UC Berkeley

CUDARE Working Papers

Title

Demand and Price Uncertainty: Rational Habits in International Gasoline Demand

Permalink

<https://escholarship.org/uc/item/25q4w08n>

Author

Scott, K. Rebecca

Publication Date

2013-03-14

University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

Year 2013

Paper 1131

Demand and Price Uncertainty:
Rational Habits in International Gasoline
Demand

K. Rebecca Scott

Demand and Price Uncertainty: Rational Habits in International Gasoline Demand

K. Rebecca Scott

14 March 2013

Key Words: gasoline demand, rational habits, price elasticity

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

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 FE 2SLS, system GMM, and pooled mean group (PMG) estimators are considered, and my preferred estimators provide evidence of rational habits in a panel of 29 countries for the years 1990-2011. Such habits may help to explain some of the cross-country and cross-time variation in 'total' price elasticity. These habits also imply that the effect of price uncertainty must be taken into account when projecting the impacts of potential policies on gasoline consumption.

1 Introduction

Environmental concerns have driven widespread policy interest in reducing gasoline consumption. With this interest in cutting demand comes a very practical interest in the consumer behavior that shapes not only demand, but the effectiveness of various policies in controlling it.

Although this interest has generated a great deal of empirical work measuring price and income elasticities, less attention has been devoted to the deeper behaviors underlying demand. Consumers may pay for fuel at the pump, but in fact they decide their purchases everywhere but the gas station—in discrete, infrequent choices about where to live in relation to work and what kind of vehicle to buy; and in nearly-continuous choices about daily life: carpool routines, non-essential travel, 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 response 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.

2 Models

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 Scott (2010). 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.

Another way to test for rational habits would be to follow the approach of Scott (2012) and test 1. whether consumers respond to anticipated future prices and 2. whether consumers respond differently to the tax and before-tax components of the gasoline price. Scott (2012) applied this type of model to US gasoline demand, estimating models of the form

$$g_{it} = \lambda g_{i,t-1} + \beta_1 y_{it} + \beta_2 p_{it} + \beta_3 p_{i,t+1} + \beta_4 t + \mu_i + \varepsilon_{it} \quad (\text{lead price model})$$

$$g_{it} = \lambda g_{i,t-1} + \beta_1 y_{it} + \beta_2 p_{it} + \beta_5 \ln(1 + \tau_{i,t+1}) + \beta_6 \ln \text{BeforeTax}_{i,t+1} + \beta_4 t + \mu_i + \varepsilon_{it} \quad (\text{lead price components model})$$

$$g_{it} = \lambda g_{i,t-1} + \beta_1 y_{it} + \beta_2 p_{it} + \beta_7 \ln(1 + \tau_{it}) + \beta_8 \ln \text{BeforeTax}_{it} + \beta_4 t + \mu_i + \varepsilon_{it} \quad (\text{price components model})$$

where g is log gasoline consumption per capita, y is log real per capita income, p is the real log gasoline price, BeforeTax is the (level) real price before taxes, τ is the ratio of the tax to BeforeTax , and i and t designate state and year. Unfortunately, these models do not apply very well to the international case. First, as in the US case, anticipated lead prices and price components are difficult to identify. Second, unlike the US case (in which nearly all gasoline taxes are specific), the tax and before-tax components of the price do not move so differently over time (see Figure 1). This is because gasoline in most OECD countries is subject to a VAT of 10 to 25% in addition to a specific tax.¹ This VAT, which on average comprises a quarter of the total price, moves in lock-step with the before-tax price, and therefore pushes the total-tax path towards the before-tax path. The more similar the behavior of taxes and the before-tax price, the less reason consumers have to respond to them differently.

Table 6 presents estimates of the (lead) price (components) models for the international data set I will be examining in this paper, performed using the same FE 2SLS approach described in Scott (2012).² In the lead price and lead price components models, the coefficients on the 'habit' variables $p_{i,t+1}$, $\ln(1 + \tau_{t+1})$, and $\ln \text{BeforeTax}_{t+1}$ are all statistically insignificant. In the price components model, the coefficients on the tax and before-tax components are statistically indistinguishable. If we imagine two 10% increases in the price, one driven by taxes and one driven by the before-tax price, their effects on consumption would be similar.

This style of model therefore does not yield evidence of rational habits in the international context—and more importantly, it is likely *incapable* of capturing rational habits even if they exist. To explore whether rational habits shape gasoline demand internationally, we will need a different model.

¹IEA (2012), Country Notes.

²This approach treats lead prices as anticipated lead by prices by instrumenting them using current information. The instrumentation strategy, including the definition of TaxIV_{it} , is identical to Scott (2012).

Figure 1: US vs. world-average gasoline prices and taxes

Fortunately, 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 Scott (2010), 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 Scott (2010), 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 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

$$g_{it} = \lambda g_{i,t-1} + \delta_1 y_{it} + \delta_2 p_{it} + \delta_3 \hat{\sigma}_{it} + \delta_4 \hat{\sigma}_{it} p_{it} + \delta_5(i)t + \mu_i + \varepsilon_{it} \quad (1)$$

where $\hat{\sigma}$ is a measure of price uncertainty or forecastability. 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.

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

$$\frac{\partial g_{it}}{\partial p_{it}} = \delta_2 + \delta_4 \hat{\sigma}_{it} \quad (2)$$

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

Although the model given by (1) assumes that coefficients are homogeneous across individuals, this may be a shaky assumption when the individuals are countries with different infrastructure, different terrain, and different socioeconomic foundations. Indeed, coefficient homogeneity seems less plausible across countries than it might be within an individual country. Since falsely imposing homogeneity may, as I discuss in Section 5, lead to biases when estimating, I will also estimate models that relax this assumption. Relaxing coefficient homogeneity is most convenient if we transform the ADL(1,0) specification of (1) into an error-correction model, equivalent to an ADL(1,1). To begin the transformation, re-write (1) to allow all coefficients to vary by country, then subtract $g_{i,t-1}$ from each side. Adding and subtracting $\delta_1 y_{i,t-1}$, $\delta_2 p_{i,t-1}$, $\delta_3 \hat{\sigma}_{i,t-1}$, and $\delta_4 (\hat{\sigma}p)_{i,t-1}$ from the right-hand side then yields

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

which can be factored into

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

Equation (4) is an ECM, and the new coefficient ϕ_i is the error-correction term. ϕ_i 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. Section (5.3) will discuss various homogeneity restrictions that can be imposed on θ and δ . For now, note that if no coefficient homogeneity restrictions are imposed, the ECM can just as easily be factored into ADL(1,1) form:

$$\begin{aligned} g_{it} = & (\phi_i + 1) g_{i,t-1} + \delta_{1i} y_{it} - (\phi_i \theta_{1i} + \delta_{1i}) y_{i,t-1} + \delta_{2i} p_{it} - (\phi_i \theta_{2i} + \delta_{2i}) p_{i,t-1} \\ & + \delta_{3i} \hat{\sigma}_{it} - (\phi_i \theta_{3i} + \delta_{3i}) \hat{\sigma}_{i,t-1} + \delta_{4i} (\hat{\sigma}p)_{it} - (\phi_i \theta_{4i} + \delta_{4i}) (\hat{\sigma}p)_{i,t-1} + \mu_i + \varepsilon_{it} \end{aligned} \quad (5)$$

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.

The difficulties in estimating (1) and (4) will be discussed in Section (5), where I weigh the advantages and disadvantages of three approaches: least squares estimation of the ADL model, system GMM estimation of the ADL model, and maximum-likelihood estimation of the ECM. My preferences are for FE 2SLS of the ADL model or a pooled mean group estimator of the ECM. To be transparent about the sensitivity of my results to various modelling and estimation choices, however—and to illustrate the potential pitfalls of certain approaches—I report the results of all these estimators. For ease of comparison

Table 1: Average elasticities in gasoline demand literature reviews

Study	Average price elasticity		Average income elasticity	
	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		

^aLagged endogenous models, annual. ^bTime series models^cPrice: time series models; income: dynamic models

with the literature, I also estimate a ‘standard’, non-habits version of both the ADL and error-correction specifications. Before concluding, finally, I use country-specific price uncertainty measurements to interpolate country-specific price elasticities, and I consider the extent to which cross-country variation in price uncertainty predicts cross-country variation in elasticity.

3 Gasoline Demand Literature Review

Rational habits models were first applied to gasoline demand in Scott (2012), which explores rational habits theory and shows that US consumers respond more strongly to tax- than market-driven price changes. A subsequent working paper by Li, Linn and Muehlegger (2012) employs a similar empirical methodology and corroborates this differential in responsiveness to tax and non-tax components of the gasoline price. Nor is this difference confined to the US: Rivers and Schaufele (2012) find that British Columbia’s carbon tax provoked a greater reduction in gasoline demand than a comparable non-tax price increase.

