Energy Economics 47 (2015) 77–86

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneco

Income elasticity of gasoline demand: A meta-analysis☆

Tomas Havranek [a,b,*], Ondrej Kokes [c]

[a] Czech National Bank, Czech Republic
[b] Charles University, Prague, Czech Republic
[c] Cambridge Econometrics, United Kingdom

ARTICLE INFO

Article history:
Received 29 April 2013
Received in revised form 24 October 2014
Accepted 2 November 2014
Available online 13 November 2014

JEL classification:
C83
Q41

Keywords:
Gasoline
Income elasticity
Publication bias
Meta-analysis

ABSTRACT

In this paper we quantitatively synthesize empirical estimates of the income elasticity of gasoline demand reported in previous studies. The studies cover many countries and report a mean elasticity of 0.28 for the short run and 0.66 for the long run. We show, however, that these mean estimates are biased upwards because of publication bias—the tendency to suppress negative and insignificant estimates of the elasticity. We employ mixed-effects multilevel meta-regression to filter out publication bias from the literature. Our results suggest that the income elasticity of gasoline demand is on average much smaller than reported in previous surveys: the mean corrected for publication bias is 0.1 for the short run and 0.23 for the long run.

1. Introduction

The income elasticity of gasoline demand is a key parameter in energy and environmental economics. It helps us understand, among other things, how emissions of greenhouse gases stemming from the consumption of gasoline will evolve in the future as developing countries get richer. Because of its policy relevance, the elasticity has been estimated by hundreds of researchers in recent decades. Nevertheless, the extensive research has not resulted in a consensus on the magnitude of the elasticity. In this paper we synthesize the estimated income elasticities of gasoline demand and try to provide a benchmark value of the elasticity based on the available empirical literature. To this end we employ meta-analysis, the set of methods designed for quantitative literature surveys.

Meta-analysis was developed in medical science to summarize the results of clinical trials; one of the first meta-analyses was Pearson (1904). Clinical trials are costly and often can only use a handful of observations; aggregation of the results of clinical trials on the same topic increases the number of degrees of freedom and improves the robustness and precision of the resulting estimated treatment effect. In the last few decades the methods of meta-analysis have spread from medical research to other fields; for example, the first meta-analysis in economics was Stanley and Jarrell (1989). The excellent survey by Nelson and Kennedy (2009) summarizes 140 meta-analyses conducted in environmental and natural resource economics since the early 1990s. Meta-analysis, we believe, is not a substitute for good narrative literature surveys, but complements them with a formal treatment of various biases potentially present in the literature.

At least since Rosenthal (1979), researchers conducting literature surveys have been concerned with the so-called file-drawer problem, or publication bias. When some results are strongly predicted by theory, researchers may treat opposite findings with suspicion. Such results are often difficult to publish, and researchers may choose to hide those counter-intuitive findings in their file drawers. The process can be unintentional and still result in publication bias; for example, if researchers use the “correct” sign of the estimated coefficient as a model selection test. The bias is particularly serious in medical research, and the best medical journals now require registration of clinical trials as a necessary condition for submission, so that the profession knows whether results end in file drawers (Krakovsky, 2004; Stanley, 2005). A well-known case of publication bias concerns the antidepressant drug Paxil, which was originally found to be effective by most published studies. When, however, unpublished results are included, the drug does not seem to outperform a sugar pill, and may have severe side effects (Turner et al., 2008).

☆ We are grateful to Zuzana Irsova, Karel Janda, seminar participants at Charles University, Prague, and two anonymous referees of Energy Economics for their helpful comments on previous versions of the manuscript. We thank Carol Dahl for providing us her data set with estimates of gasoline demand elasticities. We acknowledge support from the Czech Science Foundation (grant #P402/11/0948 and #15-02411S) and the Grant Agency of Charles University (grant #558713). An online appendix providing the data and the Stata code used in this paper is available at meta-analysis.cz/gasoline. The views expressed here are ours and not necessarily those of the Czech National Bank.
* Corresponding author at: Czech National Bank, Czech Republic. Tel.: +64 3 364 0549; fax: +64 3 364 0545.
E-mail address: tomas.havranek@ies-prague.org (T. Havranek).

http://dx.doi.org/10.1016/j.eneco.2014.11.004
0140-9883/

Motivated by the practice in medical science, the American Economic Association has been considering establishing a registry for controlled experiments because of potential publication bias (Siegfried, 2012, p. 648). But for non-experimental fields, such as the literature estimating the income elasticity of gasoline demand, meta-analysis tools remain the only way to correct for the bias. We suspect that negative estimates of the elasticity are reported less often than they should, which biases the mean estimate in the literature upwards. The reason is that negative estimates of income elasticity are counter-intuitive: it does not make much sense for gasoline demand to decrease with rising income. We expect that researchers unintentionally discard negative estimates (which imply that gasoline is an inferior good), even though they should report them from time to time because of the sampling error, especially if the true underlying elasticity is small. As discussed by Stanley and Doucouliagos (2012), such discarding of unintuitive results may paradoxically improve individual studies—it would not make much sense to build conclusions on negative estimates of the elasticity. But the literature as a whole gets biased upwards as the negative results become underrepresented.

