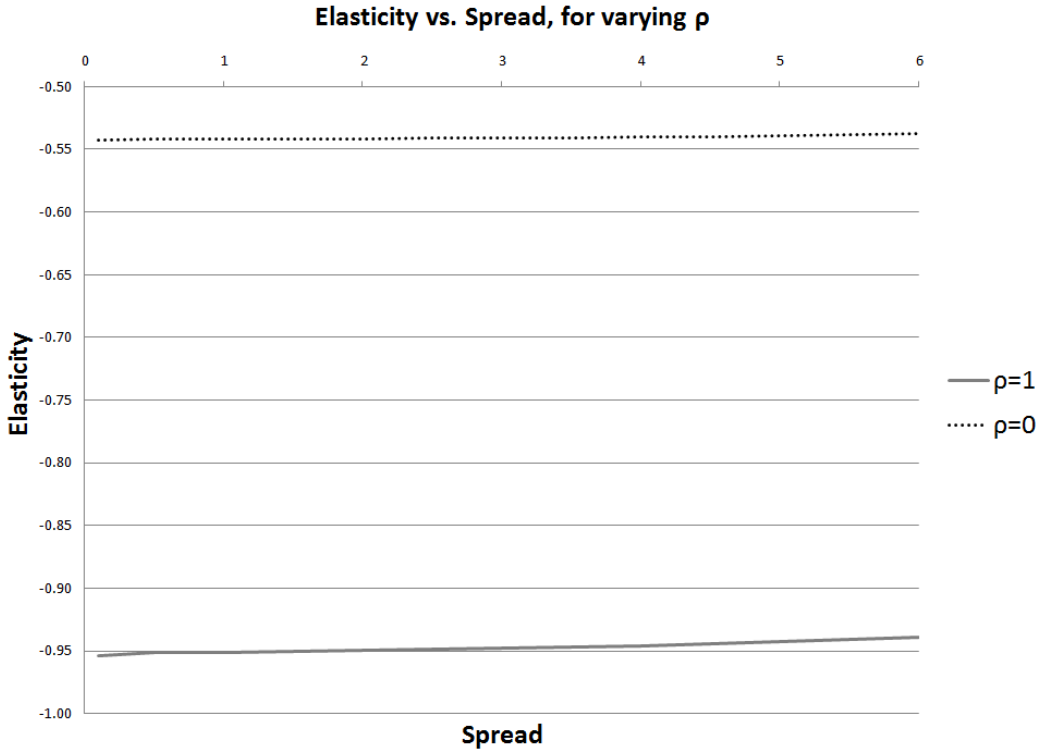


Figure 2.8: Effect of future price uncertainty, for varying price-shock durations

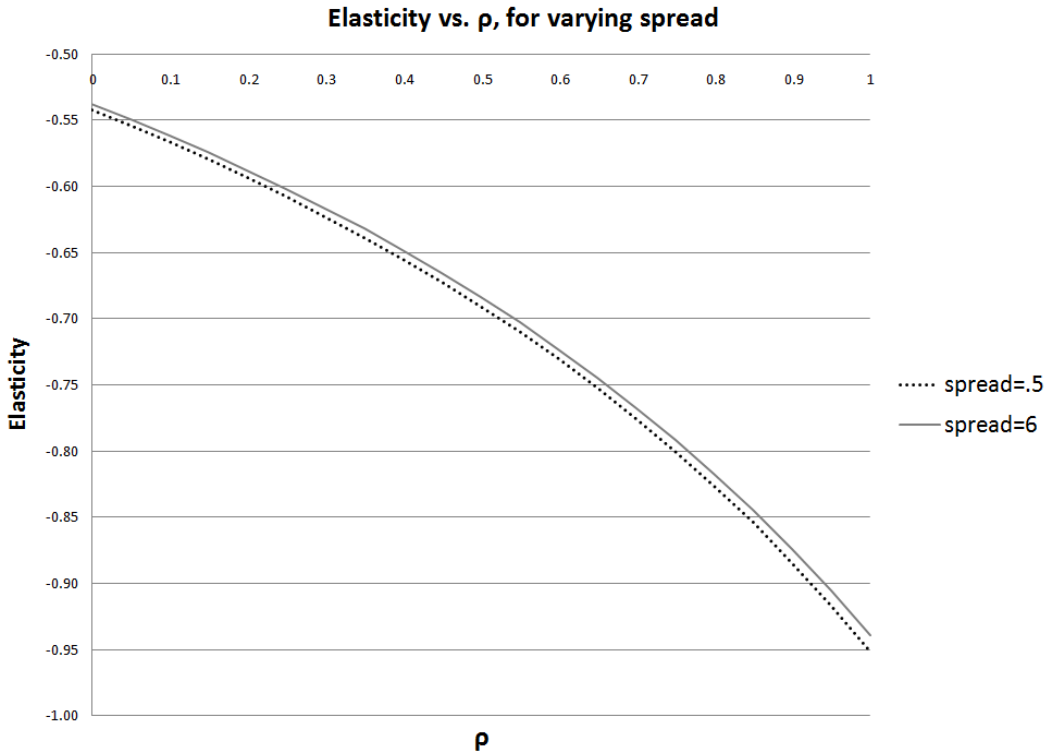


Figure 2.9: Elasticity vs. price-shock duration, for varying uncertainty

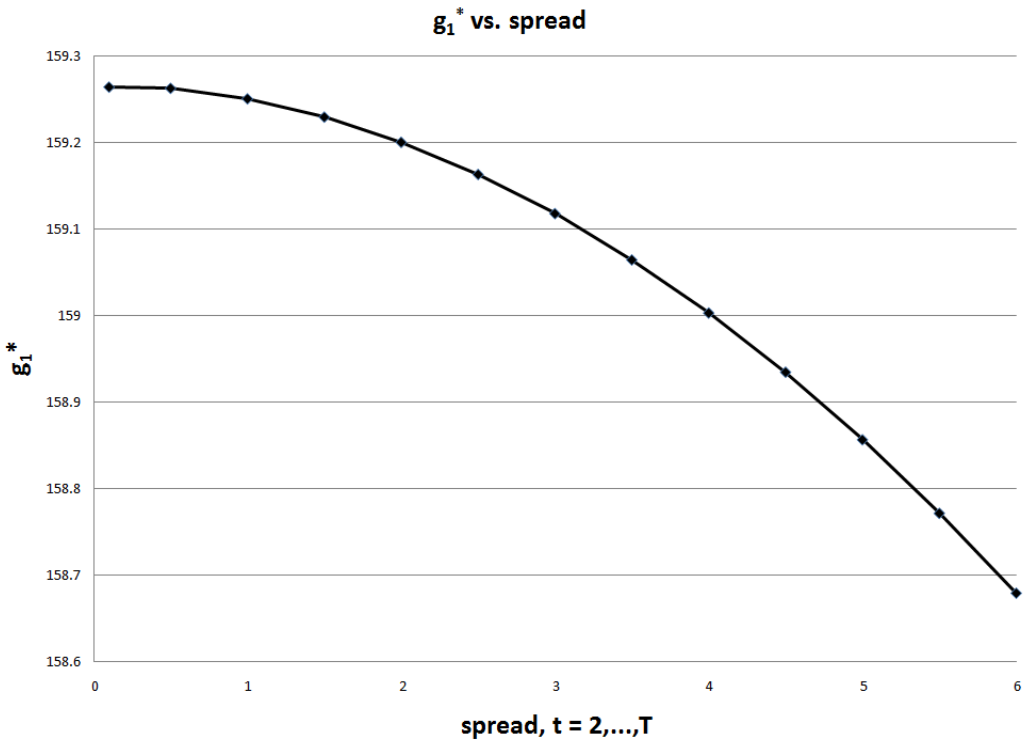


Figure 2.10: First-period consumption vs. price uncertainty

2.5.4 Habit Strength

All other things equal, stronger habits imply lower price sensitivity. This relationship is clear in Figure 2.11, which plots elasticity against the habit-strength parameter δ. When $\rho = 0.5$, increasing 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 2.11 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 2.12 plots optimal first-period consumption against habit strength: first-period consumption decreases with δ until about $\delta = 0.7$, whereupon the trend reverses.

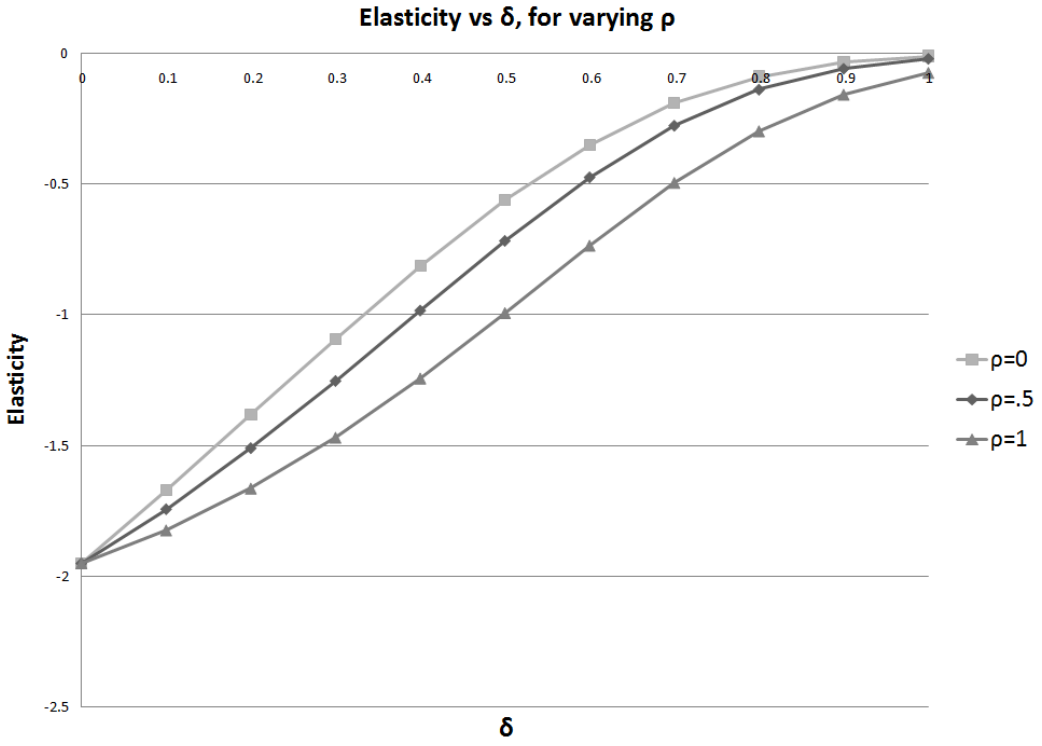


Figure 2.11: Elasticity vs. habit-strength, for varying ρ

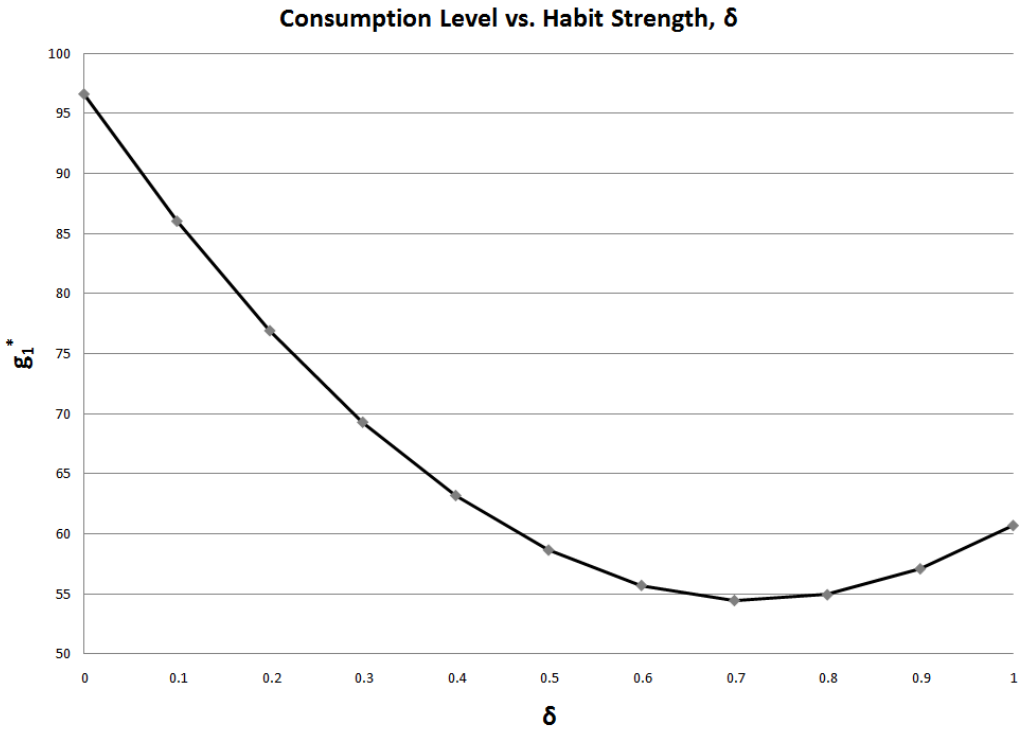


Figure 2.12: Optimal first-period consumption vs. habit strength

2.5.5 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 2.13, 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.

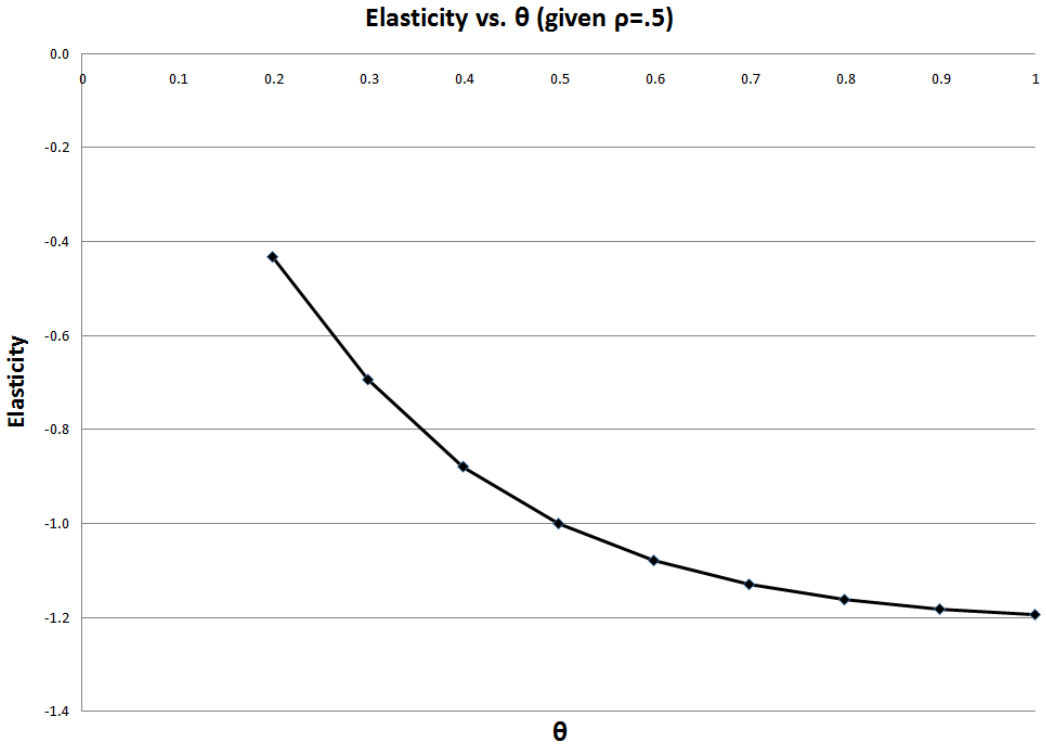


Figure 2.13: Elasticity vs. utility's weight on gasoline

Predictably, utility's weight on gasoline also affects optimal consumption levels, with first-period consumption g_1^* increasing with θ: see Figure 2.14.

Also predictably, a higher weight on gasoline in the utility function increases elasticity's sensitivity to the duration of price shocks. Figure 2.15 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.

2.5.6 Discounting of Future Utility

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

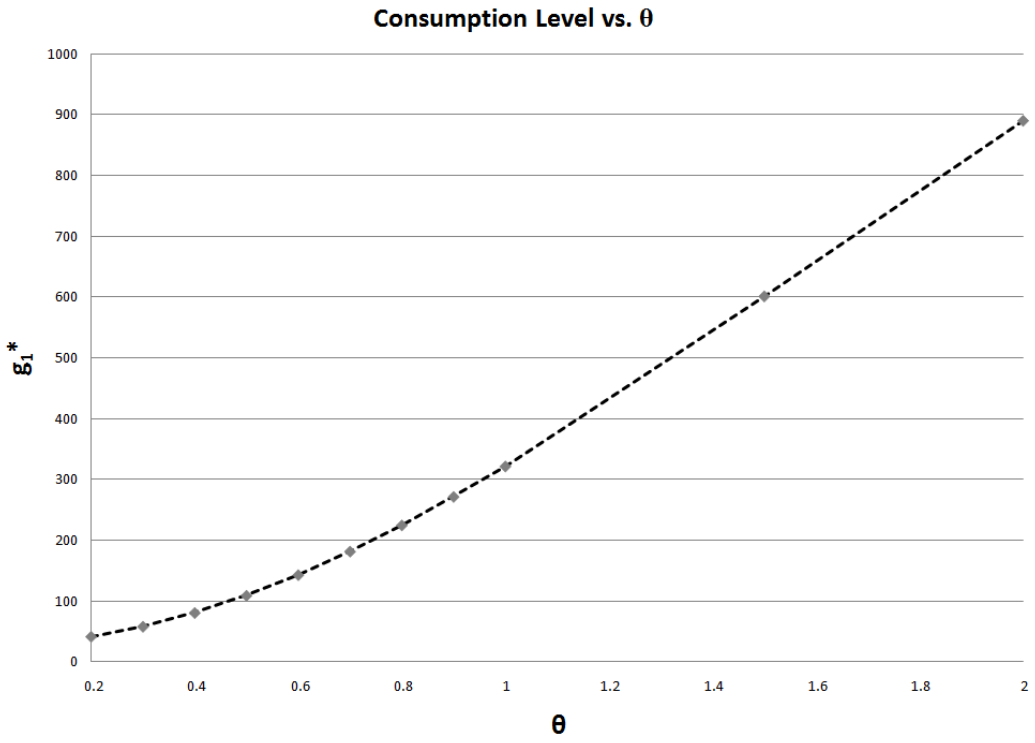


Figure 2.14: Optimal first-period consumption vs. utility's weight on gasoline

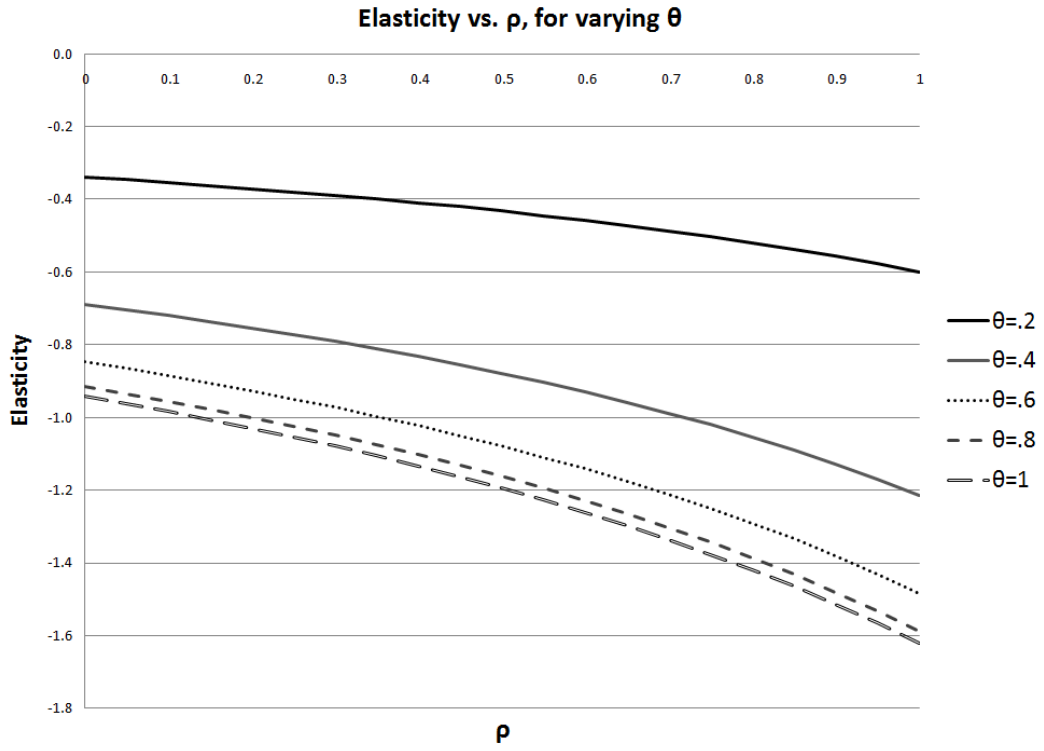


Figure 2.15: Elasticity vs. price-shock durability, ρ, for varying θ

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 2.16 and 2.17 illustrate these effects.

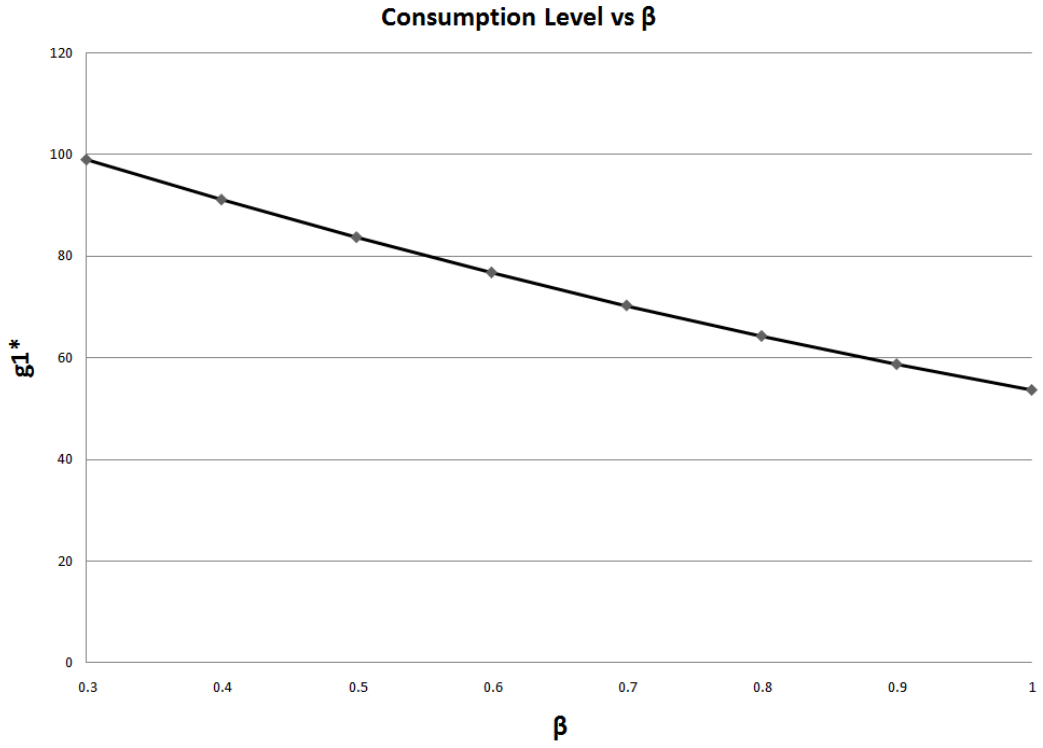


Figure 2.16: Optimal first-period consumption vs. discount factor β

As β increases, the duration of price shocks becomes more important to elasticity. This fits with intuition, and Figure 2.18 illustrates the effect.

2.6 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

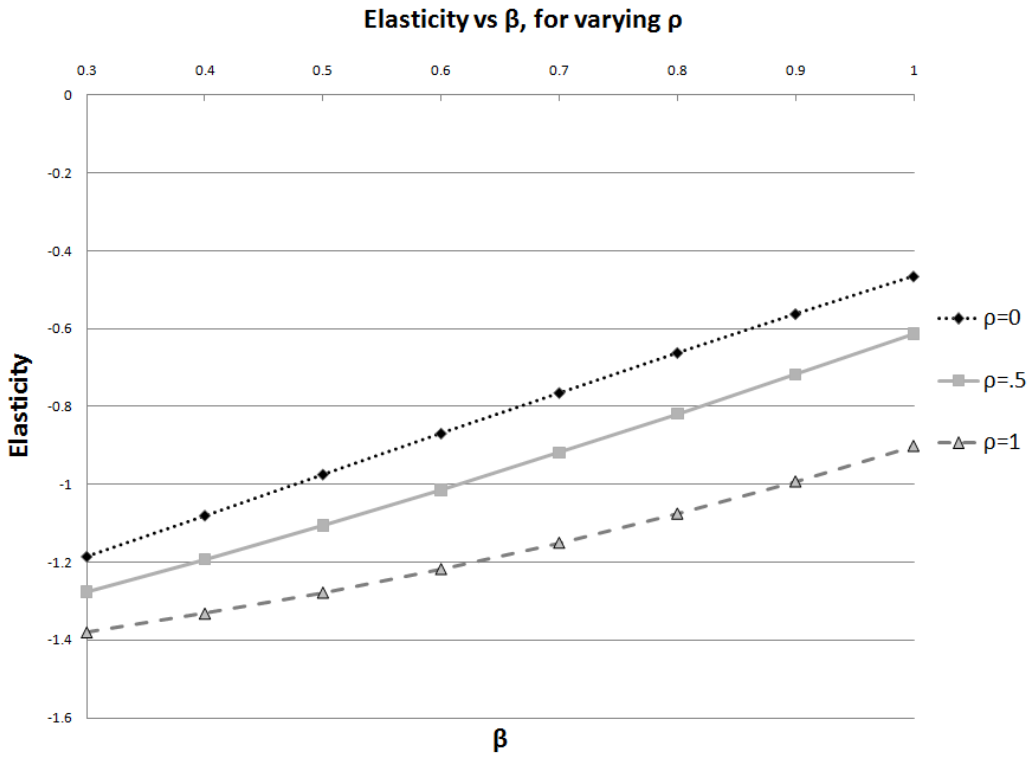


Figure 2.17: Elasticity vs. discount factor β, for varying ρ

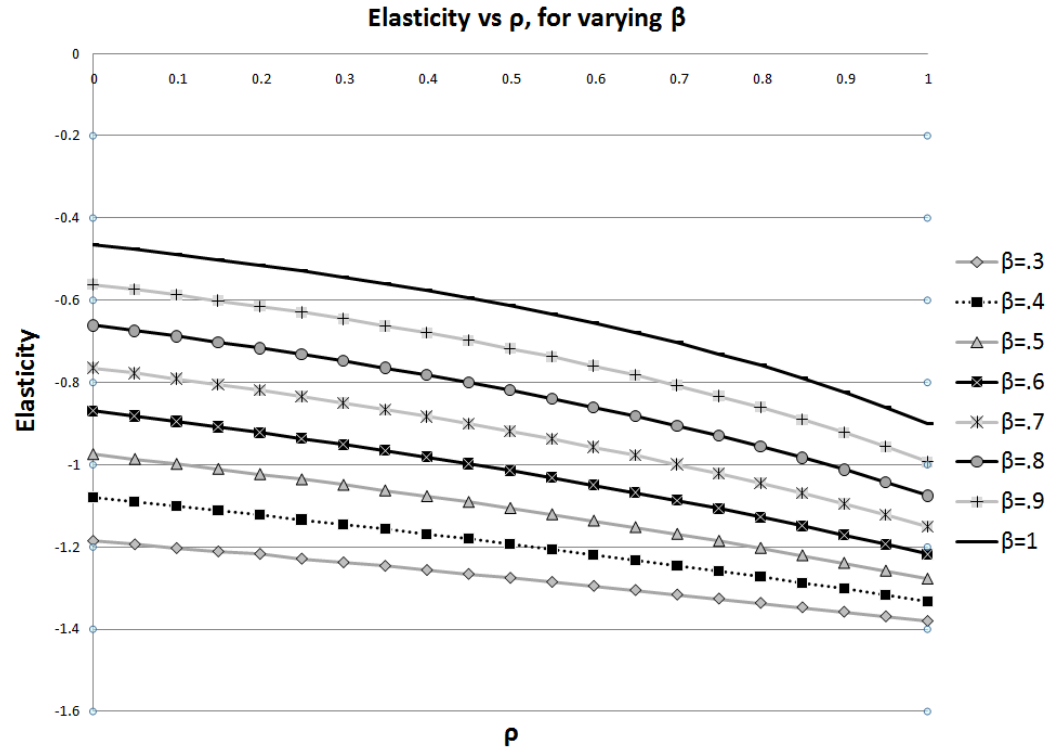


Figure 2.18: Elasticity vs. price-shock duration, for varying β

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 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 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 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.

Another potential approach is to look at a non-structural model that captures the effects implied by the structural model studied here. Given the structural model's simplifications, this more agnostic approach is appealing—and it is also more likely to be feasible, given the huge amount of data that would be necessary to identify the many parameters of the structural model. The non-structural approach will still allow us to explore how regional differences in price processes contribute to regional differences in price elasticity. Indeed, I will use a non-structural approach to examine international gasoline demand in Chapter 3.

For now, the structural model developed in this chapter identifies biases that we should beware when using traditional elasticity as a measure of consumers' responsiveness to the price of habit-forming goods.

