PRELIMINARY: PLEASE DO NOT CITE WITHOUT PERMISSION

Characterizing Intertemporal Substitution via Pre-announced VAT Increase: Evidence from Japan

David Cashin*
University of Michigan

Version: October 23, 2012

*Department of Economics, University of Michigan, email: cashin@umich.edu. I would like to thank Joel Slemrod, Chris House, James Hines, Matthew Shapiro, Mel Stephens, Daniel Silverman, Justin Wolfers, Daniel Murphy, and seminar participants at the University of Michigan and the U.S. Treasury's Office of Tax Analysis for helpful comments and suggestions. I would especially like to thank Takashi Unayama for running the regressions.

Abstract

This study measures intertemporal substitution via a pre-announced increase in Japan's VAT rate from three to five percent. Matching the time path of expenditures generated by a dynamic structural model of household consumption to empirical estimates of the intertemporal substitution response to the VAT increase, I find that expenditure is sensitive to a change in the future price level due to accelerated purchases of durables and stockpiling of storables. However, consumption is relatively insensitive. The intertemporal elasticity of substitution in consumption (IES) is 0.07, and is precisely estimated. Because the IES is small, the results suggest policies that alter the future price level will have a limited impact on the timing of household expenditure.

Keywords: Intertemporal substitution, consumption, fiscal policy, VAT

JEL Classification: D12 E21 E6 H24 H31

1. Introduction

The sensitivity of household *expenditure* to a change in the future price level is of central importance to macroeconomics and public finance. The more sensitive is expenditure, the greater will be the impact of countercyclical policies that alter intertemporal price levels, such as nominal interest rate adjustments by the monetary authority or adjustments to the tax rate on expenditure (herein referred to as a 'consumption tax') by the fiscal authority. Alternatively, in an economy operating at full employment, the deadweight loss of tax policies that alter intertemporal prices is increasing in the sensitivity of household expenditure.

Cashin and Unayama (2012) take a novel approach to measuring the sensitivity of household *consumption* to a change in the future price level, using a pre-announced increase in Japan's Value Added Tax (VAT) rate from three to five percent as a natural experiment to estimate the intertemporal elasticity of substitution in consumption (IES). With a more appropriate categorization of non-durable goods and services than in previous studies, and an econometric specification that is robust to non-separable preferences over durables and non-durables, their point estimate of the IES is 0.21. This estimate is lower than previous estimates of the IES derived from survey data (e.g. Attanasio and Weber, 1993 and 1995; Vissing-Jorgensen, 2002), but the standard errors are sufficiently large such that the previous estimates cannot be ruled out either.

This study builds on the work of Cashin and Unayama (2012) in two ways. First, it characterizes the sensitivity of both household consumption and expenditure to the increase in Japan's VAT. For storable (e.g. laundry detergent) and durable (e.g. household appliances) goods, the timing of consumption and expenditure does not necessarily coincide. Storables can be stockpiled during low price periods for consumption in high price periods, while durables can be purchased during low price periods, with most of the flow of services generated by the durable consumed during a high price period. As a result, expenditure on these goods and services should be more sensitive to a change in the future price level than expenditure on (consumption of) non-storable non-durable goods and services. Given that storables and durables comprise a significant share of household expenditure, it is important to account for this behavior in order to fully characterize intertemporal substitution.[1] Second, this study utilizes information

[1] Expenditure on durables and storables accounted for 21 and 23 percent of expenditures subject to Japan's VAT during the sample period used in this study.

on the durable and storable expenditure responses to the VAT rate increase to estimate the IES. Cashin and Unayama (2012) examine only the response of non-storable non-durable expenditure (e.g. dining out) to the tax rate increase. The additional information provided by the durable and storable responses should yield a more precise estimate of the IES.

To characterize the sensitivity of household consumption and expenditure to the VAT rate increase, I rely on empirical estimates of the expenditure response coupled with a dynamic structural model of household consumption. Specifically, using household survey data, I estimate the percentage deviation in durable, storable, and non-storable non-durable expenditures in the months surrounding implementation of the VAT rate increase relative to a base month that followed announcement of the tax rate increase, but preceded implementation. As a result, the estimates should be devoid of any income effect resulting from the tax rate increase, instead capturing only intertemporal substitution effects in the form of stockpiling, accelerated purchases of durables, and intertemporal substitution in consumption. I then employ a dynamic structural model of household consumption that simulates the durable, storable, and non-storable non-durable intertemporal substitution response to the VAT rate increase. The model is governed by parameters such as the intertemporal elasticity of substitution in consumption (IES), the elasticity of substitution between durables and non-durables, and adjustment cost parameters for storables and durables that characterize the sensitivity of household consumption and expenditure to a change in the future price level. The parameter estimates chosen are the set generating a time-path of expenditures that most closely match the empirical estimates of the intertemporal substitution response to the VAT rate increase.

I find that expenditure was sensitive to the VAT rate increase only in the months immediately prior to and following implementation. Durable expenditures were 8 and 22 percent higher in the two months prior to the tax rate increase than they would have been in its absence, dropped sharply following implementation, and returned to trend within a few months of implementation. This suggests that households did indeed accelerate purchases that would have otherwise been made after the tax rate increase, and is further confirmed by the fact that expenditure on goods with higher levels of durability was more sensitive to the VAT increase.

Expenditure on storable goods was nine percent higher in the month prior to implementation than it otherwise would have been, dropped precipitously in the month of implementation, and gradually returned to trend in the ensuing months, indicating that households engaged in a significant amount of stockpiling just prior to the price increase. This point is further confirmed by the fact that expenditure on goods with higher levels of storability was more sensitive to the tax rate increase.

Finally, non-storable non-durable expenditures were 1.51 percent higher in the month prior to the tax rate increase than they would have been in its absence, but showed little variation in other months prior to and following implementation. The lack of variation in expenditure prior to and following implementation suggests that the IES is small. The slight jump in expenditure in the month prior to implementation suggests that durables and non-durables are complements, because the reduction in the user cost of durables in that month coincided with an increase in non-storable non-durable expenditure.[2]

The structural parameter estimates confirm these conjectures. The point estimate for the IES is 0.07, with a 95 percent confidence interval given by [-0.03, 0.17]. As expected, the additional information provided by the durable and storable intertemporal substitution responses to the VAT rate increase yields a more precise estimate of the IES. The intratemporal elasticity of substitution between the durable stock and non-durables is estimated to be -0.01, with a 95 percent confidence interval given by [-0.04, 0.02]. That is, I cannot reject that durables and non-durables are perfect complements. Adjustment costs significantly reduce the accelerated purchase of durables and the stockpiling of storables in the month prior to implementation (relative to the frictionless case). For example, the marginal cost of adjustment to the durable stock is 0.36 percent of the average household's monthly income in the month prior to the tax rate increase.

To evaluate the external validity and generalizability of the elasticity and adjustment cost parameter estimates, I then compare the time path of expenditures generated by the model to the response of durable and non-durable (storable and non-storable) retail sales to the July 1989 increase in New Zealand's Goods and Services Tax (GST) rate, which is examined in Cashin (2011). This particular tax rate increase was also announced prior to its implementation, and it featured an increase in the tax rate from 10 to 12.5 percent. The time path of durable and non-

[2] One might worry that the 1.67 percent increase in non-storable non-durable expenditure in March 1997 is due to arbitrage. For example, it is plausible that a household purchased bananas on March 31, 1997 for consumption in the first week of April 1997. However, by definition, arbitrage opportunities are limited for this group. Furthermore, Cashin and Unayama (2012) reject the arbitrage explanation, as there is little evidence of an equal and offsetting reduction in April 1997 non-storable non-durable expenditures.

durable expenditures generated by the model is similar to the observed response, which suggests the estimates presented in this paper are applicable in other contexts.

Finally, using the structural parameter estimates derived in this study, I simulate the potential impact of the recently proposed VAT increase in Japan, which would gradually increase the VAT rate from five to eight percent in April 2014, and then from eight to ten percent in October 2015. I find that the deadweight loss associated with the tax rate increase is just 0.0008 percent of the present value of lifetime steady state expenditure, and is insensitive to the date of announcement. These results are not surprising given the low estimate of the IES.

The results presented in this study suggest that policies that alter the future price level will have a limited impact on the timing of household expenditure. On the other hand, the deadweight loss associated with pre-announced increases in consumption tax rates will be small. In addition, the findings indicate that the assumption of perfect complements (i.e. Leontief preferences) is a good starting point when jointly modeling durable and non-durable consumption.

The remainder of the paper is organized as follows. Section 2 provides background on Japan's April 1997 VAT rate increase, explaining why the pre-announced tax rate increase presents an ideal natural experiment to measure intertemporal substitution. Section 3 provides an overview of the data, empirical specification, and empirical results. Section 4 introduces the model, estimation strategy, and the parameter estimates. Section 5 discusses the estimate of the IES in the context of previous studies. Section 6 concludes.

2. The VAT Rate Increase: An Ideal Natural Experiment to Estimate the IES

2.1. Japan's VAT and the April 1997 Rate Increase

Japan's 'Consumption Tax' is a VAT. Unlike VAT in many other countries, it has a single flat rate with relatively few exemptions.[3] The VAT was introduced in 1989 at a rate of three percent, and the rate was increased from three to five percent in April 1997. The 1997 VAT

[3] Exemptions included transfer of lease or land, transfer of securities and transfer of means of payment, interest on loans and insurance premiums, transfer of postal and revenue stamps, fees for government services, international postal money orders, foreign exchange, medical care under the Medical Insurance Law, social welfare services specified by the Social Welfare Services Law, midwifery service, burial and crematory service, transfer or lease of goods for physically handicapped persons, tuition, entrance fees, facilities fees, and examinations fees of schools designated by the Articles of the School Education Law, transfer of school textbooks, and the lease of housing units.

rate increase, which is the focus of this study, was originally proposed as a part of the Murayama Tax Reform, which passed through the Japanese Diet in late 1994.[4] Because the primary purpose of the reform was to continue the shift from direct to indirect taxation, the future VAT rate increase was coupled with immediate cuts in income tax rates.

Although the Murayama reform package set a target date of April 1997 for the VAT rate increase, it was unclear whether the increase would actually be implemented then. This is because the reform legislation also stated that the rate increase would be imposed only if the economy had sufficiently recovered from a prolonged recession from 1991 to 1993, and feeble growth thereafter. Having judged the economy to have sufficiently recovered, the ruling Liberal Democratic Party (LDP) decided to raise the tax rate as scheduled. The bill to raise the VAT rate passed through the Upper House on June 25, 1996, and the tax rate increase was scheduled to become effective on April 1, 1997.

Even after this passage, the LDP stated that they would revisit the issue of the tax rate increase when they submitted the fiscal year 1997 budget. The VAT rate increase was the central issue in October 1996 elections to the Lower House of the Diet, with the LDP's opposition promising to shelve the tax rate increase if elected. The LDP narrowly won the election, and on December 26, 1996, the government submitted the fiscal year 1997 budget, finally deciding to increase the VAT rate to five percent on April 1, 1997.

2.2. The VAT Rate Increase as a Natural Experiment

Estimation of the IES requires variation in the real interest rate, which is the price of current consumption relative to future consumption. Because the real interest rate is defined as the nominal interest rate minus the inflation rate, a change in the inflation rate will induce the necessary variation. As a result, the April 1997 VAT rate increase, which represented an exogenous increase in the future price level during a period in which nominal interest rates were stable, presents an ideal natural experiment to estimate the IES.

First of all, the tax rate increase can be regarded as an exogenous change in the future price level. Not only is it the case that the tax system is exogenous from the perspective of individual households, but it is also true that the impact of the tax rate increase is largely independent of consumer behavior. This is because the VAT by and large applies to expenditures

[4] For a discussion of the political process, see Ishi (2001) and Takahashi (1999).

regardless of the characteristics of the consumer, the point of purchase, or the type of goods purchased. Figure 1 shows the seasonally-adjusted month-to-month percentage change in the consumer price index for goods and services subject to the VAT. While inflation was negligible in most months prior to and following implementation of the tax rate increase, the price level increased by just under two percent between March and April 1997, which is consistent with full forward shifting of the VAT onto consumers at the time of implementation. As a result, focus can be placed on a one-time price change, and I can ignore the influence of an additional factor (i.e. variation in pre-tax prices) that affects the real interest rate.

The influence of the nominal interest rate on the real interest rate can also be ruled out. Figure 2 presents the average contracted interest rates on short-term loans and discounts, which are the average interest rates applied to a contract of less than one year between a commercial bank and lender. The average interest rate fell precipitously throughout 1995, but remained relatively constant thereafter. This suggests that households would not change their nominal interest rate expectations in the months surrounding implementation of the VAT rate increase. In other words, households should not have expected any changes in nominal interest rates by the central bank that would offset or augment the intertemporal substitution incentives.

These facts imply that the tax rate increase can be regarded as an exogenous change in the real interest rate, which allows for consistent estimation of the intertemporal substitution response using ordinary least squares (OLS). Previous studies of intertemporal substitution (e.g. Hall, 1988; Attanasio and Weber, 1993 and 1995; Ogaki and Reinhart, 1998) have relied on an instrumental variables approach to address the well-documented endogeneity between the real interest rate and consumption growth.[5] However, there are several potential issues with the instruments that have been employed. First, as Yogo (2004) notes, it is notoriously difficult to predict the real interest rate, and therefore, some of the previous studies in this literature (especially those using aggregate data) suffer from the weak instrument problem. Weak instruments lead to estimates of the IES biased in the direction of OLS, which itself is likely to suffer from a downward bias.[6,7] Even if the weak instrument problem is overcome, there still

[5] For example, an increase in the real interest rate will induce an income effect in addition to the intertemporal substitution effect. If households are net savers, then failure to account for the innovation in income will lead to an OLS estimate of the IES that is biased downwards (see Appendix Figure A.1 for a simple demonstration in a two-period setting).

[6] Two stage least squares (2SLS) estimators using weak instruments are biased in the direction of OLS for the following reason. Suppose the structural equation is given by $y_i = \beta x_i + \eta_i$, and the first stage equation by

exists the potential for correlation between the instrument(s) and the expectational error term, which I discuss further in Section 5. Furthermore, Attanasio and Weber (1993, 1995) show that studies using lagged instruments and aggregate non-durable expenditure data suffer from a downward bias in estimates of the IES known as aggregation bias. [8] This study avoids these issues by using an exogenous institutional price change.