Despite this recent attention to the components of prices, however, the predominant focus in the gasoline demand literature has been on measuring simple income and price elasticities. Studies that estimate these elasticities abound: in a 1998 meta-analysis, for instance, Espey (1998) considers 363 short- or medium-run and 277 long-run price elasticity estimates. There have been so many studies that simultaneous literature reviews have even managed, as Goodwin noted in 1992, 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 agree that average price elasticity lies around -0.25 to -0.3 in the short run and -0.6 to -0.8 in the long run, while average income elasticity is about 0.4 in the short run and somewhere around unit elastic in the long run. Some of these averages are summarized in Table 1.

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 most of these studies have looked at single countries (or the world as a whole) rather than panels, and their estimates are summarized in Table 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-

Table 2: Gasoline demand studies based on cointegration and error-correction models

Study	Country	Price elasticity		Income elasticity	
		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
Liddle (2012)	14 OECD countries	-0.155	-0.429	0.283	0.344
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)	Fiji		-0.159		0.427
			to -0.244		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

^aElasticities with respect to vehicles per capita substituted for income elasticities.

^b1990q1-2002q4

^cEstimates elasticity of demand for crude oil, not gasoline. Results for 1973-1999.

^dAs reported in Wadud, Graham, and Noland (2009).

^eEnergy for transport, not just gasoline.

^fSingle-step nonlinear least squares, post-1978.

^gReported short- and long-run elasticities in fact the same.

run price elasticities (Espey 1998) might also explain some of the difference. In Section 6 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 5.

The large variation in elasticity estimates across studies cries out to be explained. Part of the explanation may be methodological. Studies based on a panel of countries, for instance, tend to yield elasticities that are similar to single-country studies' elasticities in the short run but of higher magnitude in the long run (Espey 1998). Elasticities are also affected by the inclusion of controls for vehicle ownership and/or characteristics of the vehicle stock (Dahl and Sterner 1991, Espey 1998). Another part of the explanation is that the true underlying elasticities vary across both place and time. Variation across place is noted by Espey (1998), who observes short-run price responsiveness to be relatively low in the US and relatively high in Europe. Variation over time is demonstrated by Hughes et al. (2006), who find that US short-run price elasticities have shifted since the 1970s; and by Neto (2012), who finds that price elasticity has fluctuated over time in Switzerland. The reasons *behind* these cross-sectional and cross-time variations remain to be explored, however; and one explanation may be found in price uncertainty's variations across time and space.

4 Data and Specification of Variables

4.1 Data

My data set consists of a panel of 29 countries for the period 1990-2011. As price data for some countries is limited, the actual series length ranges from twelve to twenty-one years, with an average of 20.2 years. A list of included countries is provided in Appendix 9.

In order to focus as much as possible on passenger vehicles rather than freight, data is isolated to gasoline, not diesel. Information on gasoline consumption³ is taken from the International Energy Agency (IEA)’s Oil Information (2012) and transformed into per capita terms using annual population estimates from the OECD’s Employment and Labour Market Statistics (2011). 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* (2012Q4). 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 9.

As a measure of income, I take GDP from the World Bank’s *World Development Indicators* 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* (2012Q4). I concentrate on this stream because the Brent is used extensively as a pricing benchmark.⁴ Prices and income are converted from nominal to real terms using country-specific CPIs taken from the IEA’s *Energy Prices and Taxes* (2012Q4).

Further specifics about my data and data sources are provided in Appendix 9.

4.2 Measuring Price Uncertainty

In order to begin examining models (1) and (4), I need some measure of (future) price uncertainty. I will consider two different measures of uncertainty, one that focuses on price volatility and one that focuses on price forecastability.

Price volatility may be taken as a crude proxy for consumers’ uncertainty about the future because it implies that future prices may be hard to predict. It may also suggest to consumers, moreover, that current price changes will be short-lived. The simplest measure of volatility would 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 have two drawbacks: first, it would imply that consumers had definite information about future prices; and second, it would make the identification of price volatility impossible without imposing 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 important volatility within each year—in the extreme case, countries with identical annual price series could have wildly different price paths from month to month or quarter to quarter. To mitigate this problem and more fully capture volatility, 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.

³This series (‘motor gasoline demand’) explicitly excludes aviation gasoline.

⁴The 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.

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

$$\begin{aligned}\hat{\sigma}_{iq} &= \sqrt{\frac{1}{x} \sum_{j=1}^x (p_{i,q-j} - \bar{p}_{iq})^2}, \text{ where} \\ \bar{p}_{iq} &= \frac{1}{x} \sum_{j=1}^x p_{i,q-j+1}\end{aligned}\tag{6}$$

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 $\hat{\sigma}_{iq}$ over the quarters in each year⁵ or, of course, lengthen the window of the rolling standard deviation beyond $x = 4$.

A more pertinent measure of uncertainty may be one that focuses specifically on the predictability or forecastability of prices. To construct such a measure, I turn again to quarterly price data. As with the volatility measure, I avoid implying consumers have certain future information by using a rolling measure, in this case an expanding-window forecast based on the simple model

$$p_{iq} = \beta_q p_{i,q-1} + \mu_i + \varepsilon_{iq}\tag{7}$$

β_q is estimated for every quarter q using a standard fixed-effects regression with data up to the contemporary quarter q . Note that although β_q is quarter-specific, it is not country-specific, as the price series do not extend far enough back to create meaningful estimates of β_q by 1990 without exploiting cross-sectional variation. The relatively short length of the time series also makes it expedient to use a fixed-effects (within) estimator despite its bias here, as difference/system GMM would cost us some observations early in the sample. After estimation, $\hat{\beta}_q$ is used to generate one-period-ahead forecasts:

$$\hat{p}_{i,q+1} = \hat{\beta}_q p_{iq} + \hat{\mu}_i\tag{8}$$

and these forecasts are used to calculate absolute forecast error, again denoted $\hat{\sigma}_{iq}$:

$$\hat{\sigma}_{iq} = |\hat{p}_{iq} - p_{iq}|\tag{9}$$

These quarterly absolute forecast errors are then averaged over each year to produce the annual average absolute forecast error, which I use as my measure of price (un)forecastability.

An alternative to these measures, were suitable data available, would be to model the evolution of price volatility over time and incorporate price-volatility *forecasts* into the demand model. This would be particularly worthwhile as the theory developed in Scott (2010) posits a relationship between demand and expected *future* price uncertainty, not current volatility: the rolling-uncertainty measures used at present function as proxies for expected future uncertainty. A better volatility forecast might be possible, particularly if gasoline prices follow an ARCH-type process and tend to move between periods of high and low volatility. The data at hand, unfortunately, are not suitable for such an exercise.

One exercise that *is* possible now is to use before-tax prices rather than total prices to calculate the rolling standard deviation measure. This alternative is of interest because a specific tax hike may

⁵The chief difference of this alternative method lies in the mean price, $\bar{p}_{iq}$, 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.

contribute to volatility in the year it takes effect but actually signal a *reduction* in price uncertainty, as the tax is likely to be a relatively stable component of the price. I explore this possibility by removing excise taxes (but not value-added taxes) from the log prices used to measure volatility in (6). This does not improve my results, however, so I conclude that tax-driven price volatility is not a significant issue.

4.3 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 would be to transpose all variables into a common currency. This is the approach towards which convention in the gasoline and energy demand literature leans,⁶ and it has the advantage of not introducing restrictions on the type of estimator that can be deployed. The downside of the common-currency approach is that it would not capture locally-perceived price volatility very well: quarterly prices would have to be converted using nominal exchange rates, as PPP exchange rates are not available with greater than annual frequency. Gasoline prices are driven by oil prices, moreover, and oil is generally a net import for these countries; and so the local price of gasoline will fluctuate along with the nominal exchange rate. Converting local prices back into a common currency using nominal exchange rates would strip away this exchange-rate-driven price volatility.

Provided care is taken in the choice of estimator, however, another option is to work in log real local currency. 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 can 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-measured 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 5.3 removes $-\ln r_i$ by

⁶See, 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).

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. As I will be using estimators that de-mean or difference, I opt to keep price and income in local currencies.

As an aside, it should be noted that the inclusion or exclusion of $\ln r_i$ does not influence the price volatility measure (rolling standard deviation), as standard deviation is unaffected by a constant. Similarly, the price unforecastability measure is not affected by $\ln r_i$, as $\ln r_i$ is part of the country-specific effect that disappears when the realized price and the forecast are differenced.