2.A Setting up the Ideal Problem as a Dynamic Programming Problem

We can reformulate the problem given in (2.12) as a Bellman equation:

$$V_t\left(A_t, s_t, p_t, x_{1t}, ..., x_{nt}\right) = \max_{g_t, c_t} \left\{u(\overline{g}_t, c_t) + \beta E\left[V_{t+1}\left(A_{t+1}, s_{t+1}, p_{t+1}, x_{1,t+1}, ...x_{n,t+1}\right)\right]\right\} \quad (2.21)$$

where

- $\overline{g}_t = g_t - \delta s_t$;
- A_i indicates assets in period i;
- V_i is the value in period i;
- $x_{1t}, ..., x_{nt}$ 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:

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

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

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

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:

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

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,

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

and the price-predicting state variables,

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

Substituting the equations of motion into (2.21), we have

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

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:

$$u_{\overline{g}} + \beta E\left[V_{1,t+1}\left(1+r\right)\left(-p_t\right) + V_{2,t+1}\alpha\right] = 0$$

or

$$u_{\overline{g}} = \beta\left(1+r\right)\left(p_t\right) E\left[V_{1,t+1}\right] - \beta\alpha E\left[V_{2,t+1}\right] \tag{2.22}$$

2. FOC c:

$$u_c + \beta E\left[V_{1,t+1}\left(1+r\right)\left(-1\right)\right] = 0$$

or

$$E\left[V_{1,t+1}\right] = \frac{u_c}{\beta\left(1+r\right)} \tag{2.23}$$

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

$$E\left[V_{2,t+1}\right] = \frac{1}{\beta\alpha}\left(u_c p_t - u_{\overline{g}}\right) \tag{2.24}$$

Envelope Conditions

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

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

   Combining this with (2.23), 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: $V_{s,t} = V_{2,t} = -\delta u_{\overline{g}} + \beta\left(1-\alpha\right)E\left[V_{2,t+1}\right]$

   Combining this with (2.24), we can write

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

3. Condition with respect to current price p_t:

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

   Combining this with (2.23) yields

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

4. n conditions with respect to price-prediction variables

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

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

The problem set up in Appendix 2.A now looks like a traditional (albeit complicated) stochastic programming problem, to which one could apply standard techniques.

2.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 $q = 1, ..., D$ 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:

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

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

[8] For 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).

state given by gridpoint r as π_{qr}:

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

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 $V_q^0(c_t, g_t)$, $q = 1, ..., D$. 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

$$V_q^1 = \max_{g_t, c_t} \left[u(g_t - \delta s_t, c_t) + \beta \sum_{r=1}^{D} \pi_{qr} V_r^0 \right], \; q = 1, ..., D \tag{2.25}$$

The resulting improved guesses, V_q^1, $q = 1, ..., D$, can then be substituted into the right-hand side of (2.25), and we can again solve the maximization problem to improve our guess at the value function:

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

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

$$V_q^{a+1} = \max_{g_t, c_t} \left[u(g_t - \delta s_t, c_t) + \beta \sum_{r=1}^{D} \pi_{qr} V_r^a \right], \; q = 1, ..., D \tag{2.26}$$

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 $\left|V^{a+1} - V^a\right|$ 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 (2.26)). 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.

2.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 2.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 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

$$\mathbf{T}(g(\mathbf{x})) = \mathbf{0} \tag{2.27}$$

We'd like to know the policy function $g(\mathbf{x})$ that satisfies (2.27). Although we don't know the

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

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 $g(\mathbf{x})$, $\widehat{g}(\mathbf{x})$, can therefore be written

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

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

$$\mathbf{T}(\widehat{g}(\mathbf{x})) \approx 0 \tag{2.28}$$

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

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

If we wished to place more importance on satisfying some of the conditions than on others, we could, of course, minimize a projection of $\mathbf{T}(\widehat{g}(\mathbf{x}))$ 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 $\mathbf{T}(\widehat{g}(\mathbf{x}))$ 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 2.21 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.

2.C Deriving Expected Utility for Alternative Utility Functions

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

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

and if prices are distributed uniformly,

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

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

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:

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

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

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, $u_t = c_t^{\alpha} \overline{g}_t^{\beta}$:

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

2.D Optimization Methods

To perform the constrained optimization needed to find the optimal consumption path $g_1, ..., g_T$, 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 Koleman 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.

Chapter 3

Demand and Price Volatility: Rational Habits in International Gasoline Demand

Abstract

The combination of habits and a forward outlook suggests that consumers will be sensitive not just to prices but to price *dynamics*. In particular, rational habits models suggest *1.* that price volatility and uncertainty will reduce demand for a habit-forming good and *2.* that such volatility will dampen demand's responsiveness to price. These two implications can be tested by augmenting a traditional partial-adjustment or error-correction model of demand. I apply this augmented model to data on gasoline consumption, as rational habits provide a succinct representation for the investment and behavioral decisions that determine gasoline usage. The trade-offs among 2SLS, system GMM, and pooled mean group (PMG) estimators are considered, and my preferred PMG estimator provides evidence for the two implications of rational habits in a panel of 29 countries for the years 1990-2009. The sensitivity of certain results to the choice of estimator offers a cautionary illustration of the cost of assumptions such as coefficient heterogeneity. Given the evidence uncovered in favor of rational gasoline habits, such habits may help to explain some of the cross-country variation in "total" price elasticity. These habits also imply that the effect of price volatility must be taken into account when projecting the impacts of potential policies on gasoline consumption.

3.1 Introduction

The same consumer behavior that shapes gasoline demand also shapes the effectiveness of policies for controlling gasoline demand. Reducing this demand has become a widespread policy goal, driven by environmental concerns both global and local; and the intricacies of demand behavior are now of very practical interest.

Although this interest has generated a great deal of empirical work measuring the effects of income and prices on gasoline consumption, less attention has been devoted to the deeper behaviors underlying gasoline demand. Consumers may purchase fuel at the pump, but in fact they make their gasoline-buying decisions almost everywhere but the gas station. They make these decisions in the form of discrete, infrequent choices about what type of vehicle to buy and where to live in relation to work, as well as in nearly-continuous choices about daily routines—carpooling, driving style, how much non-essential travel to undertake and whether to cycle or take the car. The investment nature of the former decisions and the habitual nature of the latter help to explain consumers' sluggish responses to changes in the gasoline price, and also suggest that models of gasoline demand should allow for the effects of long-run choices.

One way to incorporate these effects is a rational habits model, in which consumers' utility for a particular good—in this case, gasoline—is a function of how much of the good they consumed in the past. Since past investment decisions affected past consumption, the habit setup captures both investment- and habit-influenced behaviors. The model assumes consumers are 'rational', or forward-looking; and therefore when deciding their current gasoline consumption, they consider how this consumption will affect their future utility. The future burden of a gasoline habit depends upon future market conditions, and so demand in this model depends upon consumers' expectations of the future, particularly of future prices.

If rational habits do in fact shape gasoline demand, then studies that confine themselves to the effects of income and contemporary prices on demand may overlook nuances of behavior that are relevant to policy. Given rational habits, demand will be affected by the *process* by which prices are generated as well as by the current price level, and differences in price elasticity across countries may be driven by differences in price regimes as well as by differences in environment and infrastructure.

The ideal way to examine the hypothesis of rational habits would be to estimate or calibrate a structural model, such as the one introduced in Chapter 2. Unfortunately, this model can only be solved numerically, and so its parameters cannot not be estimated using traditional estimation

methods. More problematically, the model contains too many parameters to be calibrated precisely using available data—eight, plus any parameters necessary to model the gasoline price process. Even to identify each of these parameters would be a stretch.

Fortunately, however, rational habits models suggest two implications that are easily testable by augmenting a traditional demand model with extra variables. If consumers are forward-looking,

1. demand for a habit-forming good will decline with the uncertainty in its (future) price, and

2. responsiveness to price changes will be dampened by price uncertainty and the expectation that price changes will be short-lived.

The first of these implications is proved by Coppejans et al. (2007), who consider a mean-preserving spread in the distribution of future prices and find that this reduces consumption of the habit-forming good. This first implication is also demonstrated in Chapter 2 of this thesis, where the level of demand is shown to decline with the variance of future prices (Figure 2.10). To test whether price uncertainty in fact reduces demand for gasoline, I introduce a measure of price uncertainty into a traditional, non-structural dynamic model of gasoline demand. Coppejans et al. take a similar approach in examining the effect of price uncertainty on smoking behavior, augmenting a static demand model with an estimate of the expected one-period-ahead standard deviation of price.

The second implication above is demonstrated in Chapter 2 of this thesis, where the magnitude of price elasticity is shown to be negatively related to the variance of the future price distribution (Figures 2.8 and 2.9) and positively related to the expected duration of price changes (Figure 2.7). To test whether price uncertainty dampens consumers' responsiveness to prices in practice, I augment my demand model with an interaction between a price-uncertainty measure and the gasoline price. The "total" elasticity with respect to price is therefore given by a combination of the price and interaction coefficients.

To capture the sluggish adjustment associated with any habits model, rational or myopic, I can add these two rational-habits regressors to a partial-adjustment, or ADL, model. Partial-adjustment models are frequently used to estimate gasoline demand, and my augmented version takes the form

$$g_{it} = \lambda g_{i,t-1} + \delta_1 y_{it} + \delta_2 p_{it} + \delta_3 \widehat{\sigma}_{it} + \delta_4 \left(p_{it} - \overline{p}_i\right) \widehat{\sigma}_{it} + \mu_i + \varepsilon_{it} \tag{3.1}$$

where g is log gasoline consumption per capita, y_{it} is log real income or expenditure per capita, p is the log real gasoline price, and $\widehat{\sigma}$ is a measure of price uncertainty. For simplicity I will sometimes

refer to the de-meaned gasoline price, $(p_{it} - \overline{p}_i)$, as $\widetilde{p}_{it}$. (Coefficients in (3.1) are restricted to be homogeneous across countries; later I will relax this assumption.) The rational habits model implies that δ_3 should be negative, with price uncertainty discouraging gasoline consumption, and δ_4 should be positive, with the interaction term offsetting some of the negative effect of prices:

$$g_{it} = \underbrace{\lambda}_{+,<1} g_{i,t-1} + \underbrace{\delta_1}_{+} y_{it} + \underbrace{\delta_2}_{-} p_{it} + \underbrace{\delta_3}_{-} \widehat{\sigma}_{it} + \underbrace{\delta_4}_{+} (p_{it} - \overline{p}_i)\, \widehat{\sigma}_{it} + \mu_i + \varepsilon_{it} \tag{3.2}$$

Short-run income elasticity is given here by δ_1; short-run "uncertainty elasticity", by δ_3; and short-run "total" price elasticity, by

$$\frac{\partial g_{it}}{\partial p_{it}} = \delta_2 + \delta_4 \widehat{\sigma}_{it}$$

Corresponding long-run elasticities are calculated, of course, by dividing each short-run elasticity by $1 - \lambda$.

Estimating this model entails some complications, and to address these I will reparameterize (3.1) into an error-correction model. I discuss alternative formulations of the model in Section 3.4.3. For now, I begin with a brief review of the literature on estimating gasoline demand and a discussion of my data set, which is a country-level panel. I then revisit the specification of my model and consider possible estimation methods, weighing the advantages and disadvantages of three approaches: least squares and system GMM estimation of the ADL model in (3.1) and maximum-likelihood estimation of an error-correction model allowing for coefficient heterogeneity. My preference is for the error-correction model with heterogeneous short-run coefficients, but I report the results of all these estimators in order to be transparent about the sensitivity of my results to the specification of the model—and to illustrate the potential pitfalls of certain approaches. For ease of comparison with the literature, I also estimate a "standard", non-habits version of the error-correction specification. Before concluding, finally, I use country-specific price volatilities to interpolate country-specific price elasticities, and I consider the extent to which cross-country variation in price volatility predicts cross-country variation in elasticity.

3.2 Gasoline Demand Literature Review

The predominant focus in the gasoline demand literature is on measuring the effects of income and prices on consumption. Little attention has been directed at structural models in which consumers choose consumption to maximize expected utility.

Neither myopic nor rational habits have been investigated as a potentially demand-shaping

behavior. Although Breunig and Gisz (2009) use an unobservable habit stock variable in a petrol demand regression, they insert this directly into a nonstructural demand model rather than into the consumer's utility-maximization problem. The habit stock in this scenario simply substitutes for the ADL model's directly-observed lags of consumption, allowing for gradual adjustment. Breunig and Gisz focus on the econometric complications generated by this unobserved stock, which implies that the error terms have a moving-average form. This correlation of the errors across time is analogous to that implied by the inclusion of lagged consumption in the usual ADL model (see Section 3.4). Breunig and Gisz estimate an ECM-type model using maximum likelihood to account for moving-average errors, and argue that, because this approach yields less precise elasticities than OLS estimates of an ADL-type model, the latter must be spuriously precise. Yet they do not establish that the OLS estimates owe their precision to a disregard of intertemporal error correlations rather than to their basis in a model with ten fewer parameters. Across the standard ADL and Breunig and Gisz's model, moreover, point estimates of long-run elasticities are about the same: -0.25 and -0.20, respectively, for price; and 0.34 and 0.27, respectively, for income.

As a determinant of gasoline demand, therefore, habits have not been introduced into the literature; and the effects of habit-related variables such as price uncertainty and price expectations remain unexamined. The typically-acknowledged determinants of demand—income and price—have received wide attention, however; and studies measuring income and price elasticities abound. In her 1998 meta-analysis, for instance, Espey (1998) considers 363 short- or medium-run and 277 long-run price elasticity estimates, along with similar numbers of income elasticity estimates. In the past there have been so many studies that simultaneous literature reviews have even managed, as Goodwin (1992) recalls, to draw on nearly-disjoint sets of papers. More recently, simultaneous reviews by Goodwin, Dargay and Hanly (2004) and Graham and Glaister (2004), commissioned by the same source for the same purpose, have continued to draw from a pool of literature large enough to allow substantial differences in their samples.

In this large literature, gasoline demand elasticities have run a wide gamut, with own-price elasticity estimates ranging from 0 to -1.36 in the short run and 0 to -2.72 in the long run, and income elasticities ranging from 0 to 2.91 in the short run and 0 to 2.73 in the long run (Espey 1998). Overall, however, the reviews are in basic agreement about average elasticities. Average own-price elasticity, these reviews find, is around -0.25 to -0.30 in the short run and -0.6 to -0.8 in the long run; average income elasticity is around 0.4 in the short run and somewhere around unit elastic in the long run. Some of these averages are summarized in Table 3.1.

Although much of the variation in estimated elasticities has not been explained—and may, in fact, arise from rational habits and variation in price volatility—literature surveys have uncovered

Table 3.1: Average elasticities in gasoline demand literature reviews

	Average **price** elasticity		Average **income** elasticity	
Study	Short run	Long run	Short run	Long run
Dahl and Sterner (1991)[a]	-0.24	-0.80	0.45	1.31
Goodwin (1992)[b]	-0.27	-0.71		
Espey (1998)	-0.26	-0.58	0.42	0.88
Goodwin, Dargay and Hanly (2004)[c]	-0.25	-0.64	0.39	1.08
Brons et al. (2008)	-0.34	-0.84		

[a]Lagged endogenous models, annual. [b]Time series models

[c]Price: time series models; income: dynamic models

some trends. Studies based on a panel of countries, for instance, tend to produce price elasticities that are similar to single-country elasticities for the long run but of higher magnitude for the short run (Espey 1998). Short-run price responsiveness seems to be relatively low in the United States and relatively high in Europe (Espey 1998). Including a measure of vehicle ownership and/or the characteristics of the vehicle stock affects the resulting estimates (Dahl and Sterner 1991, Espey 1998). And price elasticities may be changing over time: Espey (1998) observes that short-run price elasticities appear to have decreased in magnitude, and long-run price elasticities *increased* in magnitude, since the 1970s and 1980s. Hughes, Knittel and Sperling (2006) corroborate this shift in short-run price elasticities for the U.S.

The majority of the studies summarized in these reviews are based on partial adjustment models. In recent years, increasing attention has been diverted toward error-correction models and questions of cointegration. So far it seems all these studies have looked at single countries (or the world as a whole) rather than panels, and their estimates are summarized in Table 3.2. The most striking difference between these ECM-based price elasticity estimates and the average elasticities reported in the literature reviews is that the ECM-based elasticities are generally smaller in magnitude. Whether the ECM model is actually responsible for this difference, however, is not clear: the preponderance of single-country ECM studies and the tendency for single-country studies to yield smaller-magnitude short-run price elasticities (Espey 1998) might also explain some of the difference. In Section 3.5 I will report ECM results that are consistent with a story in which the panel dimension, rather than the choice of an ECM, drives this difference. Not only do my panel-based ECM price estimates turn out to be relatively high in magnitude, but those of my estimators that exploit cross-section variation (PMG and DFE) yield elasticities that are higher in magnitude than those elasticities based on single-country regressions (MG). The differences between these estimators will be explored in Section 3.4.

Table 3.2: Gasoline demand studies based on cointegration and error-correction models

		Price elasticity		Income elasticity	
Study	Country	Short run	Long run	Short run	Long run
Akinboade, Ziramba and Kumo (2008)	South Africa		-0.47		0.36
Alves and Bueno (2003)	Brazil	-0.0919	-0.465	0.122[g]	0.122[g]
Bentzen (1994)[a]	Denmark	-0.32	-0.41	0.89	1.04
Cheung and Thomson (2004)	China	-0.19	-0.56	1.64	0.97
De Vita, Endresen, and Hunt (2006)[b]	Namibia		-0.794		0.957
Eltony and Mutairi (1995)	Kuwait	-0.37	-0.46	-0.47	0.92
Krichene (2002)[c]	World	-0.02	-0.005	1.54	1.2
Nadaud (2004)[d]	France	-0.06	-0.09	0.27	0.28
Polemis (2006)	Greece	-0.10	-0.38	0.36	0.79
Ramanathan (1999)	India	-0.209	-0.319	1.178	2.682
Ramanathan and Subramanian (2003)[d]	Oman	-0.05	-0.52	0.35	0.96
Rao and Rao (2009)	Figi		-0.159 to -0.244		0.427 to 0.462
Samimi (1995)[e]	Australia	-0.2	-0.12	0.25	0.52
Wadud, Graham, and Noland (2009)[f]	US	-0.085	-0.116	0.520	0.592

[a] Elasticities with respect to vehicles per capita substituted for income elasticities.

[b] 1990q1-2002q4

[c] Estimates elasticity of demand for crude oil, not gasoline. Results for 1973-1999.

[d] As reported in Wadud, Graham, and Noland (2009).

[e] Energy for transport, not just gasoline.

[f] Single-step nonlinear least squares, post-1978.

[g] Reported short- and long-run elasticities in fact the same.

3.3 Data and Specification of Variables

3.3.1 Data

My data set consists of a panel of 29 countries for the period 1990-2009. As price data for some countries is limited, the actual series length ranges from fifteen to twenty years, with an average of 18.7 years. A list of included countries is provided in Appendix 3.A.

In order to focus as much as possible on passenger vehicles rather than freight, data is isolated to gasoline, not diesel. Information on gasoline consumption[1] is taken from the International Energy Agency (IEA)'s *Oil Information* (2009) and transformed into per capita terms using annual population estimates from the UN's *World Population Prospects* (2009). Only annual consumption data is provided; quarterly observations are not available. Data on gasoline prices and taxes, broken down by product and grade, is taken from the IEA's *Energy Prices and Taxes* (2009Q4). International differences in product definitions and regulations mean that data availability for each product varies by country. Depending on this availability, I use either regular unleaded or 95 RON to create a price series for each country. These choices are discussed further in Appendix 3.A.

As a measure of income, I take GDP from the World Bank's *World Development Indicators* (2009) and transform this into real per capita terms. When crude oil prices are used, these are the spot prices for the Brent stream, taken again from the IEA's *Energy Prices and Taxes* (2009Q4). I concentrate on this stream because the Brent is used extensively as a pricing benchmark.[2]

For reasons explained below, I consider regressions based both in a common currency and in real local currencies. When dealing in a common currency, I convert the IEA's USD-denominated series into real 2005 USD. When dealing in local currencies, I convert all prices and income from nominal to real terms using country-specific CPIs taken from the IEA's *Energy Prices and Taxes* (2009Q4).

Further specifics about my data and data sources are provided in Appendix 3.A.

3.3.2 Measuring Price Volatility

In order to begin examining the model in (3.1), I need a measure of the price uncertainty or volatility faced by gasoline consumers. The simplest measure of volatility would, of course, be the country-specific variance or standard deviation of the (log) gasoline price. Such a measure, taken over the entire sample period and constant over time, however, would hold two disadvantages:

[1] This series ("motor gasoline demand") explicitly excludes aviation gasoline.

[2] The West Texas Intermediate (WTI) stream is widely used in North America. Using this stream instead of or alongside the Brent, however, makes little difference in the resulting estimates.

first, it would imply that consumers had information about future prices that they do in reality not have; and second, it would make identification of price volatility effects impossible without restrictive assumptions. Instead, I use a measure of the *rolling* standard deviation of log prices. This constrains consumers' information about volatility to current and past prices, and also captures the evolution in price volatility over time.

Annual price observations mask volatility within each year. In the extreme case, two countries with identical annual price series could have wildly different price paths from month to month or quarter to quarter. Relying on annual data to construct a measure of price volatility, therefore, may yield a measure that is biased toward zero, with the size of the 'bias' increasing with 'true' volatility. To mitigate this problem, I exploit the availability of quarterly price data to construct the rolling standard deviation measure. This has the added advantage of allowing me to keep several years of data at the beginning of the period that I would otherwise lose to measuring volatility.

I construct the rolling standard deviation for each quarter, denoted $\widehat{\sigma}_{iq}$, as

$$\begin{aligned} \widehat{\sigma}_{iq} &= \sqrt{\widehat{\sigma}^2_{iq}}\text{, where} \\ \widehat{\sigma}^2_{iq} &= \frac{1}{x}\sum_{j=1}^{x}\left(p_{i,q-j} - \overline{p}_{iq}\right)^2 \\ \text{and } \overline{p}_{iq} &= \frac{1}{x}\sum_{j=1}^{x} p_{i,q-j+1} \end{aligned} \tag{3.3}$$

where x is the number of quarters in the rolling window and p_{iq} is the log real price of gasoline in quarter q. As a sensible default, I choose $x = 4$. In this case, the fourth-quarter rolling standard deviation in any year is the standard deviation of the year's prices from their year-long mean, and it is this that I use as my measure of annual rolling standard deviation. As an alternative, one could average $\widehat{\sigma}_{iq}$ over the quarters in each year. (The chief difference of this alternative method lies in the mean price, $\overline{p}_{iq}$, from which each quarter's deviations are calculated: in the alternative method, prices are always compared to a *past* average. In the default measure, each quarter's price is compared to the mean price for the entire year, which for the first three quarters includes *lead* prices.) As an alternative, one could, of course, lengthen the window of the rolling standard deviation beyond $x = 4$.

A further alternative would be to consider price *predictability* rather than volatility—that is, to measure the performance of price forecasts in each country. This would entail modelling each country's price series, generating rolling forecasts, and calculating root mean squared forecasting error for each period. One could also consider modelling the evolution of price volatility over time

and incorporating price-volatility forecasts into the demand model. This would be a particularly worthwhile exercise as the theory developed in Chapter 2 posits a relationship between demand and expected *future* price volatility, not current volatility: the rolling-volatility measure used at present functions as a proxy for expected future volatility, a very simple forecast. Better volatility forecasts may be possible, particularly if gasoline prices follow an ARCH-type process and tend to move between periods of high and low volatility. If consumers are aware of better forecasts, then the current proxy may contain substantial measurement error. Even if this measurement error is white noise, it will bias estimates of volatility's effects toward 0. The rolling-volatility proxy may, therefore, render the current model susceptible to underestimating the influence of uncertainty on demand. Although I stick to the rolling-volatility measure for now, price forecastability measures and explicit volatility forecasts are well worth examining in the future.

One extension I do consider now is the use of before-tax prices rather than total prices to calculate $\widehat{\sigma}_{iq}$. This variation is of interest because "volatility" is not synonymous with "uncertainty": tax increases may, for example, contribute to price volatility within the year they come into effect but actually lead to a *reduction* in price uncertainty. If this is the case, then there is, in effect, a systematic measurement error in $\widehat{\sigma}_{iq}$ that will bias the coefficient on the interaction term downward and make consumers' price elasticity appear less sensitive to price volatility. Using before-tax prices to construct $\widehat{\sigma}_{iq}$ circumvents this danger, and so I will consider this alternative definition as a check.

3.3.3 Currency Issues

The international nature of the panel leads to an additional consideration, namely how to deal with prices, price volatility, and income denominated in different currencies. These variables must be treated in a way that allows comparisons across countries, and ideally they should capture real prices and income, and variation therein, as perceived locally.

One approach is to transpose all variables into a common currency. Indeed, this is the approach toward which convention in the gasoline- and energy demand literature leans,[3] and it is the approach I shall adopt for my main discussion. Measuring all monetary variables in the same units (in this case, real (2005) US dollars) has the advantage that it does not introduce restrictions on the type of estimator that can be deployed; and it has the side benefit, of course, of allowing easy comparisons of price and income levels across countries.

Provided care is taken in the choice of estimator, however, another option is to work in log

[3] See, for example, Angelier and Sterner (1990), Baltagi and Griffin (1983, 1997), Dahl (2011), Johanssen and Schipper (1997), Judson, Schmalensee, and Stoker (1999), Narayan and Smyth (2007), Nguyen-Van (2010), and Storchmann (2005).

real local currency. As a check on my common-currency findings, I also estimate local-currency versions, with results reported in Appendix 3.B. To see that the local-currency approach can be valid, imagine that each country i has a real gasoline price P_{it}, denominated in real local currency units, and let the real exchange rate with respect to some common currency be given by $r_t = r_i e_{it}$. The time-invariant component, r_i, represents the long-run exchange rate, and e_{it} represents fluctuations away from relative PPP. At any time, country i's gasoline price can be written in real common-currency units as

$$P_{it} r_i e_{it}$$

The log of this real common-currency price is

$$\ln P_{it} + \ln r_i + \ln e_{it}$$

and so the log real local-currency price, $\ln P_{it}$, can be thought of as

$$\ln P_{it} = \ln (\text{real common-currency price}) - \ln r_i - \ln e_{it}$$

This log *local* currency measure therefore removes $\ln e_{it}$, the fluctuations from relative PPP that might otherwise contribute noise to a measure of locally-perceived prices. The removal of this noise is the chief advantage of using log real common-currency prices. Unfortunately, in addition to $\ln e_{it}$, the log local currency measure also removes $\ln r_i$, the term that allows us to compare prices across countries. Essentially we could think of each log real local-currency price $\ln P_{it}$ as containing a measurement-error term, $-\ln r_i$, which pushes $\ln P_{it}$ away from a measure that would be comparable across countries.

Fortunately, since $-\ln r_i$ is a constant, it disappears as soon as $\ln P_{it}$ is de-meaned or differenced. Any estimator that does so, therefore, will be immune to this "measurement-error" problem. A fixed-effects (within-groups) estimator, for example, will remove $-\ln r_i$ by subtracting the country-specific mean from each variable; the pooled mean groups estimator discussed in Section 3.4.3 removes $-\ln r_i$ by differencing. In these types of estimators, prices and income become unitless, measured in percentage changes or percentage deviations. Only when an estimator exploits cross-sectional variation but fails to de-mean or difference is the "measurement-error" problem really a problem, biasing estimates of the coefficients on prices and income toward 0. This bias will be apparent in the pooled 2SLS estimators that I use to establish an upper bound for λ (Tables 3.12 and 3.22), where the local-currency versions find smaller-magnitude price and income elasticities

than their common-currency counterparts.

As an aside, it should be noted that the inclusion or exclusion of $\ln r_i$ does not influence the volatility measure $\widehat{\sigma}_{it}$, as standard deviation is unaffected by a constant. Similarly, the long-run exchange rate does not affect the interaction term $\widehat{\sigma}_{it}\widetilde{p}_{it}$, as $\widehat{\sigma}_{it}$ is not affected by a constant and $\widetilde{p}_{it} = \ln P_{it} - \overline{\ln P_i}$ removes any effect of $\ln r_i$ by de-meaning.

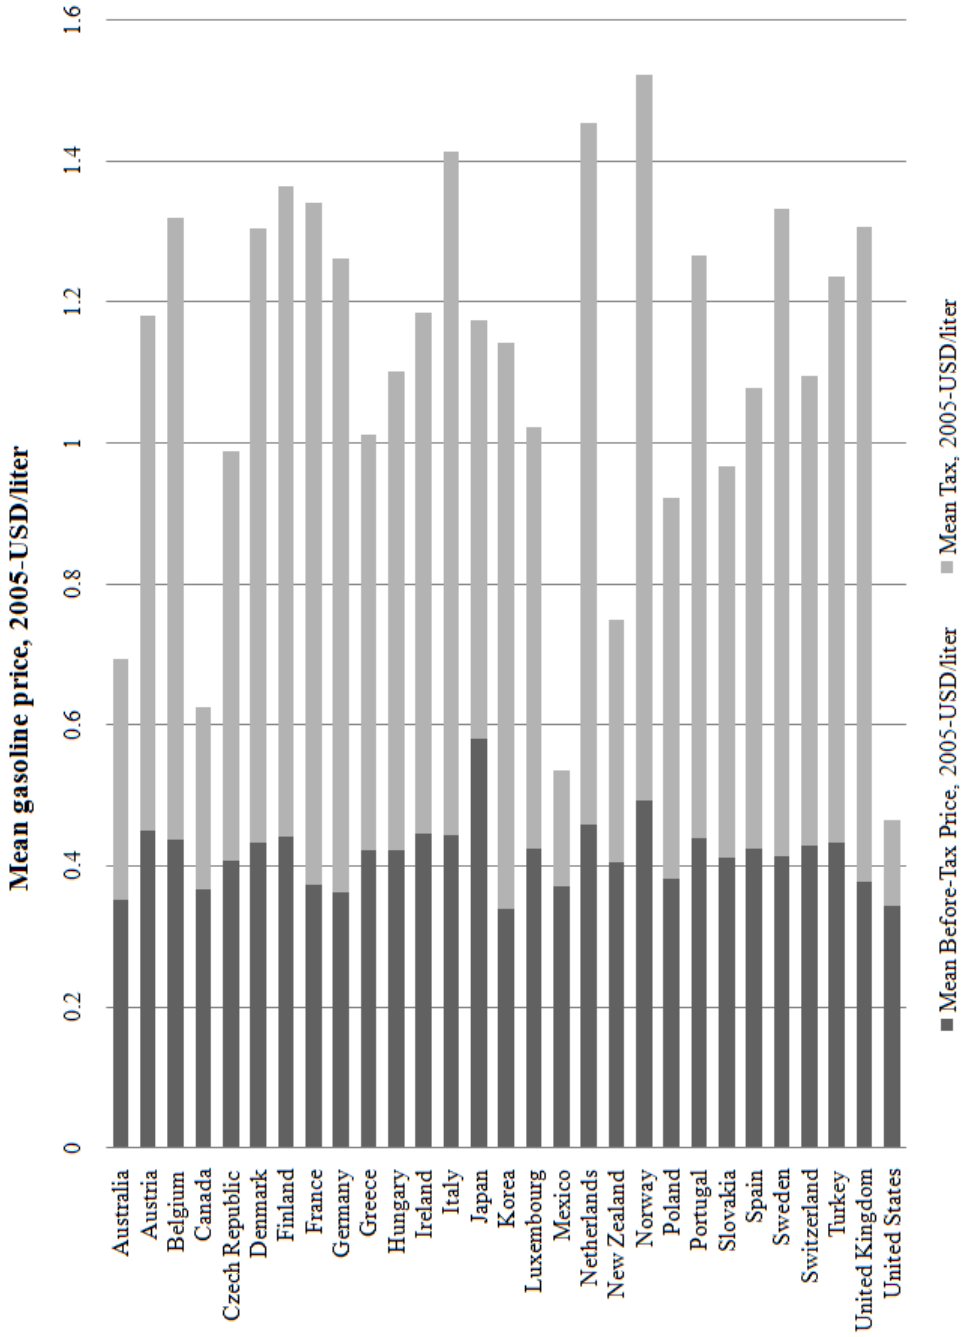
There is, of course, a trade-off when choosing which approach to take toward currency. The chief disadvantage of the common-currency approach is that my USD-denominated variables incorporate exchange rate volatility that may not be reflected in local perceptions of prices and income. On the other hand, since oil is generally imported from abroad, exchange rate fluctuations will in fact be passed along into the gasoline price. The USD-denominated price therefore actually removes some of the exchange-rate driven fluctuation in the locally-perceived price, providing a measure that focuses on price changes driven by the world crude oil price rather than by exchange-rate fluctuations. Depending on how consumers think about oil prices and price expectations, USD-denominated prices may in fact be more appropriate in this sense. More to the point, if converting to USD affected results by introducing exchange-rate noise into the price and income measures, then price and income elasticities should be of lower magnitude when estimated in a common currency than in local currencies—and this is not the case. Overall, therefore, the common-currency approach is both clear and appropriate.