While exogenous variation in the real interest rate is a necessary condition for estimating the IES, it must also be the case that households are aware of the change. While I do not provide direct evidence on household awareness of the VAT rate increase, indirect evidence is available in the form of news coverage regarding the VAT rate increase prior to its implementation. Figure 3 reports the number of articles that mention the phrase 'Consumption Tax' in the Nihon Keizai Shimbun, Japan's leading business newspaper with a circulation of over three million (in 2010), and the Yomiuri Shimbun, a leading non-business newspaper with a circulation of over 10 million (in 2010).[9] There was a steady upward trend that began just prior to enactment of the June 1996 legislation. Coverage peaked in the Yomiuri Shimbun in October 1996, which coincided with elections to the Lower House of the Diet. Overall coverage in both papers was consistently high in the months following the election but prior to the tax change, with nearly 300 articles in the Nihon Keizai Shimbun mentioning the 'Consumption Tax' in March 1997. This suggests that households were aware of the tax rate increase and might therefore engage in intertemporal substitution behavior.

The news coverage also suggests that households were aware of the effects of the Murayama reform package as a whole. Figure 3 shows that coverage initially peaked in September 1994, which coincided with the passage of the Murayama reform package. Accordingly, households may have known about the VAT rate increase well in advance of its

---

$x_i = \pi z_i + \xi_i$. If π is truly zero due to weak instruments, then any variation in the predicted value of x_i, $\hat{x}_i$, will come from ξ_i. It follows that the variation in $\hat{x}_i$ is no different from the variation in x_i, and the OLS and IV estimates are estimating the same quantity on average. For more information, see Pischke (2010).

[7] Using OLS, Gruber (2006) obtains an estimate of the IES of -0.55, which is significantly less than his estimates when instrumenting for the after-tax real interest rate. Vissing-Jorgensen (2002) finds that estimates of the IES converge towards zero as the number of instruments is increased. This is because the weak instrument problem is increasing in the degree of overidentification.

[8] Attanasio and Weber (2010) sum up aggregation bias as follows: "The aggregate consumption growth rate is computed by taking logs of the mean of individual consumption, whereas [the log-linearized Euler equation] implies that means of the logs should be taken instead…the difference between these two terms is highly serially correlated, thus invalidating lagged consumption growth as an instrument."

[9] Circulation numbers come from Japan's Audit Bureau of Circulations.

implementation, and furthermore, that it was intended to be compensated. As a result, one would expect any income effect associated with the tax reform to be small, and to have taken effect well in advance of its implementation. This conjecture is important because it will allow me to assume that deviations in expenditure around the time of the VAT rate increase are due solely to intertemporal substitution.

In addition to public awareness of the VAT rate increase, it seems likely that households expected to bear the burden of the tax rate increase in the form of higher prices at the time of implementation. For one, when VAT was imposed in April 1989 at a rate of three percent, the price of goods and services that had not previously been subject to tax increased by just under three percent upon implementation. Furthermore, the Japanese government made it clear that they expected consumers to bear the full burden of the VAT increase. [10] It is also worth noting that Carroll et al. (2010) find that full forward shifting at the time of a consumption tax rate increase is the norm across most countries, likely as a result of factor price rigidities.

A final reason that Japan's VAT rate increase presents an ideal natural experiment for estimation of the IES is that the relative pre-tax price of goods and services did not change around the time of the VAT rate increase. Figure 4 shows the price of durables and storables relative to non-storable non-durables in the months surrounding the VAT rate increase. As the figure demonstrates, there was little change in the relative pre-tax price of these goods. This finding, coupled with the fact that the nominal interest rate was constant, is important because it will greatly simplify estimation of the intratemporal elasticity of substitution between durables and non-durables, which Ogaki and Reinhart (1998) show is critical to proper estimation of the IES when preferences over durables and non-durables are non-separable. [11] The intratemporal elasticity of substitution is identified off of variation in the user cost of durables relative to non-durables. The user cost is a function of the current and future VAT rate, the current and future pre-tax price of durables relative to non-durables, and the nominal interest rate. Since the nominal interest rate and the pre-tax price of durables relative to non-durables were constant, the

10 When the VAT was introduced in 1989, the government took several steps to ensure this outcome. First, a Special Council on the Transition was formed to promote enforcement of the VAT across agencies. Second, the government carried out an extensive advertising campaign to allay the public's fear of price hikes and to restrain overcharging by traders. A telephone service was also set up so consumers could report complaints about prices. Finally, the Economic Planning Agency increased the budget for the price monitoring system. The situation was nearly identical in 1997.

11 See Footnote 24 for the intuition behind this result.

intratemporal elasticity of substitution can be identified solely off of the exogenous change in the VAT rate.

To summarize, the April 1997 VAT rate increase presents an ideal natural experiment to estimate the IES for the following reasons: the tax rate increase can be regarded as an exogenous change in the real interest rate; the real interest rate was relatively stable prior to and following implementation; the tax rate increase was predictable and consumer awareness was high; and there was little to no change in the relative pre-tax prices among goods and services around the time of the tax rate increase.

3. The Intertemporal Substitution Response to the VAT Rate Increase

3.1. Data

The Japanese Family Income and Expenditure Survey (JFIES) is used to estimate the intertemporal substitution response to the VAT rate increase.[12] The JFIES is a rotating panel survey in which households are interviewed for six consecutive months, and approximately 8,000 households are interviewed each month.[13]

The estimates make use of JFIES data from the period between April 1992 and March 2002, a symmetric five-year window around the April 1997 tax rate increase. Data from the "bubble" period (before April 1992) are excluded because household expenditures grew at a much faster pace than they did after the bursting of the economic bubble in 1991, while remaining more or less flat after that. The sample period ends in March 2002, which coincided with the beginning of another boom.

The sample is limited to households who complete all six interviews, but nearly all households can be used, as the response rate of the JFIES is quite high. Although data for agricultural households is available in the JFIES after 1999, they are excluded from the analysis to maintain consistency over the sample period. Also, the analysis restricts the sample to male-headed households and those whose head does not change his job. The latter restriction is imposed because March is the end of the fiscal year in Japan. As a result, several job changes are

[12] See Stephens and Unayama (2011, 2012) for more information regarding the JFIES design and content.
[13] Until 2002, single-person and agricultural households were excluded from the JFIES. As of the 2009 JFIES, single-person households comprised 11.8 percent of the population and were responsible for 18.1 percent of expenditures, while agricultural households accounted for 2.0 percent of the population, and 2.1 percent of expenditures.

observed, which may cause systematic changes in consumption around the time of the VAT rate increase. After imposing the sample restrictions, the dataset includes 646,900 observations from 129,380 households. Table 1 presents summary statistics for the sample.

The JFIES expenditure data is highly disaggregated by item type, which allows for an accurate categorization of goods and services. For the purposes of this study, it is critical to distinguish not only between taxable and tax-exempt goods and services, but also between durables, storables, and non-storable non-durables.

To construct the expenditure data, expenditure on goods and services that are not subject to the VAT are excluded. As shown in Table 1, expenditure on taxed items comprised 70% of total expenditure, while most tax-exempt expenditure consists of rent for housing and education (e.g. tuitions for school).[14]

The second step is to divide goods and services that are subject to the VAT into three sub-categories: durables (D), storables (S), and non-storable non-durables (N). N are defined as goods and services which are neither storable nor durable. That is, they depreciate relatively quickly over time when not in use, and when in use, are fully consumed. For example, fresh fruit, if not eaten, will spoil, and is fully consumed with use. This category also includes services such as taxi fare and dining out, which are consumed at the point of purchase. It follows that monthly expenditure on N should approximately coincide with its consumption.

S are defined as goods and services that depreciate slowly over time if not used and fully if used. For example, laundry detergent can be stored for long periods of time with little to no effect on its ability to clean clothing, but once it is put into use, whatever amount was used has been fully consumed. This category also includes public transit (rail and bus) passes, due to the fact that many Japanese households purchase passes which are good for train travel for several months after first use. Thus, one might expect that a household would purchase a pass good for several months during a low price period, and use the pass during a relatively high price period. More generally, the characteristics that define S allow for stockpiling during low price periods in order to consume in relatively high price periods. As a result, monthly expenditure on S does not necessarily coincide with consumption, and expenditure on S should be more sensitive to changes in intertemporal prices than N.

[14] I do not use tax-exempt goods and services as a control group because they are largely necessities, and as such, are unlikely to respond to changes in economic circumstances in the same manner that taxable items would.

Finally, D are defined as goods and services which depreciate relatively slowly over time if not used and do not depreciate fully with use. This category includes traditional durables such as refrigerators and automobiles, as well as goods such as clothing and footwear that are classified as semi-durables in the JFIES. In addition, this category includes a select group of services such as home repair and tailoring, which consumers derive benefits from long after the service is provided. Like S, expenditure on D should be more sensitive to changes in intertemporal prices than N. This is because D can be purchased during a low price period, with most of its service flow consumed during a relatively high price period.[15] See Appendix Table A.1 for a complete categorization of N, S, and D.

Monthly expenditures on N, S, and D are then deflated using tax-inclusive consumer price indices specific to each category.[16] The analysis thus makes use of real monthly expenditures for Japanese households from April 1992 through March 2002. Table 1 shows that more than half of taxable expenditure is on N, while expenditure on S and D is similar.

3.2. Empirical Model

Suppose that the logarithm of real monthly expenditure by household h on good-type $j \in \{D, S, N\}$ in year y and month m can be expressed as

$$\ln E^{j}_{h,y,m} = \mu^{j}_{h} + \boldsymbol{\delta^{j}_{m} Z_{m}} + \boldsymbol{\phi^{j} X_{h,y,m}} + \boldsymbol{\gamma^{j}_{y,m} D_{y,m}} + B^{j}_{y,m} + \epsilon^{j}_{h,y,m} \qquad (1)$$

where μ^{j}_{h} is a household fixed effect; $\boldsymbol{Z_{m}}$ is a vector of month dummies intended to capture seasonality effects; $\boldsymbol{X_{h,y,m}}$ is a vector of (potentially) time-varying household characteristics, including the number of household members, the number of workers, the number of household members under the age of 18, the number of household members over age 65, and interview dummies, which control for "survey fatigue", the tendency of households to report lower expenditure in later interviews; $\boldsymbol{D_{y,m}}$ is a vector of dummies for months surrounding the VAT rate increase, where $\boldsymbol{\gamma^{j}_{y,m}}$, the coefficients of interest, are intended to capture the (approximate) percentage deviation in expenditure on good j relative to some base month, which in practice is

[15] Barrell and Weale (2009) refer to this as an 'arbitrage' effect.
[16] In particular, Laspeyres price indices are constructed for each of the four categories using item-specific price indices and expenditure shares in 1990 for each of these items as the weights.

the month preceding the first $D_{y,m}$ dummy;[17] $B_{y,m}$ accounts for aggregate factors other than the tax rate increase that impact household expenditure, such as the business cycle and other policies that impact household expenditure; and $\epsilon^j_{h,y,m}$ accounts for unobservables that impact monthly household expenditures on good-type j.

Taking the first difference of (1) removes the household fixed effect, which yields

$$\Delta\ln E^j_{h,y,m} = \boldsymbol{\Delta}(\boldsymbol{\delta^j_m Z_m}) + \boldsymbol{\phi^j \Delta X_{h,y,m}} + \boldsymbol{\Delta}(\boldsymbol{\gamma^j_{y,m} D_{y,m}}) + \Delta B^j_{y,m} + \Delta\epsilon^j_{h,y,m} \qquad (2)$$

In order to separately identify the impact of the VAT rate increase on household expenditures from the impact of changes in $B^j_{y,m}$, additional restrictions must be placed on $B^j_{y,m}$. Suppose that $B^j_{y,m}$ follows either of the two conditions listed below:

1) There is no change in $B^j_{y,m}$ from one month to the next.
2) $B^j_{y,m}$ follows a linear trend.

Under condition (1), the term $\Delta B^j_{y,m}$ drops out, while under condition (2), the term $\Delta B^j_{y,m}$ becomes a constant, c^j. More generally, if there is little change in $B^j_{y,m}$ other than the linear trend, (2) can be rewritten as

$$\begin{aligned}\Delta\ln E^j_{h,y,m} &= c^j + \boldsymbol{\Delta}(\boldsymbol{\delta^j_m Z_m}) + \boldsymbol{\phi^j \Delta X_{h,y,m}} + \boldsymbol{\Delta}(\boldsymbol{\gamma^j_{y,m} D_{y,m}}) + (\Delta B^j_{y,m} - c^j + \Delta\epsilon^j_{h,y,m}) \\ &= c^j + \boldsymbol{\Delta}(\boldsymbol{\delta^j_m Z_m}) + \boldsymbol{\phi^j \Delta X_{h,y,m}} + \boldsymbol{\Delta}(\boldsymbol{\gamma^j_{y,m} D_{y,m}}) + \Delta\mathrm{v}^j_{h,y,m} \qquad (3)\end{aligned}$$

where $\Delta v^j_{h,y,m}$ is composed of $\Delta\epsilon^j_{h,y,m}$ and perturbations in $\Delta B^j_{y,m}$ from c^j.

3.2. Empirical Specification

Recall that Japan's VAT rate increase took effect in April 1997. I am therefore interested in (percentage) deviations in expenditure in the months prior to and following April 1997. These

[17] Halvorsen and Palmquist (1980) demonstrate that in regressions with a logarithmic dependent variable, it is incorrect to interpret the coefficient on a dummy variable multiplied by 100 as the percentage effect of that variable on the variable being explained. Nonetheless, when the coefficients on the dummy variables are close to zero, as is the case in this study, multiplying the coefficient by 100 provides a good approximation to the actual percentage effect of the variable on the variable being explained.

deviations will inform us of the nature of intertemporal substitution, such as whether it is driven by changes in the timing of consumption (i.e. intertemporal substitution in consumption), changes in the timing of expenditure (e.g the stockpiling of storables), or both, and will be used in conjunction with the dynamic structural model of household consumption introduced in the next section to generate the structural parameter estimates.