4.4 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 14.8 per person (Turkey, 2011) to a high of 1352.6 kg per person (Luxembourg, 1994), with an overall mean of 372.5 kg per person. Over time, per-capita consumption has been slightly decreasing in most countries, and the 1990 average of 374.8 kg per person has fallen to a 2011 average of 309.1 kg. Figure 3 plots the path of 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 2) and much of the variation over time driven by fluctuations in the crude oil price. The overall mean real price in the sample is 1.17 Y2005-USD per liter, and country-averaged prices range from 0.49 Y2005-USD per liter in the United States to 1.58 Y2005-USD per liter in Norway. The price uncertainty measures, too, have both cross-sectional and cross-time variation that should allow us to identify their effects. The price volatility measure, as defined in (6), has an average over all observations of 0.035 and a comparatively-high standard deviation of 0.029. Country-specific mean volatilities range from a low of 1.4% for Mexico to a high of 6.3% for the United States, and annual average volatilities range from a low of 1.6% in 1997 to 7.9% in 2008. Meanwhile the price unforecastability measure, as defined in (9), has country-specific means ranging from 1.7% (Poland) to 59% (Turkey) and annual means varying from 7.7% (2008) to 24% (1992). The time paths of the gasoline price, price volatility, and price unforecastability are plotted in Figures 4, 5, and 6, respectively.

Figure 2: Mean gasoline prices, by tax and non-tax components

Figure 3: Gasoline consumption per capita, by country

Figure 4: Real gasoline prices, by country

Figure 5: Gasoline price volatility, by country

Figure 6: Gasoline price forecastability, by country

5 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 and Bover (1995) and Blundell and Bond's (1998) system GMM estimator. Finally I turn to the 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 (MG); Pesaran, Shin, and Smith's (1999) pooled mean group method (PMG); and a dynamic fixed effects method (DFE). Each of these ways of specifying and estimating the model has its advantages and trade-offs, but, as I will discuss in Section 6, the weaknesses of system GMM are particularly troublesome in this application. My preferences are for the least squares estimates of the ADL and the PMG estimates of the error-correction model.

5.1 Least Squares

The first approach I consider is estimating (1) using a simple within-groups (fixed-effects) estimator with outside instruments for prices.

The major strength of this approach is that it allows us to address the potential endogeneity of prices. Prices are likely to be endogenous because the gasoline supply schedule faced by a nation is not flat: any positive demand shock will drive prices up, and vice versa. Causality, that is, runs *from* consumption to prices as well as from prices to consumption; and ignoring this endogeneity could lead us to underestimate consumers' responsiveness to price changes. Fortunately, it is straightforward to address the endogeneity by instrumenting prices with outside variables. Two obvious instrument candidates are the crude oil price and the tax level. Both should be highly relevant, as they are major determinants of the local gasoline price. The world crude oil price should be exogenous insofar as an individual country's demand does not affect the world crude oil market (though of course this may not hold quite true for large countries or for the case of demand shocks correlated across many countries). The portion of the gasoline tax levied as a specific tax should also, barring gasoline-demand-driven political pressure, be exogenous. The portion of the gasoline tax levied as an ad valorem tax suffers the same endogeneity as the before-tax price on which it is based, but this is remedied by using an instrument of the form

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

Dividing the VAT (which I take to be any tax beyond the excise tax) by the before-tax price returns a percentage that is independent of the price itself. The excise tax, meanwhile, is divided by the fitted values from a fixed-effects regression of before tax-prices on the world crude oil price⁷:

$$BeforeTax_{it} = \beta_0 + \beta_1 crude_{it} + \mu_i + e_{it}$$

Insofar as the crude oil price is exogenous, these fitted values should be exogenous as well; and altogether the tax instrument should be exogenous.

Since the price-volatility and price-unforecastability interaction terms are potentially endogenous as well, I instrument them using interactions between the uncertainty measures and each of the price instruments.

The resulting estimates are reported in Table 7. All least-squares estimates are computed in Stata using Schaffer's (2010) *xtivreg2*.

⁷Since we are working in local currencies, the world crude oil price does in fact vary across countries.

5.1.1 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.

5.1.2 Lagged Endogenous Variable and the Small-T Bias

Although the within-groups estimator of the ADL tackles the potential endogeneity of prices, it is afflicted by a second issue, namely a 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 $\varepsilon_{i,t-1}$:

$$g_{i,t-1} = \lambda g_{i,t-2} + \delta_1 y_{i,t-1} + \delta_2 p_{i,t-1} + \delta_3 \hat{\sigma}_{i,t-1} + \delta_4 \hat{\sigma}_{i,t-1} p_{i,t-1} + \delta_{5i} (t-1) + \mu_i + \varepsilon_{i,t-1}$$

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 (1) (with the trend omitted for simplicity) it becomes clear that transformed lagged consumption is correlated with the transformed error term:

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

There will be $T-1$ pairwise correlations between the terms comprising $\tilde{g}_{i,t-1}$ and $\tilde{\varepsilon}_{it}$. $T-2$ 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 $T \rightarrow \infty$, 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 (1) with all $\delta = 0$), Nickell (1981) derives the small- T bias as

$$p \lim_{N \rightarrow \infty} (\hat{\lambda} - \lambda) = \left\{ \frac{2\lambda}{1 - \lambda^2} - \left[\frac{1}{\frac{1+\lambda}{T-1} \left(1 - \frac{1}{T} \frac{1-\lambda^T}{1-\lambda} \right)} \right] \right\}^{-1} \quad (10)$$

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 $\hat{\lambda}$ should be even greater in this case, and so (10) may be used to establish the minimum size of the bias. Assuming a true λ of 0.8 and setting T equal to its average of 20 suggests a bias in $\hat{\lambda}$ of about -0.10. This puts a lower bound on the magnitude of the bias. We can also put an upper 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*. Estimating a pooled version of (1) with country-specific trends yields $\hat{\lambda} = \frac{0.894}{(0.0203)}$ when $\hat{\sigma} = \text{volatility}$ and $\hat{\lambda} = \frac{0.899}{(0.0196)}$ when $\hat{\sigma} = \text{unforecastability}$. The coefficient on lagged consumption therefore appears to be below 0.9. This should quiet any fears of a unit root, which would have troubling implications for the long-run interpretation of the model.

5.2 System GMM

To address the bias generated by the inclusion of lagged endogenous variables in models such as (1), 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 (1) 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:

$$\textbf{Difference moment conditions: } E[g_{i,t-k}\Delta\varepsilon_{it}] = 0, \quad k = 2, \dots, T \quad (11)$$

$$\textbf{Levels moment conditions: } E[(\mu_i + \varepsilon_{it})\Delta g_{i,t-1}] = 0 \quad (12)$$

$$\textbf{IV-style moment conditions: } E[\Delta w_{it}\Delta\varepsilon_{it}] = 0 \quad \forall t \text{ and } E[w_{it}\varepsilon_{it}] = 0 \quad \forall t \quad (13)$$

where w represents each of the price instruments discussed in Section 5.1. 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 Table 12. 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 6. For now, note that addressing the endogeneity of lagged consumption costs so much precision that the results are almost meaningless: the only statistically significant coefficients are λ and the trend, and even those are not significant in all variations of the instrument matrix. The estimates of λ , moreover, are all greater than 1 and in several cases statistically-significantly greater than the upper bound of 0.9. Despite the bias inherent in estimators that do not specifically address the endogeneity of $g_{i,t-1}$, therefore, these methods are more informative and possibly *less* biased than system GMM.

5.3 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 (1), 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 ($\delta_1, \dots, \delta_4$). 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.

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. The PMG

Table 3: 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: $\theta_{1i} = \theta_1, \theta_{2i} = \theta_2, \theta_{3i} = \theta_3, \theta_{4i} = \theta_4;$ $\delta_{1i} = \delta_1, \delta_{2i} = \delta_2, \delta_{3i} = \delta_3, \delta_{4i} = \delta_4; \phi_i = \phi$

estimator, which uses maximum likelihood based on the assumption that the ε_{it} are independent and normal, is constructed in terms of an error-correction model such as (4). It imposes the constraints $\{\theta_{1i} = \theta_1, \theta_{2i} = \theta_2, \theta_{3i} = \theta_3, \theta_{4i} = \theta_4\}$, while allowing the δ s to remain country-specific. The constraints on the long-run coefficients allow us to exploit *some* cross-sectional variation, which is useful given that the series are not particularly long; and meanwhile the freedom of the short-run coefficients to vary across countries mitigates the biases discussed above. To examine the value of assuming homogenous long-run coefficients (which drastically reduces the number of parameters to be estimated), I also report mean group (MG) estimates. The mean group estimator simply estimates (4) separately for each country and forms an average of the resulting coefficient estimates. To examine the cost of restricting the short-run *and* the long-run coefficients to be the same across all countries, as the ADL inherently did, I estimate a dynamic fixed-effects (DFE) version of the ECM. The restrictions imposed by these various estimators are summarized in Table 3.

Results using the PMG, MG, and DFE estimators are reported in Tables 8, 13, and 14. The MG and PMG short-run elasticities and ϕ_i s reported in these tables are the *means* of the country-specific estimates. 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*, which is further detailed in Blackburne and Frank (2007a).