3.3.4 A Brief Overview of the Data

Gasoline consumption per capita shows a reasonable amount of variation, both across countries and over time. Annual consumption in the sample ranges from a low of 31.7 kg per person (Turkey, 2008) to a high of 1352.4 kg per person (Luxembourg, 1994), with an overall mean of 295.9 kg per person. Over time, per-capita consumption has been slightly decreasing in most countries, and the 1990 average of 37.5 kg per person has fallen to a 2008 average of 34.4 kg. Figure 3.2 plots the path of log gasoline consumption over time for each of the 29 countries in the sample.

Prices, like consumption, have shown considerable variation, with much of the cross-sectional variation driven by tax differences (see Figure 3.1) and much of the variation over time driven by fluctuations in the crude oil price. The overall mean real price in the sample is 1.11 Y2005-USD per liter, and country-averaged prices range from 0.47 Y2005-USD per liter in the United States to 1.52 Y2005-USD per liter in Norway. The volatility measure, too, has both cross-sectional and cross-time variation that should allow us to identify its effect. The four-period rolling standard deviation of price, as defined in (3.3), has an average over all observations of 0.048 Y2005-USD per liter, with a comparatively-high standard deviation of 0.037. Country-specific mean volatilities

Figure 3.1: Mean gasoline prices, by tax- and non-tax components



range from a low of 0.035 for Mexico to a high of 0.063 for the United States, and annual average volatility ranges from low of 0.022 in 1998 to a high of 0.14 in 2008. The time paths of the gasoline price and the volatility measure are plotted in Figures 3.3 and 3.4, respectively.

Income per capita, finally, has generally been increasing, and this upward trend is depicted in Figure 3.6.

Figure 3.2: Gasoline consumption per capita, by country

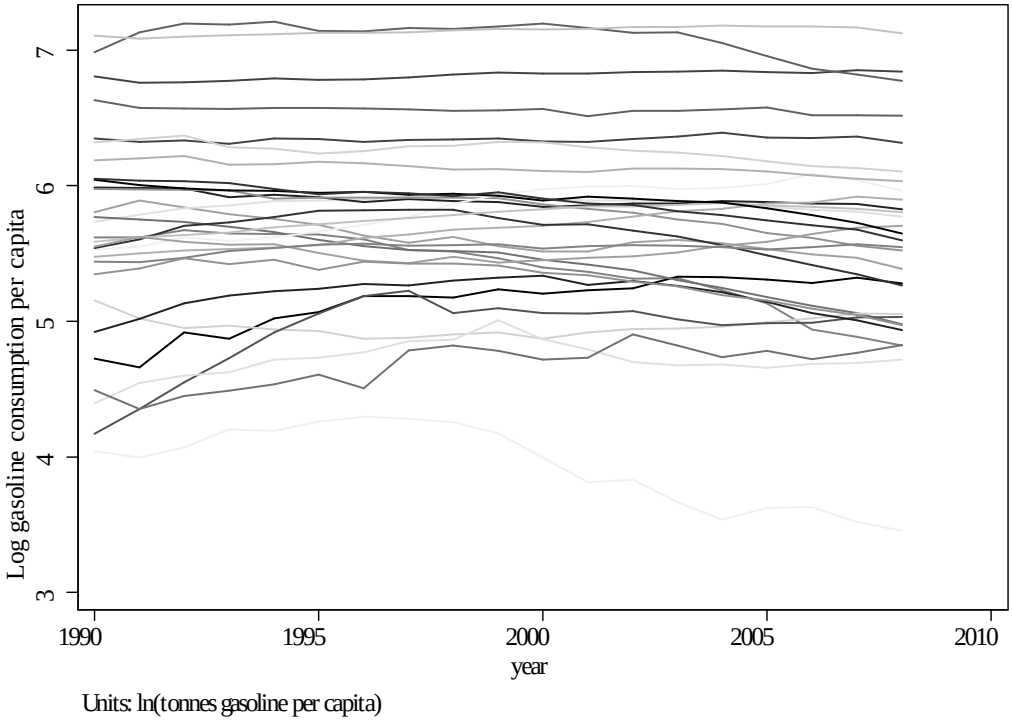


Figure 3.3: Real log gasoline prices, by country

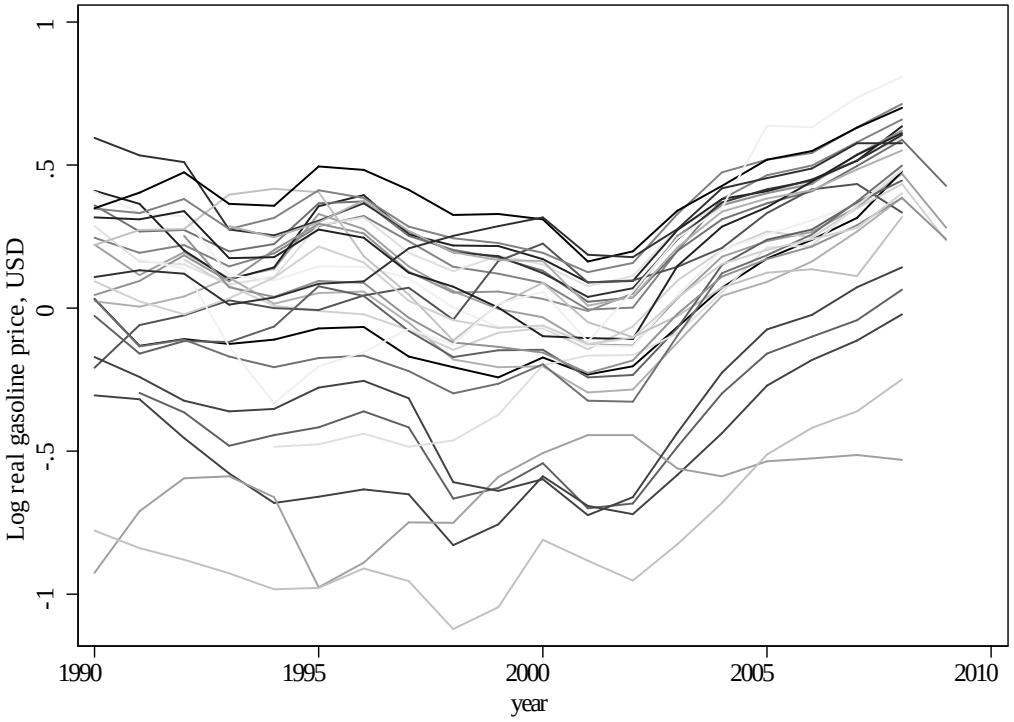
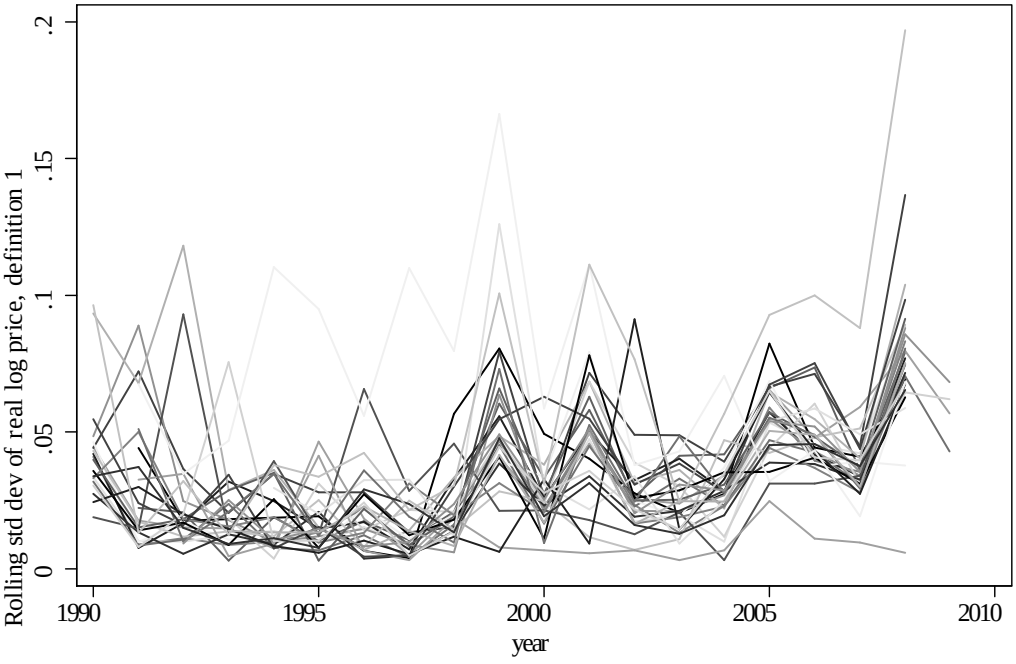


Figure 3.4: Gasoline price volatility, by country



Volatility measure: rolling standard deviation of log real (Y2005-USD) gasoline price. Omitting 1990 outliers for Czech Republic and Slovkia

Figure 3.5: Interaction between de-meaned log real price and price volatility measure (USD), by country

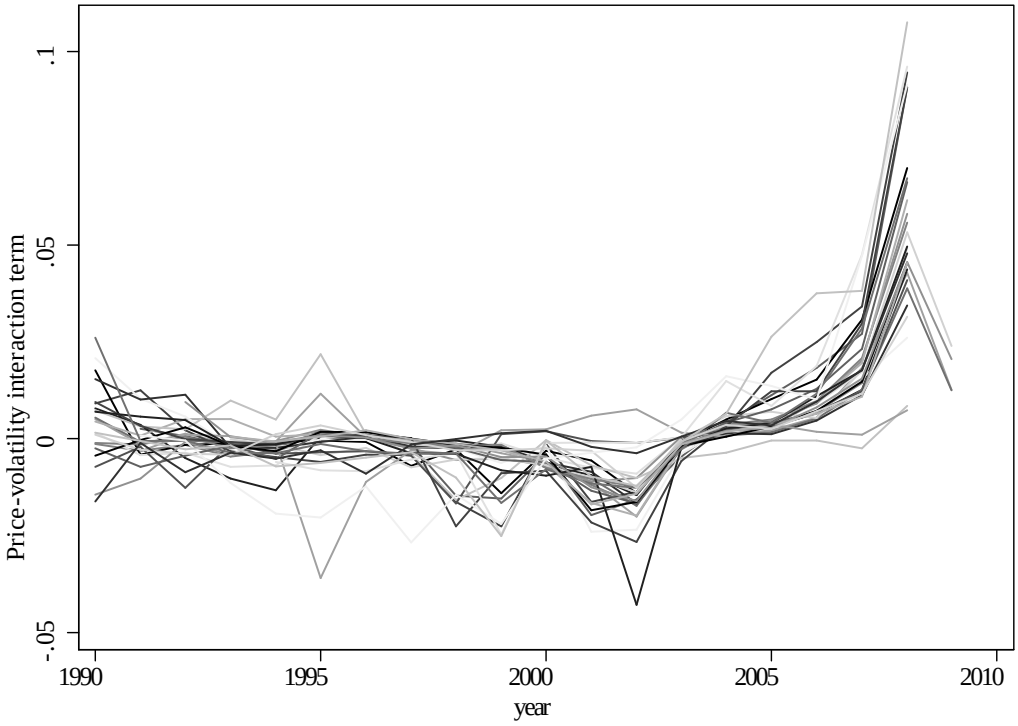
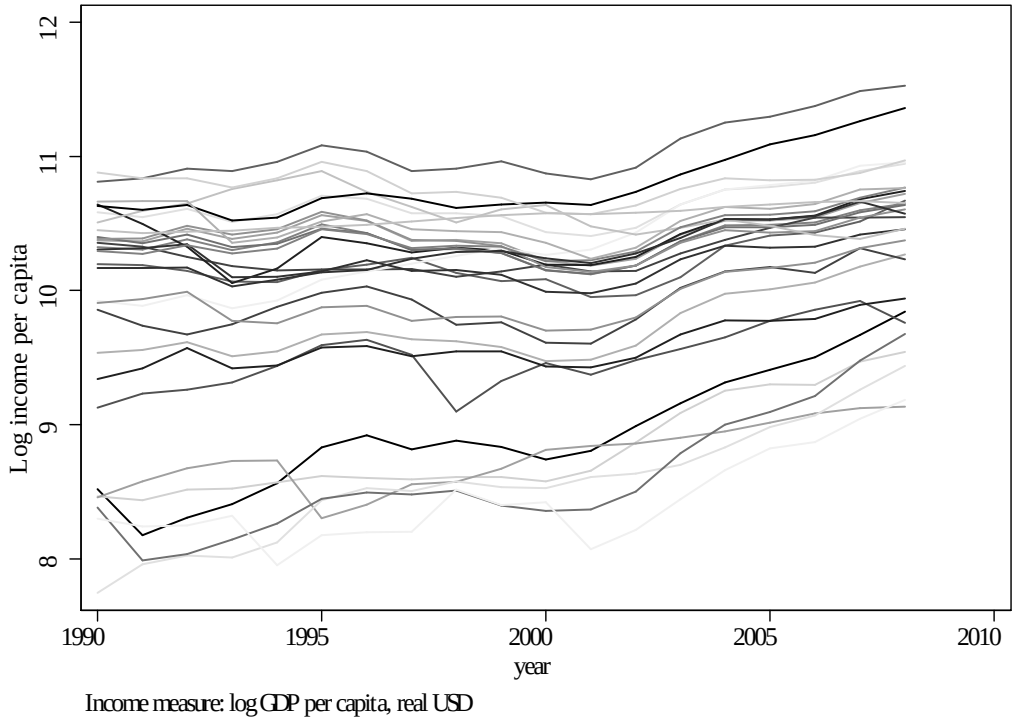


Figure 3.6: Real income per capita, by country



3.4 Models and Estimation

I consider several strategies to examine gasoline demand for the effects implied by a rational habits model. These strategies vary by model specification as well as by estimation method. To start, I estimate a partial adjustment (ADL) model using standard least-squares panel methods. Next I estimate the same model using Arellano & Bover (1995) and Blundell and Bond's (1998) system GMM estimator. Finally I turn to an error-correction version of the model, which I estimate under a range of coefficient-homogeneity restrictions using Pesaran and Smith's (1995) mean groups method; Pesaran, Shin, and Smith's (1999) pooled mean group method; and a dynamic fixed effects method. Each of these ways of specifying and estimating the model has its advantages and trade-offs, but, as I will discuss in Section 3.5, the least-squares and GMM methods' weaknesses are troublesome in this application, and I prefer the pooled mean group estimator of the ECM.

3.4.1 Least Squares

The model laid out in Section 3.1 was formulated as a partial adjustment model—specifically, as an ADL(1,**0**) model:

$$g_{it} = \lambda g_{i,t-1} + \delta_1 y_{it} + \delta_2 p_{it} + \delta_3 \widehat{\sigma}_{it} + \delta_4 \left(p_{it} - \overline{p}_i\right) \widehat{\sigma}_{it} + \mu_i + \varepsilon_{it} \tag{3.4}$$

where g is log per-capita gasoline consumption, y is log per-capita income, p is the log real gasoline price, $\widehat{\sigma}$ is a measure of volatility in the log real gasoline price, and μ_i is a time-invariant country-specific effect. I shall also consider a version of this model that includes a common time trend,

$$g_{it} = \lambda g_{i,t-1} + \delta_1 y_{it} + \delta_2 p_{it} + \delta_3 \widehat{\sigma}_{it} + \delta_4 \left(p_{it} - \overline{p}_i\right) \widehat{\sigma}_{it} + \delta_5 t + \mu_i + \varepsilon_{it} \tag{3.5}$$

and a version that includes country-specific time trends,

$$g_{it} = \lambda g_{i,t-1} + \delta_1 y_{it} + \delta_2 p_{it} + \delta_3 \widehat{\sigma}_{it} + \delta_4 \left(p_{it} - \overline{p}_i\right) \widehat{\sigma}_{it} + \delta_{5i} t + \mu_i + \varepsilon_{it} \tag{3.6}$$

Short-run elasticities in these models are given by the δ coefficients, and long-run elasticities can be calculated as $\frac{\delta}{1-\lambda}$.

On each of these models I first employ a within-groups (fixed-effects) estimator, eliminating the fixed effects μ_i by de-meaning each variable by its country-specific average. A within-groups estimator makes more sense in this situation than a GLS/random-effects estimator because, even

putting aside the issue of the lagged endogenous variable, the country-specific effects are likely to be correlated with the other regressors. Indeed, a Hausman test soundly rejects the equivalence of fixed- and random-effects estimates of a static version of (3.4); see Table 3.3. (The same non-equivalence holds for the dynamic version of the model.) Since the random-effects estimator of the static model is consistent only if μ_i is in fact random in relation to the exogenous variables, and the fixed-effects estimator is consistent either way, rejecting the equivalence of the resulting estimates confirms that μ_i is not random.

Table 3.3: Random- vs. fixed-effects estimator, static model, USD

		(1) Static Random effects	(2) Static Fixed effects	(3) Dynamic Random effects	(4) Dynamic Fixed effects
g_{t-1}	λ			0.957**	0.906**
				(0.0113)	(0.0175)
				[0.000]	[0.000]
y	δ_1	0.523**	0.372**	0.0284**	0.0428**
		(0.0329)	(0.0368)	(0.0107)	(0.0156)
		[0.000]	[0.000]	[0.008]	[0.006]
p	δ_2	-0.574**	-0.419**	-0.0717**	-0.0894**
		(0.0508)	(0.0515)	(0.0177)	(0.0207)
		[0.000]	[0.000]	[0.000]	[0.000]
$\widehat{\sigma}$	δ_3	0.0418	-0.0280	-0.118+	-0.105
		(0.182)	(0.172)	(0.0714)	(0.0729)
		[0.818]	[0.871]	[0.098]	[0.151]
$\widehat{\sigma}\widetilde{p}$	δ_4	-2.348**	-2.197**	-0.132	-0.236
		(0.479)	(0.453)	(0.172)	(0.175)
		[0.000]	[0.000]	[0.444]	[0.179]
R^2		0.8558	0.8564	0.9949	0.9947
Hausman test[a]		77.85 [0.000]		19.07 [0.002]	

Standard errors in parentheses; P-values in brackets.
** p<0.01, * p<0.05, + p<0.1

[a] Under H_0 that random- and fixed-effects estimates equivalent, $\chi^2(4)$ for static model, $\chi^2(4)$ for dynamic.

Addressing Price Endogeneity

The simple within-groups estimator is still afflicted by two problems. The first of these is the potential endogeneity of prices: unless the gasoline supply schedule is flat, any positive demand shock will drive up prices, and vice versa. The gasoline supply schedule is unlikely to be flat when

'individuals' are national aggregates rather than single households, and so causality runs *from* consumption *to* prices as well as from prices to consumption. Ignoring this endogeneity may lead us to underestimate consumers' responsiveness to price changes.

Fortunately, a strength of the current approach is that it allows us to address the endogeneity of prices by instrumenting for them using outside variables. Two obvious instrument candidates are the tax level and the crude oil price. Both should be highly relevant, as they are major determinants of the local gasoline price. Both should also be exogenous, insofar as an individual country's demand does not drive its tax level or the world crude oil price. (Exogeneity may, of course, be violated if demand-driven political pressure affects the tax level or if an individual country's gasoline consumption is great enough to affect the world oil market.) I use tax instruments and crude-price instruments, therefore, to estimate 2SLS within-groups versions of (3.4) through (3.6).

The use of log prices slightly complicates the specification of these instruments. In order to create a good first-stage fit for log prices, note that the real price level P_{it} can be decomposed into the tax level, Tax_{it}, and the before-tax price level, $BeforeTax_{it}$:

$$P_{it} = Tax_{it} + BeforeTax_{it} \text{ , or}$$

$$P_{it} = \left(1 + \frac{Tax_{it}}{BeforeTax_{it}}\right) BeforeTax_{it}$$

$$p_{it} = \ln P_{it} = \ln\left(1 + \frac{Tax_{it}}{BeforeTax_{it}}\right) + \ln BeforeTax_{it} \tag{3.7}$$

Assuming the crude oil price is linearly related to the before-tax gasoline price, the log real crude oil price is a sensible instrument for the second term in (3.7). In the first term, the appearance of the before-tax price means the tax component of log price is still endogenous—assuming, of course, that taxes are levied predominantly as level amounts rather than as percentages. To eliminate the endogeneity introduced by the before-tax price, I define the tax instrument as

$$TaxIV_{it} = \ln\left(1 + \frac{Tax_{it}}{\widehat{BeforeTax}_{it}}\right)$$

where $\widehat{BeforeTax}_{it}$ is predicted using the coefficients estimated by a fixed-effects regression of $BeforeTax_{it}$ on the crude oil price:

$$BeforeTax_{it} = \beta_0 + \beta_1 crude_{(i)t} + \mu_i + e_{it}$$

Note that when the model is estimated using local-currency data, the real crude oil price ($crude_{it}$)

differs across countries; when a common currency is used, the real crude oil price ($crude_t$) is the same for all.

As instruments for the price-volatility interaction term, finally, I construct interactions between the volatility measure $\widehat{\sigma}_{it}$ and a de-meaned version of each of the price instruments.

The resulting estimates of models (3.4) through (3.6) are reported in Table 3.11 (USD) and Table 3.21 (local currencies). All least-squares estimates are computed in Stata using Schaffer's (2010) *xtivreg2*.

An Aside on the Potential Reverse Causality of Price Volatility

At this point it is worth considering the potential for the causality between price volatility and price elasticity to run in reverse. In a closed market, if a country with an inherently low elasticity experiences a supply shock, its prices will have to undergo relatively large changes in order to re-balance supply and demand. In this sense, price elasticity could actually be driving price volatility, rather than the reverse.

Fortunately, by and large the integration of the world oil market rescues the original interpretation of causality. Oil is a globally-traded commodity: following a supply shock (or a global demand shock), it is the world crude oil price that must change in order to re-balance world supply and demand. In response to such a shock, therefore, the change in a country's gasoline price should not depend much on the country's individual price elasticity.

Of course, this only applies to shocks that are in fact global: local supply shocks, for example refinery disruptions, may feed into gasoline price changes that vary with local price elasticity. Refining shocks, however, do not seem to be the major source of gasoline price variation, at least at the annual level; Kilian (2010), for example, finds that on a 12-month horizon, only 11% of variability in the U.S. gasoline price can be attributed to refining shocks. Reverse causality, therefore, is likely to explain at most a small portion of the correlation between price volatility and elasticity.

Lagged Endogenous Variable and the Small-T Bias

The second problem afflicting the within-groups estimator of the ADL model is the small-T bias that arises because lagged gasoline consumption is included as a regressor. To see why this occurs, observe that, by definition, lagged gasoline consumption is positively correlated with the fixed effect μ_i and the idiosyncratic error term $\varepsilon_{i,t-1}$:

$$g_{i,t-1} = \lambda g_{i,t-2} + \delta_1 y_{i,t-1} + \delta_2 p_{i,t-1} + \delta_3 \widehat{\sigma}_{i,t-1} + \delta_4 \left(p_{i,t-1} - \overline{p}_i\right) \widehat{\sigma}_{i,t-1} + \mu_i + \varepsilon_{i,t-1}$$

The within-groups transformation removes μ_i from the model, but introduces all years' idiosyncratic error into the modified error term. Writing out the within-transformation of model (3.4), it becomes clear that transformed lagged consumption is correlated with the transformed error term:

$$\begin{aligned} g_{it} - \overline{g}_i &= \underbrace{\lambda\left(g_{i,t-1} - \frac{1}{T}\sum_{t=1}^{T} g_{i,t-1}\right)}_{\widetilde{g}_{i,t-1}} + \delta_1\left(y_{it} - \overline{y}_i\right) + \delta_2\left(p_{it} - \overline{p}_i\right) + \delta_3\left(\widehat{\sigma}_{it} - \overline{\widehat{\sigma}}_i\right) \\ &\quad + \delta_4\left(\left(p_{it} - \overline{p}_i\right)\widehat{\sigma}_{it} - \frac{1}{T}\sum_{t=1}^{T}\left(p_{it} - \overline{p}_i\right)\widehat{\sigma}_{it}\right) + \underbrace{\left(\varepsilon_{it} - \frac{1}{T}\sum_{t=1}^{T}\varepsilon_{it}\right)}_{\widetilde{\varepsilon}_{it}} \end{aligned}$$

$$\begin{aligned} g_{it} - \overline{g}_i &= \lambda\left(g_{i,t-1} - \frac{1}{T}g_{i0} - \frac{1}{T}g_{i1} - ... - \frac{1}{T}g_{it} - ... - \frac{1}{T}g_{i,T-1}\right) + \delta_1\left(y_{it} - \overline{y}_i\right) + \delta_2\left(p_{it} - \overline{p}_i\right) \\ &\quad + \delta_3\left(\widehat{\sigma}_{it} - \overline{\widehat{\sigma}}_i\right) + \delta_4\left(\left(p_{it} - \overline{p}_i\right)\widehat{\sigma}_{it} - \frac{1}{T}\sum_{t=1}^{T}\left(p_{it} - \overline{p}_i\right)\widehat{\sigma}_{it}\right) \\ &\quad + \left(\varepsilon_{it} - \frac{1}{T}\varepsilon_{i1} - \frac{1}{T}\varepsilon_{i2} - ... - \frac{1}{T}\varepsilon_{it} - ... - \frac{1}{T}\varepsilon_{iT}\right) \end{aligned}$$

There will be $T-1$ pairwise correlations between the terms comprising $\widetilde{g}_{i,t-1}$ and $\widetilde{\varepsilon}_{it}$. $T-2$ of them are positive and of order $\frac{1}{T^2}$ (by definition, $-\frac{1}{T}g_{i1}$ is correlated with $-\frac{1}{T}\varepsilon_{i1}$, and so on), and one of them is negative and of order $\frac{1}{T}$ ($g_{i,t-1}$ is correlated with $-\frac{1}{T}\varepsilon_{it}$). As Bond (2002) notes, the negative bias of order $\frac{1}{T}$ will be larger in magnitude than the sum of the positive biases, and altogether the within-groups estimator of λ will be biased downward.

Although the within-groups estimator will be consistent as $T \to \infty$, the size of the bias may still be significant for a panel of length $T = 15$ to 20. For an autoregressive model without outside regressors (i.e. equation (3.4) with all $\delta = 0$), Nickell (1981) derives the small-T bias as

$$\operatorname*{plim}_{N\to\infty}\left(\widehat{\lambda} - \lambda\right) = \left\{\frac{2\lambda}{1-\lambda^2} - \left[\frac{1}{\frac{1+\lambda}{T-1}\left(1 - \frac{1}{T}\frac{1-\lambda^T}{1-\lambda}\right)}\right]\right\}^{-1} \tag{3.8}$$

In the same paper Nickell derives a more opaque expression for the small-T bias in a model that includes additional regressors. He shows that the downward bias of $\widehat{\lambda}$ should be even greater in this case, and so (3.8) may be used to establish the minimum size of the bias. Given a true λ near 0.7, as suggested by the estimates in Table 3.11 (column 3), and setting T equal to its average of 18.6, (3.8) suggests the bias in $\widehat{\lambda}$ is about -0.10. Since the size of the bias increases in λ, and the true λ is suspected to be greater than the estimate $\widehat{\lambda} \approx 0.7$, the minimum small-sample bias of $\widehat{\lambda}$ may be even larger than 0.1. In other words, the bias is unlikely to be trivial. We can, however,

put a bound on it: as Bond (2002) notes, the positive correlation between the lagged endogenous variable and the fixed effect μ_i means that the OLS estimator of λ in the untransformed model will be biased *upward*. In Tables 3.12 and 3.22, therefore, I report time- and country-pooled 2SLS estimates for models (3.4) through (3.6). The coefficient on lagged consumption—around 0.9 if country-specific time trends are included and just below 1 otherwise—represents an upper bound on λ. This should quiet any fears of a unit root for λ, which would have troubling implications for the long-run interpretation of the model.

Finally, note that spurious regression—a concern that is at the forefront of much of the recent literature—is not a problem in this model. Spurious regression may indeed arise in models where the dependent variable and its regressors are nonstationary. In such a case, a regression may attribute the variables' similar but unrelated evolution over time to a relationship amongst the variables. In fact, as I'll show in section (3.4.3), both my dependent variable g and several of the explanatory variables are nonstationary in levels. Because these $I(1)$ variables are cointegrated, however (see Section (3.4.3)), the relationships estimated amongst these variables are not spurious; and as Pesaran and Shin (1999) and Bentzen and Engsted (2001) demonstrate, the ADL model remains valid.