The baseline specification used to generate the empirical estimates of the expenditure response to the VAT rate increase is

$$\Delta \ln E^{j}_{h,y,m} = c^{j} + \boldsymbol{\Delta}(\boldsymbol{\delta^{j}_{m} Z_{m}}) + \boldsymbol{\phi^{j} \Delta X_{h,y,m}} + \sum_{y=1997,m=1}^{y=1997,m=12} \Delta(\gamma^{j}_{y,m} D_{y,m}) + \Delta \nu^{j}_{h,y,m} \qquad (4)$$

where $\gamma^{j}_{y,m}$ is the average percentage deviation in household expenditures on good-type j in year y and month m relative to December 1996, after controlling for household fixed effects, the linear trend in expenditure growth, seasonality, and time-varying household characteristics. [18, 19] I choose December1996 as the base month against which expenditure in the months surrounding the VAT rate increase are compared because it coincided with passage of the fiscal year 1997 budget, which made the tax rate increase a certainty, and because news coverage of the proposed tax rate increase had been high for the previous few months. Therefore, households knew that the tax rate increase would be implemented in April 1997 as planned, and should have responded accordingly no later than December 1996. This further implies that the $\gamma^{j}_{y,m}$ will capture only the stockpiling, accelerated durable purchase, and (negative) intertemporal substitution in consumption effects associated with the VAT rate increase, as intended. Nonetheless, it is worth

[18] To avoid taking the logarithm of zero, monthly durable expenditure is set to ¥100, or approximately U.S. $1, in months that a household reported zero expenditure on durables. The results are robust to different choices of minimum durable expenditure values (e.g. ¥1 or ¥1000). Overall, 94 percent of the monthly observations report positive durable expenditures.

[19] As a robustness check, year dummies were also added to the empirical specification given in Equation (4). These dummies capture average monthly growth rates in household spending within a year relative to the omitted year (captured by c), and are important if growth rates varied considerably over the sample period. I find that inclusion of year dummies does not significantly impact the results of the baseline specification. Furthermore, it is worth noting that all standard errors are robust to arbitrary forms of serial correlation within households over time.

noting that the empirical estimates and the structural parameter estimates presented in Section 5 are robust to the choice of other base months as well. [20]

In order to make the assertion that $\hat{\gamma}_{y,m}^{j}$ is capturing only the intertemporal substitution effects resulting from the VAT rate increase, it must be that $\Delta B_{y,m}^{j} \approx c$. Otherwise, $\hat{\gamma}_{y,m}^{j}$ will be biased upwards when there is a significantly positive change in $\Delta B_{y,m}^{j}$ in a particular month, and vice versa. As a result, I do not utilize the estimates of $\gamma_{y,m}^{j}$ beyond July 1997 in the structural estimation procedure described below. There are two reasons for this. First, household expenditures on N increased significantly in the third quarter of 1997. I suspect that this is due to unseasonably warm weather that summer that lead to high cooling costs, but I do not yet have the data to test this hypothesis. Second, and more important, is the impact of the Japanese banking crisis, which began in November 1997 and likely had a negative impact on household expenditure at the end of 1997 and beyond. The structural estimation will thus make use of the estimates of $\gamma_{y,m}^{j}$ from January 1997 through July 1997.

3.3. Identification of $\gamma_{y,m}^{j}$

It may not be readily apparent how first differenced year-month dummies for the months surrounding the VAT rate increase identify $\gamma_{y,m}^{j}$. Figure 5 demonstrates how they do so. In the top portion of this hypothetical example, a household engages in stockpiling in March 1997, the month prior to implementation, which leads to an increase in storable expenditure relative to previous months, captured by γ_{Mar}. In April 1997, there is an equal and offsetting stockpiling effect, as well as the (negative) intertemporal substitution in consumption effect. The combined impact of these two effects is captured by γ_{Apr}.

When taking the first difference of expenditure, as depicted in the bottom portion of Figure 5, it is clear that inclusion of a dummy variable for April 1997 will yield a coefficient equal to $\gamma_{Apr} - \gamma_{Mar}$ rather than γ_{Apr}. The solution to this problem is to difference out the effect from the previous month. In practice, this means including the first difference of the March 1997 dummy in the empirical specification, rather than just a March 1997 dummy. That is, a dummy that takes on a value of 1 in March 1997 and a dummy that takes on a value of -1 in March 1997,

[20] In particular, I chose October 1996, the election month that made the tax rate increase a relative certainty, as an alternative base month.

with the dummies sharing a common coefficient. Doing so, the April 1997 dummy will capture γ_{Apr} as intended.[21]

3.4. Empirical Estimates of the Intertemporal Substitution Response to the VAT Rate Increase

Figure 6 presents estimates of $\gamma_{y,m}^{j}$ for durables, storables, and non-storable non-durables for January to July 1997, along with the corresponding 95 percent confidence intervals. On average, there was little change in non-storable non-durable expenditures prior to and following implementation of the VAT rate increase. Note, however, that non-storable non-durable expenditures were 1.51 percent higher in March 1997 than they otherwise would have been, which is significant at the ten percent level. Given that the user cost of durables fell in the same month, this result suggests that durables and non-durables are complements, while the lack of variation in expenditures prior to and following implementation suggests the IES is small. The intuition for both conjectures is discussed further in Section 4.4 with the aid of the model.

The top right graph in Figure 6 shows the intertemporal substitution response for storables. Expenditure in March 1997 was nine percent higher than it otherwise would have been. In April 1997, storable non-durable expenditure was 7 percent lower than it would have been in the absence of the VAT rate increase, and gradually increased over the next few months. This pattern suggests that households stockpiled goods just prior to implementation, and then consumed from their storable inventory over the next few months.

This explanation is further reinforced by comparing the intertemporal substitution response for storable non-durables that possess different levels of storability. Figure 7 examines the response of domestic household goods (e.g. laundry detergent, toilet paper), personal care items (e.g. medicine, shaving cream), beverages (alcoholic and non-alcoholic), and storable foods (e.g. butter, noodles, yogurt) to the VAT rate increase. Domestic household goods and personal care items are storable for long periods of time, while beverages and foods are storable for a relatively shorter period. As Hendel and Nevo (2006) note, this is at least in part because storability for the latter groups decreases once the container or packaging is opened. I find that expenditure is more sensitive for goods with higher levels of storability, which is consistent with the consumer inventory model of stockpiling behavior (see Hendel and Nevo, 2004 and 2006).[22]

[21] The empirical specification is robust to non-symmetric stockpiling effects as well.

[22] I also plan to examine whether there are differences in the expenditure response for goods with similar levels of storability, but different standard sizes (e.g. aluminum foil v. laundry detergent).

It is also worth noting that there was a highly significant intertemporal substitution response for public transit passes in March 1997, as hypothesized in Section 3.1.[23]

The bottom left graph in Figure 6 presents estimates of the durable intertemporal substitution response. Expenditures in the final two months prior to implementation were 8 and 23 percent higher than they otherwise would have been. Expenditure in April 1997 was 13 percent lower than it would have been in the absence of the VAT rate increase, and gradually returned to trend over the next few months. This pattern is consistent with accelerated purchases of durables that would have otherwise been bought after the tax rate increase.

Figure 8 lends further support to this conjecture. Note that the intertemporal substitution response in March 1997 was largest for furniture and household appliances, followed by consumer electronics, with almost no response for automobiles. Perhaps not coincidentally, there is some evidence of an inverse relationship between the expenditure response observed in March 1997 and the depreciation rate estimates associated with each good type in Fraumeni (1997).[24] All else equal, a dynamic model of durable consumption would predict that expenditure in the month prior to implementation would be more sensitive for goods with lower depreciation rates, and the estimates in Figure 8 are generally consistent with this prediction. Finally, it is interesting to note that the dip in durable expenditures in July 1997 is due primarily to a reduction in expenditures on household appliances, and specifically, air conditioners. This suggests that households were forward looking enough to purchase air conditioners in March 1997 that would not be used until later in the year.

In summary, the empirical results suggest that the timing of expenditure was sensitive to the VAT rate increase, but the timing of consumption was not. And while expenditure did respond to the price change, the response was confined to the months immediately preceding and following implementation.

4. Characterizing the Intertemporal Substitution Response to the VAT Rate Increase

[23] Storable non-durable expenditure does not return to its initial level following implementation. This could be due to intertemporal substitution in consumption, but there was also a general decline in storable expenditure that began in 1996, and continued beyond implementation of the VAT rate increase. Because the empirical specification assumes a linear trend in expenditure growth over the entire sample period, it may not entirely capture this decline, which was only evident for a portion of the sample period.

[24] Specifically, the annual rates of depreciation given in Fraumeni (1997) are the following: furniture (0.12), household appliances (0.15), home electronic equipment (0.18), and motor vehicles (0.17).

4.1. The Household's Problem

This section develops a dynamic model of household consumption of durables, storables, and non-storable non-durables that mimics the environment that characterized the Japanese economy around the time of the VAT rate increase. Its purpose is to allow for the estimation of structural parameters that fully characterize the intertemporal and intratemporal substitution in consumption, stockpiling, and accelerated durable purchase response to the VAT rate increase. Furthermore, unlike the standard approach to estimation of the IES, the model will allow me to incorporate the storable and durable expenditure responses into estimates of this important policy parameter, which should yield more precise estimates.

The model is constructed as follows. In each period t, taken here to be one month, a representative household chooses non-storable non-durable consumption, C_t^N; storable consumption, C_t^S; storable expenditure, X_t^S; the stock of storables, S_t, that will be carried over into period $t+1$; the durable stock, D_t, which provides a flow of consumption services; durable expenditure, X_t^D; and financial assets, A_t, to maximize the present value of expected lifetime utility subject to the budget constraint, laws of motion for S_t and D_t, and stochastic processes for the tax rate on expenditure, τ_t, and income, Y_t. Formally, the household solves

$$\max \sum_{t=0}^{\infty} \beta^t E_0 \left\{ \left(\frac{1}{1-\frac{1}{\sigma}}\right) \left(\left[(1-\psi^D)^{\frac{1}{\epsilon^D}} \left\{ \left[(1-\psi^S)^{\frac{1}{\epsilon^S}} {C_t^N}^{1-\frac{1}{\epsilon^S}} + {\psi^S}^{\frac{1}{\epsilon^S}} {C_t^S}^{1-\frac{1}{\epsilon^S}} \right]^{\frac{\epsilon^S}{\epsilon^S-1}} \right\}^{1-\frac{1}{\epsilon^D}} + {\psi^D}^{\frac{1}{\epsilon^D}} {D_t}^{1-\frac{1}{\epsilon^D}} \right]^{\frac{\epsilon^D}{\epsilon^D-1}} \right)^{1-\frac{1}{\sigma}} \right\}$$

subject to

1) $A_t = (1+i)A_{t-1} + Y_t - (1+\tau_t)(C_t^N + X_t^S + X_t^D) - \frac{\zeta^D}{2}(D_t - D_{t-1})^2 - \frac{\zeta^S}{2}(S_t - S^*)^2$

2) $S_t = S_{t-1} + X_t^S - C_t^S$

3) $D_t = (1-\delta)D_{t-1} + X_t^D$

4) $\tau_t = \tau_{t-1} + \epsilon_{t^*-31}^{\tau}$

5) $\epsilon_{t^*-31}^{\tau} = 2$

6) $Y_t = Y_{t-1} + \epsilon^Y_{t^*-31}$

7) $\epsilon^Y_{t^*-31} = 1.94$

Intertemporal preferences are assumed to be iso-elastic and governed by the IES, σ, which is one of the parameters to be estimated. Note that the value of σ will be determined by changes in C_t^N, C_t^S, and D_t in response to the VAT rate increase. As noted earlier, a potential advantage of the approach employed in this study is that information from all three goods will be used to determine σ.

The intratemporal preference specification is assumed to take a nested constant elasticity of substitution (CES) form. Preferences over the durable stock and a non-durable composite good are governed by ψ^D, a parameter measuring the overall importance of the durable stock in generating utility, and ϵ^D, the elasticity of substitution between durables and non-durables.[25] The value of ϵ^D will also be estimated using the expenditure response to the VAT rate increase. In particular, it is identified off of the change in the durable to non-durable consumption ratio resulting from the reduction in the user cost of durables prior to the VAT rate increase. It is worth noting that because I allow for non-separable preferences over durables and non-durables (i.e. ϵ^D is not restricted to be equal to σ), the estimate of σ should be free from intratemporal substitution bias (see Ogaki and Reinhart, 1998).[26] Preferences over the non-durable composite good are also assumed to be of the CES form, where ψ^S is the share of storable non-durables in non-durable consumption, and ϵ^S is the elasticity of substitution between storable and non-storable non-durables.[27, 28]

[25] As in many previous studies of durable goods, such as Bernanke (1985), I assume that the service flow derived from durables is proportional to the durable stock.

[26] Intratemporal substitution bias is a problem when preferences between durables and non-durables, assumed to be separable, are in fact non-separable, and households substitute between durables and non-durables at the same time they engage in intertemporal substitution. To illustrate, suppose the real interest rate rises. Because the user cost of durables is an increasing function of the real interest rate, the user cost of durables also rises. If durables and non-durables are substitutes, then households will substitute away from durables towards non-durables. Under the assumption of separable preferences, the IES can be identified by examining only the non-durable consumption response. Because households substituted away from durables and towards non-durables when the interest rate is high, the resulting estimate of the IES will be biased downwards.

[27] ψ^S represents a share only if preferences are Cobb-Douglas, or if prices are assumed to be one, as in this study.

[28] The nested CES form restricts the intratemporal elasticity of substitution between storables and durables and non-storables and durables to be the same. To test the validity of this assumption, I used a quadratic specification that allowed for an interaction term between storables and durables, as well as non-storables and durables. After

Note that preferences over the durable stock and the non-durable composite good are assumed to be homothetic. In contrast, Pakos (2004) provides evidence that durables are luxuries and non-durables are necessities, and that the assumption of homotheticity when preferences are actually non-homothetic biases estimates of ϵ^D upward. However, given the relatively short period of time this study concerns itself with, the modest increase in the VAT rate, the fact that the VAT rate increase was compensated, and that any income effect associated with the tax reform should have occurred prior to the period I am concerned with, the assumption of homotheticity seems innocuous.