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. I must therefore check whether any of my variables are nonstationary, and if so whether they are cointegrated.

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

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

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 $\hat{\rho}$ estimated for an $AR(1)$ to be used 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.⁸ One complication is whether to

⁸For a thorough overview of these tests, see Baltagi and Kao (2000).

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 (2011) *xtunitroot*, are reported in Table 4. Where the test does not reject a variable's nonstationarity in levels, I repeat the test on first differences. All variables except g appear to be trend stationary, while without trend g , y , and p are I(1). Since all the outside regressors are trend stationary, and g is clearly cointegrated with its lag, there should be no trouble with spurious regression.

Table 4: Fisher-style Phillips-Perron unit root tests

Variable	Without trends				
	Level		First Difference		Order of Integration
	Test Statistic*	P-Value	Test Statistic ^a	P-Value	
g	52.0571	0.6947	288.9463	0.0000	I(1)
y	28.7177	0.9996	468.0812	0.0000	I(1)
p	58.8670	0.4435	495.8439	0.0000	I(1)
$\hat{\sigma}$ =volatility	389.7647	0.0000			I(0)
$\hat{\sigma}$ =unforecastability	83.6054	0.0155			I(0)
$\hat{\sigma}p$, volatility*price	363.8242	0.0000			I(0)
$\hat{\sigma}p$, unforecastability*price	102.1594	0.0003			I(0)
Variable	With time trends				
	Level		First Difference		Order of Integration
	Test Statistic*	P-Value	Test Statistic ^a	P-Value	
g	45.1560	0.8908	413.0863	0.0000	I(1)
y	151.1144	0.0000			I(0)
p	95.6477	0.0014			I(0)
$\hat{\sigma}$ =volatility	362.5181	0.0000			I(0)
$\hat{\sigma}$ =unforecastability	274.8653	0.0000			I(0)
$\hat{\sigma}p$, volatility*price	379.2350	0.0000			I(0)
$\hat{\sigma}p$, unforecastability*price	190.3141	0.0000			I(0)

*Inverse χ^2 with 58 degrees of freedom.

Phillips-Perron tests conducted using 3 lags for the Newey-West standard errors.

6 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 5), not all the advantages are equal, and indeed not all the advantages are fully realized. After weighing each approach’s advantages and disadvantages, I prefer the fixed-effects 2SLS estimator of the partial adjustment model or the PMG estimator of the error-correction model.

Table 5: 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	✓		✓

6.1 Least-Squares Estimates of the ADL Model

A first glance at the least-squares results in Table 9 reveals the striking value of allowing for country-specific trends. With a homogenous trend, the coefficient on lagged consumption—which we already suspect to be biased downward—is dangerously close to 1. With heterogeneous trends capturing some of the country-specific evolution in consumption behavior, the coefficient on lagged consumption falls to more realistic values of between 0.7 and 0.8. I focus therefore on the models with heterogenous trends, which are summarized in Table 7.

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, this seems to be of moderate value when uncertainty is measured as unforecastability and little value when uncertainty is measured as volatility. Compared to results with no instrumenting (Table 10), that is, the instrumented results for the uncertainty-as-unforecastability model have higher-magnitude coefficients on prices and the price-uncertainty interaction—exactly the effect we expect if the instrumenting corrects for endogeneity in those variables. For the volatility model, however, instrumenting dulls the effect of the price-uncertainty interaction. In both cases, there is in fact endogeneity to address, as tests for endogeneity⁹ strongly reject that prices and the price-uncertainty interactions are exogenous (though they do fail to reject that prices alone are exogenous).

The ineffectiveness of instrumenting in the uncertainty-as-volatility model may be due to weakness of the instruments, as the Kleibergen-Paap rk Wald F statistic suggests a maximal IV relative bias of over 30% in that case (see Stock and Yogo 2005), versus under 5% for the uncertainty-as-unforecastability model. The instruments’ exogeneity is not contradicted by overidentifying restrictions tests, which in both cases fail to reject exogeneity. We might remain slightly suspicious, as the overidentifying restrictions test *does* reject exogeneity of the instruments in the traditional model without habits variables, suggesting that the crude oil price might be somewhat endogenous. Unfortunately, omitting the crude-oil instruments and basing identification solely on the tax instruments yields an insignificant coefficient on prices in the uncertainty-as-volatility model; so even if crude-oil endogeneity *were* a problem, omitting the instrument would not be a solution.

⁹That 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.

The main weakness of the least-squares estimator is, of course, that it does not address the endogeneity of lagged consumption, and so its estimate of λ may be biased downward. A simple way to check whether this is problematic is to set λ over a range of values, then estimate the rest of the model given the specified λ . Table 11 shows the results of this exercise for the model using unforecastability as its measure of uncertainty. Increasing λ decreases the magnitude of the short-run coefficients while increasing the magnitude of the implied long-run effects. Although a difference of 0.1 in λ has a considerable effect on long-run elasticities, note that even pushing λ to 0.95 does not affect the sign or the significance of the coefficients: the small- T bias in $\hat{\lambda}$ does not drive findings of habits effects. This holds for the uncertainty-as-volatility model as well.

Note, moreover, that $\hat{\lambda}$ is potentially subjected to two biases: an upward bias because lagged consumption is endogenous, and a downward bias because parameters are constrained to be homogenous across countries. These biases counteract each other, and so depending on their relative sizes, addressing one without addressing the other could do more harm than good.

Overall, the fixed-effects 2SLS estimator seems very well-specified in the case of the unforecastability model, slightly less so in the case of the volatility model. The specification issues in the latter case may contribute to the failure of the uncertainty-as-volatility model to capture evidence of habits effects: the coefficients on volatility and the price-volatility interaction are both insignificant. Also contributing may be the use of the price-volatility measure as a proxy for price uncertainty. The uncertainty-as-unforecastability model, by contrast, captures strong evidence of habits effects: the coefficients on $\hat{\sigma}$ and $\hat{\sigma}p$ are negative and positive, respectively, with high statistical significance. Higher unforecastability is associated with lower demand and lower-magnitude price elasticity. For every 10% increase in average price-forecast error, the estimate suggests, the demand level decreases by nearly 7% in the short run and 24% in the long run; while the magnitude of price elasticity decreases by 1.5 percentage points in the short run and 5 percentage points in the long run.

6.2 System GMM Estimates of the ADL Model

The unique strength 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 downward, and therefore make adjustment seem faster than it truly is. As expected, the system GMM estimates of λ are higher than the least-squares estimates—troublingly, in fact, they are all higher than 1.

Several shortcomings could be contributing to these overly-high λ s. One 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 $\hat{\lambda}$ (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, which in the least-squares case reduced $\hat{\lambda}$ considerably. System GMM is likely, therefore, to be producing estimates that are considerably higher than the true λ .

Just as troublingly, the precision of the system GMM estimates is devastatingly low: the only coefficients anywhere near statistical significance are lagged consumption and the trend. In this particular application, with a medium-length T and a strong call for individual-specific trends, the system GMM estimates are not just flawed, but useless.

6.3 Error-Correction Model

As with the least-squares estimates of the ADL model, a first glance at the PMG estimates of the ECM model (Table 13) reveals the cost of imposing homogeneity on the trend: the error-correction term $\hat{\phi}$ is dangerously and implausibly close to 0. Again I opt for the heterogeneous-trend models, summarized in Table 8.

Recall that the chief advantage of the PMG estimator is that it mitigates biases induced by falsely imposing coefficient homogeneity when regressors are autocorrelated. Given that the regressors in this application are positively autocorrelated, such bias on $\hat{\lambda}$ would be upward, and correcting that bias should lead to lower λ and higher-magnitude ϕ , i.e. a faster speed of adjustment. Surprisingly, however, adjustment is slightly slower according to the PMG estimates than according to FE 2SLS estimates: $\hat{\phi}_{PMG}$ is in the region of -0.2 , corresponding to a λ of 0.8 , versus $\hat{\lambda}_{FE} = 0.71$ to 0.76 . This surprise can't be attributed entirely to differences in the underlying models. Focusing on the ECM, we'd expect to see $|\phi|$ fall as the coefficient-homogeneity restriction were relaxed from DFE to PMG to MG (Table 14), but this is not entirely the case: although $|\phi_{MG}|$ is highest for all variants of the model and $|\hat{\phi}_{PMG}| > |\hat{\phi}_{DFE}|$ for the unforecastability version, $|\hat{\phi}_{PMG}| < |\hat{\phi}_{DFE}|$ for the volatility version as well as the traditional model. The difference between $\hat{\phi}_{PMG}$ and $\hat{\phi}_{DFE}$ is in all instances quite small, so the surprise is not too troubling; but it does suggest that heterogeneity of the short-run coefficients is not driving much of a bias in the homogenous-coefficient estimators (including FE 2SLS). Heterogeneity of *long-run* coefficients may contribute more to such a bias, given the larger gap between $\hat{\phi}_{PMG}$ and $\hat{\phi}_{MG}$. Unfortunately, relaxing all coefficient-homogeneity restrictions would be unwise, as the MG estimator tends to be sensitive to outliers and less precise than PMG (Pesaran, Shin & Smith 1999), and these tendencies are clearly problematic in the MG estimates of the volatility model.¹⁰