3.4.2 System GMM

To address the bias generated by the inclusion of lagged endogenous variables in models such as (3.4), the dynamic panel literature has come up with several alternative estimators. The first of these was Anderson and Hsiao's (1981) estimator, which in this case calls for differencing the model and then using the second lag of the *level* of consumption as an instrument for the included endogenous variable, which has been transformed into the first lag of *differenced* consumption. This instrument should be both relevant and exogenous, as $g_{i,t-2}$ is correlated with $\Delta g_{i,t-1} = g_{i,t-1} - g_{i,t-2}$ by definition but uncorrelated (provided the model is properly specified) with the contemporary error term $\Delta\varepsilon_{it} = \varepsilon_{it} - \varepsilon_{i,t-1}$. Subsequent papers (e.g., Holtz-Eakin, Newey and Rosen 1988 and Arellano and Bond 1991) pointed out that all lags of the endogenous variable beyond the first should be independent of the contemporary differenced error term, and that this independence could be exploited using GMM.

Even more moment conditions were suggested by Ahn and Schmidt (1995) and Arellano and Bover (1995)—nonlinear moment conditions in the case of Ahn and Schmidt, and "levels" moment conditions in the case of Arellano and Bover. These latter "levels" moment conditions are particularly useful in cases where the coefficient on the lagged endogenous variable is near 1, which is exactly the case at hand. In such a situation, Blundell and Bond (1998) point out, the included

endogenous variable's *level* provides little information about its future *evolution*, and so lags of the endogenous variable are weak instruments. By contrast, lagged differences of the included endogenous variable *will* be informative about its future level, and this relationship can be used to construct a set of moment conditions for the non-differenced model. Combining these "levels" moment conditions with Arellano and Bond's (1991) "difference" moment conditions yields the "system GMM" estimator. Blundell, Bond, and Windmeijer (2000) show using Monte Carlo simulations that system GMM is more precise and less biased than difference GMM when λ is high and the model includes outside regressors.

I therefore construct a system GMM estimator for (3.4) based on these difference and level moment conditions, as well as on additional "IV-style" moment conditions based upon the outside instruments for the gasoline price and price-volatility interaction term:

$$\textbf{Difference moment conditions: } E\left[g_{i,t-k}\Delta\varepsilon_{it}\right] = 0,\ k = 2, ..., T \tag{3.9}$$

$$\textbf{Levels moment conditions: } E\left[\left(\mu_i + \varepsilon_{it}\right)\Delta g_{i,t-1}\right] = 0 \tag{3.10}$$

$$\textbf{IV-style moment conditions: } E\left[\Delta w_{it}\Delta\varepsilon_{it}\right] = 0\ \forall\ t \text{ and } E\left[w_{it}\varepsilon_{it}\right] = 0\ \forall\ t \tag{3.11}$$

where w represents each of the price instruments discussed in Section 3.4.1. The validity of the difference moment conditions requires that the ε_{it} are not serially correlated; the levels moment conditions further require that the relationship between g_{it} and the fixed effect μ_i is constant over time; and the IV-style moment conditions require that each w fulfills the standard conditions for 2SLS instruments.

Estimates based on this system GMM estimator, computed in Stata using Roodman's (2011) *xtabond2*, are reported in Tables 3.13 (USD) and 3.24 (local currencies). The two-step version of the system GMM estimator is known to produce downward-biased standard errors (Arellano and Bond 1991, Blundell and Bond 1998), so I apply Windmeijer's (2005) correction. To address the problem of overfitting, which arises in difference and system GMM when moment conditions proliferate (Roodman 2006, 2007), I also report a version in which the difference moment conditions are restricted to $k = 2, 3, 4$ and a version based on a 'collapsed' instrument matrix. The meaning and implications of collapsing the instrument matrix are discussed in Roodman (2007).

Discussion of the system GMM results is postponed to Section 3.5. For now, note that addressing the endogeneity of lagged consumption has cost us considerable precision. Few of the coefficients in the system GMM estimates differ significantly from 0, even those coefficients that the least-squares estimates found to be highly significant. Despite the bias inherent in estimates that do not

specifically address the endogeneity of $g_{i,t-1}$, therefore, these methods are more informative—and probably preferable, as the nature of the bias is known.

3.4.3 Heterogeneous Coefficients and the Error-Correction Model

Neither the least-squares nor the GMM methods discussed above allow for cross-country heterogeneity in the coefficients. If coefficients indeed vary by country and regressors are autocorrelated, then, as Pesaran and Smith (1995) prove, estimators that assume homogeneity will be biased and inconsistent. In particular, fixed-effects models such as (3.4), which constrain all coefficients except the country-fixed effects to be the same across countries, will produce downward-biased estimates for the coefficients on the outside regressors $(\delta_1, ..., \delta_4)$. The fixed-effects model will also yield a biased estimate of the coefficient on the lagged endogenous variable (λ), the direction of which will be positive if the autocorrelations of the other regressors are positive. These biases do not go away as the sample size and sample period increase, and their practical effect will be to exaggerate the difference between short- and long-run responses. The intuition for this practical effect is simple: if the coefficient on lagged consumption is biased upward, the speed of adjustment will appear slower, and this will magnify artificially-small short-run responses into artificially-large long-run responses.

If regressors are I(1), moreover, these biases are potentially severe. As the autocorrelation coefficient of an outside regressor approaches a unit root, Pesaran and Smith (1995) show, the estimator for the coefficient on the lagged endogenous variable converges in probability to 1, while the estimators for the coefficients on the outside regressors converge to 0. Many of my regressors are in fact I(1), as I'll show shortly, and so the consequences of coefficient heterogeneity are potentially grave.

To address the problem of coefficient heterogeneity, Pesaran, Shin and Smith (1999) suggest a pooled-mean group (PMG) estimator that allows short-run responses to vary across individuals. Since the PMG estimator and its alternatives are derived and discussed in the literature in terms of an error-correction model, I turn to an ECM model of gasoline demand—the unrestricted version of which is equivalent to an ADL(1,**1**). To demonstrate the relationship between the former ADL(1,**0**) specification and the ECM and the ECM's equivalence to an ADL(1,**1**), I first re-write the existing model (3.4) to allow coefficients to vary by country:

$$g_{it} = \lambda_i g_{i,t-1} + \delta_{1i} y_{it} + \delta_{2i} p_{it} + \delta_{3i} \widehat{\sigma}_{it} + \delta_{4i} \left(\widehat{\sigma}\widetilde{p}\right)_{it} + \mu_i + \varepsilon_{it} \tag{3.12}$$

(For notational simplicity, define $\widetilde{p}_{it} = p_{it} - \overline{p}_i$.) Next, subtracting $g_{i,t-1}$ from each side of (3.12),

then adding and subtracting $\delta_{1i}y_{i,t-1}$, $\delta_{2i}p_{i,t-1}$, $\delta_{3i}\widehat{\sigma}_{i,t-1}$, and $\delta_{4i}\widehat{\sigma}\widetilde{p}_{i,t-1}$ from the right-hand side, yields

$$\begin{aligned}\Delta g_{it} &= (\lambda_i - 1)\, g_{i,t-1} + \delta_{1i}y_{i,t-1} + \delta_{2i}p_{i,t-1} + \delta_{3i}\widehat{\sigma}_{i,t-1} + \delta_{4i}\left(\widehat{\sigma}\widetilde{p}\right)_{i,t-1} \\ &\quad +\delta_{1i}\Delta y_{it} + \delta_{2i}\Delta p_{it} + \delta_{3i}\Delta\widehat{\sigma}_{it} + \delta_{4i}\Delta\left(\widehat{\sigma}\widetilde{p}\right)_{it} + \mu_i + \varepsilon_{it}\end{aligned}$$

which can be factored into

$$\begin{aligned}\Delta g_{it} &= \underbrace{-(1-\lambda_{it})}_{\phi_i}\left[g_{i,t-1} - \underbrace{\frac{\delta_{1i}}{1-\lambda_i}}_{\theta_{1i}}y_{i,t-1} - \underbrace{\frac{\delta_{2i}}{1-\lambda_i}}_{\theta_{2i}}p_{i,t-1} - \underbrace{\frac{\delta_{3i}}{1-\lambda_i}}_{\theta_{3i}}\widehat{\sigma}_{i,t-1} - \underbrace{\frac{\delta_{4i}}{1-\lambda_i}}_{\theta_{4i}}\left(\widehat{\sigma}\widetilde{p}\right)_{i,t-1}\right] \\ &\quad +\delta_{1i}\Delta y_{it} + \delta_{2i}\Delta p_{it} + \delta_{3i}\Delta\widehat{\sigma}_{it} + \delta_{4i}\Delta\left(\widehat{\sigma}\widetilde{p}\right)_{it} + \mu_i + \varepsilon_{it}\end{aligned} \tag{3.13}$$

This is an ECM with no coefficient-homogeneity restrictions, but it can be factored just as easily into ADL(1,**1**) form:

$$\begin{aligned}g_{it} &= (\phi_i + 1)\, g_{i,t-1} + \delta_{1i}y_{it} - (\phi_i\theta_{1i} + \delta_{1i})\, y_{i,t-1} + \delta_{2i}p_{it} - (\phi_i\theta_{2i} + \delta_{2i})\, p_{i,t-1} \\ &\quad +\delta_{3i}\widehat{\sigma}_{it} - (\phi_i\theta_{3i} + \delta_{3i})\,\widehat{\sigma}_{i,t-1} + \delta_{4i}\left(\widehat{\sigma}\widetilde{p}\right)_{it} - (\phi_i\theta_{4i} + \delta_{4i})\left(\widehat{\sigma}\widetilde{p}\right)_{i,t-1} + \mu_i + \varepsilon_{it}\end{aligned} \tag{3.14}$$

Estimates based on the ECM differ from those based on ADL(1,**0**) by virtue of slightly different underlying models, therefore, as well as by different coefficient-homogeneity restrictions.

In the unrestricted ECM given by (3.13), the new coefficient ϕ_i is the error-correction term. It should be negative, and its magnitude is the speed of adjustment, or the portion of long-run adjustment that takes place during the first period after a change in one of the regressors. The new coefficients θ_{1i} through θ_{4i} are long-run elasticities, and as before the coefficients δ_{1i} through δ_{4i} are short-run elasticities.

Although Pesaran, Shin and Smith's (1999) estimator allows regressors to be I(0) or I(1), it does require a cointegrating relationship among I(1) variables—that is, it requires a stable long-run relationship between the dependent and explanatory variables. Before checking for cointegration, I check the order of integration of each of my variables.

The classical tests for a unit root in a *single* time series are the Dickey-Fuller and Augmented Dickey-Fuller, which test for the stationarity of an AR(1) and an AR(p) process, respectively. The Dickey-Fuller tests $H_0 : \rho - 1 = 0$ (nonstationarity) vs. $H_a : |(\rho - 1)| < 0$ (stationarity) in the

transformed model

$$\Delta x_t = (\rho - 1)\, x_{t-1} + (\alpha_0 + \alpha_1 t) + \varepsilon_t$$

whereas the Augmented Dickey-Fuller uses the same null and alternative hypotheses for the transformed model

$$\Delta x_t = (\rho - 1)\, x_{t-1} + \sum_{i=j}^{p-1} \beta_i \Delta x_{t-j} + (\alpha_0 + \alpha_1 t) + \varepsilon_t$$

Alternately, the Phillips-Perron (1988) test makes it possible to test for nonstationarity in an AR(p) process without knowing p. In this test, corrections for serial correlation in ε_t allow $\widehat{\rho}$ estimated for an AR(1) to be used to to test $\rho = 1$ for any AR(p).

When x is a panel variable, the cross-section dimension introduces several complications, which have prompted the development of a variety of panel unit root tests.[4] One complication is whether to consider a single autocorrelation coefficient for the panel as a whole or to consider each individual i as a separate time series with its own coefficient ρ_i. Tests taking the former route include Levin, Lin, and Chu (2002) and Harris and Tzavalis (1999). Tests taking the latter route include Im, Pesaran and Shin (2003), Maddala and Wu (1999) and Choi (2001), all of which involve performing a separate test for each individual and aggregating the results. Im, Pesaran and Shin (2003) perform this aggregation by averaging the *test statistics* for individual (A)DF tests; Maddala and Wu (1999) and Choi (2001) perform the aggregation over the *p-values* associated with (A)DF or other individual unit root tests. These tests based on p-values are known collectively as Fisher-type tests, and weigh a null hypothesis that all panels are *nonstationary* against the alternative that one or more panels are stationary.

To explore the order of integration of my variables, I use a Fisher-type test based on individual Phillips-Perron tests, with and without a trend. The choice of the Phillips-Perron test protects my results from the ADF's sensitivity to the choice of lag length p. The choice of a Fisher-type test has the advantages of high power compared to tests based on ADF test statistics (Choi 2001) and, crucially, the flexibility to deal with series whose length varies across individuals. Results of these tests, performed using Stata's (2009) *xtunitroot*, are reported in Tables 3.4 and 3.5. Where the test does not reject a variable's nonstationarity in levels, I repeat the test on first differences. In each instance the second test does reject nonstationarity, and so for the USD-denominated variables I am able to conclude that volatility is I(0) with or without trend and the remaining variables are I(1) with or without trend. For the common-currency variables, results are somewhat different: g and p are I(1) with or without trend; y is I(1) without trend and I(0) with; and $\widehat{\sigma}$ and $\widehat{\sigma}\widetilde{\rho}$ are I(0) with or without trend.

[4] For a thorough overview of these tests, see Baltagi and Kao (2000).

Table 3.4: Fisher-style Phillips-Perron unit root tests, currency in USD

Without trends

	Level		First Difference		
Variable	Test Statistic*	P-Value	Test Statistic*	P-Value	Order of Integration
g	68.0062	0.1732	364.1834	0.0000	$I(1)$
p	17.8696	1.0000	148.7426	0.0000	$I(1)$
y	11.8702	1.0000	220.1072	0.0000	$I(1)$
$\widehat{\sigma}$	190.6238	0.0000			$I(0)$
$\widehat{\sigma}\widetilde{p}$	19.1043	1.0000	199.4135	0.0000	$I(1)$

With time trends

	Level		First Difference		
Variable	Test Statistic*	P-Value	Test Statistic*	P-Value	Order of Integration
g	71.6704	0.1071	507.0169	0.0000	$I(1)$
p	14.6156	1.0000	136.2582	0.0000	$I(1)$
y	24.2946	1.0000	156.7488	0.0000	$I(1)$
$\widehat{\sigma}$	195.1116	0.0000			$I(0)$
$\widehat{\sigma}\widetilde{p}$	12.0569	1.0000	211.9994	0.0000	$I(1)$

*Inverse $\chi 2$ with 58 degrees of freedom.
Phillips-Perron tests conducted using 3 lags for the Newey-West standard errors.

Since a regression involving nonstationary variables may be spurious, and since Pesaran, Shin and Smith's (1999) estimator requires cointegration of any I(1) variables, I turn now to checking for cointegration. One approach to cointegration testing involves estimating a static model and then checking the residuals for nonstationarity. This was in fact the first approach to cointegration testing, introduced by Engle and Granger (1987) for single time series and later adapted to panels by Kao (1999) and others. Pedroni (2004) expanded this two-step, residual-based approach to allow for cross-sectional heterogeneity of coefficients. Meanwhile, a second approach to cointegration testing has been to check whether the error-correction term in an ECM—e.g., ϕ_i in (3.13), or α_i in (3.15)—is zero. Banerjee, Dolado and Mestre (1998) develop this approach for single time series, and Westerlund (2007) adapts it to panels, basing his tests on a model of the form

$$\Delta y_{it} = \alpha_i y_{i,t-1} + \lambda_i^{'}\mathbf{x}_{i,t-1} + \sum_{j=1}^{q_i} \alpha_{ij}\Delta y_{i,t-j} + \sum_{j=0}^{q_i} \gamma_{ij}^{'}\Delta \mathbf{x}_{i,t-j} + \mu_i + \pi_i t + \varepsilon_{it} \tag{3.15}$$

where $\mathbf{x}$ and y comprise the variables to be examined for possible cointegration. Westerlund's approach has the advantages of high power and flexibility: not only does it allow cross-sectional heterogeneity in coefficients, but, unlike the two-step residual-based approaches, it does not impose a common factor restriction on the relationship between short- and long-run adjustment. For these reasons I choose Westerlund's test to check for cointegration of my I(1) variables.

Westerlund's (2007) test actually has several varieties. The "group mean" variety (G) allows

Table 3.5: Fisher-style Phillips-Perron unit root tests, local currencies

Without trends

	Level		First Difference		
Variable	Test Statistic*	P-Value	Test Statistic*	P-Value	Order of Integration
g	68.0062	0.1732	364.1834	0.0000	$I(1)$
p	72.0861	0.1010	343.6998	0.0000	$I(1)$
y	21.1883	1.0000	412.3704	0.0000	$I(1)$
$\widehat{\sigma}$	313.8033	0.0000			$I(0)$
$\widehat{\sigma}\widetilde{p}$	183.9754	0.0000			$I(0)$

Fisher-style unit root tests on models **With time trends**

	Level		First Difference		
Variable	Test Statistic*	P-Value	Test Statistic*	P-Value	Order of Integration
g	71.6704	0.1071	507.0169	0.0000	$I(1)$
p	81.6818	0.0219	316.9540	0.0000	$I(1)$
y	165.6083	0.0000			$I(0)$
$\widehat{\sigma}$	450.9815	0.0000			$I(0)$
$\widehat{\sigma}\widetilde{p}$	229.8273	0.0000			$I(0)$

*Inverse $\chi 2$ with 58 degrees of freedom.
Phillips-Perron tests conducted using 3 lags for the Newey-West standard errors.

the error-correction term to vary across individuals, and tests against the alternative hypothesis that the error-correction term is smaller than zero for at least one individual; the "panel" variety (P), by contrast, constrains the error-correction term to be the same across individuals, testing against the alternative hypothesis that this single term is smaller than zero. In both cases, the null hypothesis is that the error-correction term is zero, and that the variables in question are not cointegrated. Given my small sample size, the extra precision afforded by constraining the error-correction term to be constant across countries appears to be important, and so I focus on this variety. For each of these varieties, Westerlund suggests two different test statistics: a t-statistic version, denoted with the subscript τ, which is simply the estimated error-correction term divided by its standard error (or, for the group mean version, the average of these quotients); and a less intuitive version, denoted with the subscript α, that controls for the panel length $T_{(i)}$. I report both of these test statistics.

Tests are performed in Stata with Persyn and Westerlund's (2008) *xtwest*, using standard errors that have been bootstrapped to allow cross-sectional dependence (see Westerlund 2007 and Persyn and Westerlund 2008). Results are reported in Tables 3.6 and 3.7. Note that some countries' series are insufficiently long to be used in certain versions of the tests; these omissions are listed in the tables. Using the panel-variety tests, I can clearly reject the hypothesis of non-cointegration of the I(1) variables in the local-currency version of the model. In the common-currency version, the panel-variety tests reject the hypothesis of non-cointegration provided that the lag q_i in (3.15),

chosen for each country using the AIC, is allowed to be as high as 2.[5,6]

Table 3.6: Cointegration tests, USD

Variables	Lag length q_i	Omitted Countries	Test	Value	P-Value	Bootstrapped[a] p-value
g, y, p, $\widehat{\sigma}\widetilde{p}$	Chosen by AIC from 0 and 1	*Czech Republic*	G_τ	-0.036	1.000	1.000
			G_α	-0.203	1.000	1.000
			P_τ	-1.431	1.000	0.720
			P_α	-0.427	0.999	0.730
g, y, p, $\widehat{\sigma}\widetilde{p}$	Chosen by AIC from 0-2	*Australia, Czech Republic, Hungary, Mexico, Poland, Spain, Sweden, Turkey*	G_τ	-1.991	0.103	0.390
			G_α	-0.057	1.000	0.630
			P_τ	-0.144	1.000	0.010
			P_α	-0.051	0.999	0.010
g, y, p, $\widehat{\sigma}\widetilde{p}$	Chosen by AIC from 0 and 1	*Australia, Czech Republic, Hungary, Mexico, Poland, Spain, Sweden, Turkey*	G_τ	-0.071	1.000	0.990
			G_α	-0.091	1.000	1.000
			P_τ	-1.708	1.000	0.580
			P_α	-0.566	0.996	0.560

[a]100 draws

Table 3.7: Cointegration tests, local currencies

Variables	Lag length q_i	Omitted Countries	Test	Value	P-Value	Bootstrapped[a] p-value
g, y, p	Chosen by AIC from 0 and 1	none	G_τ	-1.460	0.344	0.340
			G_α	-3.833	0.975	0.320
			P_τ	-9.074	0.001	0.000
			P_α	-3.501	0.136	0.050
g, y, p	Chosen by AIC from 0-2	none	G_τ	-1.540	0.207	
			G_α	-3.244	0.995	
			P_τ	-9.074	0.001	
			P_α	-3.501	0.136	
g, y, p	Chosen by AIC from 0-3	Australia, Hungary, Poland, Spain, Sweden, Turkey	G_τ	-1.311	0.627	0.420
			G_α	-0.961	1.000	0.260
			P_τ	-6.129	0.109	0.000
			P_α	-2.520	0.492	0.000

[a]100 draws

Given this evidence that a long-run relationship amongst the variables of my model does exist, I move on to estimating the ECM. Pesaran, Shin and Smith's (1999) pooled mean group estimator uses maximum likelihood, based on the assumption that the ε_{it} are independent and normal, to

[5]Allowing q_i to be as high as 2, unfortunately, forces me to drop several countries from the sample. It does not appear to be the omission of these countries that is responsible for the rejection of non-cointegration, however, as running the test for $q_i \leq 1$ with the same countries omitted does not lead to rejection.

[6]Persyn and Westerlund (2008) observe that the Westerlund tests are sensitive to their specification in "small" T situations. As my T is only two thirds that of Persyn and Westerlund's example, and I am attempting to test for cointegration amongst a larger number of variables, I suspect that my inability to reject non-cointegration in certain specifications (e.g., $q_i \leq 1$) arises from this sensitivity and relatively low power, and not from a lack of a long-run relationship amongst the variables. The strong rejection of non-cointegration amongst the local-currency variables, moreover, suggests that a cointegrating relationship does exist.

Table 3.8: Restrictions on the error-correction model

Model	**Restrictions**
Pooled mean-group (PMG)	Long-run parameters homogenous: $\theta_{1i} = \theta_1, \theta_{2i} = \theta_2, \theta_{3i} = \theta_3, \theta_{4i} = \theta_4$
Mean-Group (MG)	None
Dynamic Fixed-Effect (DFE)	All parameters homogenous: $\theta_{1i} = \theta_1, \theta_{2i} = \theta_2, \theta_{3i} = \theta_3, \theta_{4i} = \theta_4$; $\delta_{1i} = \delta_1, \delta_{2i} = \delta_2, \delta_{3i} = \delta_3, \delta_{4i} = \delta_4; \phi_i = \phi$

estimate

$$\begin{aligned} \Delta y_{it} &= \phi_i \left[g_{i,t-1} - \theta_1 y_{i,t-1} - \theta_2 p_{i,t-1} - \theta_3 \widehat{\sigma}_{i,t-1} - \theta_4 \left(\widehat{\sigma} \widetilde{p} \right)_{i,t-1} \right] \\ &\quad + \delta_{1i} \Delta y_{it} + \delta_{2i} \Delta p_{it} + \delta_{3i} \Delta \widehat{\sigma}_{it} + \delta_{4i} \Delta \left(\widehat{\sigma} \widetilde{p} \right)_{it} + \mu_i + \varepsilon_{it} \end{aligned} \tag{3.16}$$

Note that (3.16) is simply the general ECM from (3.13) with the long-run coefficients θ_{1i}, θ_{2i}, θ_{3i}, and θ_{4i} constrained to be equal across individuals.

The resulting pooled mean group estimates are reported in Tables 3.14, 3.15, and 3.25. Table 3.14 reports results for all countries; Table 3.15 reports results when the four countries with the worst fit from the ADL are omitted. The short run elasticities ($\widehat{\delta}_1$ through $\widehat{\delta}_4$) reported in these tables are the *means* of the country-specific short-run elasticities; individual short-run elasticities are reported in Appendix 3.D.

To examine the value of the assumption that the long-run coefficients are homogenous (which drastically reduces the number of parameters to be estimated), I also report mean group estimates of the ECM model. The mean group estimator simply estimates (3.13) separately for each country and forms an average of the resulting coefficient estimates. To examine the value and cost of restricting both the long- *and* the short-run coefficients to be homogenous, as the ADL inherently did, I also estimate a dynamic fixed-effects version of the ECM. The restrictions imposed by these various estimators are summarized in Table 3.8. Results are reported alongside the PMG results.

All estimates of the ECM—pooled mean groups, mean groups, and dynamic fixed-effects—are computed in Stata using Blackburne and Frank's (2007b) *xtpmg*.

3.5 Results and Discussion

Results, particularly with respect to the habits-related parameters, vary considerably across these model specifications and estimation methods. Although each approach does have its advantages (summarized in Table 3.9), not all the advantages are equal; and after weighing each approach's

advantages and disadvantages, my preference is for the pooled mean group estimates of the error-correction model.

Table 3.9: Advantages of various estimation approaches

Advantage	Fixed-effects 2SLS with instrumenting for price	GMM	Pooled mean group estimator of ECM
Corrects for endogeneity of prices	✓		
Addresses endogeneity of g_{t-1}		✓	
Addresses coefficient heterogeneity			✓
Precision	✓		✓

3.5.1 Least-Squares Estimates of the ADL Model

The unique advantage of the fixed-effects estimator of the ADL model is that it allows me to use outside instruments to address the endogeneity of the price variables. In practice, however, it is not clear that instrumenting for prices is very valuable. Tests for endogeneity[7] do, by and large, reject that prices are exogenous, and may or may not, depending on the specification of the trend, reject that price and the price-volatility interaction are jointly exogenous. The instruments used to address the apparent endogeneities, moreover, do appear to be relevant, as underidentification is rejected[8] no matter how the trend is specified. The problems arise when it comes to the instruments' exogeneity: overidentifying restrictions tests strongly reject that all the instruments are exogenous. The endogeneity, it becomes clear when the regressions are re-run using subsets of the instruments, lies in the instruments based on crude oil prices. Unfortunately, omitting the crude-oil instruments and basing price identification solely on the tax instruments yields estimates that are nonsensical and uninformative, with most coefficients' p-values around 0.9.

In theory, it's possible that an instrument that is not *perfectly* exogenous may nonetheless remove some of the instrumented variable's endogeneity, and therefore retain some value. Alas, this is not such a case. Any endogeneity in the price should bias its coefficient upwards towards 0, and yet the price coefficients in the non-instrumented ADL models (Table 3.10) are more negative than their instrumented counterparts. The implied long-run price elasticity from the instrumented version of the individual-trend model, moreover, is of a smaller magnitude than the long-run elasticity estimated using PMG on the error-correction model. Though price endogeneity is an issue, it appears we are better off ignoring it than addressing it—or, rather, better off merely

[7] That is, difference-in-Sargan/Hansen tests comparing a regression in which the variable(s) in question are treated as endogenous to a regression in which they are assumed to be exogenous.

[8] Kleibergen-Paap (2006) tests are used to test for underidentification.

acknowledging it than addressing it using the instruments at hand.

The actual fixed-effects 2SLS estimates, it should be noted, are nonsensical unless country-specific trends are included. With no trend or a homogenous trend, prices have no statistically-significant effect; and with no trend, even income does not have a statistically-significant effect. Moreover, *none* of the fixed-effects 2SLS results—common or local currency, with or without any type of trend—finds the rational-habits variables, $\widehat{\sigma}$ and $\widehat{\sigma}\widetilde{p}$, to have a statistically-significant effect. Although this might ordinarily be taken as a sign that $\widehat{\sigma}$ and $\widehat{\sigma}\widetilde{p}$ are not relevant to gasoline demand, the weakness of the price and income effects suggests that the estimated effects of $\widehat{\sigma}$ and $\widehat{\sigma}\widetilde{p}$ should not be trusted, either. Overall, I discount the findings of the fixed-effects 2SLS estimator of the ADL.

Table 3.10: Within-groups (fixed-effects) estimates, USD

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}	λ	0.911**	0.888**	0.727**
		(0.0249)	(0.0230)	(0.0372)
		[0.000]	[0.000]	[0.000]
y	δ_1	0.0404+	0.105**	0.183**
		(0.0217)	(0.0233)	(0.0318)
		[0.064]	[0.000]	[0.000]
p	δ_2	-0.117**	-0.135**	-0.198**
		(0.0312)	(0.0299)	(0.0331)
		[0.000]	[0.000]	[0.000]
$\widehat{\sigma}$	δ_3	-0.178+	-0.0470	-0.0592
		(0.0931)	(0.0904)	(0.0792)
		[0.056]	[0.604]	[0.455]
$\widehat{\sigma}\widetilde{p}$	δ_4	0.260	0.263	-0.00520
		(0.288)	(0.275)	(0.244)
		[0.367]	[0.340]	[0.983]
t			-0.00426**	
			(0.000524)	
			[0.000]	
R^2		0.891	0.904	0.938

Robust standard errors in parentheses;
P-values in brackets.
** p<0.01, * p<0.05, + p<0.1

3.5.2 System GMM Estimates of the ADL Model

Table 3.11: Within-groups (fixed effects) estimates with 2SLS instrumenting for price, USD

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}	λ	0.941**	0.922**	0.681**
		(0.0288)	(0.0274)	(0.0417)
		[0.000]	[0.000]	[0.000]
y	δ_1	0.0347	0.0770**	0.196**
		(0.0251)	(0.0265)	(0.0368)
		[0.167]	[0.004]	[0.000]
p	δ_2	0.0479	0.0361	-0.168**
		(0.101)	(0.0928)	(0.0520)
		[0.635]	[0.698]	[0.001]
$\widehat{\sigma}$	δ_3	0.283	0.342	0.0906
		(0.318)	(0.295)	(0.142)
		[0.374]	[0.248]	[0.525]
$\widehat{\sigma}\widetilde{p}$	δ_4	-2.313	-2.062	-0.868
		(1.572)	(1.442)	(0.680)
		[0.142]	[0.153]	[0.202]
t			-0.00364**	[country-specific
			(0.000550)	trends]
			[0.000]	
R^2		0.871	0.886	0.934
Implied long-run elasticities				
y	$\frac{\delta_1}{1-\lambda}$	0.585	0.983*	0.614**
		(0.462)	(0.422)	(0.125)
		[0.207]	[0.020]	[0.000]
p	$\frac{\delta_2}{1-\lambda}$	0 .807	0.460	-0.528**
		(1.762)	(1.208)	(0.165)
		[0.647]	[0.703]	[0.001]
Tests				
Endogeneity: price[a]		4.081	0.297	7.396
		[0.0434]	[0.5859]	[0.0065]
Endogeneity: price & interaction[b]		4.230	4.477	7.675
		[0.1206]	[0.1066]	[0.0215]
Underidentification[c]		8.159	8.341	19.689
		[0.0428]	[0.0395]	[0.0002]
Exogeneity[d]		22.997	7.175	7.919
		[0.000]	[0.0277]	[0.0191]

Robust standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

[a] Difference-in-Sargan/Hansen $\chi^2(1)$ under H_0.