The following assumptions are made with respect to prices. The nominal interest rate, i, is constant, since Japan's benchmark nominal interest rate was constant in the years prior to and following the VAT rate increase. The model abstracts from time-varying pre-tax prices on the three composite goods, as the price ratios for these goods were stable during the period of interest. Finally, the burden of the VAT is assumed to fall entirely on the representative consumer in the form of higher prices, which is consistent with Japan's experience and the experience of other countries with a VAT.

The budget constraint also includes durable and storable non-durable adjustment cost functions, both of which take quadratic functional forms.[29] The durable adjustment cost function is intended to capture the time cost of shopping for durable goods and services. This is because a durable purchase is an infrequent event requiring more effort than a non-durable purchase. The adjustment cost is increasing and convex in net expenditure, reflecting the fact that the time devoted to shopping for a durable is likely increasing in expenditure, and the opportunity cost of one's time is an increasing and convex function. The parameter associated with the adjustment cost function, ζ^D, is another parameter which I will estimate based on the expenditure response to the VAT rate increase. It is identified by the difference between the durable expenditure response that would be observed in March 1997 (holding the other parameters fixed) in the absence of frictions and the observed response, and is increasing in that difference.

doing so, the implied intratemporal elasticities of substitution between storables and durables and non-storables and durables were similar.

[29] Including the adjustment cost functions in the budget constraint rather than the preference specification does not impact the time path of expenditures generated by the model, though it will impact the point estimate for the parameters ζ^D and ζ^S.

The adjustment cost function for storable non-durables depends on two parameters, S^* and ζ^S. S^*, which is assumed to be greater than zero, is a storable inventory bliss point. $S_t > S^*$ generates a cost to the household due to space constraints and the time cost associated with stockpiling. $S_t < S^*$ also generates a cost due to the inconvenience of holding too few storables. For example, there is a time cost associated with having to make a shopping trip to pick up a new tube of toothpaste after the previous tube runs empty. ζ^S is a parameter to be estimated based on the expenditure response to the VAT rate increase. In particular, it will be identified primarily by the difference between the storable non-durable expenditure response in March 1997 that would be observed in the absence of frictions (again holding the other parameters fixed) and the observed response, and is increasing in that difference.

The law of motion for the stock of storable non-durables is the same as that used by Hendel and Nevo (2006). Note that the stock of storable non-durables carried over from one period to the next does not depreciate. This seems plausible for highly storable items like laundry detergent, but perhaps less so for storable foods that have been opened. In effect, I assume that these foods are fully consumed before they perish. The law of motion for durables is standard and depends on the durable depreciation rate, δ.

Recall that Japan's VAT rate increase was part of a compensated tax reform package introduced in September 1994, 31 months prior to its implementation in period t^*. The model abstracts away from the income tax cuts that took effect immediately upon passage of the reform package, instead compensating households for the two percentage increase in the tax rate on expenditure with a 1.94 percent increase in income. The tax rate increase and the offsetting compensation are known to the representative household well before t^*, and thus any change in expenditure around the time of the tax rate increase that is generated by the model will be attributable to intertemporal substitution.

The model does not account for the labor/leisure decision. This is a simplifying assumption, as I do not have access to monthly labor supply data. In addition, the time period examined is short, and I find it unlikely that households immediately adjusted their labor supply in response to such a modest, and compensated, change in the VAT rate. Finally, recall that households whose job status changed around the time of the VAT rate increase are excluded from the empirical estimates, which further mitigates any impact the VAT rate increase may have had on the labor/leisure decision.

4.2. Econometric Methodology

To estimate the model, the parameters are separated into three groups. The first group includes i, β, ψ^S, δ, ϵ^S, and S^*, which are fixed prior to estimation based on available data. The value of i is set to 0.0015, which corresponds to an annual interest rate of 0.018, or 1.8 percent. This was the average annual interest rate on short-term loans and discounts prior to and following the VAT rate increase (see Figure 2). β is set such that $\beta(1+i)=1$, because the model begins in steady state. ψ^S is set to 0.29, which was storable non-durable expenditure as a share of non-durable expenditure in the JFIES in 1996. The value of δ is set to 0.022, which corresponds to an annual depreciation rate of 0.23. This value was computed by combining good-specific annual depreciation rates from Fraumeni (1997) with good-specific expenditure shares on durables from the JFIES. The value of ϵ^S has no impact on the time path of expenditures generated by the model, because there is no change in the relative price of storable and non-storable non-durables. The output generated by the model is also completely insensitive to the choice of S^*.[30]

The second group of parameters are given by the P x 1 vector $\theta = [\sigma\ \ \epsilon^D\ \ \zeta^D\ \ \zeta^S\ \ ]^T$, where $P=4$ is the number of parameters to be estimated. These parameters will be estimated by minimizing a measure of the distance between the time path of expenditures generated by the model and the empirical estimates presented in Section 3.4., a procedure which I discuss further below.

The third group consists of just one parameter, ψ^D, which can be written as a function of parameters from the first two groups and the ratio of steady state durable expenditure to non-storable non-durable expenditure, $\frac{X^D}{C^N}$. I set this ratio to 0.42, which was the ratio in the JFIES in 1996.[31] Finally, initial income, $Y_0 + rA_0$, is normalized to 1, and the initial tax rate on expenditure is set to $\tau_0 = 0.03$.

To generate a time path of expenditures from the model, the following method is used. Given a full set of model parameters, I first solve for the model's steady state. I then log-

[30] The difference between S_t and S^*, which is what generates the storage cost, is independent of the magnitude of S^*. Rather, this difference is a function of τ_t, i, and ζ^S. For example, in steady state, $S - S^* = -(1+\tau)\left(\frac{r}{1+r}\right)\left(\frac{1}{\zeta^S}\right)$.

[31] Specifically, $\psi^D = \frac{(1-\psi^S)\left(\frac{1}{\delta}\right)\left(\frac{r+\delta}{1+r}\right)^{\epsilon^D}\left(\frac{X^D}{C^N}\right)}{1+(1-\psi^S)\left(\frac{1}{\delta}\right)\left(\frac{r+\delta}{1+r}\right)^{\epsilon^D}\left(\frac{X^D}{C^N}\right)}$. This expression is obtained by solving for steady state durable expenditure in terms of non-storable non-durable expenditure, and then rearranging and solving for ψ^D.

linearize the model around its steady state. The shocks to the tax rate on expenditure and income are then introduced. They propagate through the system of equations, generating a time path of percentage deviations in durable, storable, and non-storable non-durable expenditures from their steady state values.

Recall that $\hat{\gamma}$, the vector of $\hat{\gamma}_{y,m}^{j}$'s to which the time path of expenditures generated by the model will be matched, are percentage deviations in expenditure relative to December 1996, four months prior to implementation of the VAT increase. It follows that in order to make the output generated by the model consistent with the empirical estimates, I must convert the output from percentage deviations in expenditure relative to the steady state to percentage deviations in expenditure relative to expenditure four periods prior to the tax rate increase.[32]

To estimate θ, I use an econometric procedure similar to that employed by Christiano et al. (2005). I conduct a grid search over combinations of θ. $\hat{\theta}$ is the vector of parameter values that minimizes a weighted sum of the squared deviations between the M x 1 vector of time-series output generated by the model, $\gamma(\theta)$, and the M x 1 vector of empirical estimates, $\hat{\gamma}$, depicted in Figure 6. The estimates are chosen to match the durable, storable, and non-storable non-durable empirical estimates from January through July 1997. Thus, M = 7 months x 3 goods = 21. Formally,

$$\hat{\theta} = \arg\min_{\theta}[\hat{\gamma} - \gamma(\theta)]^{T} V^{-1} [\hat{\gamma} - \gamma(\theta)] \qquad (5),$$

where V is an M x M diagonal matrix that contains the sample variances of the $\hat{\gamma}_{y,m}^{j}{}'s$ along the diagonal, and T is the transpose operator.[33] The sample variances are the basis of the confidence intervals reported in Figure 6. With this choice of V, θ is chosen so that $\gamma(\theta)$ lies as much as possible inside the confidence intervals. Standard errors for the parameter estimates are computed using the delta method, which is documented in Appendix B.

[32] In practice, this conversion has little impact on the results, because the compensation to the household for the tax rate increase removes the income effect, and the intertemporal elasticity of substitution in consumption is found to be small. Thus, there is little difference between steady state expenditure and expenditure four periods prior to the tax rate increase.

[33] It is worth noting that matching the model to the empirical estimates derived from the JFIES data is equivalent to matching the model to the actual JFIES data.

4.4. Identifying the Structural Parameters

This section provides intuition for the identification and magnitudes of the structural parameter estimates presented in the next section. Identification of σ and ϵ^D are closely linked. Recall that the user cost of durables fell relative to non-durables prior to implementation of the VAT rate increase. As illustrated in Figure 9, if $\sigma < \epsilon^D$, one would observe non-storable non-durable expenditures trending downwards in the months prior to implementation, as households substitute away from non-durables to durables, followed by an upward trend thereafter. If instead $\sigma = \epsilon^D$ (i.e. preferences are separable over durables and non-durables), one would observe no change in non-storable non-durable expenditure in the months prior to implementation, and a fall in expenditure upon implementation that remains constant thereafter. Finally, if $\sigma > \epsilon^D$, one would observe an upward trend in non-storable non-durable expenditure in the months prior to implementation, a fall in expenditure upon implementation, and a slight decline thereafter. Because the empirical estimates of the expenditure response to the VAT increase exhibit an upward trend in non-storable non-durable expenditures in the months prior to implementation, and a decline upon implementation, I expect to find that $\sigma > \epsilon^D$.

Assuming that $\sigma > \epsilon^D$, what can be said about the magnitude of σ? Figure 10 illustrates the non-storable non-durable expenditure response for $\sigma = 0.2$, 0.5, and 2 when $\epsilon^D = 0.1$. Note that the larger is σ, the larger is the fall in non-storable non-durable expenditure following implementation. In addition, larger values of σ imply that the fall in durable expenditure upon implementation greatly exceeds the spike in expenditure in the month prior to implementation, as illustrated in Figure 11. The empirical estimates in Figure 6 show that non-storable non-durable expenditure was not noticeably lower following implementation than it was before, though storable expenditures were somewhat lower. Durable expenditures fell upon implementation, but not by an amount greater than the spike in expenditure in the month prior to implementation. Together, these facts suggest that σ is small, and likely close to zero.

Taking as given that σ and ϵ^D are small, Figure 12 demonstrates that ζ^D is identified primarily by the spike and trough in durable expenditure in the months prior to and including the VAT rate increase. Simply, the larger is ζ^D, the smaller will be the spike and trough in durable expenditure in March and April 1997, respectively. Similarly, the larger is ζ^S, the lower will be the amount of stockpiling in the month prior to implementation, and consequently, the lower will be the jump in storable expenditures.

4.3. Structural Parameter Estimates

Table 2 presents point estimates and 95 percent confidence intervals for the model parameters that comprise θ. The estimate of σ is 0.07, with a 95 percent confidence interval given by [-0.03, 0.17]. The point estimate is similar to, though slightly lower than, the estimate in Cashin and Unayama (2012). However, the confidence interval is much tighter, as expected. The null hypothesis that the IES is zero cannot be rejected at conventional levels of significance. The results imply that consumption growth is insensitive to a change in the future price level. [34]

The point estimate of ϵ^D is -0.01, with a 95 percent confidence interval given by [-0.04, 0.02]. The null hypothesis that durables and non-durables are perfect complements (i.e. $\epsilon^D = 0$) cannot be rejected, which suggests the use of Leontief preferences when jointly modeling durable and non-durable consumption, and is consistent with the findings of Pakos (2004). That said, the null hypothesis that preferences over durables and non-durables are separable (i.e. $\sigma = \epsilon^D$) also cannot be rejected at conventional significance levels.[35]

The point estimate for ζ^D, the durable adjustment cost parameter, is 0.10, and is significant at the five percent level. To get a better sense of what this value implies, note that a household will increase its durable stock prior to the VAT increase so long as the marginal benefit of adjustment exceeds the marginal cost. As demonstrated in Figure 12, the marginal benefit of adjustment is decreasing in D_t, which is due to the fact that the marginal utility of the contemporaneous service flow derived from durables is decreasing in D_t, and also because accelerated purchases of durables today implies additional costly adjustments in the future.[36] The marginal adjustment cost is given by $\zeta^D(D_t - D_{t-1})$, and increases linearly in D_t. Substituting the relevant values into the marginal adjustment cost function for March 1997, the month in

[34] There exists a literature (e.g. Watanabe et al., 2001; Poterba, 1988) which finds that the income effect associated with a tax change does not become evident until implementation. If this is true, then the estimate of σ will be biased upwards, as the fall in expenditure upon implementation will be attributed solely to the intertemporal substitution effect rather than the income effect.

[35] The Wald test statistic, W, for the null hypothesis $H_0\colon \sigma = \epsilon^D$ is $W = 1.41$, which falls below the critical value of 3.84 when the degrees of freedom equals one, and the significance level is five percent. It also falls short of the critical value at a ten percent significance level.

[36] Specifically, the marginal benefit of adjustment is given by $(1+\tau_t)\frac{U_{D_t}}{U^N_{C_t}} + \left(\frac{1-\delta}{1+r}\right)(1+\tau_{t+1}) - (1+\tau_t) + \frac{1}{1+r}\zeta^D(D_{t+1} - D_t)$. It is decreasing in D_t for two reasons. First, the marginal utility of the service flow provided by durables is decreasing in the durable stock, which can be seen in the first term. Second, increasing D_t implies additional costly adjustment in the future, as shown in the final term.

which durable expenditure was most sensitive to the VAT rate increase, yields a value of 0.0036. That is, the marginal adjustment cost is 0.36 percent of monthly income.