Although the PMG estimator does—at least in theory—mitigate coefficient-heterogeneity bias, it is still vulnerable to the endogeneity of lagged consumption. It would be ideal to correct the resulting upward bias on the speed of adjustment, but it does not seem at the moment that appropriate corrections have been derived. Pesaran and Zhao (1999) have worked out bias corrections for the (*not*-pooled) mean group estimator, which perform well when $\lambda < 0.8$ but not at all when $\lambda \geq 0.8$; and recent work by Choi, Mark and Sul (2010) offers bias corrections for homogenous-coefficient models; but there do not seem to be corrections for the case where λ is moderately large and/or coefficient homogeneity is relaxed. I will have to accept the PMG estimator as it is, noting for reassurance that, given the bound on λ worked out in Section 5.1.2, $|\phi|$ *must* truly be greater than 0.1 .

In interpreting the PMG results, my chief focus is the long-run coefficients. The short-run coefficient estimates reported in the tables are in fact means of country-specific estimates, and are prone to the same weaknesses as the MG estimator. The long-run coefficients do, however, provide evidence of rational habits. In the case of the uncertainty-as-volatility model, the coefficient on price volatility $\hat{\sigma}$ is not statistically-significantly negative ($p = 0.1$), but the coefficient on the interaction term $\hat{\sigma}p$ is positive with high significance: a 0.1 increase in the rolling standard deviation of prices reduces the magnitude of long-run price elasticity by 4.6 percentage points. This finding contrasts the FE 2SLS estimates, according to which price volatility did *not* have a statistically-significant effect on price elasticity. Consistent with the FE 2SLS estimates, though, evidence of habits is strong in the uncertainty-as-unforecastability model: the coefficients on $\hat{\sigma}$ and $\hat{\sigma}p$ are negative and positive, respectively, with high statistical significance. In practical terms, the estimates suggest that for every 10% increase in average price-forecast error, demand in the long run decreases by 23% and the magnitude of long-run price elasticity decreases by 2 percentage points. Compared to the FE 2SLS estimate, the PMG estimate suggests a larger effect of

¹⁰Baltagi and Griffin (1997) apply a variety of estimators to gasoline demand and are similarly unimpressed with the MG estimator.

price unforecastability on demand level and a similar effect of unforecastability on price elasticity.

6.4 Discussion

Depending on how price uncertainty is characterized, we have seen across the various models and estimators here that consumers' behavior is consistent with rational habits for gasoline. But how important are rational habits effects in explaining variations in demand behavior across countries?

Although both the FE 2SLS and the PMG estimator constrain long-run coefficients to be homogenous, we can use country-specific average price uncertainties, $\bar{\sigma}_i$, to interpolate the variation in price elasticities predicted by variation in price uncertainty. Recall that 'total' long-run price elasticity is given by $\frac{\delta_2 + \delta_4 \bar{\sigma}_{it}}{1 - \lambda}$ in the partial-adjustment model and $\theta_2 + \theta_4 \hat{\sigma}_{it}$ in the error-correction model. Figure 7 plots total long-run price elasticities ($\hat{\theta}_2 + \hat{\theta}_4 \bar{\sigma}_i$) predicted by the PMG estimates of the uncertainty-as-volatility ECM, along with the range of elasticities implied by the minimum and maximum volatilities experienced by each country. Figures 8 and 9 do the same for the uncertainty-as-unforecastability models. Overall, neither price volatility nor price unforecastability predicts very big variations in price elasticity. How successful the models are in *explaining* cross-sectional variations in price elasticity depends, of course, on what those variations are. Unfortunately, precise country-specific elasticity estimates are out of reach given the data at hand; but I make two sets of *rough* country-specific estimates, one using 2SLS on the traditional partial-adjustment model for each country individually, the other using the country-specific elasticities that would be averaged to produce MG estimates of the traditional ECM. Ignoring country-specific elasticities that are clear outliers,¹¹ only 10 to 13% of the variation in the country-specific price elasticity estimates can be explained by the 'total' long-run price elasticities predicted for the uncertainty-as-unforecastability model. Even less variation in the country-specific price elasticities—1 to 7%—can be explained by the 'total' long-run price elasticities predicted for the uncertainty-as-volatility model. Whether this low explanatory power actually reflects a weakness of price uncertainty in explaining price-elasticity variation, or whether it merely reflects the imprecision of country-specific elasticities, is not clear. Even for those countries not considered outliers, the standard errors of the price elasticities are generally huge.

The power of price uncertainty to explain cross-sectional variation in the demand *level* is far more pronounced, or at least clearer to see. The long-run effect of uncertainty implied by the countries' mean unforecastability, $\frac{\hat{\sigma}_3 \bar{\sigma}_i}{1 - \lambda}$ or $\hat{\theta}_3 \bar{\sigma}_i$, explains¹² 37% of the variation in countries' mean log per-capita consumption, $\bar{g}_i$. (Price volatility's effect on the demand level is, as previously noted, not statistically significant.) This cross-sectional explanatory power is markedly greater than unforecastability's partial R^2 of 6.7% (FE 2SLS model); and although it's possible that price unforecastability proxies for other cross-sectional differences, it could indeed be that price unforecastability is a major driver of the variation in gasoline demand across countries.

Of course, price uncertainty may also help to explain variations in price elasticity and demand *within* countries over time. From a policy perspective, this may be the more immediately interesting dimension, since taxes and other interventions can be used to manipulate consumers' price uncertainty. To examine rational habits' explanatory power across this dimension, we would need information on the evolution of country-specific elasticities over time—and for that, we would require more data, preferably micro data.

In addition to examining the rational habit effects within my models, it is worth checking how my estimates of price and income elasticity fit with the existing literature. To facilitate this, I report estimates of 'traditional' models without rational-habits variables (see Tables 7 and 8). Long-run income elasticity in both the partial-adjustment and the error-correction model is about 1, which fits neatly

¹¹ Outliers were taken to be long-run price elasticities greater than 0 or smaller than -5 . Four of the ADL- and nine of the ECM-based estimates were outliers by these criteria.

¹² After omitting three countries with particularly high price unforecastability.

Figure 7: Long-run price elasticities implied by PMG estimates of uncertainty-as-volatility ECM

Figure 8: Long-run price elasticities implied by FE 2SLS estimates of uncertainty-as-unforecastability ADL

Figure 9: Long-run price elasticities implied by PMG estimates of uncertainty-as-unforecastability ECM

alongside the unit-elastic averages common in the literature. Long-run price elasticity is about -0.9 in the partial-adjustment model and -0.6 in the error-correction model, spanning the mid-to-upper range (in magnitude) of estimates in the literature. In the models that do include habits variables, meanwhile, income elasticities continue to look reasonable, again around 1 in the volatility models and falling to a still-plausible 0.6-0.8 in the unforecastability models. The ‘total’ long-run price elasticities ($\frac{\hat{\delta}_2 + \hat{\delta}_4 \hat{\sigma}}{1 - \hat{\lambda}}$ and $\hat{\theta}_2 + \hat{\theta}_4 \hat{\sigma}$), like the traditional long-run price elasticities, fall on the mid-to-upper end of literature, ranging from -0.5 (ECM, volatility) to -1.0 (ECM, unforecastability), with the partial-adjustment model results sandwiched in between (-0.8 for volatility, -0.7 for unforecastability).

That the price elasticities here are relatively high-magnitude suggests that the ECM literature’s tendency towards low price elasticities is due not to the choice of an error-correction model, but to the lack of a panel dimension in most ECM studies to date. This story is bolstered by a comparison of the MG and DFE estimates of price elasticity in the traditional ECM: the DFE estimate, which exploits cross-sectional variation, is three times the MG estimate, which does not. Although long-run price elasticity is higher according to the DFE (and PMG) estimate of the traditional ECM than the FE 2SLS estimate of the ADL, there does not seem to be a systemic difference in ‘total’ price elasticities between the ECM and ADL models that would make a strong case for an alternative story.

7 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 in shaping the demand of forward-looking consumers.

The rational habits framework suggests two major implications for consumer behavior. This paper 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 both implications is strong when price uncertainty is proxied using price unforecastability, and there is mixed evidence for the first implication when price uncertainty is proxied using price volatility. The effects are not unsizeable: according to my preferred model and estimator specifications, an increase in absolute price-forecasting error of 10% is associated in the long run with a 22 to 23% decrease in demand and 5 to 7 percentage-point decrease in price elasticity.