[b] Difference-in-Sargan/Hansen $\chi^2(2)$ under H_0.

[c] Kleibergen-Paap (2006) rk statistic; $\chi^2(3)$ under H_0.

[d] Overidentifying restrictions test (Hansen J statistic); $\chi^2(2)$ under H_0.

Table 3.12: Pooled estimates with 2SLS instrumenting for prices, USD

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}	λ	0.997**	0.991**	0.924**
		(0.00957)	(0.00910)	(0.0143)
		[0.000]	[0.000]	[0.000]
y	δ_1	-0.00216	0.00318	0.0569**
		(0.00879)	(0.00834)	(0.0122)
		[0.806]	[0.703]	[0.000]
p	δ_2	-0.0292*	-0.0341*	-0.0806**
		(0.0140)	(0.0134)	(0.0181)
		[0.037]	[0.011]	[0.000]
$\widehat{\sigma}$	δ_3	-0.0575	-0.0309	-0.0247
		(0.0957)	(0.0953)	(0.0889)
		[0.548]	[0.746]	[0.781]
$\widehat{\sigma}\widetilde{p}$	δ_4	-0.260	0.0969	0.201
		(0.209)	(0.224)	(0.228)
		[0.213]	[0.666]	[0.379]
t			-0.00223**	[country-specific
			(0.000503)	trends]
			[0.000]	
R^2		0.996	0.996	0.997

Robust standard errors in parentheses;
p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

The unique advantage of the system GMM estimator is that it addresses the endogeneity of lagged consumption. That endogeneity should, as previously discussed, bias the coefficient on lagged consumption upward, and therefore make adjustment seem slower than it truly is. As expected, the system GMM estimates of λ are higher than the least-squares estimates. Troublingly, in fact, the GMM estimates of λ are in most cases higher than 1. But although these estimates have been pushed upward by the removal of a downward bias, they have also been pushed upward by the GMM estimator's own shortcomings. One of these shortcomings is that the system GMM estimator, like the least-squares estimator, imposes homogeneity of coefficients across countries. Given that the other regressors are positively autocorrelated, this will lead to an upward bias on $\widehat{\lambda}$ (Pesaran and Smith 1995). Another shortcoming is that, unlike the least-squares estimator, the system GMM estimator does not allow me to identify a model with country-specific trends,[9] which in the least-squares case reduced $\widehat{\lambda}$ considerably. The GMM estimator merely trades one bias for another, therefore, and it does so at an incredibly high cost of precision: in nearly all variants of the system GMM estimator, in fact, $\widehat{\lambda}$ is the *only* coefficient that's statistically significantly different from 0. Overall, the system GMM estimates are useless in this application.

3.5.3 Error-Correction Model

Turning to the error correction model, note that pooled mean group, mean group, and dynamic fixed effects estimators are all affected by the endogeneity of lagged consumption, just as the least squares estimators of the ADL were. It would be ideal to correct for the resulting biases, and this may become possible as relevant bias corrections are developed in the literature. At the moment, however, it does not seem that appropriate corrections have been derived. Bias corrections for the (*not*-pooled) mean group estimator are worked out in Pesaran and Zhao (1999); these corrections perform well when $\lambda < 0.8$ but not at all when $\lambda \geq 0.8$, which appears likely in the case of gasoline demand. Moreover, the assumption of homogeneity of the long-run coefficients appears to be important in this application (as will be discussed shortly), and so a mean group estimator that allows heterogeneity of the long-run coefficients is inappropriate. Recent work by Choi, Mark and Sul (2010) offers bias corrections that constrain the long-run coefficients to be homogenous, but unfortunately they also constrain the *short-run* coefficients and speed of adjustment to be homogenous. For now, I will have to accept estimates of the ECM (and ADL) as they are.

Despite the endogeneity of g_{t-1}, pooled mean group estimates of the ECM yield reasonable

[9] Country-specific trends require the estimation of an extra 28 parameters, which system GMM cannot handle given my sample size.

Table 3.13: System GMM, USD

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				No trend	With trend	No trend	With trend	No trend	With trend
				Collapsed	Collapsed	Reduced	Reduced	No outside	No outside
		No trend	With trend	IV matrix	IV matrix	IV count	IV count	IVs	IVs
g_{t-1}	λ	1.031**	1.010**	1.085**	1.055**	1.042**	1.032**	0.995**	0.966**
		(0.0331)	(0.0240)	(0.0532)	(0.0693)	(0.0402)	(0.0326)	(0.0359)	(0.0337)
		[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
y	δ_1	-0.0394+	-0.0177	-0.0451	-0.0116	-0.0445	-0.0158	-0.000187	0.0354
		(0.0220)	(0.0168)	(0.0432)	(0.0325)	(0.0262)	(0.0190)	(0.0302)	(0.0327)
		[0.084]	[0.301]	[0.306]	[0.725]	[0.101]	[0.414]	[0.995]	[0.288]
p	δ_2	0.0233	-0.00381	0.0740	0.0435	0.0345	0.0139	-0.0758	-0.108**
		(0.0383)	(0.0282)	(0.0624)	(0.0633)	(0.0478)	(0.0328)	(0.0485)	(0.0380)
		[0.548]	[0.894]	[0.246]	[0.498]	[0.476]	[0.675]	[0.129]	[0.008]
$\widehat{\sigma}$	δ_3	-0.145	-0.0776	0.0843	0.160	-0.0232	0.0671	-0.133	-0.0168
		(0.123)	(0.121)	(0.204)	(0.225)	(0.146)	(0.155)	(0.0890)	(0.108)
		[0.248]	[0.527]	[0.683]	[0.483]	[0.875]	[0.668]	[0.145]	[0.877]
$\widehat{\sigma}\widetilde{p}$	δ_4	-0.201	0.119	-0.824	-0.450	-0.421	-0.103	0.237	0.563
		(0.296)	(0.263)	(0.592)	(0.651)	(0.318)	(0.341)	(0.346)	(0.340)
		[0.503]	[0.653]	[0.175]	[0.496]	[0.196]	[0.765]	[0.498]	[0.109]
t			-0.00215*		-0.00280		-0.00276*		-0.00367*
			(0.00103)		(0.00175)		(0.00106)		(0.00142)
			[0.046]		[0.120]		[0.014]		[0.015]
Sargan		377.2	368.4	97.30	86.25	267.3	253.3	427.4	391.3
Sargan df		169	168	17	16	64	63	165	164
Sargan P		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Hansen		25.39	23.91	22.65	23.51	25.89	22.26	26.10	25.42
Hansen df		169	168	17	16	64	63	165	164
Hansen P		1	1	0.161	0.101	1.000	1.000	1	1
AR(2) test		-0.522	-0.495	-0.596	-0.556	-0.550	-0.528	-0.283	-0.304
AR(2) P		0.602	0.621	0.551	0.578	0.582	0.597	0.778	0.761

Standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

results. As expected, the speed of adjustment implied by PMG estimates based on all countries (with or without trend) is slightly faster than the speed of adjustment implied by 2SLS on the ADL. This is because the PMG approach is not susceptible to the false assumption of coefficient homogeneity that biases the 2SLS estimate of λ upward. Similarly, the PMG estimates of the error-correction model imply slightly faster adjustment than their DFE (dynamic fixed effects) counterparts. (That is, the PMG estimates of the error-correction term ϕ are more negative than the DFE estimates.) Although the PMG estimates imply slightly faster adjustment, however, the difference is not so great as to beggar belief. The pooled 2SLS estimates of the ADL (Table 3.12) establish lower bounds for the speed of adjustment of 0.3 to 7.6 percent in the first year, and we know that these estimates are biased toward 0 in *two* ways. In that light, the speeds of adjustment estimated by PMG—5.4 to 33.7 percent of adjustment occurring in the first year—do not seem suspiciously fast.

My preferred specification is, therefore, the pooled mean group estimator of the ECM. The PMG estimator's distinct advantage, as has been mentioned, is that by allowing for heterogeneity in the short-run coefficients, it avoids some of the heterogeneity-induced bias of least squares or system GMM. The PMG estimator does, of course, restrict long-run coefficients to be the same across countries; and another matter to consider is whether this restriction is acceptable. A Hausman test comparing the common-currency PMG and MG estimates fails to reject the hypothesis that the estimates are equivalent for the model without trend ($p = 0.71$) but squarely rejects the equivalence of the models with trend ($p = 0.00$), so it is not clear whether the restrictions are entirely true. The efficiency gains offered by the long-run restrictions if they are true, however, are huge; and without them, the MG estimator has trouble even establishing statistically-significant price effects. This is unsurprising, as the MG estimator is based on individual regressions in which time series of 15 to 19 observations are used to estimate the short- and long-run effects of four different variables, plus intercept and time trend. The relatively short length of this panel makes the restrictions necessary, whether or not the long-run coefficients are purely homogenous. On a positive note, PMG does at least introduce a considerable relaxation of homogeneity in comparison to estimators that assume homogeneity of *all* coefficients; and as Pesaran, Shin and Smith (1999) observe, PMG is not as sensitive to outliers as MG.

Although PMG is not as sensitive to outliers as MG, it is not entirely insensitive to them, either; and so in addition to estimates based on all countries in my data set (Table 3.14), I report estimates omitting the four countries that consistently suffered the worst fit in least-squares estimates—the Czech Republic, Poland, Slovakia, and Turkey (Table 3.15). The latter three coun-

tries are, incidentally, the three poorest in my sample,[10] and the Czech Republic is the sixth-poorest. The first three, moreover, were beginning their political and economic transformations away from communism at the beginning of the sample period in 1990, and so it is not surprising that their gasoline-consumption patterns would be atypical of the broader sample. In response to the removal of these four countries, the PMG estimates do in fact change a bit, with the long-run income, price, and price-volatility interaction elasticities increasing in magnitude. Unless otherwise noted, the following discussion refers to the PMG estimates with these outlier countries removed.

By and large, my long-run coefficients are not out of line with the literature. To facilitate comparisons, I report estimates of an ECM in the form of (3.16) with price volatility and the price-volatility interaction term omitted (Tables 3.16, 3.17, and 3.26). The resulting model contains only the variables standardly included in gasoline demand models. The PMG estimates of this model, omitting the four worst-fitting countries (Table 3.17), suggest a long-run income elasticity of 0.7 to 1. These nearly-unit-elastic estimates fit neatly within the values usually estimated in the literature. Meanwhile, the estimated long-run price elasticities, ranging from -0.6 to -1.3, span the middle to upper end (in magnitude) of estimates in the literature. Although the PMG estimator is more reliable in its long-run than its short-run estimates, the short-run elasticities in Table 3.17 are plausible, with short-run income elasticities somewhat lower than the literature's averages and short-run price elasticities ranging from low (in magnitude) to average.

Returning to the model that includes the rational-habits variables, long-run results with respect to income and prices appear, again, reasonable. In the no-trend and homogenous-trend models, long-run income elasticities of 0.8 and 0.9 closely match the literature. In the same models, long-run "total" price elasticities are again on the high-magnitude end, hovering around unit elastic. Since these high-magnitude price elasticities are also found in the "standard" model without the rational-habits variables, the difference from the literature can be ascribed to my panel data set and the choice of an ECM with heterogeneous short-run coefficients rather than to the inclusion of $\widehat{\sigma}$ and $\widehat{\sigma}\widetilde{p}$. In the model with heterogeneous trends, results look somewhat different but still plausible, with long-run income elasticity on the low end and long-run price elasticity near the literature average.

The parameters of primary interest are, of course, the coefficients on price volatility and the price-volatility interaction term, which allow us to test for the behaviors implied by a rational habits model. The rational habits model implies, first, that price uncertainty should reduce demand for the habit-forming good, and this is indeed borne out by the regressions: trend or no trend, all

[10] As measured by mean real GDP per capita, USD, over the course of the sample.

Table 3.14: Estimates of the error-correction model, all countries, USD

		(1) PMG No trend	(2) PMG Homogenous trend	(3) PMG Heterogenous trends	(4) MG No trend	(5) MG (Heterogeneous) trends	(6) DFE No trend	(7) DFE (Homogenous) trend
t			-0.0473**	-0.00507*		-0.0129*		-0.0481*
			(0.00303)	(0.00228)		(0.00651)		(0.0219)
			[0.000]	[0.026]		[0.048]		[0.028]
Long run								
y	θ_1	0.555**	0.774**	0.388**	0.752*	0.313	0.0812	0.746*
		(0.0306)	(0.0905)	(0.0412)	(0.335)	(0.209)	(0.579)	(0.372)
		[0.000]	[0.000]	[0.000]	[0.025]	[0.134]	[0.888]	[0.045]
p	θ_2	-0.693**	-0.981**	-0.626**	-0.945+	-0.466	-2.046	-1.585*
		(0.0755)	(0.122)	(0.0398)	(0.495)	(0.400)	(1.488)	(0.744)
		[0.000]	[0.000]	[0.000]	[0.056]	[0.244]	[0.169]	[0.033]
$\widehat{\sigma}$	θ_3	-1.249**	-1.464**	-0.881**	-0.653	-0.411	-4.456	-2.551
		(0.249)	(0.353)	(0.100)	(1.067)	(0.678)	(4.013)	(1.727)
		[0.000]	[0.000]	[0.000]	[0.540]	[0.545]	[0.267]	[0.140]
$\widehat{\sigma}\widetilde{p}$	θ_4	0.816	-0.307	-0.0718	6.402	-1.075	22.64	12.63
		(1.205)	(1.872)	(0.520)	(10.63)	(6.308)	(22.87)	(10.01)
		[0.498]	[0.870]	[0.890]	[0.547]	[0.865]	[0.322]	[0.207]
Short run								
EC term	ϕ	-0.142**	-0.133**	-0.337**	-0.300**	-0.604**	-0.0534	-0.0789+
		(0.0530)	(0.0326)	(0.0733)	(0.0934)	(0.0918)	(0.0501)	(0.0415)
		[0.007]	[0.000]	[0.000]	[0.001]	[0.000]	[0.287]	[0.057]
Δy	δ_1	0.184**	0.162**	0.0901*	0.130*	0.0173	0.225**	0.177**
		(0.0472)	(0.0344)	(0.0420)	(0.0543)	(0.0526)	(0.0396)	(0.0317)
		[0.000]	[0.000]	[0.032]	[0.016]	[0.742]	[0.000]	[0.000]
Δp	δ_2	-0.177**	-0.108**	-0.0415	-0.102+	0.0217	-0.131+	-0.0683
		(0.0543)	(0.0389)	(0.0356)	(0.0553)	(0.0556)	(0.0690)	(0.0607)
		[0.001]	[0.006]	[0.245]	[0.065]	[0.697]	[0.058]	[0.261]
$\Delta\widehat{\sigma}$	δ_3	0.0180	0.148*	0.169**	0.0864	0.224*	0.130	0.215*
		(0.0654)	(0.0579)	(0.0640)	(0.124)	(0.108)	(0.109)	(0.0911)
		[0.783]	[0.011]	[0.008]	[0.486]	[0.038]	[0.232]	[0.019]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	δ_4	-0.318	0.151	0.0938	0.275	0.351	-1.258	-0.903
		(0.228)	(0.231)	(0.235)	(1.014)	(0.862)	(0.785)	(0.706)
		[0.163]	[0.513]	[0.690]	[0.786]	[0.684]	[0.109]	[0.200]

Standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

Table 3.15: Estimates of the error-correction model, USD, omitting 4 worst-fit countries*

		(1) PMG No trend	(2) PMG Homo-genous trend	(3) PMG Hetero-genous trends	(4) MG No trend	(5) MG (Hetero-geneous) trends	(6) DFE No trend	(7) DFE (Homo-geneous) trend
t			-0.0491**	-0.00586*		-0.0190**		-0.0525
			(0.00374)	(0.00241)		(0.00587)		(0.0340)
			[0.000]	[0.015]		[0.001]		[0.122]
Long run								
y	θ_1	0.893**	0.782**	0.350**	0.803*	0.401*	0.363	1.042+
		(0.0767)	(0.110)	(0.0516)	(0.385)	(0.202)	(0.625)	(0.564)
		[0.000]	[0.000]	[0.000]	[0.037]	[0.047]	[0.562]	[0.065]
p	θ_2	-1.295**	-0.971**	-0.593**	-1.021+	-0.535	-2.949	-2.233+
		(0.118)	(0.154)	(0.0458)	(0.572)	(0.453)	(2.963)	(1.222)
		[0.000]	[0.000]	[0.000]	[0.074]	[0.237]	[0.320]	[0.068]
$\widehat{\sigma}$	θ_3	-0.731**	-1.904**	-0.942**	-1.090	-0.985+	-1.564	-1.071
		(0.261)	(0.410)	(0.106)	(1.203)	(0.589)	(2.433)	(1.763)
		[0.005]	[0.000]	[0.000]	[0.365]	[0.095]	[0.520]	[0.544]
$\widehat{\sigma}\widetilde{p}$	θ_4	7.879**	-2.120	-0.213	8.537	-0.0767	36.01	21.10
		(1.913)	(2.526)	(0.631)	(12.31)	(7.251)	(41.90)	(14.48)
		[0.000]	[0.401]	[0.736]	[0.488]	[0.992]	[0.390]	[0.145]
Short run								
EC term	ϕ	-0.0542	-0.100**	-0.283**	-0.283**	-0.597**	-0.0416	-0.0617
		(0.0354)	(0.0296)	(0.0767)	(0.0918)	(0.0933)	(0.0600)	(0.0508)
		[0.126]	[0.001]	[0.000]	[0.002]	[0.000]	[0.488]	[0.225]
Δy	δ_1	0.228**	0.182**	0.133**	0.168**	0.0587	0.214**	0.177**
		(0.0488)	(0.0369)	(0.0412)	(0.0585)	(0.0565)	(0.0425)	(0.0327)
		[0.000]	[0.000]	[0.001]	[0.004]	[0.299]	[0.000]	[0.000]
Δp	δ_2	-0.214**	-0.148**	-0.0897**	-0.145*	-0.0402	-0.139*	-0.0838
		(0.0567)	(0.0327)	(0.0297)	(0.0588)	(0.0510)	(0.0643)	(0.0571)
		[0.000]	[0.000]	[0.003]	[0.014]	[0.430]	[0.030]	[0.142]
$\Delta\widehat{\sigma}$	δ_3	-0.154**	0.113*	0.107+	0.0521	0.170	-0.156+	-0.0230
		(0.0506)	(0.0534)	(0.0596)	(0.137)	(0.117)	(0.0806)	(0.0780)
		[0.002]	[0.034]	[0.074]	[0.703]	[0.146]	[0.053]	[0.768]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	δ_4	-0.537	0.374	0.129	0.555	0.807	-1.394*	-1.047**
		(0.344)	(0.299)	(0.273)	(1.002)	(0.897)	(0.607)	(0.389)
		[0.118]	[0.211]	[0.635]	[0.580]	[0.368]	[0.022]	[0.007]

Standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

* *Czech Republic, Poland, Slovakia, and Turkey*

Table 3.16: Estimates of error-correction model with only the standard regressors, USD

		(1) PMG No trend	(2) PMG Homo-genous trend	(3) PMG Hetero-geneous trends	(4) MG No trend	(5) MG (Hetero-geneous) trends	(6) DFE No trend	(7) DFE (Homo-genous) trend
t			-0.0523**	-0.00558**		0.0261		-0.0446*
			(0.00432)	(0.00169)		(0.0465)		(0.0200)
			[0.000]	[0.001]		[0.575]		[0.025]
Long run								
y	θ_1	-0.908**	0.863**	0.647**	0.0733	-0.712	0.0468	0.786**
		(0.140)	(0.117)	(0.0658)	(0.429)	(0.851)	(0.526)	(0.277)
		[0.000]	[0.000]	[0.000]	[0.864]	[0.403]	[0.929]	[0.005]
p	θ_2	0.139	-1.075**	-0.834**	-0.302	0.0632	-1.039*	-0.985**
		(0.129)	(0.0961)	(0.0567)	(0.336)	(0.363)	(0.454)	(0.261)
		[0.282]	[0.000]	[0.000]	[0.369]	[0.862]	[0.022]	[0.000]
Short run								
EC term	ϕ	-0.0505+	-0.116**	-0.286**	-0.288**	-0.480**	-0.0557	-0.0822+
		(0.0282)	(0.0263)	(0.0527)	(0.0649)	(0.0746)	(0.0523)	(0.0435)
		[0.073]	[0.000]	[0.000]	[0.000]	[0.000]	[0.287]	[0.059]
Δy	δ_1	0.272**	0.164**	0.0819*	0.137*	0.0260	0.222**	0.173**
		(0.0365)	(0.0339)	(0.0404)	(0.0547)	(0.0537)	(0.0382)	(0.0340)
		[0.000]	[0.000]	[0.043]	[0.012]	[0.628]	[0.000]	[0.000]
Δp	δ_2	-0.229**	-0.0990*	-0.0353	-0.0759	0.0126	-0.196**	-0.126**
		(0.0442)	(0.0402)	(0.0403)	(0.0501)	(0.0493)	(0.0483)	(0.0424)
		[0.000]	[0.014]	[0.381]	[0.129]	[0.799]	[0.000]	[0.003]

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

Table 3.17: Estimates of error-correction model with only the standard regressors, USD, omitting 4 worst-fitting countries*

		(1) PMG No trend	(2) PMG Homo-genous trend	(3) PMG Hetero-geneous trends	(4) MG No trend	(5) MG (Hetero–geneous) trends	(6) DFE No trend	(7) DFE (Homo-genous) trend
t			-0.0530**	-0.00486**		0.0331		-0.0551
			(0.00501)	(0.00106)		(0.0539)		(0.0393)
			[0.000]	[0.000]		[0.540]		[0.161]
Long run								
y	θ_1	0.710**	0.769**	1.020**	-0.270	-0.929	-0.143	0.936+
		(0.0695)	(0.148)	(0.0871)	(0.410)	(0.983)	(1.097)	(0.538)
		[0.000]	[0.000]	[0.000]	[0.509]	[0.344]	[0.896]	[0.082]
p	θ_2	-0.866**	-1.046**	-0.982**	-0.0222	0.197	-1.283	-1.080*
		(0.0632)	(0.122)	(0.0681)	(0.324)	(0.416)	(1.494)	(0.550)
		[0.000]	[0.000]	[0.000]	[0.945]	[0.635]	[0.390]	[0.050]
Short run								
EC term	ϕ	-0.0645+	-0.0845**	-0.177**	-0.238**	-0.427**	-0.0336	-0.0600
		(0.0387)	(0.0209)	(0.0285)	(0.0616)	(0.0758)	(0.0630)	(0.0540)
		[0.095]	[0.000]	[0.000]	[0.000]	[0.000]	[0.593]	[0.267]
Δy	δ_1	0.243**	0.184**	0.109**	0.168**	0.0586	0.224**	0.173**
		(0.0497)	(0.0346)	(0.0395)	(0.0586)	(0.0589)	(0.0424)	(0.0312)
		[0.000]	[0.000]	[0.006]	[0.004]	[0.320]	[0.000]	[0.000]
Δp	δ_2	-0.237**	-0.136**	-0.107**	-0.124**	-0.0458	-0.200**	-0.133**
		(0.0491)	(0.0298)	(0.0316)	(0.0462)	(0.0470)	(0.0528)	(0.0434)
		[0.000]	[0.000]	[0.001]	[0.008]	[0.329]	[0.000]	[0.002]

Standard errors in parentheses; p-values in brackets.

** p<0.01, * p<0.05, + p<0.1

* *Czech Republic, Poland, Slovakia, and Turkey*

countries or all-but-four, the PMG estimate of the long-run effect of price volatility, θ_3, is negative and statistically significant. Depending on the specification, values of $\widehat{\theta}_3$ range from -0.7 to -1.9, suggesting that a 0.1 increase in $\widehat{\sigma}$ is associated with a 7 to 19% reduction in gasoline consumption.

Evidence on rational habits' second implication, that the magnitude of price elasticity should decline with price uncertainty, is not clear-cut. When a time trend is included in the USD-denominated model, the PMG estimate of the long-run coefficient on the price-volatility interaction term, $\widehat{\theta}_4$, is not statistically significant. Absent the trend, $\widehat{\theta}_4$ is positive and highly significant when the four outlier-countries are omitted but not statistically significant when they are included.[11] Given that these four countries experienced some of the highest price volatilities in the sample at a time when their economies were rapidly expanding, it is easy to imagine that they might skew the effect of volatility on price elasticity. It seems reasonable, therefore, to put greater faith on the estimates of $\widehat{\theta}_4$ that omit these countries.

Recall the possibility that tax changes might contribute to measured price volatility while actually decreasing price uncertainty. To check whether this problem is hiding some of the effect of price volatility on price elasticity, I estimate a version of the model in which before-tax prices are used to construct an alternative measure of price volatility, $\widehat{\sigma}_{BT}$. In Table 3.18 I report results using the before-tax volatility term *a.* both in the interaction term and on its own and *b.* only in the interaction term, with total-price volatility used on its own. Continuing to use total-price volatility on its own seems reasonable, as rational habits imply that it is *overall* price volatility that dampens consumption—that is, high volatility that arises from low tax rates will still dampen consumption. Either way, however, the use of before-tax volatility actually *decreases* the coefficient on the interaction term, so it does not seem as if tax-driven price volatility is masking the effect of "true" price volatility. The models based on total-price volatility therefore remain my preferred specifications.

If, as the version of the preferred model *with* trend suggests, people's responsiveness to prices is immune to the volatility of prices, then θ_2 is a perfectly serviceable measure of price elasticity. On the other hand, if price elasticity does vary with price uncertainty—as the model *without* trend implies—then total long-run price elasticity is not θ_2, but $\theta_2 + \theta_4\widehat{\sigma}_{it}$. This is analogous to the total

[11] In the local-currency models (Table 3.25), as in the USD-denominated models, the coefficient on the interaction term is sensitive to the specification of the trend. With no trend, $\widehat{\theta}_4$ in the local-currency model is not statistically significant; with a homogenous trend, it is very negative; and with heterogenous trends, it is positive. This pattern is, confusingly enough, nearly opposite to that in the USD-denominated models, where $\widehat{\theta}_4$ is only positive and statistically significant if the trend is excluded.