Although not reported here, the interpretation of ζ^D as capturing a time cost appears accurate, as estimates of $\zeta^D$are in general lower (higher) for groups with arguably more (less) free time (e.g. retired v. working, non-working wife with no children v. family with small children). Suppose then that one treats the adjustment cost as a time cost, and that monthly household income is actually U.S. $4,000, which is approximately the average monthly wage income in the U.S. as of 2011.[37] Then the marginal adjustment cost is $14.40, which is approximately two-thirds of the mean hourly wage in the U.S. If one thinks of the hourly wage as the opportunity cost of leisure, then the marginal adjustment cost is equivalent to roughly 40 minutes of leisure for the average American worker.

The point estimate for the storage cost parameter, ζ^S, is 1.18, and is significant at the one percent level. A household will continue stockpiling so long as the marginal benefit of doing so exceeds the marginal cost. As shown in Figure 13, the marginal benefit of stockpiling in March 1997 is constant at $\left(\frac{1+\tau_{t+1}}{1+i}\right) - (1+\tau_t) \approx 0.02$, since i is small.[38] That is, for every purchase of a storable made in March 1997 rather than April 1997, a household saves approximately two percent on the storable's purchase price. The marginal cost of stockpiling is given by $\zeta^S(S_t - S^*)$. The lower is ζ^S, the greater the amount of stockpiling that will be observed in March 1997. A household will continue to accumulate storable inventory until the marginal adjustment cost is approximately two percent of monthly income.

4.4. Sensitivity tests

Recall that several of the model parameters were fixed prior to estimation. While their values were fixed based on the available data, it seems reasonable to test the sensitivity of the structural parameter estimates to different values of the fixed parameters. Table 3 presents the results of the sensitivity analysis. In general, the structural parameter estimates are robust to alternative fixed parameter values. Given a change in a Group 1 parameter, each new parameter estimate lies within its confidence interval shown in Table 2. Furthermore, there is little variation in the value of σ regardless of the value at which the fixed parameters are set.

[37] Source: Bureau of Labor Statistics' Occupational Employment Statistics. Retrieved from http://www.bls.gov/oes/current/oes_nat.htm

[38] This expression holds under the model assumptions that $\beta(1+i) = 1$ and relative pre-tax prices are constant.

4.5. Evaluating the Model

Figure 15 plots the time path of expenditures generated by the model against the empirical estimates previously shown in Figure 6. Overall, the model matches the empirical estimates reasonably well, though the validity of the model is rejected by the test of overidentifying restrictions.[39] Non-storable non-durable expenditures lie entirely within the confidence interval. The model underpredicts non-storable non-durable expenditure in the month prior to implementation, and overpredicts expenditure in the month following implementation, which could be due to limited arbitrage opportunities for non-storable non-durables (e.g. buying bananas on March 31, 1997, and consuming them in early April) that the model does not allow.

The model closely matches storable expenditure in the months prior to and including implementation, but the match is rather poor in the months following implementation. There are likely two reasons for this. First, the model considers only one storable composite good that is costly to store. It follows that the representative household will fully consume the stockpiled storable good in the month of implementation before making any additional purchases. Storable expenditure will then return to a new steady state in the months following implementation. In reality, some households stockpile, while others don't. Among those that do, some stockpile a lot, and others stockpile a little. As a result, there is a gradual return to trend in storable expenditure following implementation that the model is not flexible enough to match. Second, storable expenditure was trending downward throughout the latter half of the 1990s, but this was not so in the first half of the decade. The constant in Equation (4) assumes a linear trend in consumption growth during the sample period. As a result, the constant may fail to appropriately capture the downward trend in storable non-durable expenditure, instead influencing the estimates of $\gamma_{y,m}^{S}$ that the model attempts to match.

As with non-storable non-durable expenditure, the time path of durable expenditure generated by the model stays within the 95 percent confidence interval in all months used in the estimation of the structural parameters. Furthermore, the pattern of durable expenditure predicted by the model is generally consistent the empirical estimates. However, the model underpredicts the sensitivity of durable expenditure to the VAT rate increase in February and March 1997. This is a result of the choice of a quadratic adjustment cost specification, which requires symmetric expenditure responses on either side of the tax rate increase. The model could roughly match the

[39] The J-statistic equals 51.95, which exceeds the critical value of 27.59 (df = 17, $\alpha = 0.05$).

expenditure response in February and March 1997, but if it did, there would have to be a much larger reduction in expenditure in April 1997 than was observed. The estimate of ζ^D is thus a compromise between matching the expenditure response in March without overshooting the decline in expenditure in April 1997. The model also overpredicts expenditure in July 1997, though this can be explained by the fact that many households purchased air conditioning units prior to the VAT rate increase. As a result, there was a significant decrease in purchases in July 1997, which coincides with the end of the rainy, relatively cool portion of the summer in Japan.

4.6. External validity test

One might question the external validity of the parameter estimates, given that the estimates are derived from one event. To address this concern, I simulate the expenditure response to New Zealand's Goods and Services Tax (GST) rate increase from 10 to 12.5 percent, which was announced in March 1989 and implemented in July 1989. Because the GST rate increase was uncompensated, the simulated time path of expenditures generated by the model allows for an income effect. However, the vector of parameter estimates characterizing intertemporal substitution, θ, remains unchanged.

Figure 16 compares the time path of expenditures generated by the model to the empirical estimates of the retail sales response to the GST rate increase, which is documented in Cashin (2011). The model performs very well for non-durable expenditures (storables and non-storables are summed to match the available retail sales data from New Zealand), not only staying within the confidence intervals, but matching quite closely with the point estimates. The model overpredicts the arbitrage response for durable goods, which I suspect is at least in part due to the fact that the elderly comprise a larger share of the population in Japan than in New Zealand.[40] Although not reported here, durable expenditure was significantly more sensitive to the VAT rate increase for retired households in Japan than working households, likely as a result of their lower opportunity cost of shopping.

Overall, however, the simulation lends additional support to the main finding in this paper, which is that expenditure is sensitive to a change in the future price level, albeit over a short period immediately preceding and following the price change, while consumption growth is not.

[40] As of 2011, 30 percent of the population in Japan was over 60, while just over 18 percent of the population was over the age of 60 in New Zealand (The Economist Pocket World in Figures, 2011 edition).

4.7. Application: The Efficiency Cost of Japan's Proposed VAT Rate Increase

In August 2012, Japan's Prime Minister, Yoshihiko Noda, successfully pushed a bill through Japan's Diet (legislature) to gradually increase the VAT rate from five to ten percent. Specifically, the bill would increase the VAT rate from five to eight percent in April 2014, and from eight to ten percent in October 2015. The measure is intended to rein in public debt, which now exceeds 200 percent of GDP. However, there is some uncertainty as to whether the government will proceed with the tax rate hike. This is because the law contains a provision stating that the government should consider the overall economic situation before proceeding with the VAT increase.[41]

To evaluate the deadweight loss of the proposed tax rate increase and whether the uncertainty associated with its implementation will significantly reduce the marginal excess burden (by reducing intertemporal substitution), I simulate the (Hicks) compensated expenditure response to the VAT rate increase under two scenarios, using the parameter estimates in Table 2. The first assumes that passage of the August 2012 bill constitutes announcement of the tax rate increase. The second assumes that haggling over the aforementioned provision prevents the tax rate increase from becoming a certainty until December 2013, when the fiscal year 2014 budget is submitted. Using a compensating variation measure, I then calculate the deadweight loss of the tax rate increase under each scenario. See Appendix C for details on the estimation.

Figure 16 presents the time paths of expenditure generated under the two scenarios. What is immediately obvious is that the announcement date has very little impact on the amount of intertemporal substitution, stockpiling, and arbitrage that households engage in, though it does change the timing of the initial durable response. The negligible difference in the magnitude of the expenditure response follows from the low value of σ, and suggests that the deadweight loss of the VAT rate increase is largely independent of the delay between announcement and implementation of the rate increase. The low value of σ also suggests that the deadweight loss will be quite small.

Indeed, I find that the deadweight loss is only 0.0008 and 0.0007 percent of the present value of steady state lifetime expenditure for the August 2012 and December 2013

[41] "…the provision says that before implementing the higher levy, the government will 'give holistic consideration to the economy and will take appropriate measures, which could include halting [the tax's] enactment' if economic conditions are languishing." Warnock, Eleanor. "Loophole Threatens Implementation of Controversial Japan Sales Tax." Wall Street Journal Online August 14, 2012. http://online.wsj.com/article/BT-CO-20120814-702666.html

announcement scenarios, respectively.[42] In contrast, suppose that $\sigma = 0.9$, a value consistent with previous intertemporal substitution studies using survey data. In that case, the deadweight loss associated with the phased-in rate increase is 0.007 and 0.004 percent for the August 2012 and December 2013 scenarios, respectively. It is clear that as σ increases, the delay between announcement and implementation has a larger impact on the deadweight loss associated with the tax rate increase. However, the small estimate of σ in this study implies that the delay between announcement and implementation is of little consequence, as the distortions resulting from the tax change are largely confined to the months immediately preceding the change, and some delay between announcement and implementation is inevitable.

5. Discussion of the IES Estimate

The estimate of σ in this study resides at the low end of previous estimates using survey data. Not surprisingly, it is most similar to the estimate of Cashin and Unayama (2012). However, the estimate is significantly lower than the estimates of Attanasio and Weber (1993, 1995), Vissing-Jorgensen (2002), and Gruber (2006), whose baseline point estimates are 0.8, 0.8-1, and 2, respectively.

There exist several potential explanations for the difference between the estimate of σ in this study and Cashin and Unayama (2012), and those in other studies using survey data. One possibility is that the estimate in this study does not account for borrowing constraints, and thus yields an estimate of σ that is biased downwards. However, when Cashin and Unayama (2012) split the sample between groups that are more and less likely to be borrowing constrained, there is little difference in the estimate of σ. For example, the difference in σ for retired households, who are unlikely to experience future income growth and are in the asset decumulation stage, and working households, is not significant.

While the previous studies utilizing survey data do not suffer from aggregation bias, they may face other methodological issues that would bias the estimate of σ upwards. For one, the analyses include storable and some durable expenditures (e.g. apparel). If households stockpile in response to an increase in the future price level, these studies would incorrectly attribute the response to increased consumption during the low-price period, and thus a larger estimate of σ.

[42] This assumes that $i = 0.018$ annually, as in the estimation presented in Section 4. The magnitude of the deadweight loss is decreasing in the interest rate.

Working against this explanation, however, is the fact that the previous studies use quarterly or semiannual data, where stockpiling behavior is more difficult to observe. In addition, the changes in the real interest rate that are the subject of those studies are not necessarily anticipated like the change in this study, further mitigating the amount of stockpiling that is likely to occur.

Another possible issue with all but one of the previous studies is the potential for a positive correlation between the variables used to instrument for the endogenous real interest rate and the expectational error term (i.e. innovations to a household's permanent income). Lagged stock-market and Treasury bill return data are often used as instruments for the contemporaneous real interest rate, under the assumption that they can predict the real interest rate, and are uncorrelated with the expectational error term. However, stock and Treasury bill returns are generally considered to be leading economic indicators, while wage growth is a lagging indicator. Unless a household equates increases in market returns with future wage growth, and responds accordingly, the potential exists for an upward bias in estimates of σ due to a positive correlation between lagged market returns and contemporaneous wage growth, the latter of which appears as an innovation to permanent income. Gruber (2006) addresses this issue by relying on cross-sectional variation in capital income tax rates to identify the IES. This strategy yields an estimate of the IES similar to the one he obtains using the standard instrumental variables approach, which suggests that endogeneity between the instruments normally used in the literature and the expectational error term are not responsible for the large difference between the estimate of σ presented in this study, and those found in previous studies.

A final possible issue with the previous studies is their assumption of separable preferences over durables and non-durables, which allows them to focus solely on non-durable consumption growth to estimate σ. Suppose durables and non-durables are instead complements, as found in this study. Then a reduction in the real interest rate will lead to a fall in the user cost of durables. This will in turn lead to even greater non-durable consumption growth, because the two goods are complements. It follows that the estimate of σ will be biased upwards, because some of the non-durable consumption growth, which should be attributed to complementarities between durables and non-durables, is instead attributed to the reduction in the real interest rate. Indeed, Cashin and Unayama (2012) obtain an estimate of σ of 0.91 under the assumption of separable preferences, which is significantly larger than their baseline estimate of 0.21.

6. Conclusion

This study uses a pre-announced increase in Japan's VAT rate from three to five percent to measure intertemporal substitution. The main finding is that households accelerated purchases of durable goods and stockpiled storable goods just prior to the tax rate increase, but did not change their consumption patterns.

To the extent that the results in this study can be applied to other contexts, they suggest that policies that alter the future price level will have a limited impact on household expenditure, with changes in the timing of expenditure confined to the months just prior to and following the price change. Consequently, one might question the efficacy of countercyclical policies that alter the real interest rate with the intention of yielding a more persistent change in the timing of household expenditure.

Even if the intertemporal substitution effects are small, however, it should not be forgotten that these policies also have income effects, which were not addressed in this study, and could be significant for, say, a borrowing constrained household in a country where the consumption tax rate has been temporarily reduced to combat a recession.[43] Without empirical evidence on the income effects, it seems premature to rule out the effectiveness of countercyclical policies that alter the real interest rate, though it does lead one to wonder whether countercyclical policies should directly target income rather than intertemporal substitution effects.

[43] Crossley and Low (2009) address this issue in the context of the 2009 temporary VAT cut in the U.K.

References

Attanasio, O.P. and Weber, G. (1993). 'Consumption Growth, the Interest Rate and Aggregation', *Review of Economic Studies,* Vol. 60, No. 3 (July), pp. 631-649.

Attanasio, O.P. and Weber, G. (1995). 'Is Consumption Growth Consistent with Intertemporal Optimization?', *Journal of Political Economy,* Vol. 103, No. 6 (December), pp. 1121-1157.

Attanasio, O.P. and Weber, G. (2010). 'Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy', *Journal of Economic Literature,* Vol. 48, pp. 693-751.

Barrell, R. and Weale, M. (2009). 'The Economics of a Reduction in VAT', *Fiscal Studies,* Vol. 30, no. 1, pp. 17-30.

Bernanke, B. (1985). 'Adjustment Costs, Durables, and Aggregate Consumption', *Journal of Monetary Economics,* Vol. 15, pp. 41-68.