The sensitivity of my results to the choice of price-uncertainty proxy suggests that this may be a fertile area for further work. Explicit modelling of price volatility and/or an exploration of how consumers form perceptions of price uncertainty might bring evidence of rational habits into starker relief. For now, however, the volatility and unforecastability proxies offer convincing evidence that consumers around the world exhibit behaviors associated with rational habits. In light of these habits, policymakers should heed the role of price uncertainty in shaping gasoline demand. Policies that stabilize gasoline prices will have the undesired effect of stoking demand, but they will also increase the effectiveness of price instruments. Given rational habits, indeed, stable price increases are a longer policy lever than traditionally-measured price elasticity would suggest.

8 Tables

Table 6: FE 2SLS estimates of the (lead) price (components) models

	(1) Lead price	(2) Lead price components	(3) Price components
g_{t-1}	0.720** (0.0249) [0.000]	0.721** (0.0254) [0.000]	0.746** (0.0270) [0.000]
y_t	0.342** (0.0529) [0.000]	0.323** (0.0596) [0.000]	0.258** (0.0459) [0.000]
p_t	-0.296** (0.0696) [0.000]	-0.327** (0.0822) [0.000]	
p_{t+1}	0.213 (0.158) [0.177]		
$\ln(1 + \tau_{t+1})$		0.168 (0.171) [0.328]	
$\ln(BeforeTax_{t+1})$		0.206 (0.160) [0.200]	
$\ln(1 + \tau_t)$			-0.220* (0.104) [0.034]
$\ln(BeforeTax_t)$			-0.222** (0.0699) [0.001]
t_i	[indiv. trends]	[indiv. trends]	[indiv. trends]
R^2	0.944	0.942	0.942

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

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

N=29, T=10-21.

Table 7: FE 2SLS estimates of the ADL model

	(1)	(2)	(3)
	Traditional model	Volatility het. trend	Unforecastability ^Ψ het. trend
g_{t-1}	0.761** (0.0457) [0.000]	0.762** (0.0484) [0.000]	0.712** (0.0320) [0.000]
y	0.248** (0.0481) [0.000]	0.240** (0.0515) [0.000]	0.229** (0.0324) [0.000]
p	-0.214** (0.0368) [0.000]	-0.185** (0.0443) [0.000]	-0.209** (0.0276) [0.000]
$\hat{\sigma}$		0.303 (0.238) [0.204]	-0.686** (0.152) [0.000]
$\hat{\sigma}p$		-0.304 (0.239) [0.205]	0.147** (0.0358) [0.000]
t_i	[indiv. trends]	[indiv. trends]	[indiv. trends]
Implied long-run effects			
y	1.0378 (0.149) [0.000]	1.010 (0.173) [0.000]	0.797 (0.110) [0.000]
p	-0.898 (0.204) [0.000]	-0.780 (0.194) [0.000]	-0.727 (0.098) [0.000]
$\hat{\sigma}$		1.276 (1.0125) [0.208]	-2.382 (0.473) [0.000]
$\hat{\sigma}p$		-1.279 (1.011) [0.206]	0.511 (0.114) [0.000]
R^2	0.942	0.937	0.968
Tests			
Endogeneity: price ^a	0.741 [0.3893]	1.148 [0.2840]	0.068 [0.7949]
Endogeneity: price and interaction ^b		8.505 [0.0142]	9.869 [0.0072]
Underidentification ^c	127.541 (2 df) [0.0000]	8.809 (3 df) [0.0319]	35.788 (3 df) [0.0000]
Weak instruments ^d	157.185	2.544	20.889
Exogeneity ^e	5.459 (1 df) [0.0195]	2.479 (2 df) [0.2896]	1.244 (2 df) [0.5370]

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

^ΨOmitting 3 outlier countries. | N=26-29, T=11-21.^aDifference-in-Sargan/Hansen χ^2 (1) under H_0 .^bDifference-in-Sargan/Hansen χ^2 (2) under H_0 .^cKleibergen-Paap (2006) rk statistic; χ^2 (df) under H_0 .^dKleibergen-Paap rk Wald F statistic.^eOveridentifying restrictions test (Hansen J statistic); χ^2 (df) under H_0 .

Table 8: PMG estimates of the ECM with heterogeneous trend

	(1) Traditional model het. trend	(2) Volatility het. trend	(3) Unforecastability ^Ψ het. trend
ϕ_i	-0.181** (0.0359) [0.000]	-0.174** (0.0430) [0.000]	-0.223** (0.0402) [0.000]
t_i	-0.00739** (0.00129) [0.000]	-0.00671** (0.00139) [0.000]	-0.00848** (0.00190) [0.000]
Long run			
y	1.091** (0.106) [0.000]	0.965** (0.0936) [0.000]	0.548** (0.0772) [0.000]
p	-0.556** (0.0801) [0.000]	-0.538** (0.0611) [0.000]	-1.044** (0.0764) [0.000]
$\hat{\sigma}$		-0.448 (0.284) [0.114]	-2.306** (0.290) [0.000]
$\hat{\sigma}p$		0.462** (0.107) [0.000]	0.211** (0.0650) [0.001]
Short run			
Δy	0.250** (0.0750) [0.001]	0.270** (0.0708) [0.000]	0.308** (0.0878) [0.000]
Δp	-0.0849* (0.0371) [0.022]	-0.0948** (0.0315) [0.003]	0.131 (0.0934) [0.161]
$\Delta \hat{\sigma}$		0.597 (0.797) [0.454]	3.931+ (2.125) [0.064]
$\Delta \hat{\sigma}p$		-0.0982 (0.368) [0.789]	-2.690* (1.293) [0.037]

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

^ΨOmitting 3 outlier countries.

Table 9: Homogeneous and heterogeneous-trend versions of the ADL model, FE 2SLS

	(1) Volatility hom. trend	(2) Volatility het. trend	(3) Unforecastability ^Ψ hom. trend	(4) Unforecastability ^Ψ het. trend
$gt-1$	0.966** (0.0346) [0.000]	0.762** (0.0484) [0.000]	0.960** (0.0181) [0.000]	0.712** (0.0320) [0.000]
y	0.0571 (0.0413) [0.167]	0.240** (0.0515) [0.000]	0.0872** (0.0251) [0.001]	0.229** (0.0324) [0.000]
p	-0.0965** (0.0342) [0.005]	-0.185** (0.0443) [0.000]	-0.0831** (0.0306) [0.007]	-0.209** (0.0276) [0.000]
$\hat{\sigma}$	0.0978 (0.225) [0.665]	0.303 (0.238) [0.204]	-0.0760 (0.120) [0.526]	-0.686** (0.152) [0.000]
$\hat{\sigma}p$	-0.0591 (0.143) [0.678]	-0.304 (0.239) [0.205]	-0.00216 (0.0261) [0.934]	0.147** (0.0358) [0.000]
$t_{(i)}$	-0.00394** (0.00110) [0.000]	[indiv. trends]	-0.00466** (0.000972) [0.000]	[indiv. trends]
R^2	0.921	0.937	0.952	0.968
Tests				
Endogeneity:	0.053	1.148	0.413	0.068
price ^a	[0.8184]	[0.2840]	[0.5206]	[0.7949]
Endogeneity:	1.910	8.505	0.920	9.869
price and interaction ^b	[0.3849]	[0.0142]	[0.6313]	[0.0072]
Underidentification ^c	15.655 (3 df) [0.0013]	8.809 (3 df) [0.0319]	73.685 (3 df) [0.0000]	35.788 (3 df) [0.0000]
Weak instruments ^d	5.555	2.544	43.783	20.889
Exogeneity ^e	0.720 [0.6978]	2.479 (2 df) [0.2896]	4.568 [0.1019]	1.244 (2 df) [0.5370]

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

N=26-29, T=11-21. | ^ΨOmitting 3 outlier countries.^aDifference-in-Sargan/Hansen χ^2 (1) under H_0 .^bDifference-in-Sargan/Hansen χ^2 (2) under H_0 .^cKleibergen-Paap (2006) rk statistic; χ^2 (df) under H_0 .^dKleibergen-Paap rk Wald F statistic.^eOveridentifying restrictions test (Hansen J statistic); χ^2 (df) under H_0 .