To our knowledge, there has been one meta-analysis on the income elasticity of gasoline demand. Espey (1998) examines the heterogeneity in the estimates and reports mean elasticities of 0.47 for the short run and 0.88 for the long run, but she does not take publication bias into account. Additionally, two meta-analyses have been conducted on the price elasticity of gasoline demand: Brons et al. (2008) and Havranek et al. (2012). Similarly to Espey (1998), Brons et al. (2008) focus on the heterogeneity stemming from the different methods used by the authors estimating the elasticity, and do not control for publication bias. Havranek et al. (2012) show there is substantial publication bias in the literature on the price elasticity of gasoline demand: the mean estimate of the price elasticity seems to be exaggerated twofold because of publication selection.

We employ a large data set of gasoline demand elasticities collected and described by Dahl (2012). Because modern meta-analysis methods require information concerning the precision of the estimates of elasticities, we only use estimates for which standard errors or t-statistics are reported. The average reported elasticity for the short run is 0.28; for the long run it is 0.66. We find strong publication bias in the literature, especially for the estimates corresponding to the short run. To correct for publication bias we use mixed-effects multilevel meta-regression methods. The mixed-effects approach allows for between-study differences in the underlying elasticity, which is important because the studies in the data set estimate the elasticity for different countries. The method also assigns each study approximately the same weight, which is desirable because otherwise studies reporting many estimates would dominate the meta-analysis. Our results suggest that the corrected income elasticity of gasoline demand is, on average, only 0.1 for the short run and 0.23 for the long run. For the short run, for example, this is one-fifth the size of the number reported by the previous meta-analysis of Espey (1998); the difference is in part due to newer data and in part due to the correction for publication bias.

The remainder of the paper is structured as follows. In Section 2 we outline the basic models used for the estimation of the income elasticity of gasoline demand. In Section 3 we describe the meta-analysis techniques that we employ in this paper. Section 4 presents the results of our meta-analysis. Section 6 concludes the paper. The data and Stata code used for the estimation are available in an online appendix at meta-analysis.cz/gasoline.

2. Estimating the elasticity

In this section we briefly outline the econometric methods used for the estimation of gasoline demand elasticities. Energy demand exhibits unique features that do not allow researchers to treat it in the same way as demand for other consumer products. The main problem is that consumers do not demand energy directly; they demand transportation for which gasoline serves as an input, so researchers often work with demand for gasoline in the same way as with derived demand. While gasoline is a non-durable good, the dependence on durable goods complicates estimation. For example, as people demand certain amounts of travel, their gasoline consumption depends on the efficiency and price of vehicles. Over the last 40 years many potential approaches for the estimation of demand elasticities have been suggested, but no consensus on the best practice has been reached in the literature, as different researchers prefer different methodologies.

2.1. Static models

The models discussed over the decades have one thing in common—gasoline demand is modeled as a function of the price of gasoline and real income. Other explanatory variables may include the stock of vehicles, average vehicle efficiency, and prices of other inputs. The main difference between the models used in the literature is the way how the adjustment of gasoline demand to shocks in prices and income is laid out in time.

The so-called static models do not consider short-run adjustment, but only focus on the overall response in the long run. Dahl (2012) notes, however, that results from static models could be treated as estimates for the "intermediate run" because they often yield lower estimates compared with dynamic models. The benchmark static model can be specified as follows:

$$log\, G_t = \alpha + \beta_1\, log\, P_t + \beta_2\, log\, Y_t + \sum_{k=1}^{K} \beta_{k+2} Z_{kt} + u_t, \quad (1)$$

where G represents gasoline demanded, Y per capita income, P real prices, and Z_k other relevant explanatory variables, while the betas denote the corresponding elasticities. When estimating these types of regressions, of course, researchers have to make sure that the time series entering the model are stationary.

2.2. Dynamic models

The class of dynamic models, described in detail by Kennedy (1974) and Houthakker et al. (1974), assumes different consumer adaptation for the short run and long run. The demand function takes the following general form:

$$G^* = f_2(P, Y) = \alpha Y^{\beta} P^{\gamma}. \quad (2)$$

Given that the desired level in the short run may not match the actual demand for gasoline, demand adjusts over time toward the long-run level:

$$\frac{G_t}{G_{t-1}} = \left(\frac{G_t^*}{G_{t-1}}\right)^{1-\lambda}. \quad (3)$$

After substituting Eq. (2) into Eq. (3), taking the logarithm of both sides of the equation, and adding a disturbance term, we arrive at

$$log\, G_t = log\, \alpha + (1-\lambda)\beta\, log\, Y_t + (1-\lambda)\gamma\, log\, P_t + \lambda\, log\, G_{t-1} + u_t. \quad (4)$$

The regression coefficients corresponding to log Y_t and log P_t in Eq. (4) denote the short-run estimates of the income and price elasticities, respectively. Dividing them by $1-\lambda$, thus obtaining β and γ, we get the long-run estimates. Such an elegant combination of short- and long-run elasticities within one equation has made this model very popular.