Campbell, John Y., and Luis M. Viceira (1999). 'Consumption and Portfolio Decisions When Expected Returns Are Time Varying', *Quarterly Journal of Economics*, vol. 114, no. 2, pp. 433–495.

Carroll, R.J., Cline, R.J., Diamond, J.W., Neubig, T.S., and Zodrow, G.R. (2010). 'Implementing a VAT in the United States: Price Effects and Revenue Offsets', forthcoming in the 2010 *Proceedings of the Annual Conference on Taxation*.

Cashin, D. (2011). 'The Intertemporal Substitution and Income Effects of a Consumption Tax Rate Increase: Evidence from New Zealand', Working Paper, University of Michigan.

Cashin, D. and Unayama, T. (2012). 'Measuring Intertemporal Substitution: Evidence from a Consumption Tax Rate Increase in Japan', RIETI Working Paper.

Christiano, L.J., Eichenbaum, M., and Evans, C.L. (2005). 'Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy', *Journal of Political Economy*, vol. 113, no. 1, pp. 1-45.

Correia, I., Farhi, E., Nicolini, J.P., and Teles, P. (2011). 'Unconventional Fiscal Policy at the Zero Bound', NBER Working Paper no. 16758.

Crossley, T. Low H., and Wakefield, M. (2009). 'The Economics of a Temporary VAT Cut', *Fiscal Studies*, vol. 30, no. 1, pp. 17-30.

Fraumeni, B.M. (1997). 'The Measurement of Depreciation in the U.S. National Income and Product Accounts', *Survey of Current Business,* July, pp. 7-23.

Gruber, J. (2006). ‘A Tax-Based Estimate of the Elasticity of Intertemporal Substitution’, *NBER Working Paper No. 11945*.

Hall, R.E. (1988). ‘Intertemporal Substitution in Consumption’, *The Journal of Political Economy*, vol. 96, no. 2 (April), pp. 339-357.

Halvorsen, R. and Palmquist, R. (1980). ‘The Interpretation of Dummy Variables in Semilogarithmic Equations’, *The American Economic Review*, Vol. 70, No. 3, pp. 474-475.

Hendel, I. and Nevo, A. (2004). ‘Intertemporal Substitution and Storable Products’, *Journal of the European Economic Association*, vol. 2, no. 2/3, pp. 536-547.

Hendel, I. and Nevo, A. (2006). ‘Sales and Consumer Inventory’, *RAND Journal of Economics*, Vol. 37, no. 3 (Autumn), pp. 543-561.

Ishi, H. (2001). *The Japanese Tax System: Third Edition*, Oxford University Press.

Ogaki, M. and Reinhart, C.M. (1998). ‘Measuring Intertemporal Substitution: The Role of Durable Goods’, *Journal of Political Economy*, vol. 106, no. 5, pp. 1078-1098.

Pakos, M. (2004). ‘Asset Pricing with Durable Goods and Nonhomothetic Preferences’, *Munich Personal RePEc Archive Paper No. 26167*.

Pischke, J. (2010). *Weak Instruments* [PDF document]. Retrieved from http://econ.lse.ac.uk/staff/spischke/ec533/Weak%20IV.pdf

Poterba, J. (1988). ‘Are consumers forward looking? Evidence from fiscal experiments’, *AEA Papers and Proceedings*, Vol. 78, pp. 413-418.

Stephens Jr., Melvin and Unayama, T. (2011). ‘The Consumption Response to Seasonal Income: Evidence from Japanese Public Pension Benefits’, Forthcoming in *American Economic Journal: Applied Economics*.

Stephens Jr., Melvin and Unayama, T. (2012). ‘The impact of retirement on household consumption in Japan’, *Journal of the Japanese and International Economies* vol.26 pp.62-83.

Takahashi, F. (1999). ‘Manipulations Behind the Consumption Tax Increase: The Ministry of Finance Prolongs Japan’s Recession’, *Journal of Japanese Studies*, Vol. 25, No. 1 (Winter), pp. 91-106.

Vissing-Jorgensen, A. (2002). ‘Limited Asset Market Participation and the Elasticity of Intertemporal Substitution', *Journal of Political Economy*, Vol. 110, No. 4, pp. 825-853.

Watanabe, K., Watanabe, T., and Watanabe, T. (2001). ‘Tax Policy and Consumer Spending: Evidence from Japanese Fiscal Experiments’, *Journal of International Economics*, vol. 53, pp.

261-281.

Yogo, M. (2004). 'Estimating the Elasticity of Intertemporal Substitution When Instruments are Weak', *Review of Economics and Statistics*, Vol. 86, No. 3 (August), pp. 797-810.

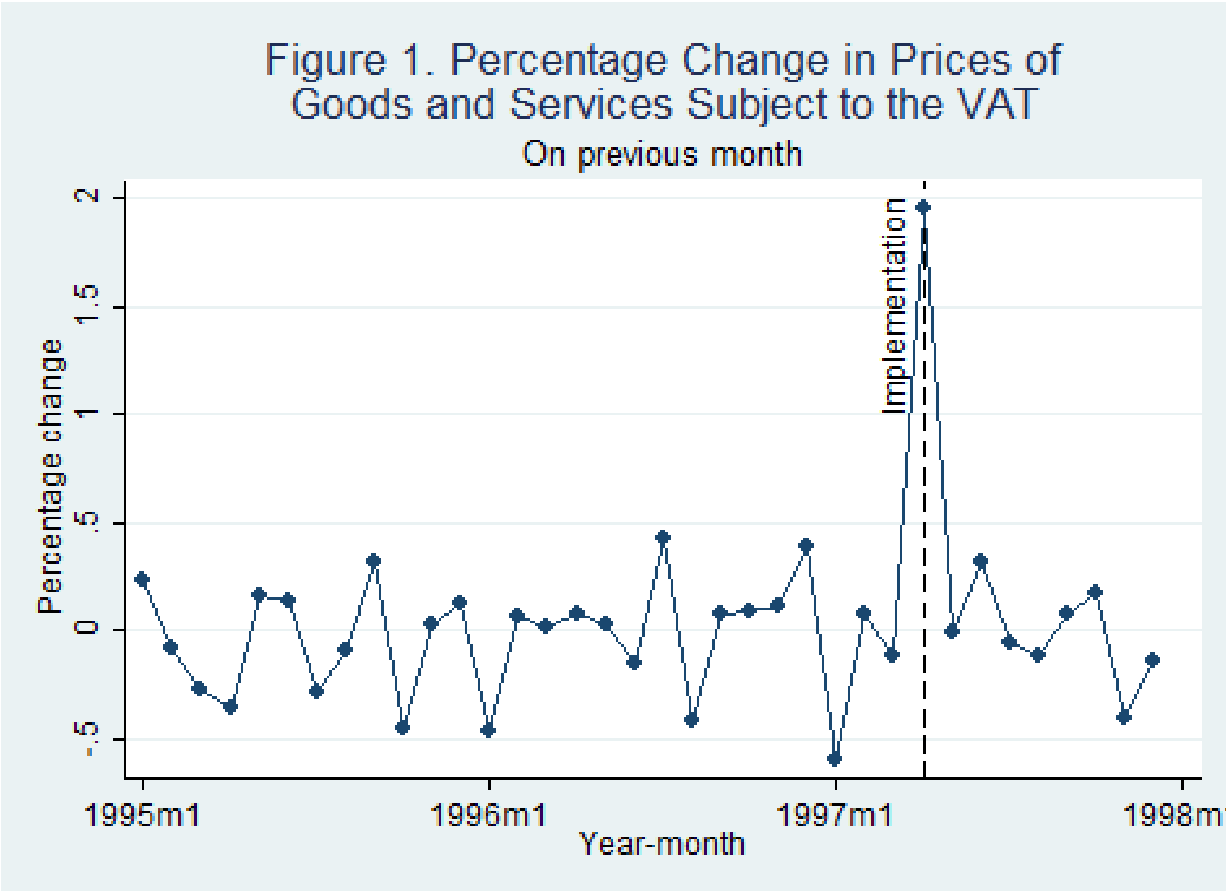

Figure 1 presents seasonally-adjusted month-to-month price changes on goods and services that were subject to the VAT. The vertical dashed line represents April 1997, the month of the VAT rate increase from three to five percent. To generate these estimates, a price index for all goods and services subject to the VAT is regressed on month dummies. The residuals from this regression are added to the constant from the regression to yield a seasonally-adjusted price index. I then compute month-to-month percentage changes in the price index.

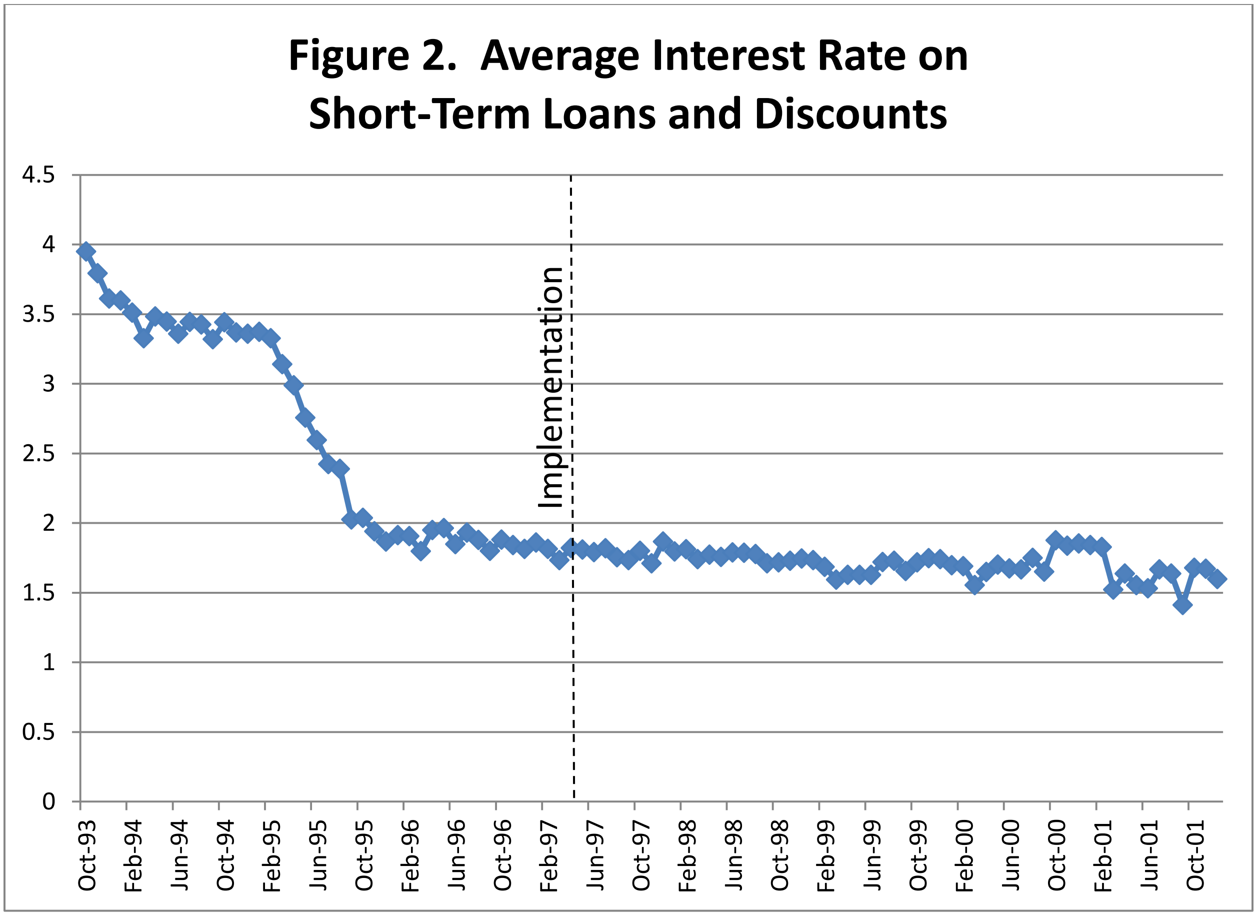

Figure 2 presents the average contracted interest rate on short-term loans and discounts. These are the average interest rates applied to a contract of less than one year between commercial banks and lenders. The data comes from the Bank of Japan.

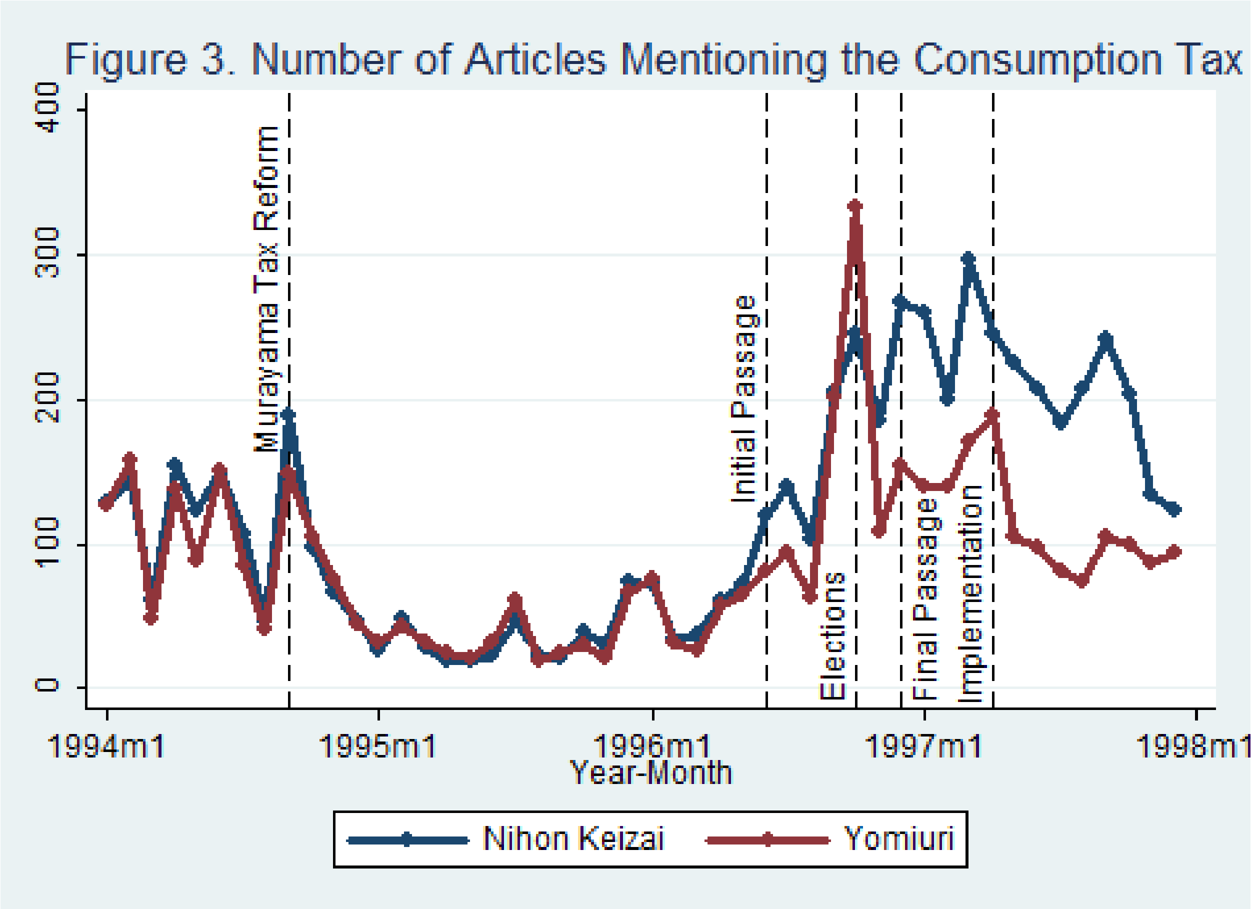

Source: Author's calculations. Circulation numbers come from Japan's Audit Bureau of Circulations.