Table 10: Fixed effects estimates of the ADL model, no price instrumenting

	(1)	(2)	(3)
	Traditional model	Volatility het. trend	Unforecastability ^Ψ het. trend
g_{t-1}	0.765** (0.0395) [0.000]	0.769** (0.0396) [0.000]	0.757** (0.0304) [0.000]
y	0.251** (0.0477) [0.000]	0.250** (0.0494) [0.000]	0.250** (0.0349) [0.000]
p	-0.187** (0.0278) [0.000]	-0.186** (0.0288) [0.000]	-0.171** (0.0190) [0.000]
$\hat{\sigma}$		-0.0635 (0.0900) [0.481]	-0.246** (0.0904) [0.007]
$\hat{\sigma}p$		0.0818* (0.0385) [0.034]	0.0493** (0.0182) [0.007]
t_i	[indiv. trends]	[indiv. trends]	[indiv. trends]
R^2	0.942	0.943	0.968

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

^ΨOmitting 3 outlier countries. | N=26-29, T=11-21.

Table 11: Sensitivity of FE 2SLS estimator to λ , unforecastability model

	(1) $\lambda \equiv 0.7$	(2) $\lambda \equiv 0.8$	(3) $\lambda \equiv 0.9$	(4) $\lambda \equiv 0.95$
g_{t-1}	0.7 (as defined)	0.8 (as defined)	0.9 (as defined)	0.95 (as defined)
y	0.235** (0.0294) [0.000]	0.188** (0.0301) [0.000]	0.141** (0.0317) [0.000]	0.117** (0.0327) [0.000]
p	-0.212** (0.0255) [0.000]	-0.185** (0.0249) [0.000]	-0.157** (0.0254) [0.000]	-0.144** (0.0260) [0.000]
$\hat{\sigma}$	-0.704** (0.137) [0.000]	-0.540** (0.137) [0.000]	-0.376** (0.141) [0.008]	-0.294* (0.146) [0.044]
$\hat{\sigma}p$	0.151** (0.0331) [0.000]	0.118** (0.0326) [0.000]	0.0862* (0.0334) [0.010]	0.0701* (0.0342) [0.041]
t_i	[indiv. trends]	[indiv. trends]	[indiv. trends]	[indiv. trends]
Implied long run elasticities				
y	0.391** (0.109) [0.000]	0.940** (0.151) [0.000]	1.410** (0.317) [0.000]	2.348** (0.654) [0.000]
p	-0.479** (0.0868) [0.000]	-0.925** (0.124) [0.000]	-1.575** (0.254) [0.000]	-2.875** (0.521) [0.000]
$\hat{\sigma}$	-0.979* (0.486) [0.044]	-2.700** (0.683) [0.000]	-3.759 ** (1.415) [0.008]	-5.876* (2.914) [0.044]
$\hat{\sigma}p$	0.234* (0.114) [0.041]	0.592** (0.163) [0.000]	0.862* (0.334) [0.010]	1.401* (0.685) [0.041]

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

Omitting 3 outlier countries.

Table 12: System GMM estimates of the ADL model

	(1) All IVs Volatility	(2) Collapsed IV matrix Volatility	(3) Reduced IV count Volatility	(4) All IVs Unforecastability Ψ	(5) Collapsed IV matrix Unforecastability Ψ	(6) Reduced IV count Unforecastability Ψ
g_{t-1}	1.016** (0.0253) [0.000]	1.004** (0.0734) [0.000]	1.026** (0.317) [0.003]	1.019** (0.0165) [0.000]	1.060** (0.0535) [0.000]	1.014** (0.207) [0.000]
y	-0.00294 (0.0458) [0.949]	-0.0429 (0.0700) [0.545]	-0.196 (0.133) [0.151]	0.0609 (0.0517) [0.250]	0.0251 (0.0594) [0.677]	-0.143 (0.152) [0.355]
σp	0.00493 (0.0425) [0.908]	0.0484 (0.0668) [0.475]	0.236 (0.178) [0.195]	-0.0456 (0.0525) [0.394]	-0.0133 (0.0571) [0.818]	0.184 (0.178) [0.312]
$\hat{\sigma}$	0.0774 (0.162) [0.636]	0.423 (0.467) [0.373]	2.395 (2.307) [0.309]	-0.191 (0.163) [0.251]	0.0926 (0.272) [0.737]	0.863 (1.142) [0.457]
$\hat{\sigma} p$	0.101 (0.0803) [0.221]	-0.0690 (0.300) [0.820]	-1.744 (2.989) [0.564]	0.0148 (0.0350) [0.677]	-0.00710 (0.0389) [0.857]	-0.224 (0.254) [0.385]
t	-0.00370** (0.000906) [0.000]	-0.00442* (0.00170) [0.015]	-0.00299 (0.00347) [0.397]	-0.00512** (0.00110) [0.000]	-0.00334 (0.00211) [0.126]	-0.00296 (0.00628) [0.642]

Robust standard errors in parentheses; p-values in brackets. | ** p<0.01, * p<0.05, + p<0.1
 Ψ Omitting 3 outlier countries.

Table 13: Homogeneous and heterogeneous-trend versions of the ECM model, PMG

	(1) Volatility hom. trend	(2) Volatility het. trend	(3) Unforecastability ^Ψ hom. trend	(4) Unforecastability ^Ψ het. trend
ϕ_i	-0.0705** (0.0122) [0.000]	-0.174** (0.0430) [0.000]	-0.0207 (0.0330) [0.529]	-0.223** (0.0402) [0.000]
$t_{(i)}$	-0.0404** (0.00627) [0.000]	-0.00671** (0.00139) [0.000]	-0.0188** (0.00414) [0.000]	-0.00848** (0.00190) [0.000]
Long run				
y	-0.0269 (0.311) [0.931]	0.965** (0.0936) [0.000]	0.876** (0.110) [0.000]	0.548** (0.0772) [0.000]
p	-1.904** (0.247) [0.000]	-0.538** (0.0611) [0.000]	-0.642** (0.0968) [0.000]	-1.044** (0.0764) [0.000]
$\hat{\sigma}$	0.0574 (1.168) [0.961]	-0.448 (0.284) [0.114]	-10.65** (1.086) [0.000]	-2.306** (0.290) [0.000]
$\hat{\sigma}p$	2.484** (0.599) [0.000]	0.462** (0.107) [0.000]	1.919** (0.197) [0.000]	0.211** (0.0650) [0.001]
Short run				
Δy	0.474** (0.0837) [0.000]	0.270** (0.0708) [0.000]	0.418** (0.0919) [0.000]	0.308** (0.0878) [0.000]
Δp	-0.121** (0.0273) [0.000]	-0.0948** (0.0315) [0.003]	-0.0923 (0.0728) [0.205]	0.131 (0.0934) [0.161]
$\Delta \hat{\sigma}$	2.334 (1.525) [0.126]	0.597 (0.797) [0.454]	2.124+ (1.227) [0.083]	3.931+ (2.125) [0.064]
$\Delta \hat{\sigma}p$	-0.577 (0.512) [0.259]	-0.0982 (0.368) [0.789]	-1.739+ (0.995) [0.081]	-2.690* (1.293) [0.037]

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

^ΨOmitting 3 outlier countries.

Table 14: PMG, MG, and DFE estimates of the ECM

	(1) PMG Traditional ^Ψ	(2) MG Traditional ^Ψ	(3) DFE Traditional ^Ψ	(4) PMG Volatility	(5) MG Volatility	(6) DFE Volatility	(7) PMG Unforecastability ^Ψ	(8) MG Unforecastability ^Ψ	(9) DFE Unforecastability ^Ψ
$\phi_{(i)}$	-0.149** (0.0320) [0.000]	-0.318** (0.0625) [0.000]	-0.198** (0.0235) [0.000]	-0.174** (0.0430) [0.000]	-0.375** (0.119) [0.002]	-0.218** (0.0247) [0.000]	-0.223** (0.0402) [0.000]	-0.434** (0.0941) [0.000]	-0.216** (0.0262) [0.000]
$t_{(i)}$	-0.00638** (0.00126) [0.000]	-0.00724* (0.00285) [0.011]	[indiv. trends]	-0.00671** (0.00139) [0.000]	-0.00545** (0.00209) [0.009]	[indiv. trends]	-0.00848** (0.00190) [0.000]	-0.0119** (0.00357) [0.001]	[indiv. trends]
Long run									
y	1.191** (0.121) [0.000]	0.0500 (0.465) [0.914]	0.946** (0.166) [0.000]	0.965** (0.0936) [0.000]	1.050* (0.483) [0.030]	0.810** (0.133) [0.000]	0.548** (0.0772) [0.000]	1.188+ (0.637) [0.062]	0.774** (0.158) [0.000]
p	-0.586** (0.0967) [0.000]	-0.240 (0.359) [0.504]	-0.734** (0.168) [0.000]	-0.538** (0.0611) [0.000]	-0.390 (0.652) [0.550]	-0.641** (0.163) [0.000]	-1.044** (0.0764) [0.000]	-1.351** (0.338) [0.000]	-0.739** (0.172) [0.000]
$\hat{\sigma}$				-0.448 (0.284) [0.114]	31.69 (32.87) [0.335]	-0.243 (0.525) [0.643]	-2.306** (0.290) [0.000]	-6.147+ (3.466) [0.076]	-1.120+ (0.613) [0.068]
$\hat{\sigma}p$				0.462** (0.107) [0.000]	-10.69 (11.14) [0.337]	0.403 (0.375) [0.283]	0.211** (0.0650) [0.001]	6.906+ (3.530) [0.050]	0.228+ (0.129) [0.078]
Short run									
Δy	0.253** (0.0795) [0.001]	0.273* (0.124) [0.028]	0.238** (0.0639) [0.000]	0.270** (0.0708) [0.000]	0.319* (0.148) [0.031]	0.251* (0.115) [0.029]	0.308** (0.0878) [0.000]	0.190 (0.151) [0.207]	0.231** (0.0706) [0.001]
Δp	-0.0873* (0.0381) [0.022]	-0.0592 (0.0366) [0.106]	-0.0522+ (0.0287) [0.068]	-0.0948** (0.0315) [0.003]	-0.0717 (0.0481) [0.136]	-0.0607* (0.0268) [0.024]	0.131 (0.0934) [0.161]	0.0910 (0.0937) [0.332]	-0.0470 (0.0335) [0.161]
$\Delta \hat{\sigma}$				0.597 (0.797) [0.454]	0.299 (1.241) [0.810]	0.0644 (0.0557) [0.247]	3.931+ (2.126) [0.064]	2.759 (1.749) [0.115]	0.0927 (0.106) [0.383]
$\Delta \hat{\sigma}p$				-0.0982 (0.368) [0.789]	0.858 (0.820) [0.296]	0.00298 (0.0673) [0.965]	-2.690* (1.293) [0.037]	-1.759+ (1.002) [0.079]	-0.0279 (0.0202) [0.167]