Table 3.18: Estimates of error-correction model with before-tax volatilities, USD

		(1) PMG Homogenous trend	(2) PMG Heterogeneous trends	(3) PMG Homogeneous trend, $\widehat{\sigma}_{BT}$ only in interaction	(4) PMG Heterogeneous trends, $\widehat{\sigma}_{BT}$ only in interaction
t		-0.0660**	-0.00316**	-0.0506**	-0.00346*
		(0.00995)	(0.00122)	(0.00408)	(0.00164)
		[0.000]	[0.009]	[0.000]	[0.035]
Long run					
y	θ_1	0.623**	0.460**	0.707**	0.220**
		(0.193)	(0.0377)	(0.121)	(0.0225)
		[0.001]	[0.000]	[0.000]	[0.000]
p	θ_2	-0.428+	-0.188**	-0.834**	-0.0738**
		(0.260)	(0.0591)	(0.155)	(0.0247)
		[0.100]	[0.001]	[0.000]	[0.003]
$\widehat{\sigma}_{BT}$	θ_3	0.200	-0.298**		
		(0.413)	(0.101)		
		[0.629]	[0.003]		
$\widehat{\sigma}$	θ_3			-1.367**	0.181**
				(0.422)	(0.0380)
				[0.001]	[0.000]
$\widehat{\sigma}_{BT}\widetilde{p}$	θ_4	-5.018*	-3.609**	-1.309	-1.160**
		(2.031)	(0.521)	(1.198)	(0.237)
		[0.013]	[0.000]	[0.274]	[0.000]
Short run					
EC term	ϕ	-0.0567**	-0.275**	-0.102**	-0.317**
		(0.0207)	(0.0429)	(0.0276)	(0.0622)
		[0.006]	[0.000]	[0.000]	[0.000]
Δy	δ_1	0.195**	0.108**	0.187**	0.171**
		(0.0296)	(0.0370)	(0.0320)	(0.0379)
		[0.000]	[0.004]	[0.000]	[0.000]
Δp	δ_2	-0.155**	-0.178**	-0.138**	-0.176**
		(0.0340)	(0.0362)	(0.0390)	(0.0476)
		[0.000]	[0.000]	[0.000]	[0.000]
$\Delta\widehat{\sigma}_{BT}$	δ_3	0.0421	0.0923*		
		(0.0365)	(0.0362)		
		[0.249]	[0.011]		
$\Delta\widehat{\sigma}$	δ_3			0.125*	-0.0185
				(0.0509)	(0.0528)
				[0.014]	[0.726]
$\Delta\left(\widehat{\sigma}_{BT}p\right)$	δ_4	0.0224	0.647**	0.181	0.157
		(0.257)	(0.202)	(0.184)	(0.188)
		[0.931]	[0.001]	[0.324]	[0.404]

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

price elasticities derived for the ADL:

$$\text{Short run price elasticity:} \qquad \frac{\partial g_{it}}{\partial p_{it}} = \delta_2 + \delta_4 \widehat{\sigma}_{it}$$

$$\text{Long run price elasticity:} \qquad \frac{\delta_2 + \delta_4 \widehat{\sigma}_{it}}{1 - \lambda}$$

The model without trend implies, therefore, that the "total" long-run elasticity with respect to price is

$$\widehat{\theta}_2 + \widehat{\theta}_4 \overline{\widehat{\sigma}} = -0.918 \tag{3.17}$$

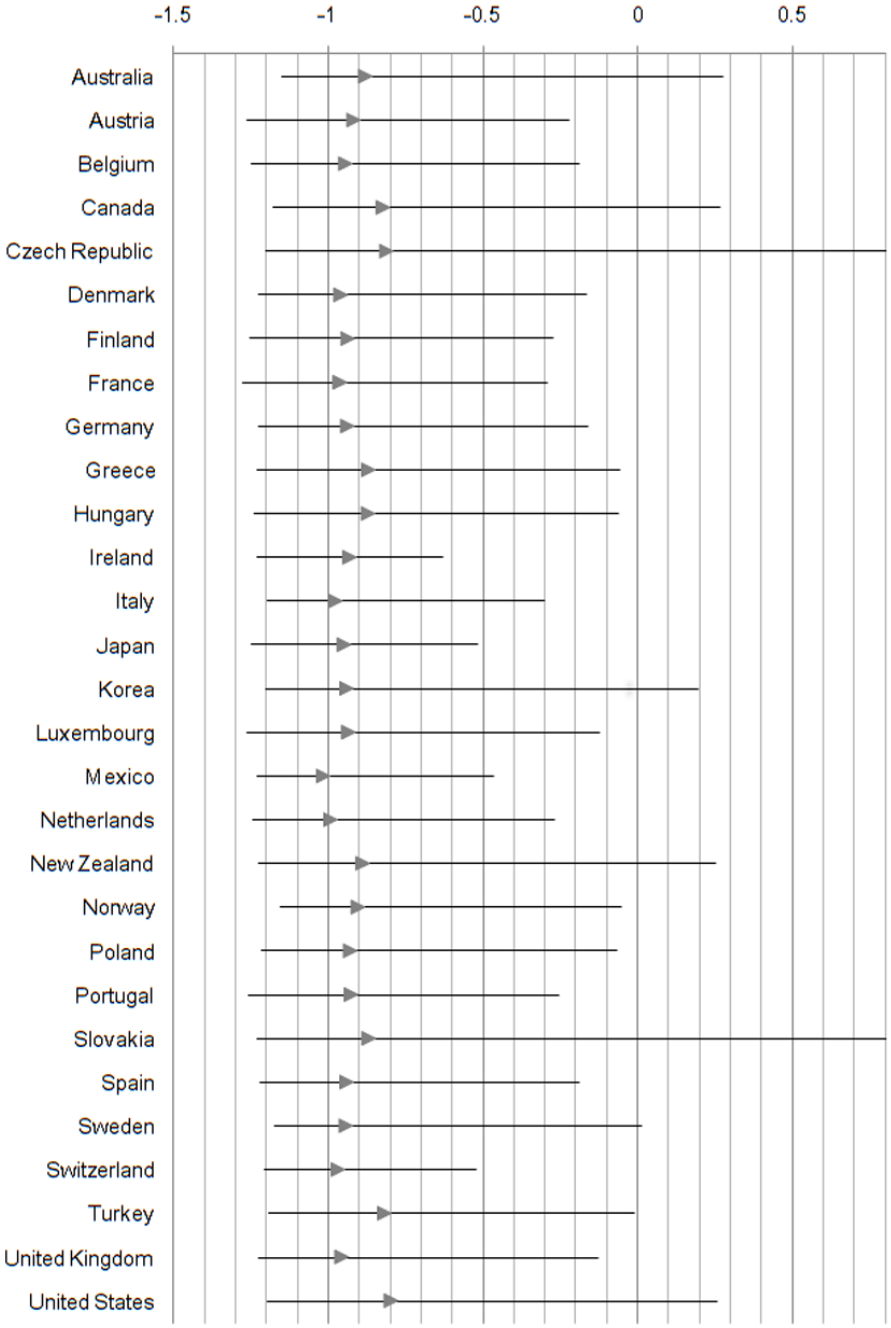
where $\overline{\widehat{\sigma}}$ is the mean rolling price standard deviation across all countries and time periods. This total elasticity is, of course, smaller in magnitude than $\widehat{\theta}_2$, but still on the relatively high end of price elasticity estimates when compared to the literature.

Although the PMG estimator constrains long-run coefficients to be homogenous, we can substitute country-specific averages $\overline{\widehat{\sigma}}_i$ into (3.17) to interpolate the variation in "total" price elasticity due to volatility differences. These country-specific total price elasticities are plotted in Figure 3.7. Over time, volatility levels within a country evolve and may thereby alter total price elasticities; and so Figure 3.7 also depicts the range of total price elasticities implied by each country's range of price volatilities $\widehat{\sigma}_{it}$. (Some of these ranges extend into positive price elasticities, but the worst offenders, the Czech Republic and Slovakia, are actually omitted from the regression, and so their implied wrong-sign elasticities should not cast doubt on the model as a whole.) It should be observed that the United States, whose price responsiveness is generally found to be low in comparison to other countries', has the highest average price volatility in the sample, and so by construction its total elasticity is of comparatively low magnitude. Although this does not establish that rational habits are responsible for the United States' low price responsiveness, the model does at least correctly predict this phenomenon.

The remaining long-term parameter, the time trend, is about -0.05 and statistically significantly. All other things equal, this implies, per-capita gasoline consumption has been declining about 5% per year. This negative trend seems to have been widespread: in the version of the ECM with heterogeneous trends, the majority of country-specific trends are negative (as are nearly all the country-specific trends estimated using 2SLS on the ADL). There are variety of potential explanations for these downward trends. Demographic shifts, for instance, might have reduced the number of drivers in ageing populations or pushed populations towards locations that demanded less driving. The more probable explanation is that technological innovations, tightening fuel-economy standards, and individual investments in fuel efficiency have dampened gasoline demand.

Figure 3.7: Country-specific "total" long-run price elasticities.

Triangle indicates country-specific price elasticity implied by mean volatility $\overline{\widehat{\sigma}}_i$; line indicates range of elasticities implied by country's range of price volatilities $[\widehat{\sigma}_{i,\min}, \widehat{\sigma}_{i,\max}]$.



The short-run coefficients reported in Table 3.15, which are the averages of the country-specific short-run coefficients, are (like MG estimates) not as reliable as the long-run estimates, since each country's short-run coefficients have been estimated using relatively little data. Nonetheless, estimates of δ_1 and δ_2 appear to be very reasonable: the short-run income elasticity of around 0.2 is approximately a quarter the long-run income elasticity; the short-run price elasticity (non-"total") of -0.15 to -0.21 is roughly fifteen to twenty percent of the long-run price elasticity.

The unreliability of the short-run estimates does show up in the remaining coefficients, however: in the short run, the effect of price volatility is estimated to be negative if a trend is excluded and *positive* if it is included; the effect of the price-volatility interaction term is statistically insignificant in either case. The relative unreliability of the short-run estimates seems to be a worthwhile trade-off, however, for the heterogeneity that improves the long-run estimates of primary interest.

3.6 Conclusion

Consumers' gasoline consumption is decided by a multitude of choices—some large, some small; some in the form of investment decisions, others in the form of everyday routines. All of these choices can be swept into a habits framework, which highlights the importance of price uncertainty and volatility in shaping the demand of forward-looking consumers.

The rational habits framework suggests two major implications for consumer behavior. This chapter has tested for both: first, that price uncertainty will depress demand; and second, that price uncertainty will dampen the price *responsiveness* of demand. Evidence for the first implication is strong. In the long run, price volatility *does* depress gasoline demand; the effect of the rolling-standard-deviation measure on demand is not only statistically significant but of a similar magnitude to price elasticity itself.

Evidence for the second implication is not as unequivocal. In some of my preferred specifications, it does indeed appear that total price responsiveness is greater when prices are less volatile. The estimated effect of price volatility on price elasticity, however, is quite sensitive to the model specification and estimation method. Future work on consumer habits should keep this sensitivity in mind and give careful thought to choosing an appropriate estimator.

It should be noted again—particularly in light of the fragile estimates of price volatility's effect on price elasticity—that this chapter has based all its tests for rational habits on a price-uncertainty measure that uses the rolling standard deviation of prices as a proxy for anticipated future price volatility. The appropriateness of this proxy is therefore key to the results. Imperfections in this proxy, even imperfections that amount to well-behaved measurement error, introduce a downward

bias (in magnitude) to estimates of uncertainty's effects, and may therefore disguise habits' influence on price elasticity. Future work on consumer habits should keep this in mind as well: thoughtful, explicit modelling of consumers' expectations may bring evidence of rational habits into starker relief.

Altogether, there is compelling evidence that consumers around the world exhibit at least some of the behaviors associated with rational gasoline habits. In light of these habits, policymakers should heed the role of price uncertainty in shaping gasoline demand. Policies to stabilize gasoline prices may have the undesired effect of increasing demand, and to offset this increase there may be a need for further intervention. This intervention may be made easier, however, by the increased effectiveness of price instruments when price volatility is low. Given rational habits, in fact, stable price increases are a longer policy lever than traditionally-measured price elasticity suggests. Not only may gasoline taxes explain some of the variation in price elasticity around the world, but their influence on price volatility may allow them to re-shape one country's demand to look like another's.

3.A Further Data Exposition

3.A.1 Geographic Scope

The twenty-nine countries in my sample are Australia, Austria, Belgium, Canada, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, the Netherlands, New Zealand, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and the United States.

3.A.2 Grades and Prices

The IEA reports price and tax data by product, i.e. by petrol type and grade. The availability of data on any product naturally depends on the product's availability to consumers, which in turn depends upon local regulations. For a number of countries—including Finland, France, Ireland, Italy, Luxembourg, and the U.K.—data on unleaded gasoline is not available until 1990. To avoid the problems associated with switching from a measure of price based solely on leaded gasoline to a measure based on unleaded gasoline (or some average of leaded and unleaded), I concentrate on the period 1990 onwards.

Even in this later period, however, data availability varies by product and country, with this variation driven by differences in regulation and definition. In Australia, for example, "95 RON" is considered premium, and regular unleaded gasoline is 91 RON (MotorMouth 2010); in the UK, "95 RON" is regular unleaded (Boulter and Latham 2009). Such differences in reported octane ratings may exaggerate differences in the actual content of the fuels, as the definition of the Research Octane Number (RON) varies by country, and its definition tends to be lower-skewing in countries that define "regular unleaded" by a lower RON. Whatever the differences in definition across countries, I need a price series for each country that is measured by a constant definition, since varying the definition over time would introduce false price variation (unless, of course, such variations reflected actual changes in the type of gasoline in use). Where annual data for "regular unleaded" is available from 1990 on—however "regular unleaded" may be defined locally—I use this as the product. Where annual data on regular unleaded is not available contiguously from 1990 but data on 95 RON is, I use 95 RON. Where data on neither regular unleaded nor 95 RON is available contiguously from 1990, I use whichever product has the longer contiguous availability. After determining which product to use for annual data, I use the same product for quarterly data. Table 3.19 shows the product for which each country's price measures are taken.

Table 3.19: Gasoline product used for price and tax data

Regular Unleaded	95 RON	
Australia	Belgium	Norway
Austria	Czech Republic	Poland
Canada	Finland	Portugal
Denmark	France	Slovak Republic
Germany	Greece	Spain
Japan	Hungary	Sweden
Korea	Ireland	Switzerland
Mexico	Italy	Turkey
New Zealand	Luxembourg	United Kingdom
United States	Netherlands	

3.B Local-currency Estimates

This appendix reports local-currency results corresponding to the common-currency results presented in the main text. As noted in Section 3.3.3, not all the estimators applied to the common-currency data can be sensibly applied to the local-currency data: in particular, the random-effects estimator of Table 3.20 and the pooled 2SLS estimator of Table 3.22 are both biased because of their comparison of non-differenced log prices and log income across different currencies.

That said, the remaining results should not be subject to the currency-generated bias. The within-groups (fixed-effects) 2SLS results (Table 3.21), in fact, make a little bit more sense than their common-currency counterparts, with negative and statistically-significant price effects no matter the specification of the trend. The results with respect to income and the rational-habits variables, however, show a pattern similar to the common-currency results, with income's effect becoming more pronounced as a homogenous and then a heterogeneous trend is introduced, and the rational-habits variables never showing a statistically-significant effect.

Turning to system GMM (Table 3.24), the estimators are just as uninformative when applied to local-currency data as they were when applied to common-currency data.

Moving on to the preferred estimator, namely the pooled mean group estimator applied to the ECM, there are some interesting contrasts to the common-currency results. Generally speaking, I find lower-magnitude long-run income and price elasticities in the local-currency data. In the "standard" model without rational-habits variables, the long-run income elasticity is about 0.3, which is considerably below both the literature's averages and most of the common-currency estimates. Similarly, the long-run price elasticities from the "standard" model with no trend or heterogeneous trends (-0.65 and -0.14, respectively) are smaller than the corresponding common-currency results.

There is a confounding blip in the homogenous-trend versions of the local-currency estimates,

with price elasticity in the "standard" model shooting to -1.28 and "total" price elasticity in the rational-habits model shooting to -1.08. The latter is driven almost entirely by an implausibly-sized and incorrectly-signed coefficient on $\widehat{\sigma}\widetilde{p}$, and so I am inclined to be suspicious of these homogenous-trend results.

With respect to the rational habits variables, the local-currency results are less conclusive than their common-currency counterparts. Whereas the coefficient on $\widehat{\sigma}$ was negative and highly statistically significant in all common-currency estimates of the ECM, when using local currencies the coefficient on $\widehat{\sigma}$ is nowhere near statistically significant unless heterogeneous trends are included—and even then the p-value is 0.065. With respect to the interaction term $\widehat{\sigma}\widetilde{p}$, the local-currency data produces the expected positive coefficient ($p = 0.07$) only when heterogeneous trends are included. Without a trend, the coefficient on $\widehat{\sigma}\widetilde{p}$ is indistinguishable from 0; and with a homogenous trend, as has been noted, the coefficient on $\widehat{\sigma}\widetilde{p}$ is hugely *negative*.

Overall, it seems that in practice the local-currency approach may be less reliable than the common-currency approach.

Table 3.20: Random- vs. fixed-effects estimator, static model, local currencies

		(1) Static Random effects	(2) Static Fixed effects	(3) Dynamic Random effects	(4) Dynamic Fixed effects
g_{t-1}	λ			0.992**	0.905**
				(0.00660)	(0.0160)
				[0.000]	[0.000]
y	δ_1	0.298**	0.265**	0.00672	-0.0178
		(0.0413)	(0.0429)	(0.00609)	(0.0160)
		[0.000]	[0.000]	[0.270]	[0.267]
p	δ_2	-0.494**	-0.586**	-0.00261	-0.169**
		(0.0508)	(0.0635)	(0.00601)	(0.0282)
		[0.000]	[0.000]	[0.664]	[0.000]
$\widehat{\sigma}$	δ_3	-0.216	-0.169	-0.185*	-0.0772
		(0.267)	(0.266)	(0.0937)	(0.0944)
		[0.419]	[0.527]	[0.048]	[0.414]
$\widehat{\sigma}\widetilde{p}$	δ_4	-0.410	0.335	-1.469**	0.331
		(0.693)	(0.752)	(0.275)	(0.407)
		[0.554]	[0.656]	[0.000]	[0.417]
R^2		0.2892	0.2274	0.9951	0.8288
Hausman test[a]		11.21 [0.024]		74.79 [0.000]	

Standard errors in parentheses; P-values in brackets.
** p<0.01, * p<0.05, + p<0.1

[a] Under H_0 that random- and fixed-effects estimates equivalent, $\chi^2(4)$ for static model, $\chi^2(4)$ for dynamic.

Table 3.21: Within-groups (fixed effects) estimates with 2SLS instrumenting for price, local currencies

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}		0.906**	0.908**	0.697**
		(0.0287)	(0.0279)	(0.0418)
		[0.000]	[0.000]	[0.000]
y		-0.00128	0.0566*	0.235**
		(0.0217)	(0.0272)	(0.0595)
		[0.953]	[0.038]	[0.000]
p		-0.203**	-0.158**	-0.210**
		(0.0459)	(0.0524)	(0.0482)
		[0.000]	[0.003]	[0.000]
$\widehat{\sigma}$		-0.0667	-0.0322	-0.000603
		(0.125)	(0.127)	(0.125)
		[0.595]	[0.800]	[0.996]
$\widehat{\sigma}\widetilde{p}$		0.419	0.305	-0.454
		(0.513)	(0.513)	(0.523)
		[0.415]	[0.552]	[0.386]
t			-0.00245*	[country-specific
			(0.00103)	trends]
			[0.017]	
R^2		0.899	0.902	0.935
Implied long-run elasticities				
y	$\frac{\delta_1}{1-\lambda}$	-.0136	.616	.776
		(0.232)	(0.292)	(0.186)
		[0.953]	[0.036]	[0.000]
p	$\frac{\delta_2}{1-\lambda}$	-2.165	-1.715	-0.695
		(0.656)	(0.653)	(0.161)
		[0.001]	[0.009]	[0.000]
Tests				
Endogeneity: price[a]		9.769	4.228	6.198
		[0.0018]	[0.0398]	[0.0128]
Endogeneity: price & interaction[b]		11.735	6.557	8.041
		[0.0028]	[0.0377]	[0.0179]
Underidentification[c]		125.510	92.181	98.878
		[0.000]	[0.0000]	[0.000]
Exogeneity[d]		5.653	8.344	4.362
		[0.0592]	[0.0154]	[0.113]

Robust standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

[a] Difference-in-Sargan/Hansen $\chi^2(1)$ under H_0.
[b] Difference-in-Sargan/Hansen $\chi^2(2)$ under H_0.
[c] Kleibergen-Paap (2006) rk statistic; $\chi^2(3)$ under H_0.
[d] Overidentifying restrictions test (Hansen J statistic); $\chi^2(2)$ under H_0.

Table 3.22: Pooled estimates with 2SLS instrumenting for prices, local currencies

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}	λ	1.002**	1.001**	0.959**
		(0.00662)	(0.00675)	(0.0168)
		[0.000]	[0.000]	[0.000]
y	δ_1	0.00454	0.00443	0.00922
		(0.00675)	(0.00676)	(0.0156)
		[0.501]	[0.513]	[0.555]
p	δ_2	-0.00118	-0.000166	-0.0142
		(0.00673)	(0.00670)	(0.0166)
		[0.861]	[0.980]	[0.393]
$\widehat{\sigma}$	δ_3	-0.0981	-0.0139	-0.194
		(0.117)	(0.117)	(0.124)
		[0.401]	[0.906]	[0.119]
$\widehat{\sigma}\widetilde{p}$	δ_4	-1.800**	-0.929*	-0.866*
		(0.469)	(0.434)	(0.390)
		[0.000]	[0.033]	[0.027]
t			-0.00216**	[country-specific
			(0.000620)	trends]
			[0.001]	
R^2		0.996	0.996	0.997

Robust standard errors in parentheses;
p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

Table 3.23: Within-groups (fixed-effects) estimates, local currencies

		(1) No trend	(2) Common trend	(3) Individual trends
g_{t-1}	λ	0.905**	0.893**	0.729**
		(0.0218)	(0.0231)	(0.0358)
		[0.000]	[0.000]	[0.000]
y	δ_1	-0.0178	0.0559+	0.255**
		(0.0215)	(0.0287)	(0.0606)
		[0.408]	[0.052]	[0.000]
p	δ_2	-0.169**	-0.156**	-0.168**
		(0.0377)	(0.0397)	(0.0392)
		[0.000]	[0.000]	[0.000]
$\widehat{\sigma}$	δ_3	-0.0772	-0.00558	-0.0446
		(0.117)	(0.119)	(0.114)
		[0.510]	[0.963]	[0.695]
$\widehat{\sigma}\widetilde{p}$	δ_4	0.331	0.386	-0.348
		(0.422)	(0.402)	(0.432)
		[0.433]	[0.337]	[0.420]
t			-0.00280**	
			(0.00103)	
			[0.007]	
R^2		0.903	0.905	0.938

Robust standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

Table 3.24: System GMM, local currencies

		(1) No trend	(2) With trend	(3) No trend Collapsed IV matrix	(4) With trend Collapsed IV matrix	(5) No trend Reduced IV count	(6) With trend Reduced IV count	(7) No trend No outside IVs	(8) With trend No outside IVs
g_{t-1}	λ	1.005**	0.991**	1.035**	0.928**	1.006**	0.724	1.002**	0.977**
		(0.00733)	(0.0143)	(0.0271)	(0.136)	(0.149)	(0.575)	(0.0180)	(0.0410)
		[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.218]	[0.000]	[0.000]
y	δ_1	0.0100	0.0569	0.0311	0.234	0.167	0.602	0.0151+	0.0423
		(0.00814)	(0.0364)	(0.0378)	(0.189)	(0.297)	(0.981)	(0.00863)	(0.0349)
		[0.228]	[0.129]	[0.417]	[0.227]	[0.579]	[0.545]	[0.091]	[0.236]
p	δ_2	-0.00597	-0.0451	-0.0184	-0.201	-0.133	-0.535	-0.0157	-0.0298
		(0.00701)	(0.0319)	(0.0320)	(0.170)	(0.259)	(0.892)	(0.0206)	(0.0312)
		[0.402]	[0.168]	[0.571]	[0.248]	[0.610]	[0.554]	[0.454]	[0.348]
$\widehat{\sigma}$	δ_3	-0.146	0.0971	-0.191	0.0145	-1.463	-1.224	-0.106	0.160
		(0.191)	(0.203)	(0.257)	(0.232)	(2.323)	(2.302)	(0.240)	(0.252)
		[0.451]	[0.636]	[0.464]	[0.951]	[0.534]	[0.599]	[0.660]	[0.530]
$\widehat{\sigma}\widetilde{p}$	δ_4	-1.649*	-0.421	-1.995*	-0.00918	-0.455	3.255	-2.904*	-1.773+
		(0.655)	(0.678)	(0.955)	(1.367)	(4.162)	(7.916)	(1.144)	(0.881)
		[0.018]	[0.540]	[0.046]	[0.995]	[0.914]	[0.684]	[0.017]	[0.054]
t			-0.00399**		-0.00640		-0.0125		-0.00352*
			(0.00139)		(0.00479)		(0.0190)		(0.00135)
			[0.008]		[0.192]		[0.514]		[0.014]
Sargan		420.4	380.5	90.69	59.02	10.38	10.65	411.6	378.2
Sargan df		169	168	17	16	3	2	165	164
Sargan P		0	0	0	7.65e-07	0.0156	0.00487	0	0
Hansen		26.77	23.08	20.45	22.96	7.894	5.770	25.96	23.87
Hansen df		169	168	17	16	3	2	165	164
Hansen P		1	1	0.252	0.115	0.0482	0.0558	1	1
AR(2) test		-0.415	-0.503	-0.371	-0.311	0.399	0.436	-0.297	-0.419
AR(2) P		0.678	0.615	0.711	0.756	0.690	0.663	0.767	0.675

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

Table 3.25: Estimates of the error-correction model, all countries, local currencies

		(1) PMG No trend	(2) PMG (Homo-genous) trend	(3) PMG (Hetero-genous) trends	(4) MG No trend	(5) MG (Hetero-geneous) trends	(6) DFE No trend	(7) DFE (Homo-geneous) trend
t			-0.0499**	-0.00140		-0.0126		-0.0376+
			(0.00648)	(0.00136)		(0.00872)		(0.0212)
			[0.000]	[0.302]		[0.149]		[0.076]
Long run								
y	θ_1	0.258**	0.560*	0.0594	-0.286	0.190	-0.815	0.300
		(0.0343)	(0.268)	(0.0746)	(0.402)	(0.483)	(0.646)	(0.545)
		[0.000]	[0.037]	[0.426]	[0.476]	[0.694]	[0.207]	[0.582]
p	θ_2	-0.501**	-0.0955	-0.440**	-0.0252	0.111	-1.666	-1.112
		(0.0809)	(0.231)	(0.0640)	(0.433)	(0.549)	(1.013)	(0.728)
		[0.000]	[0.679]	[0.000]	[0.954]	[0.839]	[0.100]	[0.126]
$\widehat{\sigma}$	θ_3	-0.359	-0.376	-0.354+	-0.0753	1.458	1.214	1.791
		(0.279)	(0.841)	(0.192)	(1.978)	(1.111)	(2.166)	(1.687)
		[0.199]	[0.655]	[0.065]	[0.970]	[0.189]	[0.575]	[0.288]
$\widehat{\sigma}\widetilde{p}$	θ_4	-1.457	-27.64**	2.258+	-35.85+	-15.25	4.331	3.088
		(2.102)	(5.587)	(1.244)	(19.67)	(21.12)	(9.114)	(7.007)
		[0.488]	[0.000]	[0.070]	[0.068]	[0.470]	[0.635]	[0.659]
Short run								
EC term	ϕ	-0.130**	-0.0890**	-0.252**	-0.452**	-0.515**	-0.0663+	-0.0820*
		(0.0479)	(0.0215)	(0.0483)	(0.0726)	(0.0938)	(0.0342)	(0.0359)
		[0.007]	[0.000]	[0.000]	[0.000]	[0.000]	[0.052]	[0.022]
Δy	δ_1	0.267**	0.407**	0.272**	0.272+	0.271	0.379**	0.379**
		(0.0625)	(0.0711)	(0.0836)	(0.154)	(0.169)	(0.0794)	(0.0822)
		[0.000]	[0.000]	[0.001]	[0.077]	[0.110]	[0.000]	[0.000]
Δp	δ_2	-0.139**	-0.147**	-0.0670+	-0.111+	-0.102	-0.117*	-0.105*
		(0.0515)	(0.0316)	(0.0394)	(0.0643)	(0.0698)	(0.0498)	(0.0477)
		[0.007]	[0.000]	[0.089]	[0.085]	[0.143]	[0.019]	[0.028]
$\Delta\widehat{\sigma}$	δ_3	-0.00355	0.0783	0.103*	-0.110	-0.0395	-0.000811	0.00461
		(0.0599)	(0.0759)	(0.0520)	(0.123)	(0.135)	(0.111)	(0.115)
		[0.953]	[0.303]	[0.047]	[0.371]	[0.769]	[0.994]	[0.968]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	δ_4	-1.045*	0.720	-0.739+	1.655	1.784	-0.144	-0.00948
		(0.501)	(0.644)	(0.419)	(1.843)	(2.013)	(0.347)	(0.352)
		[0.037]	[0.264]	[0.078]	[0.369]	[0.375]	[0.679]	[0.979]

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

Table 3.26: Estimates of error-correction model with only the standard regressors, local currencies

		(1) PMG No trend	(2) PMG Homogenous trend	(3) PMG Heterogeneous trends	(4) MG No trend	(5) MG (Heterogeneous) trends	(6) DFE No trend	(7) DFE (Homogenous) trend
t			-0.0463**	-0.00383*		0.0620		-0.0342
			(0.00817)	(0.00163)		(0.0968)		(0.0212)
			[0.000]	[0.019]		[0.522]		[0.106]
Long run								
y	θ_1	0.301**	0.283	0.275**	-0.246	-1.565	-0.741	0.307
		(0.0321)	(0.320)	(0.0436)	(0.625)	(1.815)	(0.637)	(0.543)
		[0.000]	[0.375]	[0.000]	[0.694]	[0.388]	[0.245]	[0.571]
p	θ_2	-0.648**	-1.283**	-0.139**	-0.660**	-1.569	-1.435*	-0.981*
		(0.0389)	(0.165)	(0.0237)	(0.217)	(1.311)	(0.662)	(0.468)
		[0.000]	[0.000]	[0.000]	[0.002]	[0.231]	[0.030]	[0.036]
Short run								
EC term	ϕ	-0.142**	-0.0803**	-0.295**	-0.376**	-0.507**	-0.0681+	-0.0820*
		(0.0508)	(0.0175)	(0.0642)	(0.0638)	(0.0713)	(0.0362)	(0.0377)
		[0.005]	[0.000]	[0.000]	[0.000]	[0.000]	[0.060]	[0.030]
Δy	δ_1	0.273**	0.384**	0.286**	0.330**	0.239**	0.350**	0.358**
		(0.0600)	(0.0575)	(0.0726)	(0.0659)	(0.0880)	(0.0518)	(0.0578)
		[0.000]	[0.000]	[0.000]	[0.000]	[0.007]	[0.000]	[0.000]
Δp	δ_2	-0.138*	-0.0992**	-0.0875*	-0.0109	-0.00316	-0.109**	-0.0880*
		(0.0569)	(0.0371)	(0.0425)	(0.0456)	(0.0422)	(0.0411)	(0.0385)
		[0.016]	[0.007]	[0.039]	[0.812]	[0.940]	[0.008]	[0.022]

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

3.C Implementation of Estimation and Tests

All estimation has been performed in Stata 11 (2009). The panel unit root tests reported in Tables 3.4 and 3.5 are implemented using the Stata (2009) program *xtunitroot*. The cointegration tests reported in Tables 3.6 and 3.7 are implemented using Persyn and Westerlund's program *xtwest*, detailed in Persyn and Westerlund (2008). The PMG, MG, and DFE estimators of the error-correction models are implemented using Blackburne and Frank's *xtpmg* (2007b), detailed in Blackburne and Frank (2007a). GMM estimates of the ADL are computed using Roodman's (2011) *xtabond2*, and least-squares estimates are computed using Schaffer's *xtivreg2*.

3.D PMG Estimates with Individual Country-Specific Short-Run Coefficients

Table 3.27 (pp. 158-172) reports the *country-specific* short-run ECM coefficients estimated by PMG, averages of which are reported in Tables 3.14 and 3.25.