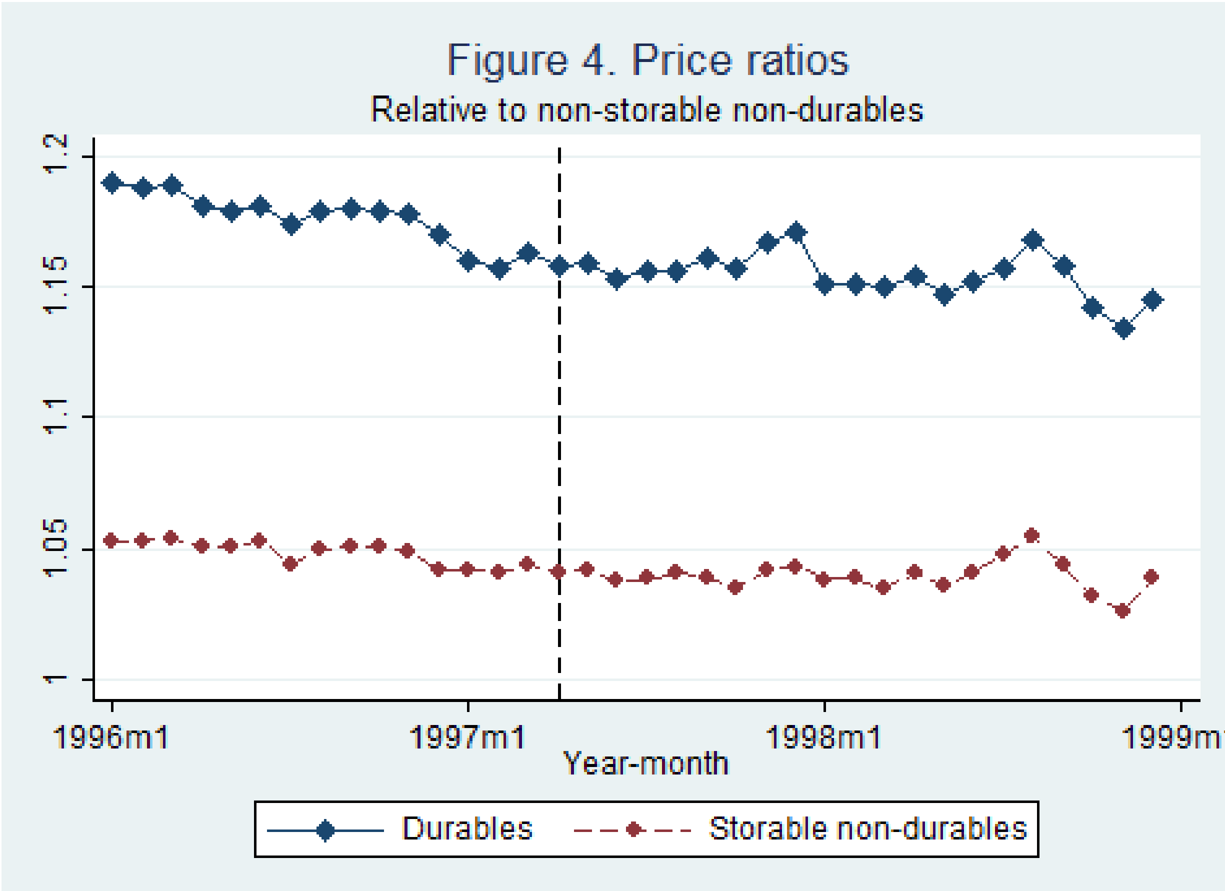

Figure 4 presents the ratio of seasonally-adjusted durable and storable non-durable prices to non-storable non-durable prices. To remove seasonality, the consumer price indices for durables, storable non-durable, and non-storable non-durable goods and services are regressed on month dummies. The residuals are added to the constant in the regression to obtain seasonally-adjusted price indices. To calculate the ratios, the seasonally-adjusted durable and storable non-durable prices are divided by the seasonally-adjusted non-storable non-durable price in each month. The dashed vertical line in the figure is April 1997, the month of implementation. The base year for the index is 2005.

Figure 5. Identifying the Intertemporal Substitution Response to the VAT Rate Increase

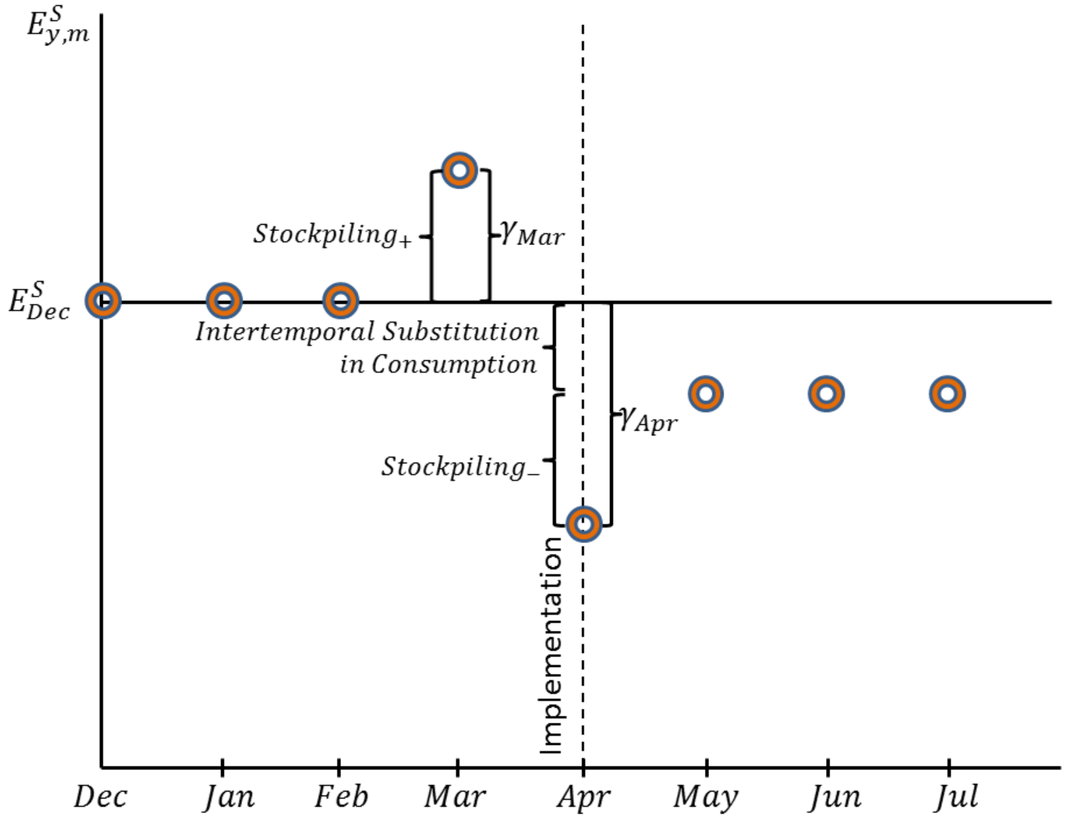

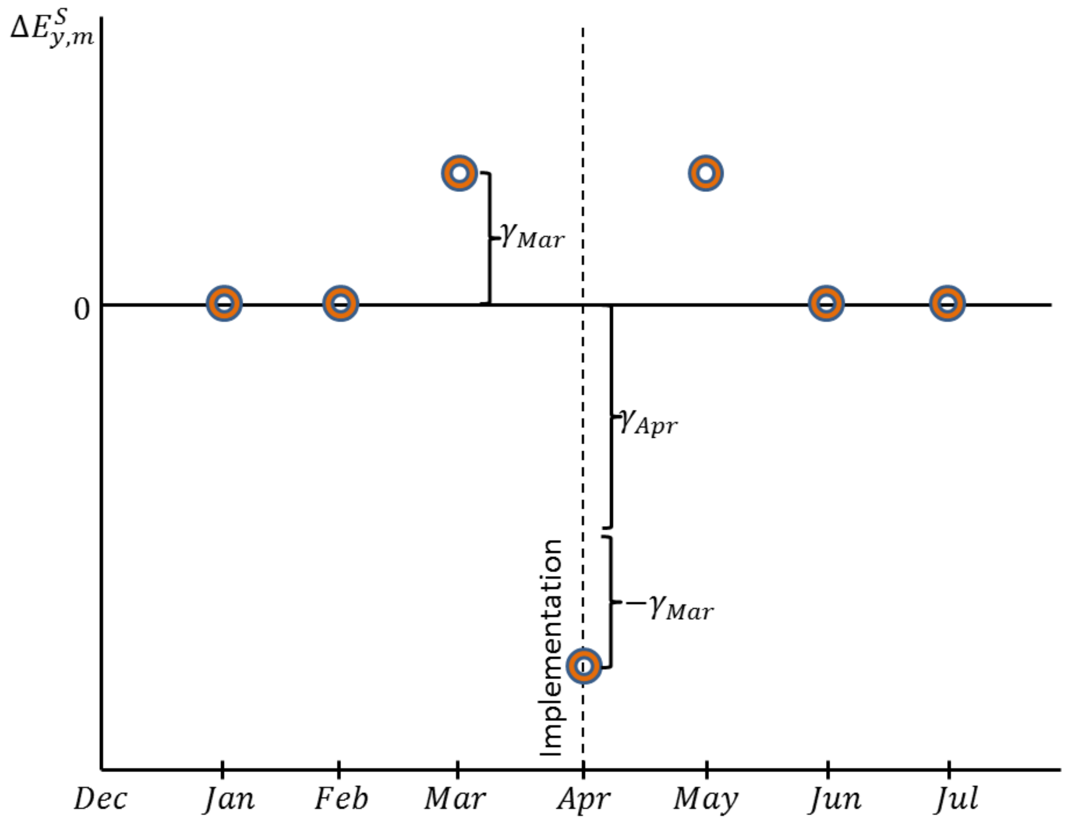

Figure 6. The Intertemporal Substitution Responseto the VAT Rate Increase

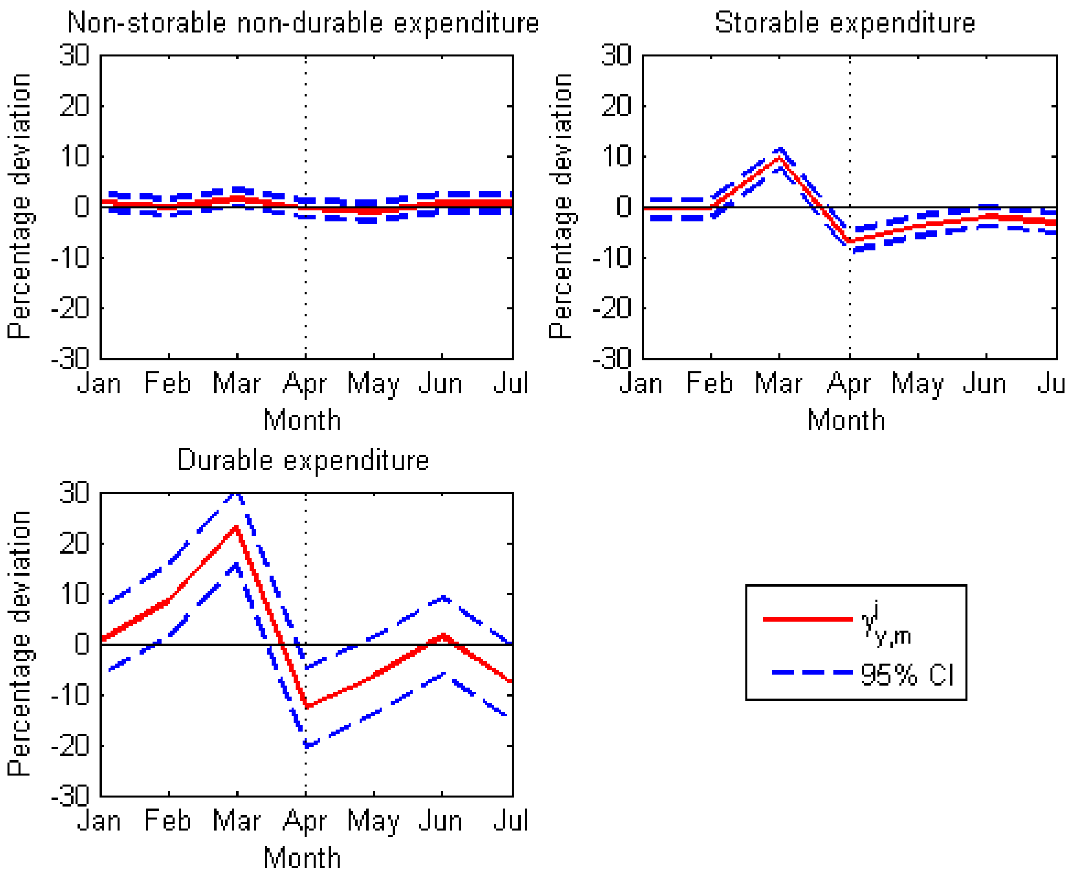

The figure above shows the percentage deviation in durable, storable, and non-storable non-durable expenditures relative to expenditure in these categories in December 1996, controlling for household fixed effects, a linear trend in consumption growth, seasonality, and time-varying household characteristics. I interpret the results as yielding the intertemporal substitution effects associated with the VAT rate increase. The red line gives the point estimates in each month. The dashed blue lines give the 95 percent confidence intervals. The dashed vertical line represents April 1997, the month of implementation. The results are based on the specification given in Equation (4). Standard errors are robust to serial correlation within households over time. If monthly durable expenditure for a household is reported as zero, it is set to ¥100 to avoid taking the logarithm of zero.