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

^ΨOmitting 3 outlier countries.

9 Further Data Exposition

9.1 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. For some estimates, the Czech Republic, Slovakia, and Turkey are excluded as outliers because the model fit them particularly poorly. All three of these countries had unusually high price volatilities, particularly early in the sample period; and the first two underwent huge political and economic change in the early 1990s.

9.2 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 Denmark, for example, regular unleaded is 92 RON and premium unleaded is 98 RON; in Germany, regular is 91 RON and premium is 95 RON; in Austria, regular is 91 RON and premium 95 or 98; and in France and the UK, the standard option is 95 RON or above. Meanwhile the definition of the Research Octane Number (RON) varies by country, and some countries use alternative measures such as the Road Octane Number (RdON) or the Anti-Knock Index (AKI). (IEA 2012, Country Notes) 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 15 shows the product for which each country's price measures are taken.

Table 15: Gasoline product used for price and tax data

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

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] Brons, Martijn, et al. 2008. A meta-analysis of the price elasticity of gasoline demand. A SUR approach. *Energy Economics* 30: 2105-2122.
- [20] Cheung, Kui-yin, and Elspeth Thomson. 2004. The Demand for Gasoline in China: A Cointegration Analysis. *Journal of Applied Statistics* 31(5): 533-544.
- [21] 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.
- [22] Choi, In. 2001. Unit root tests for panel data. *Journal of International Money and Finance* 20: 249-272.
- [23] 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.
- [24] Dahl, Carol, and Thomas Sterner. 1991. Analysing gasoline demand elasticities: A survey. *Energy Economics* 13(3): 203-210.
- [25] Dahl, Carol. 2011. Measuring global gasoline and diesel price and income elasticities. *Energy Policy*: corrected proof.
- [26] De Vita, G., K. Endresen, and L.C. Hunt. 2006. An empirical analysis of energy demand in Namibia. *Energy Policy* 34: 3447-3463.
- [27] 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.
- [28] Espey, Molly. 1998. Gasoline demand revisited: An international meta-analysis of elasticities. *Energy Economics* 20(3): 273-95.
- [29] 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.
- [30] 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.
- [31] Graham, Daniel J., and Stephen Glaister. 2004. Road Traffic Demand Elasticity Estimates: A Review. *Transport Reviews* 24(3): 261-274.
- [32] 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.
- [33] Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. 1988. Estimating vector autoregressions with panel data. *Econometrica* 56(6): 1371-95.
- [34] 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.

- [35] Im, Kyung So, M. Hashem Pesaran, and Yongcheol Shin. 2003. Testing for unit roots in heterogeneous panels. *Journal of Econometrics* 115: 53-74.
- [36] International Energy Agency (IEA). 2012. Country Notes. Energy Prices and Taxes (2012Q4). ESDS International, (Mimas) University of Manchester.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] Kleibergen, Frank, and Richard Paap. 2006. Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133: 97-126
- [41] Krichene, Noureddine. 2002. World crude oil and natural gas: a demand and supply model. *Energy Economics* 24: 557-576.
- [42] 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.
- [43] Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2012. Gasoline taxes and consumer behavior. NBER Working Paper 17891.
- [44] Liddle, Brantley. 2012. The systemic, long-run relation among gasoline demand, gasoline price, income, and vehicle ownership in OECD countries: Evidence from panel cointegration and causality modeling. *Transportation Research Part D* 17: 327-331.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] Neto, David. 2012. Testing and estimating time-varying elasticities of Swiss gasoline demand. *Energy Economics* 34: 1755-1762.
- [49] Nguyen-Van, Phu. 2010. Energy consumption and income: A semiparametric panel data analysis. *Energy Economics* 32: 557-563.
- [50] Nickell, Stephen. 1981. Biases in Dynamic Models with Fixed Effects. *Econometrica* 49(6): 1417-1426.
- [51] 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). Cambridge, UK: Cambridge UP. Accessed version revised January 1997 at <<http://www.econ.cam.ac.uk/faculty/pesaran/ardl.pdf>>, 24 February 2011.

- [52] 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.
- [53] Pesaran, M. Hashem, and Ron Smith. 1995. Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics* 68: 79-113.
- [54] 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.
- [55] Phillips, Peter C. B., and Pierre Perron. 1988. Testing for a Unit Root in Time Series Regression. *Biometrika* 75(2): 335-346.
- [56] Polemis, Michael L. 2006. Empirical assessment of the determinants of road energy demand in Greece. *Energy Economics* 28: 385-403.
- [57] Ramanathan, R. 1999. Short- and long-run elasticities of gasoline demand in India: An empirical analysis using cointegration techniques. *Energy Economics* 21: 321-330.
- [58] 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.
- [59] Rao, B. Bhaskara, and Gyaneshwar Rao. 2009. Cointegration and the demand for gasoline. *Energy Policy* 37: 3978-3983.
- [60] Rivers, Nicholas, and Brandon Schaufele. 2012. Carbon tax salience and gasoline demand. Working Paper #1211E, Department of Economics, University of Ottawa.
- [61] 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_HowtoDxtabond6_12_1_06.pdf.
- [62] 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>.
- [63] 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>.
- [64] Samimi, Rodney. 1995. Road transport energy demand in Australia: A cointegration approach. *Energy Economics* 17(4): 329-339.
- [65] 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>.
- [66] Scott, K. Rebecca. 2012. Rational habits and gasoline demand. *Energy Economics* 34(5): 1713-1723.
- [67] Scott, K. Rebecca. 2010. Rational Habits and Uncertain Relative Prices: Simulating Petrol Consumption Behavior. In *Essays in Consumption Habits and the Environment - Thesis (D.Phil.)*, University of Oxford, 2011. Also available as Economics Series Working Papers 596, University of Oxford, Department of Economics, 2012.
- [68] StataCorp. 2011. Stata Statistical Software: Release 12. College Station, TX: StataCorp LP.

- [69] Stock, J. H., and M. Yogo. 2005. Testing for weak instruments in linear IV regression. In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, ed. D. W. K. Andrews and J. H. Stock, 80-108. Cambridge: Cambridge University Press.
- [70] Storchmann, Karl. 2005. Long-run gasoline demand for passenger cars: the role of income distribution. *Energy Economics* 27(1): 25– 58.
- [71] 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.
- [72] 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

- [73] International Energy Agency (IEA). 2012. *Energy Prices and Taxes (2012Q4)*. ESDS International, (Mimas) University of Manchester.
- [74] International Energy Agency (IEA). 2012. *Oil Information (2012)*. ESDS International, (Mimas) University of Manchester. DOI: <<http://dx.doi.org/10.5257/iea/oil/2012>>.
- [75] 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>. DOI: <<http://dx.doi.org/10.5257/iea/ept/2012q4>>.
- [76] OECD. 2011. ‘Labour Force Statistics: Population projections’, OECD Employment and Labour Market Statistics (database). DOI: <<http://dx.doi.org/10.1787/data-00538-en>>.
- [77] World Bank. *World Development Indicators (WDI)*. World dataBank. <<http://databank.worldbank.org>>. Accessed 30 Jan 2013.