Table 3.27: Pooled mean group estimates of error-correction model with reporting of individual short-run parameters

	(1) USD No trend	(2) USD Homogenous trend	(3) USD Heterogeneous trends	(4) Local currency No trend	(5) Local currency Homogenous trend	(6) Local currency Heterogeneous trends
Long Run						
y	0.555**	0.774**	0.388**	0.258**	0.560*	0.0594
	(0.0306)	(0.0905)	(0.0412)	(0.0343)	(0.268)	(0.0746)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.037]	[0.426]
p	-0.693**	-0.981**	-0.626**	-0.501**	-0.0955	-0.440**
	(0.0755)	(0.122)	(0.0398)	(0.0809)	(0.231)	(0.0640)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.679]	[0.000]
$\widehat{\sigma}$	-1.249**	-1.464**	-0.881**	-0.359	-0.376	-0.354+
	(0.249)	(0.353)	(0.100)	(0.279)	(0.841)	(0.192)
	[0.000]	[0.000]	[0.000]	[0.199]	[0.655]	[0.065]
$\widehat{\sigma}\widetilde{p}$	0.816	-0.307	-0.0718	-1.457	-27.64**	2.258+
	(1.205)	(1.872)	(0.520)	(2.102)	(5.587)	(1.244)
	[0.498]	[0.870]	[0.890]	[0.488]	[0.000]	[0.070]
t		-0.0473**			-0.0499**	
		(0.00303)			(0.00648)	
		[0.000]			[0.000]	
			Individual Short-Run Estimates			
Australia						
EC term	-0.220*	-0.0444+	-0.175**	-0.384**	-0.0367+	-0.636**
	(0.0903)	(0.0227)	(0.0645)	(0.101)	(0.0212)	(0.136)
	[0.015]	[0.051]	[0.007]	[0.000]	[0.083]	[0.000]
Δy	0.0463	0.0671	0.0842	0.0927	0.578	-0.00186
	(0.0814)	(0.0854)	(0.0767)	(0.309)	(0.471)	(0.341)
	[0.570]	[0.432]	[0.272]	[0.764]	[0.220]	[0.996]
Δp	0.0162	0.00870	0.0151	0.129	0.0933	0.241*
	(0.0745)	(0.0777)	(0.0690)	(0.102)	(0.127)	(0.0988)
	[0.828]	[0.911]	[0.827]	[0.206]	[0.464]	[0.015]
$\Delta\widehat{\sigma}$	-0.292	-0.368	-0.256	-0.325+	-0.336	-0.219
	(0.222)	(0.224)	(0.207)	(0.185)	(0.230)	(0.173)
	[0.189]	[0.100]	[0.218]	[0.079]	[0.144]	[0.205]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.942+	1.048*	1.123*	1.571	1.032	-0.0404
	(0.514)	(0.501)	(0.452)	(1.391)	(1.785)	(1.238)
	[0.067]	[0.036]	[0.013]	[0.259]	[0.563]	[0.974]
t			-0.00148			0.00101
			(0.000985)			(0.00151)
			[0.134]			[0.504]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Austria						
EC term	-0.0802	-0.0468	-0.202*	-0.106	-0.00346	-0.202
	(0.0574)	(0.0675)	(0.0963)	(0.0708)	(0.0634)	(0.158)
	[0.162]	[0.488]	[0.036]	[0.136]	[0.956]	[0.199]
Δy	0.640**	0.624**	0.579**	0.522	0.625	0.544
	(0.176)	(0.188)	(0.169)	(0.607)	(0.645)	(0.603)
	[0.000]	[0.001]	[0.001]	[0.389]	[0.332]	[0.366]
Δp	-0.458*	-0.446*	-0.362*	-0.443+	-0.512*	-0.422+
	(0.178)	(0.194)	(0.178)	(0.227)	(0.245)	(0.227)
	[0.010]	[0.021]	[0.043]	[0.050]	[0.037]	[0.063]
$\Delta\widehat{\sigma}$	-0.0487	0.0784	0.0216	-0.310	-0.208	-0.242
	(0.280)	(0.308)	(0.272)	(0.414)	(0.450)	(0.421)
	[0.862]	[0.799]	[0.937]	[0.454]	[0.644]	[0.566]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-1.349	-0.743	-0.514	-3.304	-1.122	-1.937
	(1.316)	(1.515)	(1.337)	(5.616)	(8.261)	(6.326)
	[0.306]	[0.624]	[0.701]	[0.556]	[0.892]	[0.759]
t			-0.00426			-0.00264
			(0.00308)			(0.00390)
			[0.167]			[0.497]
Belgium						
EC term	0.202**	-0.486**	-0.270+	0.151**	-0.373**	-0.122
	(0.0624)	(0.117)	(0.160)	(0.0513)	(0.0803)	(0.124)
	[0.001]	[0.000]	[0.091]	[0.003]	[0.000]	[0.326]
Δy	0.497**	-0.0812	0.153	-0.380	-0.412	-0.210
	(0.162)	(0.159)	(0.179)	(0.828)	(0.506)	(0.717)
	[0.002]	[0.610]	[0.393]	[0.646]	[0.416]	[0.769]
Δp	-0.390*	0.163	0.0120	-0.524**	-0.117	-0.208
	(0.165)	(0.180)	(0.196)	(0.194)	(0.140)	(0.216)
	[0.018]	[0.367]	[0.951]	[0.007]	[0.401]	[0.336]
$\Delta\widehat{\sigma}$	-0.319	0.161	-0.00887	0.0339	0.826*	0.464
	(0.477)	(0.419)	(0.421)	(0.485)	(0.343)	(0.456)
	[0.504]	[0.701]	[0.983]	[0.944]	[0.016]	[0.309]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	2.305	1.726	1.072	2.536	5.328*	-0.426
	(1.417)	(1.512)	(1.218)	(3.846)	(2.711)	(3.676)
	[0.104]	[0.254]	[0.379]	[0.510]	[0.049]	[0.908]
t			-0.0170**			-0.0108*
			(0.00564)			(0.00458)
			[0.002]			[0.019]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Canada						
EC term	-0.0606	0.0251	-0.101**	0.0719	0.0244**	-0.365**
	(0.0585)	(0.0163)	(0.0375)	(0.0635)	(0.00845)	(0.0843)
	[0.301]	[0.123]	[0.007]	[0.258]	[0.004]	[0.000]
Δy	0.0939	0.121	-0.141	0.431*	0.423**	-0.0604
	(0.158)	(0.145)	(0.124)	(0.172)	(0.114)	(0.117)
	[0.553]	[0.404]	[0.256]	[0.013]	[0.000]	[0.606]
Δp	-0.0606	-0.117	-0.00868	-0.122	-0.120*	0.0478
	(0.102)	(0.0910)	(0.0717)	(0.0757)	(0.0526)	(0.0547)
	[0.552]	[0.198]	[0.904]	[0.107]	[0.022]	[0.382]
$\Delta\widehat{\sigma}$	0.0399	-0.132	-0.0586	-0.0929	-0.114	-0.0130
	(0.120)	(0.109)	(0.0810)	(0.0990)	(0.0827)	(0.0721)
	[0.740]	[0.226]	[0.469]	[0.348]	[0.167]	[0.857]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.0155	-0.419	-0.116	-0.926+	-1.462**	-0.797+
	(0.358)	(0.348)	(0.263)	(0.558)	(0.489)	(0.470)
	[0.966]	[0.228]	[0.658]	[0.097]	[0.003]	[0.090]
t			0.00339**			0.00502**
			(0.000785)			(0.000941)
			[0.000]			[0.000]
Czech Republic						
EC term	-1.225**	-0.338**	-1.092**	-0.888**	-0.269**	-0.847**
	(0.148)	(0.0622)	(0.211)	(0.0921)	(0.0530)	(0.0790)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Δy	-0.358**	0.177	-0.227	0.732**	1.351**	0.763**
	(0.127)	(0.145)	(0.150)	(0.211)	(0.413)	(0.165)
	[0.005]	[0.223]	[0.130]	[0.001]	[0.001]	[0.000]
Δp	0.635**	0.353	0.454**	0.300**	-0.00226	0.255**
	(0.193)	(0.251)	(0.158)	(0.0924)	(0.159)	(0.0691)
	[0.001]	[0.160]	[0.004]	[0.001]	[0.989]	[0.000]
$\Delta\widehat{\sigma}$	0.825**	0.308	0.776**	-0.110	-0.755*	0.254+
	(0.155)	(0.219)	(0.164)	(0.195)	(0.329)	(0.143)
	[0.000]	[0.159]	[0.000]	[0.572]	[0.022]	[0.077]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.528	0.698	1.371	-0.508	-1.172	-0.703
	(1.733)	(1.454)	(1.061)	(0.647)	(1.310)	(0.489)
	[0.761]	[0.631]	[0.196]	[0.433]	[0.371]	[0.150]
t			0.00975+			0.00692**
			(0.00580)			(0.00256)
			[0.093]			[0.007]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Denmark						
EC term	-0.292**	-0.0722**	-0.117+	-0.247**	-0.0617**	-0.0891**
	(0.0573)	(0.00847)	(0.0670)	(0.0403)	(0.00622)	(0.0298)
	[0.000]	[0.000]	[0.081]	[0.000]	[0.000]	[0.003]
Δy	-0.00514	0.150**	0.126*	0.0107	0.0316	0.0812
	(0.0760)	(0.0423)	(0.0618)	(0.177)	(0.1000)	(0.0919)
	[0.946]	[0.000]	[0.041]	[0.952]	[0.752]	[0.376]
Δp	-0.0439	-0.135**	-0.111*	0.00923	-0.155**	-0.103**
	(0.0764)	(0.0414)	(0.0550)	(0.0664)	(0.0330)	(0.0385)
	[0.566]	[0.001]	[0.044]	[0.890]	[0.000]	[0.007]
$\Delta\widehat{\sigma}$	-0.0448	0.0979	0.0550	-0.0526	-0.223*	-0.143
	(0.190)	(0.120)	(0.115)	(0.197)	(0.108)	(0.105)
	[0.814]	[0.414]	[0.633]	[0.789]	[0.040]	[0.175]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.249	0.119	0.128	-3.058+	1.348	-1.064
	(0.677)	(0.410)	(0.386)	(1.675)	(0.976)	(0.955)
	[0.713]	[0.772]	[0.741]	[0.068]	[0.167]	[0.265]
t			-0.00308**			-0.00281**
			(0.000782)			(0.000417)
			[0.000]			[0.000]
Finland						
EC term	-0.0424	-0.0163	-0.0364	-0.425**	-0.00120	-0.477*
	(0.0603)	(0.0242)	(0.0919)	(0.154)	(0.0107)	(0.197)
	[0.482]	[0.500]	[0.692]	[0.006]	[0.911]	[0.015]
Δy	0.229**	0.258**	0.244**	0.0924	0.152*	0.0637
	(0.0590)	(0.0461)	(0.0630)	(0.0658)	(0.0709)	(0.0731)
	[0.000]	[0.000]	[0.000]	[0.160]	[0.032]	[0.384]
Δp	-0.179*	-0.191**	-0.183*	-0.152+	-0.345**	-0.156+
	(0.0754)	(0.0715)	(0.0810)	(0.0853)	(0.0497)	(0.0914)
	[0.017]	[0.008]	[0.024]	[0.075]	[0.000]	[0.088]
$\Delta\widehat{\sigma}$	0.0557	0.108	0.0884	-0.146	-0.270*	-0.223+
	(0.106)	(0.141)	(0.144)	(0.133)	(0.135)	(0.129)
	[0.598]	[0.444]	[0.539]	[0.270]	[0.046]	[0.083]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.877+	-0.731	-0.726	1.192	0.792	0.305
	(0.528)	(0.588)	(0.666)	(0.870)	(0.931)	(0.862)
	[0.097]	[0.214]	[0.276]	[0.171]	[0.395]	[0.724]
t			-0.000520			0.00115
			(0.00119)			(0.00131)
			[0.663]			[0.379]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
France						
EC term	0.0308+	-0.142*	-0.190	0.0763**	-0.184**	-0.0171
	(0.0160)	(0.0592)	(0.122)	(0.0140)	(0.0590)	(0.0660)
	[0.054]	[0.016]	[0.119]	[0.000]	[0.002]	[0.796]
Δy	0.174**	0.111+	0.112+	0.784**	0.821**	0.864**
	(0.0594)	(0.0595)	(0.0648)	(0.196)	(0.222)	(0.195)
	[0.003]	[0.061]	[0.085]	[0.000]	[0.000]	[0.000]
Δp	-0.253**	-0.211**	-0.185**	-0.158**	-0.153*	-0.133*
	(0.0623)	(0.0623)	(0.0682)	(0.0602)	(0.0690)	(0.0600)
	[0.000]	[0.001]	[0.007]	[0.009]	[0.027]	[0.027]
$\Delta\widehat{\sigma}$	-0.0110	0.170	0.143	0.205	0.249	0.231+
	(0.0952)	(0.123)	(0.117)	(0.133)	(0.164)	(0.128)
	[0.908]	[0.167]	[0.222]	[0.123]	[0.129]	[0.071]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.585	-0.192	-0.278	-1.630	1.784	-1.781
	(0.423)	(0.501)	(0.428)	(1.355)	(2.105)	(1.339)
	[0.167]	[0.701]	[0.516]	[0.229]	[0.397]	[0.184]
t			-0.00884+			-0.00347
			(0.00479)			(0.00237)
			[0.065]			[0.143]
Germany						
EC term	0.136+	-0.0553**	-1.005**	0.215**	-0.0499**	-0.0106
	(0.0782)	(0.0185)	(0.330)	(0.0824)	(0.0132)	(0.105)
	[0.081]	[0.003]	[0.002]	[0.009]	[0.000]	[0.920]
Δy	0.297**	0.196**	-0.147	0.564	0.112	0.202
	(0.0907)	(0.0726)	(0.133)	(0.374)	(0.322)	(0.345)
	[0.001]	[0.007]	[0.268]	[0.132]	[0.729]	[0.558]
Δp	-0.490**	-0.331**	0.104	-0.405**	-0.270**	-0.256*
	(0.0952)	(0.0874)	(0.171)	(0.105)	(0.0828)	(0.109)
	[0.000]	[0.000]	[0.543]	[0.000]	[0.001]	[0.019]
$\Delta\widehat{\sigma}$	-0.375*	-0.120	0.214	0.147	0.175	0.156
	(0.176)	(0.182)	(0.207)	(0.143)	(0.121)	(0.119)
	[0.033]	[0.510]	[0.300]	[0.306]	[0.148]	[0.192]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.971	0.964+	1.649*	-0.0582	0.673	-0.206
	(0.601)	(0.524)	(0.760)	(0.983)	(0.860)	(0.906)
	[0.106]	[0.066]	[0.030]	[0.953]	[0.434]	[0.820]
t			-0.0128**			-0.00297+
			(0.00378)			(0.00155)
			[0.001]			[0.055]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Greece						
EC term	-0.0373	0.0110	-0.0536	0.0120	0.0114	-0.232**
	(0.0496)	(0.0170)	(0.0485)	(0.0369)	(0.0128)	(0.0877)
	[0.453]	[0.518]	[0.269]	[0.744]	[0.374]	[0.008]
Δy	0.0741	0.0927*	0.0576	0.156	0.0927	-0.263
	(0.0467)	(0.0425)	(0.0448)	(0.159)	(0.171)	(0.184)
	[0.113]	[0.029]	[0.199]	[0.329]	[0.587]	[0.153]
Δp	0.0198	-0.0193	-0.000402	0.1000	0.115	0.157*
	(0.0577)	(0.0573)	(0.0551)	(0.0740)	(0.0750)	(0.0623)
	[0.731]	[0.737]	[0.994]	[0.177]	[0.124]	[0.012]
$\Delta\widehat{\sigma}$	0.00219	-0.0914	-0.0867	-0.0539	-0.0626	-0.0281
	(0.108)	(0.118)	(0.109)	(0.0979)	(0.0957)	(0.0836)
	[0.984]	[0.438]	[0.424]	[0.582]	[0.513]	[0.736]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-0.819+	-1.039**	-0.705	-4.121**	-4.807**	-4.518**
	(0.425)	(0.401)	(0.431)	(1.563)	(1.683)	(1.199)
	[0.054]	[0.009]	[0.102]	[0.008]	[0.004]	[0.000]
t			0.00159+			0.00736**
			(0.000905)			(0.00239)
			[0.078]			[0.002]
Hungary						
EC term	-0.330**	0.0995	-0.184*	-0.372**	0.104*	-0.263**
	(0.0640)	(0.0684)	(0.0788)	(0.0958)	(0.0448)	(0.0926)
	[0.000]	[0.146]	[0.020]	[0.000]	[0.020]	[0.004]
Δy	0.0910	0.239*	0.135+	0.187	0.486*	0.280+
	(0.0699)	(0.108)	(0.0700)	(0.193)	(0.225)	(0.166)
	[0.193]	[0.027]	[0.054]	[0.331]	[0.031]	[0.093]
Δp	-0.179*	-0.148	-0.223*	-0.0187	-0.274+	-0.174
	(0.0886)	(0.144)	(0.0928)	(0.116)	(0.148)	(0.110)
	[0.043]	[0.306]	[0.016]	[0.872]	[0.064]	[0.113]
$\Delta\widehat{\sigma}$	0.335**	-0.0608	0.132	0.508**	0.503*	0.459**
	(0.121)	(0.208)	(0.128)	(0.195)	(0.230)	(0.156)
	[0.006]	[0.770]	[0.302]	[0.009]	[0.029]	[0.003]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	0.240	-1.359	-0.233	-4.514	0.138	-0.478
	(0.618)	(1.020)	(0.679)	(4.393)	(5.460)	(4.016)
	[0.698]	[0.183]	[0.731]	[0.304]	[0.980]	[0.905]
t			0.00519**			0.00496**
			(0.00147)			(0.00124)
			[0.000]			[0.000]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Ireland						
EC term	-0.407**	-0.110+	-0.385**	-0.0739	-0.0327	-0.0219
	(0.107)	(0.0630)	(0.118)	(0.0823)	(0.0360)	(0.147)
	[0.000]	[0.080]	[0.001]	[0.369]	[0.364]	[0.882]
Δy	-0.0688	0.00939	-0.0484	0.634**	0.621**	0.657**
	(0.139)	(0.182)	(0.154)	(0.208)	(0.214)	(0.195)
	[0.620]	[0.959]	[0.753]	[0.002]	[0.004]	[0.001]
Δp	0.116	0.0914	0.0895	0.134	0.127	0.125
	(0.137)	(0.192)	(0.162)	(0.210)	(0.208)	(0.209)
	[0.395]	[0.634]	[0.580]	[0.524]	[0.543]	[0.548]
$\Delta\widehat{\sigma}$	-0.350	-0.123	-0.551	0.00360	0.000763	0.00220
	(0.319)	(0.480)	(0.433)	(0.324)	(0.323)	(0.324)
	[0.273]	[0.798]	[0.203]	[0.991]	[0.998]	[0.995]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-3.333*	-1.547	-3.303*	-2.308	-1.633	-2.008
	(1.375)	(1.805)	(1.655)	(5.267)	(5.404)	(5.453)
	[0.015]	[0.391]	[0.046]	[0.661]	[0.763]	[0.713]
t			0.00472			-0.000825
			(0.00298)			(0.00508)
			[0.114]			[0.871]
Italy						
EC term	0.176+	-0.302**	-0.117*	0.149	-0.222**	-0.0801+
	(0.0910)	(0.0460)	(0.0559)	(0.0912)	(0.0342)	(0.0416)
	[0.054]	[0.000]	[0.037]	[0.102]	[0.000]	[0.055]
Δy	0.587**	0.160	0.239*	-0.0250	0.847*	0.407
	(0.198)	(0.121)	(0.0965)	(0.918)	(0.432)	(0.372)
	[0.003]	[0.188]	[0.013]	[0.978]	[0.050]	[0.274]
Δp	-0.718**	-0.198	-0.227*	-0.775*	-0.383*	-0.379*
	(0.223)	(0.142)	(0.114)	(0.354)	(0.176)	(0.149)
	[0.001]	[0.164]	[0.046]	[0.028]	[0.029]	[0.011]
$\Delta\widehat{\sigma}$	-0.200	0.311	0.181	-0.728	0.296	0.0596
	(0.358)	(0.229)	(0.165)	(0.742)	(0.397)	(0.310)
	[0.576]	[0.175]	[0.272]	[0.326]	[0.457]	[0.848]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	0.715	1.162	0.932	3.215	6.720	4.213
	(1.627)	(1.004)	(0.709)	(10.24)	(5.062)	(4.197)
	[0.660]	[0.247]	[0.189]	[0.753]	[0.184]	[0.315]
t			-0.0114**			-0.00953**
			(0.00135)			(0.000960)
			[0.000]			[0.000]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Japan						
EC term	-0.125**	-0.0384**	-0.132**	-0.126**	-0.0291**	-0.340**
	(0.0234)	(0.00934)	(0.0408)	(0.0266)	(0.00785)	(0.120)
	[0.000]	[0.000]	[0.001]	[0.000]	[0.000]	[0.005]
Δy	0.0723	0.0523	0.110	0.0529	-0.0450	0.168
	(0.0517)	(0.0659)	(0.0686)	(0.130)	(0.142)	(0.146)
	[0.162]	[0.427]	[0.109]	[0.685]	[0.752]	[0.249]
Δp	-0.0138	-0.0207	-0.0567	0.0162	-0.0330	0.0796
	(0.0525)	(0.0629)	(0.0611)	(0.0524)	(0.0530)	(0.0599)
	[0.793]	[0.742]	[0.354]	[0.758]	[0.534]	[0.184]
$\Delta\widehat{\sigma}$	0.0889	0.0528	0.0538	0.111	0.110	0.0798
	(0.0594)	(0.0642)	(0.0618)	(0.113)	(0.125)	(0.112)
	[0.135]	[0.411]	[0.384]	[0.330]	[0.379]	[0.476]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.435	-0.327	-0.204	-0.505	-0.428	-1.206
	(0.348)	(0.386)	(0.371)	(0.725)	(0.809)	(0.757)
	[0.212]	[0.396]	[0.581]	[0.486]	[0.597]	[0.111]
t			0.000127			0.00389+
			(0.000905)			(0.00215)
			[0.888]			[0.071]
Korea						
EC term	-0.265**	-0.173**	-0.136**	-0.196**	-0.125**	-0.156**
	(0.0448)	(0.0186)	(0.0457)	(0.0390)	(0.0256)	(0.0460)
	[0.000]	[0.000]	[0.003]	[0.000]	[0.000]	[0.001]
Δy	0.165	0.232**	0.318**	0.935**	0.802*	0.841*
	(0.123)	(0.0828)	(0.0912)	(0.360)	(0.357)	(0.379)
	[0.180]	[0.005]	[0.000]	[0.009]	[0.025]	[0.026]
Δp	-0.192	-0.196	-0.218	-0.0285	-0.261	-0.171
	(0.212)	(0.147)	(0.155)	(0.162)	(0.175)	(0.199)
	[0.365]	[0.183]	[0.159]	[0.860]	[0.135]	[0.389]
$\Delta\widehat{\sigma}$	0.175	0.409*	0.487*	0.577	0.993*	0.731
	(0.291)	(0.202)	(0.224)	(0.497)	(0.464)	(0.505)
	[0.548]	[0.043]	[0.030]	[0.245]	[0.032]	[0.148]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.971	2.484*	2.343*	1.900	5.086**	2.209
	(1.388)	(1.026)	(1.060)	(1.854)	(1.829)	(1.911)
	[0.484]	[0.015]	[0.027]	[0.305]	[0.005]	[0.248]
t			-0.00981**			-0.00384
			(0.00264)			(0.00330)
			[0.000]			[0.245]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Luxembourg						
EC term	0.103	-0.250**	-0.504**	0.00355	-0.177**	-0.167
	(0.155)	(0.0485)	(0.137)	(0.160)	(0.0403)	(0.106)
	[0.505]	[0.000]	[0.000]	[0.982]	[0.000]	[0.117]
Δy	0.276	0.0307	-0.114	0.522	0.346	0.392
	(0.237)	(0.134)	(0.153)	(0.486)	(0.285)	(0.317)
	[0.245]	[0.820]	[0.454]	[0.283]	[0.226]	[0.217]
Δp	-0.535*	-0.144	0.0589	-0.440	-0.297*	-0.191
	(0.209)	(0.132)	(0.158)	(0.268)	(0.144)	(0.185)
	[0.010]	[0.275]	[0.709]	[0.101]	[0.039]	[0.303]
$\Delta\widehat{\sigma}$	-0.225	0.378	0.309	-0.258	0.0763	-0.0114
	(0.342)	(0.248)	(0.216)	(0.478)	(0.309)	(0.320)
	[0.511]	[0.128]	[0.153]	[0.589]	[0.805]	[0.972]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	0.637	0.975	-0.126	0.388	5.536*	0.971
	(1.476)	(0.916)	(0.843)	(3.924)	(2.459)	(2.733)
	[0.666]	[0.287]	[0.881]	[0.921]	[0.024]	[0.722]
t			-0.0121**			-0.00825**
			(0.00194)			(0.00177)
			[0.000]			[0.000]
Mexico						
EC term	-0.319**	0.0943**	-0.102	-0.0364	0.0853**	-0.0792
	(0.0792)	(0.0334)	(0.0782)	(0.116)	(0.0293)	(0.0717)
	[0.000]	[0.005]	[0.191]	[0.755]	[0.004]	[0.269]
Δy	0.228+	0.445**	0.336**	0.488**	0.381**	0.210
	(0.123)	(0.112)	(0.0980)	(0.188)	(0.141)	(0.149)
	[0.065]	[0.000]	[0.001]	[0.009]	[0.007]	[0.159]
Δp	-0.133	-0.272*	-0.175+	-0.243	-0.291	-0.387+
	(0.109)	(0.113)	(0.0983)	(0.284)	(0.226)	(0.209)
	[0.219]	[0.017]	[0.074]	[0.393]	[0.198]	[0.063]
$\Delta\widehat{\sigma}$	0.473*	0.0422	0.195	0.561	0.530	0.481
	(0.204)	(0.260)	(0.224)	(0.554)	(0.443)	(0.398)
	[0.020]	[0.871]	[0.384]	[0.312]	[0.231]	[0.226]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-1.005	-1.928*	-1.673*	-9.593	-7.624	-7.228
	(0.874)	(0.979)	(0.812)	(11.63)	(9.268)	(8.368)
	[0.250]	[0.049]	[0.039]	[0.409]	[0.411]	[0.388]
t			0.00395**			0.00469**
			(0.00106)			(0.00120)
			[0.000]			[0.000]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Netherlands						
EC term	-0.0308	-0.0205	0.0739	-0.108	-0.00854	-0.499**
	(0.171)	(0.0180)	(0.117)	(0.0794)	(0.0108)	(0.166)
	[0.857]	[0.255]	[0.526]	[0.174]	[0.428]	[0.003]
Δy	0.364**	0.361**	0.407**	0.0638	-0.0502	0.520+
	(0.119)	(0.0978)	(0.108)	(0.253)	(0.243)	(0.276)
	[0.002]	[0.000]	[0.000]	[0.801]	[0.836]	[0.060]
Δp	-0.322**	-0.291**	-0.329**	-0.251**	-0.273**	-0.186*
	(0.113)	(0.0980)	(0.106)	(0.0870)	(0.0876)	(0.0794)
	[0.005]	[0.003]	[0.002]	[0.004]	[0.002]	[0.019]
$\Delta\widehat{\sigma}$	-0.0572	0.0410	-0.0279	-0.152	-0.186	-0.0388
	(0.276)	(0.217)	(0.221)	(0.174)	(0.176)	(0.158)
	[0.836]	[0.850]	[0.900]	[0.382]	[0.291]	[0.807]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-0.472	-0.426	-0.849	-0.678	-1.019	-1.541
	(0.784)	(0.653)	(0.794)	(1.620)	(1.670)	(1.379)
	[0.548]	[0.514]	[0.285]	[0.676]	[0.542]	[0.264]
t			-0.00130			0.00385*
			(0.000846)			(0.00179)
			[0.126]			[0.031]
New Zealand						
EC term	-0.0779	-0.00966	-0.0642	-0.126*	-0.00385	-0.231**
	(0.0546)	(0.0236)	(0.0500)	(0.0578)	(0.0129)	(0.0848)
	[0.154]	[0.682]	[0.199]	[0.030]	[0.765]	[0.007]
Δy	0.0755	0.103+	0.0833	0.261	0.358+	0.191
	(0.0526)	(0.0532)	(0.0526)	(0.169)	(0.185)	(0.166)
	[0.151]	[0.054]	[0.113]	[0.122]	[0.053]	[0.250]
Δp	-0.0194	-0.0181	-0.0183	-0.0264	-0.0542	0.00947
	(0.0597)	(0.0708)	(0.0687)	(0.0794)	(0.0874)	(0.0799)
	[0.746]	[0.798]	[0.790]	[0.739]	[0.535]	[0.906]
$\Delta\widehat{\sigma}$	0.00488	-0.0185	-0.0130	-0.489**	-0.442*	-0.493**
	(0.141)	(0.170)	(0.160)	(0.178)	(0.198)	(0.166)
	[0.972]	[0.914]	[0.936]	[0.006]	[0.025]	[0.003]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-0.186	-0.307	-0.173	0.858	0.307	0.0701
	(0.391)	(0.401)	(0.397)	(1.068)	(1.257)	(0.971)
	[0.634]	[0.444]	[0.663]	[0.422]	[0.807]	[0.942]
t			0.000182			0.00146
			(0.00110)			(0.000900)
			[0.869]			[0.104]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Norway						
EC term	0.0901**	-0.277**	-0.424*	0.170**	-0.133**	0.371**
	(0.0318)	(0.0827)	(0.173)	(0.0342)	(0.0500)	(0.120)
	[0.005]	[0.001]	[0.014]	[0.000]	[0.008]	[0.002]
Δy	0.464**	0.480**	0.388**	0.111	0.203	0.00589
	(0.127)	(0.108)	(0.115)	(0.111)	(0.140)	(0.122)
	[0.000]	[0.000]	[0.001]	[0.319]	[0.147]	[0.962]
Δp	-0.504**	-0.486**	-0.352**	-0.257*	-0.241+	-0.235*
	(0.134)	(0.110)	(0.130)	(0.119)	(0.134)	(0.109)
	[0.000]	[0.000]	[0.007]	[0.030]	[0.073]	[0.032]
$\Delta\widehat{\sigma}$	-0.392**	-0.104	-0.110	0.289*	0.330*	0.344**
	(0.149)	(0.166)	(0.163)	(0.138)	(0.156)	(0.127)
	[0.008]	[0.530]	[0.501]	[0.036]	[0.034]	[0.007]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.131	0.532	0.0883	-2.880	-2.949	-2.015
	(0.663)	(0.762)	(0.585)	(2.110)	(2.412)	(2.026)
	[0.843]	[0.485]	[0.880]	[0.172]	[0.222]	[0.320]
t			-0.0127**			0.00153
			(0.00402)			(0.00189)
			[0.002]			[0.420]
Poland						
EC term	-0.416**	-0.133*	-0.429**	-0.342**	-0.0370	-0.312*
	(0.118)	(0.0600)	(0.136)	(0.114)	(0.0358)	(0.131)
	[0.000]	[0.026]	[0.002]	[0.003]	[0.301]	[0.017]
Δy	-0.289*	-0.108	-0.256*	0.339	0.650*	0.358
	(0.130)	(0.131)	(0.130)	(0.274)	(0.308)	(0.302)
	[0.026]	[0.407]	[0.050]	[0.216]	[0.035]	[0.235]
Δp	0.306	0.192	0.344	0.399**	0.181	0.406**
	(0.200)	(0.232)	(0.214)	(0.138)	(0.151)	(0.157)
	[0.127]	[0.409]	[0.109]	[0.004]	[0.232]	[0.010]
$\Delta\widehat{\sigma}$	0.988**	1.140**	1.052**	0.587+	0.849*	0.643+
	(0.208)	(0.253)	(0.221)	(0.309)	(0.372)	(0.337)
	[0.000]	[0.000]	[0.000]	[0.057]	[0.022]	[0.056]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.774	0.0465	0.215	-3.091	-2.628	-3.598
	(0.783)	(0.923)	(0.822)	(2.469)	(3.249)	(2.674)
	[0.323]	[0.960]	[0.793]	[0.210]	[0.419]	[0.178]
t			-0.00127			-0.00145
			(0.00339)			(0.00239)
			[0.709]			[0.544]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Portugal						
EC term	-0.0402	-0.209**	-0.254*	-0.0706	-0.175**	-0.0600
	(0.172)	(0.0487)	(0.114)	(0.135)	(0.0272)	(0.0590)
	[0.815]	[0.000]	[0.026]	[0.600]	[0.000]	[0.309]
Δy	0.713**	0.223	0.245+	0.539	0.215	0.0249
	(0.157)	(0.150)	(0.137)	(0.404)	(0.210)	(0.246)
	[0.000]	[0.137]	[0.075]	[0.182]	[0.307]	[0.919]
Δp	-0.803**	-0.145	-0.204	-0.832**	-0.218	-0.323+
	(0.158)	(0.185)	(0.159)	(0.266)	(0.138)	(0.170)
	[0.000]	[0.435]	[0.198]	[0.002]	[0.115]	[0.058]
$\Delta\widehat{\sigma}$	-0.512**	0.283	0.0209	-0.184	0.260	0.147
	(0.188)	(0.235)	(0.172)	(0.299)	(0.175)	(0.183)
	[0.007]	[0.229]	[0.903]	[0.539]	[0.138]	[0.421]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	-0.523	0.295	-0.229	0.0261	2.656+	-0.646
	(0.717)	(0.604)	(0.465)	(2.971)	(1.553)	(1.777)
	[0.465]	[0.625]	[0.623]	[0.993]	[0.087]	[0.716]
t			-0.00728**			-0.00735**
			(0.00165)			(0.00127)
			[0.000]			[0.000]
Slovakia						
EC term	-0.575**	-0.147	-0.917**	-0.669**	-0.189+	-0.887**
	(0.197)	(0.106)	(0.217)	(0.214)	(0.0985)	(0.280)
	[0.004]	[0.166]	[0.000]	[0.002]	[0.055]	[0.002]
Δy	-0.0991	0.155	-0.260	-0.538	0.590	-1.205
	(0.233)	(0.243)	(0.210)	(0.645)	(0.647)	(0.925)
	[0.671]	[0.525]	[0.216]	[0.404]	[0.361]	[0.193]
Δp	-0.156	-0.468	0.0658	-0.0584	-0.146	-0.155
	(0.302)	(0.327)	(0.280)	(0.270)	(0.315)	(0.284)
	[0.605]	[0.153]	[0.814]	[0.829]	[0.643]	[0.584]
$\Delta\widehat{\sigma}$	0.577+	0.404	0.495+	-0.181	-0.477	-0.0233
	(0.328)	(0.369)	(0.277)	(0.480)	(0.537)	(0.480)
	[0.079]	[0.272]	[0.074]	[0.707]	[0.374]	[0.961]
$\Delta\left(\widehat{\sigma}\widehat{p}\right)$	0.236	2.592	1.035	2.791	0.186	4.486
	(2.080)	(2.430)	(1.921)	(2.131)	(2.359)	(2.762)
	[0.910]	[0.286]	[0.590]	[0.190]	[0.937]	[0.104]
t			0.0115*			0.0151+
			(0.00523)			(0.00783)
			[0.029]			[0.054]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Spain						
EC term	0.0182	-0.643**	-1.720**	0.102	-0.300*	-0.386*
	(0.0766)	(0.219)	(0.200)	(0.0631)	(0.134)	(0.194)
	[0.812]	[0.003]	[0.000]	[0.105]	[0.025]	[0.046]
Δy	0.169	-0.137	-0.302**	0.0863	0.623	1.089+
	(0.160)	(0.159)	(0.0897)	(0.543)	(0.518)	(0.604)
	[0.288]	[0.388]	[0.001]	[0.874]	[0.229]	[0.071]
Δp	-0.311	0.0458	0.204+	-0.157	-0.303	-0.0451
	(0.214)	(0.199)	(0.113)	(0.312)	(0.281)	(0.267)
	[0.147]	[0.818]	[0.071]	[0.613]	[0.280]	[0.866]
$\Delta\widehat{\sigma}$	-0.331	0.820+	1.069**	-0.0823	-0.0307	0.0330
	(0.320)	(0.447)	(0.242)	(0.443)	(0.410)	(0.378)
	[0.302]	[0.067]	[0.000]	[0.853]	[0.940]	[0.930]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.0726	2.363	2.587*	-1.094	9.539	0.601
	(1.207)	(1.701)	(1.101)	(6.998)	(7.807)	(6.003)
	[0.952]	[0.165]	[0.019]	[0.876]	[0.222]	[0.920]
t			-0.0558**			-0.0137**
			(0.00646)			(0.00523)
			[0.000]			[0.009]
Sweden						
EC term	0.0105	-0.0518*	-0.203**	-0.332+	-0.0323+	-0.466*
	(0.0921)	(0.0230)	(0.0710)	(0.169)	(0.0182)	(0.210)
	[0.909]	[0.024]	[0.004]	[0.050]	[0.076]	[0.026]
Δy	0.185+	0.141	0.0914	0.00603	-0.223	-0.216
	(0.102)	(0.0865)	(0.0776)	(0.248)	(0.267)	(0.249)
	[0.068]	[0.102]	[0.239]	[0.981]	[0.405]	[0.384]
Δp	-0.106	-0.0316	0.0177	-0.0587	-0.0968	0.0772
	(0.0922)	(0.0860)	(0.0768)	(0.107)	(0.100)	(0.121)
	[0.252]	[0.714]	[0.817]	[0.582]	[0.334]	[0.524]
$\Delta\widehat{\sigma}$	0.0889	0.256	0.265	-0.0162	-0.0551	-0.0289
	(0.216)	(0.201)	(0.170)	(0.183)	(0.179)	(0.171)
	[0.681]	[0.203]	[0.120]	[0.930]	[0.758]	[0.866]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.466	-0.236	0.153	0.385	0.455	-1.142
	(0.794)	(0.665)	(0.607)	(1.761)	(1.733)	(1.585)
	[0.557]	[0.723]	[0.801]	[0.827]	[0.793]	[0.471]
t			-0.00309**			-0.000878
			(0.000945)			(0.00136)
			[0.001]			[0.518]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
Switzerland						
EC term	-0.338**	-0.00812	-0.589**	-0.316+	-0.0652*	-0.267+
	(0.123)	(0.0239)	(0.160)	(0.177)	(0.0313)	(0.153)
	[0.006]	[0.734]	[0.000]	[0.074]	[0.037]	[0.081]
Δy	-0.0475	0.0874	-0.0963	0.0168	0.904+	0.741
	(0.0921)	(0.0924)	(0.0866)	(0.385)	(0.516)	(0.504)
	[0.606]	[0.345]	[0.266]	[0.965]	[0.080]	[0.142]
Δp	-0.186	-0.375**	-0.0702	-0.0467	-0.257+	-0.121
	(0.117)	(0.111)	(0.119)	(0.153)	(0.134)	(0.143)
	[0.113]	[0.001]	[0.555]	[0.759]	[0.054]	[0.399]
$\Delta\widehat{\sigma}$	-0.107	-0.296+	0.0892	-0.235	0.0279	-0.00775
	(0.150)	(0.177)	(0.170)	(0.252)	(0.278)	(0.259)
	[0.474]	[0.096]	[0.600]	[0.352]	[0.920]	[0.976]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	0.909	1.621	2.073*	-0.298	3.189	0.606
	(0.882)	(1.062)	(0.844)	(2.444)	(2.756)	(2.452)
	[0.303]	[0.127]	[0.014]	[0.903]	[0.247]	[0.805]
t			-0.00265*			-0.00391*
			(0.00118)			(0.00176)
			[0.024]			[0.026]
Turkey						
EC term	-0.0712	-0.441**	-0.345**	0.0170	-0.221*	-0.333**
	(0.0936)	(0.138)	(0.129)	(0.0986)	(0.0878)	(0.115)
	[0.447]	[0.001]	[0.008]	[0.863]	[0.012]	[0.004]
Δy	0.246+	0.0189	0.119	0.0535	0.0720	0.341
	(0.128)	(0.128)	(0.114)	(0.323)	(0.266)	(0.250)
	[0.054]	[0.882]	[0.297]	[0.868]	[0.787]	[0.174]
Δp	-0.183	0.212	0.104	-0.0538	0.0464	0.295
	(0.159)	(0.176)	(0.164)	(0.283)	(0.237)	(0.230)
	[0.249]	[0.227]	[0.523]	[0.849]	[0.845]	[0.199]
$\Delta\widehat{\sigma}$	0.124	0.427	0.265	-0.197	-0.233	-0.0379
	(0.358)	(0.315)	(0.298)	(0.446)	(0.375)	(0.335)
	[0.730]	[0.176]	[0.375]	[0.658]	[0.534]	[0.910]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-4.037**	-1.636	-1.535	-4.459	0.623	-1.650
	(1.427)	(1.562)	(1.381)	(3.419)	(3.361)	(2.657)
	[0.005]	[0.295]	[0.267]	[0.192]	[0.853]	[0.535]
t			-0.0193**			-0.0228**
			(0.00658)			(0.00619)
			[0.003]			[0.000]

	(1) USD No t	(2) USD Homogenous t	(3) USD Heterogeneous trends	(4) Local currency No t	(5) Local currency With t	(6) Local currency Heterogeneous trends
United Kingdom						
EC term	0.0921	-0.0638+	-0.0655	0.112+	-0.0711**	-0.0634
	(0.0637)	(0.0335)	(0.153)	(0.0587)	(0.0233)	(0.122)
	[0.148]	[0.057]	[0.669]	[0.056]	[0.002]	[0.605]
Δy	0.0532	-0.0441	-0.0272	0.515	0.707+	0.665
	(0.145)	(0.124)	(0.156)	(0.464)	(0.400)	(0.433)
	[0.713]	[0.723]	[0.862]	[0.266]	[0.077]	[0.125]
Δp	0.00828	0.0698	0.0653	-0.0459	-0.00489	0.00366
	(0.140)	(0.134)	(0.147)	(0.138)	(0.118)	(0.129)
	[0.953]	[0.601]	[0.657]	[0.739]	[0.967]	[0.977]
$\Delta\widehat{\sigma}$	-0.00265	0.107	0.0853	0.480	0.424	0.426
	(0.124)	(0.142)	(0.136)	(0.328)	(0.277)	(0.302)
	[0.983]	[0.451]	[0.531]	[0.144]	[0.126]	[0.158]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-1.158	-1.031	-1.087	-1.867	0.512	-1.430
	(0.773)	(0.749)	(0.735)	(3.101)	(2.827)	(2.847)
	[0.134]	[0.169]	[0.139]	[0.547]	[0.856]	[0.616]
t			-0.00301			-0.00322+
			(0.00228)			(0.00184)
			[0.186]			[0.080]
United States						
EC term	-0.0211	-0.00496	-0.0354+	-0.0235	-0.00488	-0.0756
	(0.0148)	(0.00519)	(0.0201)	(0.0153)	(0.00309)	(0.0532)
	[0.155]	[0.339]	[0.079]	[0.124]	[0.114]	[0.155]
Δy	0.466**	0.530**	0.405**	0.499**	0.534**	0.422*
	(0.145)	(0.139)	(0.155)	(0.140)	(0.136)	(0.166)
	[0.001]	[0.000]	[0.009]	[0.000]	[0.000]	[0.011]
Δp	-0.0108	-0.0171	-0.0126	-0.0129	-0.0226	0.00303
	(0.0227)	(0.0225)	(0.0216)	(0.0215)	(0.0198)	(0.0298)
	[0.634]	[0.447]	[0.560]	[0.547]	[0.255]	[0.919]
$\Delta\widehat{\sigma}$	0.0119	0.00159	0.00873	0.00738	0.0116	-0.00607
	(0.0376)	(0.0373)	(0.0355)	(0.0351)	(0.0358)	(0.0364)
	[0.751]	[0.966]	[0.806]	[0.833]	[0.745]	[0.868]
$\Delta\left(\widehat{\sigma}\widetilde{p}\right)$	-0.320*	-0.330*	-0.300*	-0.266+	-0.174	-0.485*
	(0.133)	(0.137)	(0.131)	(0.146)	(0.176)	(0.193)
	[0.016]	[0.016]	[0.022]	[0.069]	[0.324]	[0.012]
t			0.000553			0.000910
			(0.000555)			(0.000786)
			[0.319]			[0.247]

Standard errors in parentheses; p-values in brackets.
** p<0.01, * p<0.05, + p<0.1

References

[1] Ahn, Seung C., and Peter Schmidt. 1995. Efficient estimation of models for dynamic panel data. Journal of Econometrics 68: 5-27.

[2] Akinboade, Olude A., Emmanuel Ziramba, and Wolassa L. Kumo. 2008. The demand for gasoline in South Africa: An empirical analysis using co-integration techniques. Energy Economics 30: 3222-3229.

[3] Alves, Denisard C.O., and Rodrigo De Losso da Silveira Bueno. 2003. Short-run, long-run and cross elasticities of gasoline demand in Brazil. Energy Economics 25:191–199.

[4] Anderson, T. W., and Cheng Hsiao. 1981. Estimation of dynamic models with error components. Journal of the American Statistical Association 76(375): 598-606.

[5] Angelier, Jean Pierre, and Thomas Sterner. 1990. Tax harmonization for petroleum products in the EC. Energy Policy 18(6): 500-505.

[6] Arellano, Manuel, and Steven Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. The Review of Economic Studies 58(2): 277-297.

[7] Arellano, Manuel, and Olympia Bover. 1995. Another look at the instrumental variable estimation of error-components models. Journal of Econometrics 68 : 29-51.

[8] Baltagi, Badi H., and James M. Griffin. 1983. Gasoline demand in the OECD application of pooling and testing procedures. European Economic Review 22(2): 117-137.

[9] Baltagi, Badi H., and James M. Griffin. 1997. Pooled estimators vs. their heterogeneous counterparts in the context of dynamic demand for gasoline. Journal of Econometrics 77(2): 303-327.

[10] Baltagi, Badi H., and Chihwa Kao. Nonstationary panels, cointegration in panels and dynamic panels: A Survey. 2000. Advances in Econometrics, Vol. 25: Nonstationary Panels, Panel Cointegration, and Dynamic Panels (ed. Badi H. Baltagi).New York: Elsevier.

[11] Banerjee, Anindya, Juan J. Dolado, and Ricardo Mestre. 1998. Error-Correction Mechanism Tests for Cointegration in a Single-Equation Framework. Journal of Time Series Analysis 19(3): 267-283.

[12] Bentzen, Jan. 1994. An empirical analysis of gasoline demand in Denmark using cointegration techniques. Energy Economics 16(2): 139-143.

[13] Bentzen, Jan, and Tom Engsted. 2001. A revival of the autoregressive distributed lag model in estimating energy demand relationships. Energy 26: 45-55.

[14] Blackburne, Edward F., and Mark W. Frank. 2007a. Estimation of nonstationary heterogeneous panels. Stata Journal 7(2): 197-208.

[15] Blackburne, Edward F., and Mark W. Frank. 2007b. XTPMG: Stata module for estimation of nonstationary heterogeneous panels. Statistical Software Components S456868, Boston College Department of Economics.

[16] Blundell, Richard, Steve Bond, and Frank Windmeijer. 2000. Estimation in dynamic panel data models: Improving on the performance of the standard GMM estimator. IFS Working Papers W00/12, Institute for Fiscal Studies.

[17] Blundell, Richard, and Steven Bond. 1998. Initial conditions and moment restrictions in dynamic panel data models. Journal of Econometrics 87: 115-143.

[18] Bond, Stephen R. 2002. Dynamic panel data models: a guide to micro data methods and practice. Portuguese Economic Journal 1: 141-162.

[19] Boulter, P.G., and S. Latham. 2009. Emission factors 2009: Report 5 - a review of the effects of fuel properties on road vehicle emissions. Published Project Report 358. TRL Limited. Prepared for UK Dept. for Transport, Cleaner Fuels, & Vehicles. 20 March 2011 <http://www.dft.gov.uk/pgr/roads/environment/emissions/report-5.pdf>.

[20] Breunig, Robert V., and Carol Gisz. 2009. An Exploration of Australian Petrol Demand: Unobservable Habits, Irreversibility and Some Updated Estimates. Economic Record 85(268): 73-91.

[21] Brons, Martijn, et al. 2008. A meta-analysis of the price elasticity of gasoline demand. A SUR approach. Energy Economics 30: 2105–2122.

[22] Cheung, Kui-yin, and Elspeth Thomson. 2004. The Demand for Gasoline in China: A Cointegration Analysis. Journal of Applied Statistics 31(5): 533-544.

[23] Choi, Chi-Young, Nelson C. Mark, and Donggyu Sul. 2010. Bias Reduction in Dynamic Panel Data Models by Common Recursive Mean Adjustment. Oxford Bulletin of Economics and Statistics 72(5): 567–599.

[24] Choi, In. 2001. Unit root tests for panel data. Journal of International Money and Finance 20: 249-272.

[25] Coppejans, Mark, Donna Gilleskie, Holger Sieg, and Koleman Strumpf. 2007. Consumer demand under price uncertainty: Empirical evidence from the market for cigarettes. The Review of Economics and Statistics 89(3): 510-521.

[26] Dahl, Carol, and Thomas Sterner. 1991. Analysing gasoline demand elasticities: A survey. Energy Economics 13(3): 203-210.

[27] Dahl, Carol. 2011. Measuring global gasoline and diesel price and income elasticities. Energy Policy: corrected proof.

[28] De Vita, G., K. Endresen, and L.C. Hunt. 2006. An empirical analysis of energy demand in Namibia. Energy Policy 34: 3447-3463.

[29] Eltony, M.N., and N.H. Al-Mutairi. 1995. Demand for gasoline in Kuwait: An empirical analysis using cointegration techniques. Energy Economics 17(3): 249-253.

[30] Engle, Robert F., and C. W. J. Granger. 1987. Co-Integration and Error Correction: Representation, Estimation, and Testing. Econometrica 55(2): 251-276.

[31] Espey, Molly. 1998. Gasoline demand revisited: An international meta-analysis of elasticities. Energy Economics 20(3): 273-95.

[32] Goodwin, P. B. 1992. A review of new demand elasticities with special reference to short and long run effects of price changes. Journal of Transport Economics and Policy (May 1992): 155-169.

[33] Goodwin, Phil, Joyce Dargay, and Mark Hanly. 2004. Elasticities of Road Traffic and Fuel Consumption with Respect to Price and Income: A Review. Transport Reviews 24(3): 275-292.

[34] Graham, Daniel J., and Stephen Glaister. 2004. Road Traffic Demand Elasticity Estimates: A Review. Transport Reviews 24(3): 261–274.

[35] Harris, Richard D.F., and Elias Tzavalis. 1999. Inference for unit roots in dynamic panels where the time dimension is fixed. Journal of Econometrics 91: 201-226.

[36] Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. 1988. Estimating vector autoregressions with panel data. Econometrica 56(6): 1371-95.

[37] 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.

[38] Im, Kyung So, M. Hashem Pesaran, and Yongcheol Shin. 2003. Testing for unit roots in heterogeneous panels. Journal of Econometrics 115: 53-74.

[39] Johansson, Olof, and Lee Schipper. 1997. Measuring the Long-Run Fuel Demand of Cars: Separate Estimations of Vehicle Stock, Mean Fuel Intensity, and Mean Annual Driving Distance. Journal of Transport Economics and Policy 31(3): 277-292.

[40] Judson, Ruth A., Richard Schmalensee, and Thomas M. Stoker. 1999. Economic development and the structure of the demand for commercial energy. Energy Journal 20(2): 29-57.

[41] Kao, Chihwa. 1999. Spurious regression and residual-based tests for cointegration in panel data. Journal of Econometrics 90: 1-44.

[42] Kilian, Lutz. 2010. Explaining Fluctuations in Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market. The Energy Journal 31(2): 103-128.

[43] Kleibergen, Frank, and Richard Paap. 2006. Generalized reduced rank tests using the singular value decomposition. Journal of Econometrics 133: 97–126

[44] Krichene, Noureddine. 2002. World crude oil and natural gas: a demand and supply model. Energy Economics 24: 557-576.

[45] Levin, Andrew, Chien-Fu Lin, and Chia-Shang James Chu. 2002. Unit root tests in panel data:asymptotic and finite-sample properties. Journal of Econometrics 108: 1-24.

[46] Liu, Gang. 2004. Estimating Energy Demand Elasticities for OECD Countries: A Dynamic Panel Approach. Statistics Norway Resarch Department Discussion Papers No. 373, March 2004.

[47] Maddala, G. S. , and Shaowen Wu. 1999. A comparative study of unit root tests with panel data and a new simple test. Oxford Bulletin of Economics and Statistics, Special Issue, 1999: 631-652.

[48] MotorMouth. 2010. "Fuel Types." 8 Dec 2010 <http://motormouth.com.au/myresources/fueltypes.aspx>.

[49] Nadaud, F. 2004. The demand for car gasoline in France: a long run econometric perspective, 6th IAEE European Energy Conference, 2–3 September, Zurich, Switzerland.

[50] Narayan, Paresh Kumar, and Russell Smyth. 2007. A panel cointegration analysis of the demand for oil in the Middle East. Energy Policy 35: 6258–6265.

[51] Nguyen-Van, Phu. 2010. Energy consumption and income: A semiparametric panel data analysis. Energy Economics 32: 557–563.

[52] Nickell, Stephen. 1981. Biases in Dynamic Models with Fixed Effects. Econometrica 49(6): 1417-1426.

[53] Pedroni, Peter. 2004. Panel cointegration: Asymptotic and finite sample properties of pooled time series tests with an application to the PPP hypothesis. Econometric Theory 20: 597–625.

[54] Persyn, Damiaan, and Joakim Westerlund. 2008. Error Correction Based Cointegration Tests for Panel Data. Stata Journal 8(2): 232-241.

[55] Pesaran, M. Hashem, and Yongcheol Shin. 1999. An autoregressive distributed lag modelling approach to cointegration analysis. Chapter 11 in Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial Symposium, Strom S (ed). Cambidge, UK: Cambridge UP. Accessed version revised January 1997 at <http://www.econ.cam.ac.uk/faculty/pesaran/ardl.pdf>, 24 February 2011.

[56] Pesaran, M. Hashem, Yongcheol Shin, and Richard J. Smith. 1999. Pooled Mean Group Estimation of Dynamic Heterogeneous Panels. Journal of the American Statistical Association 94(446): 621-634.

[57] Pesaran, M. Hashem, and Ron Smith. 1995. Estimating long-run relationships from dynamic heterogenous panels. Journal of Econometrics 68: 79-113.

[58] Pesaran, M. Hashem, and Zhongyun Zhao. 1999. Bias Reduction in Estimating Long- Run Relationships From Dynamic Heterogeneous Panels. In Analysis of Panels and Limited Dependent Variables: A Volume in Honour of G. S. Maddala, eds. Cheng Hsiao et al. Cambridge, UK: Cambridge UP.

[59] Phillips, Peter C. B., and Pierre Perron. 1988. Testing for a Unit Root in Time Series Regression. Biometrika 75(2): 335-346.

[60] Polemis, Michael L. 2006. Empirical assessment of the determinants of road energy demand in Greece. Energy Economics 28: 385-403.

[61] Ramanathan, R. 1999. Short- and long-run elasticities of gasoline demand in India: An empirical analysis using cointegration techniques. Energy Economics 21: 321-330.

[62] Ramanathan, R. and G. Subramanian, G. 2003. An empirical analysis of gasoline demand in the Sultanate of Oman using cointegration techniques. Pacific and Asian Journal of Energy 13: 33–41.

[63] Rao, B. Bhaskara, and Gyaneshwar Rao. 2009. Cointegration and the demand for gasoline. Energy Policy 37: 3978-3983.

[64] Roodman, David. 2006. How to do xtabond2: An introduction to "Difference" and "System" GMM in Stata. Center for Global Development, Working Paper No. 103. Available from http://www.cgdev.org/files/11619_file_HowtoDoxtabond6_12_1_06.pdf.

[65] Roodman, David. 2007. A short note on the theme of too many instruments. Center for Global Development, Working Paper No. 125. Available from http://www.cgdev.org/content/publications/detail/14256.

[66] Roodman, David. 2011 (first version 2003). XTABOND2: Stata module to extend xtabond dynamic panel data estimator. Center for Global Development, Washington. http://econpapers.repec.org/software/bocbocode/s435901.htm.

[67] Samimi, Rodney. 1995. Road transport energy demand in Australia: A cointegration approach. Energy Economics 17(4): 329-339.

[68] Schaffer, Mark E. 2010. xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. Available from http://ideas.repec.org/c/boc/bocode/s456501.html.

[69] Stata. 2009. Stata Longitudinal-Data/Panel-Data Reference Manual, Stata User's Guide, Release 11. College Station, TX: Stata Press. 523-522.

[70] StataCorp. 2009. Stata Statistical Software: Release 11. College Station, TX: StataCorp LP.

[71] Storchmann, Karl. 2005. Long-run gasoline demand for passenger cars: the role of income distribution. Energy Economics 27(1): 25– 58.

[72] Wadud, Zia, Daniel J. Graham, and Robert B. Noland. 2009. A cointegration analysis of gasoline demand in the United States. Applied Economics 41(26): 3327-3336.

[73] Westerlund, Joakim. 2007. Testing for Error Correction in Panel Data. Oxford Bulletin of Economics and Statistics 69(6): 709-748.

[74] Windmeijer, Frank. 2005. A finite sample correction for the variance of linear efficient two-step GMM estimators. Journal of Econometrics 126(1): 25-51.

Data

[75] International Energy Agency (IEA). 2009. Energy Prices and Taxes (2009Q4). ESDS International, (Mimas) University of Manchester.

[76] International Energy Agency (IEA). 2009. Oil Information (2009). ESDS International, (Mimas) University of Manchester.

[77] United Nations, Department of Economic and Social Affairs, Population Division. 2009. World Population Prospects: The 2008 Revision, CD-ROM Edition. 9 March 2010 <http://esa.un.org/unpd/wpp2008/peps_stock-indicators.htm>.

[78] World Bank. 2009. World Development Indicators (WDI) (2009). ESDS International, (Mimas) University of Manchester.

Conclusion

Gasoline demand is sluggish: on this point, study after study agrees. This thesis has asked whether some of the observed sluggishness is not an inherent property of gasoline demand itself, but a product of the environment in which demand is estimated—an environment in which prices are uncertain and price changes are short-lived. The distinction is important because policy instruments have the potential to change that environment. When governments intervene in the market, gasoline demand may not, in fact, be so sluggish.

That the market environment matters to demand behavior is a major implication of rational habits models. I have shown that the price process of a habit-forming good is particularly important when consumers are forward-looking: given rational habits, consumers will respond *now* to their expectations of future prices, and they will allow price volatility to temper their demand and their price responsiveness.

Empirically, I have exploited both the future-price and the price-volatility implications of rational habits—but I have treated the implications separately. In the U.S. data of Chapter 1, I examined consumers' reactions to their price expectations; in the international data of Chapter 3, I examined consumers' sensitivity to price volatility. In the future, if data of the required scope and detail could be obtained, it might be interesting to combine these empirical approaches and test for both implications simultaneously. One way to combine the approaches would be to augment a traditional demand regression with (expected) future prices, a measure of price uncertainty or volatility, and interactions between this measure and both present and (expected) future prices. Even given a very large and detailed data set, however, it would likely be difficult to identify the separate effects of all these regressors. A more workable tactic might be to build a separate model of gasoline prices, using this to generate explicit (quasi-)forecasts that could proxy for consumers' expectations of gasoline prices. These quasi-forecasts could then be added to a demand regression, along with a measure of average or expected forecasting error and interactions between this forecasting error and current and forecasted prices. This approach has two advantages: the separate price-forecasting model would reduce the identification burden of the demand regression, and the

use of forecasting error in lieu of price volatility might provide a better gauge of how well consumers are able to anticipate future prices.

For now, however, the future-price and price-volatility implications have provided me with two separate empirical models, and each has turned up some evidence of rational habits in gasoline demand. In the U.S., current gasoline consumption appears to respond to anticipated future prices, particularly if future prices are broken down into tax and before-tax components. Around the world, gasoline consumption and consumption's price-responsiveness both appear to drop with increasing price volatility. All these findings fit neatly into a story of rational habits.

These findings also suggest that the apparent sluggishness of gasoline demand is the product, in part, of a fluctuating and uncertain gasoline market. Policies that change this market will also change the apparent shape of gasoline demand, and traditionally-measured price elasticity will underproject responses to policy interventions. Gasoline consumption and its accompanying emissions may have a huge inertia, but credibly-permanent price increases are a powerful brake.