Figure 7. The Storable Intertemporal Substitution Response to the VAT Rate Increase

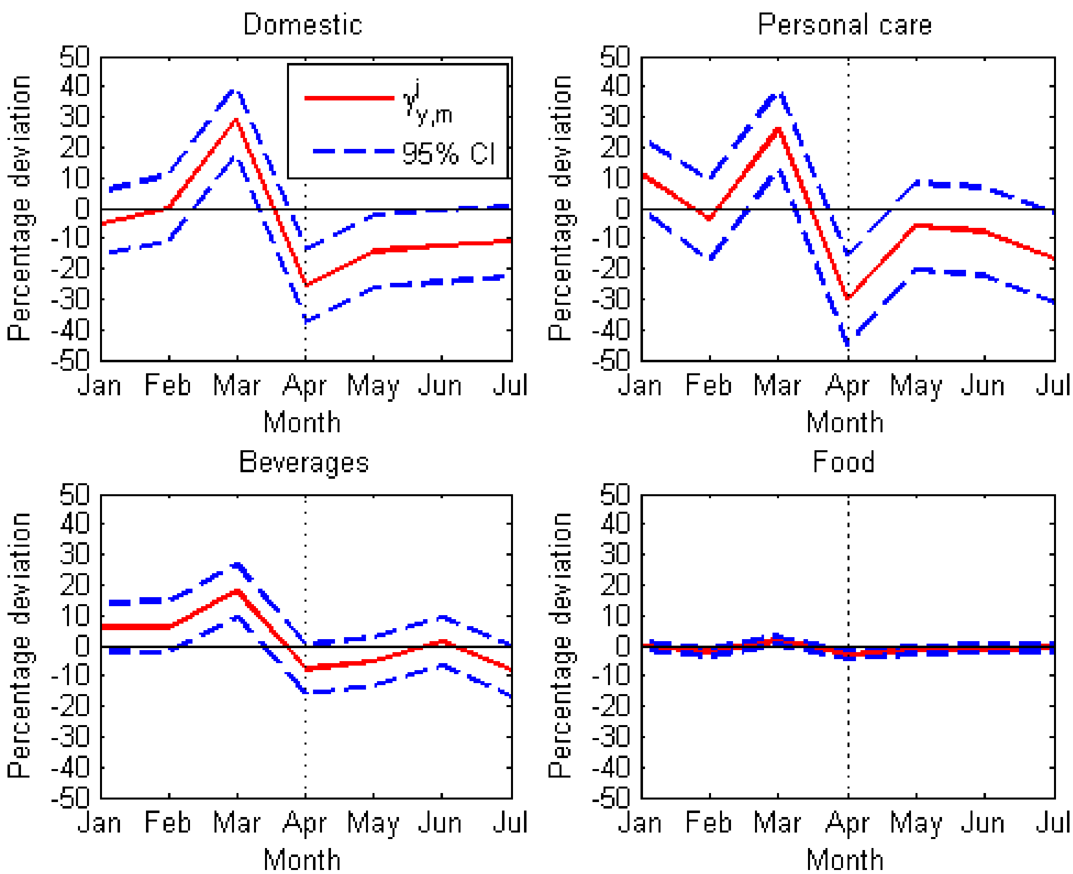

See Figure 6 for an explanation of how these plots were generated.

Figure 8. The Durable Intertemporal Substitution Response to the VAT Rate Increase

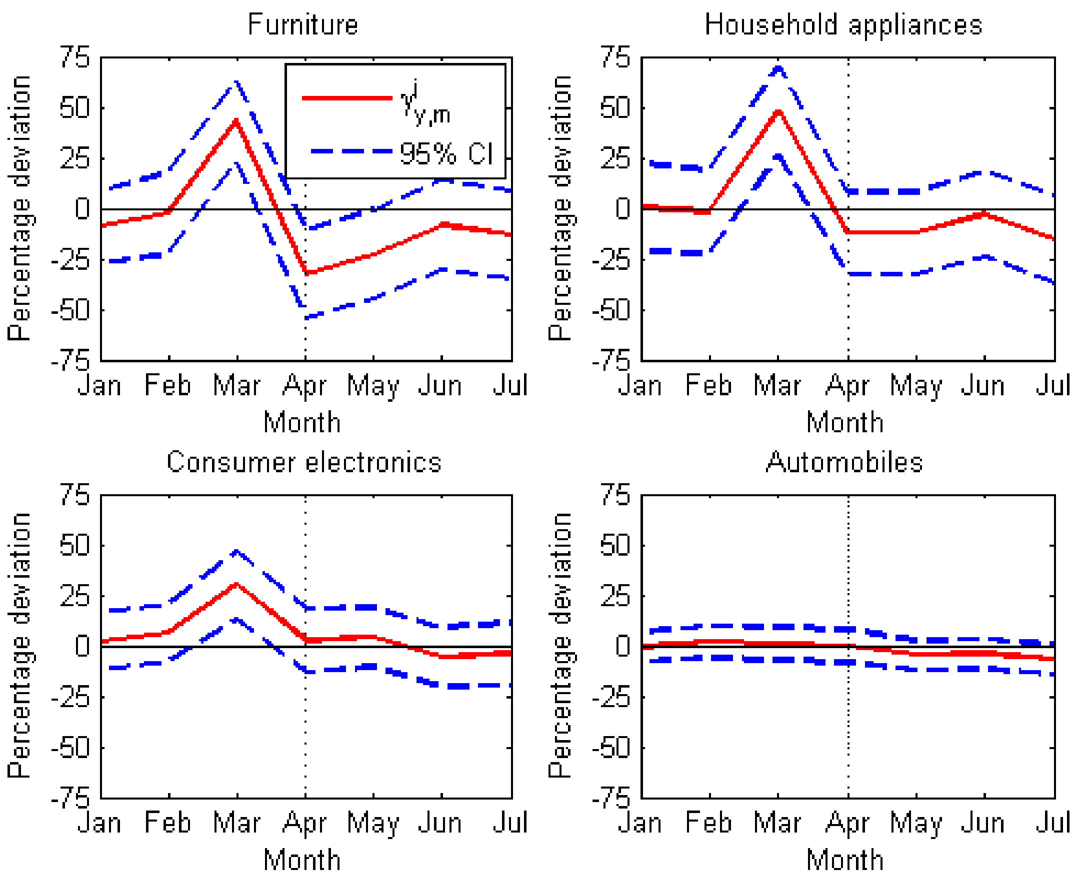

See Figure 6 for an explanation of how these plots were generated.

Figure 9. Simulated Non-storable Non-durable Expenditure Patterns for Different Values of σ

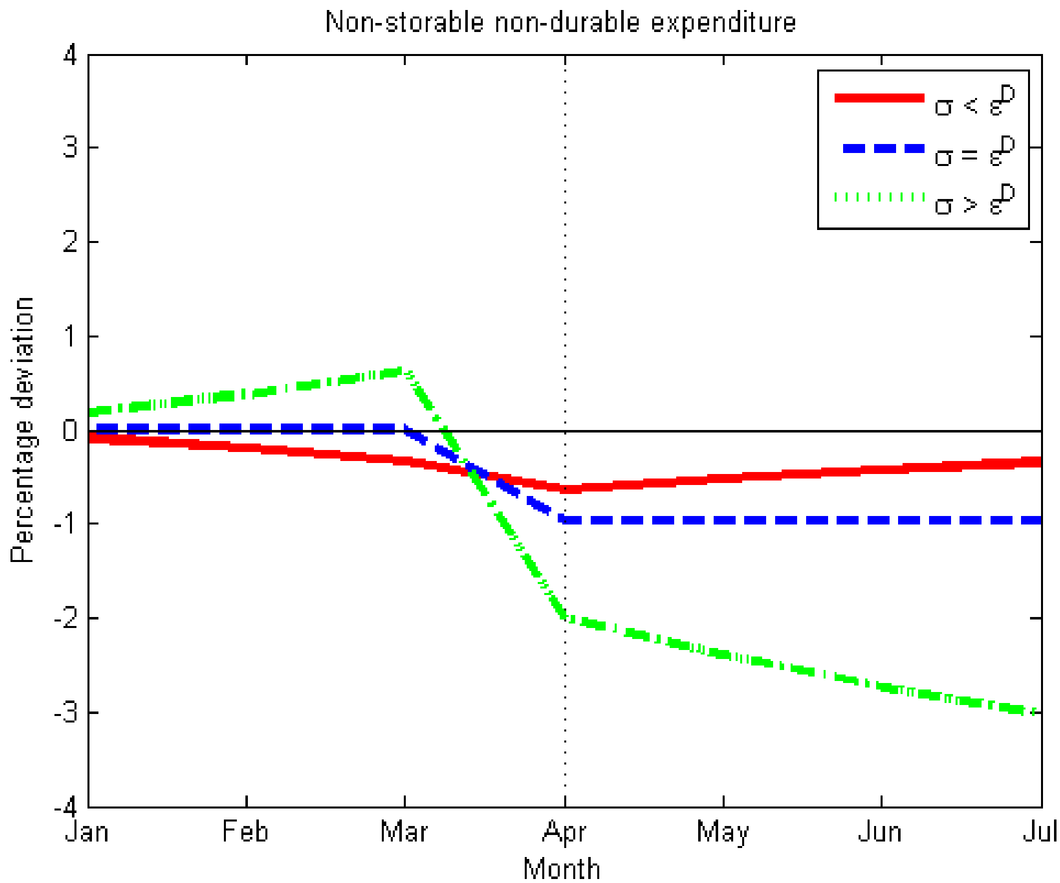

Figure 9 presents simulations of non-storable non-durable expenditure in the months prior to and following implementation of the VAT increase. For all three simulations, $\epsilon^D = 0.5$. The solid red, dashed blue, and dashed-dot green line show expenditure when $\sigma = 0.2, 0.5$, and 2, respectively.

Figure 10. Simulated Non-storable Non-durable Expenditure Patterns for Different Values of σ When $\sigma > \epsilon^D$

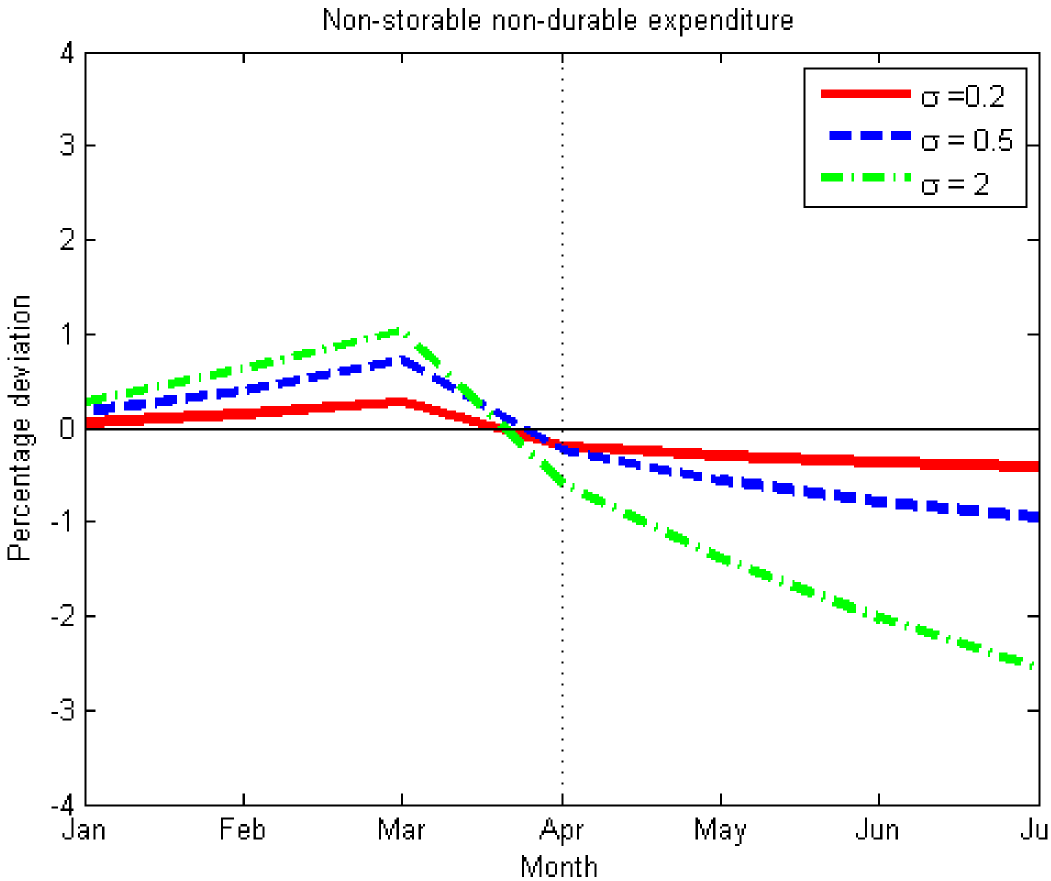

Figure 10 presents simulations of non-storable non-durable expenditure in the months prior to and following implementation of the VAT increase. For all three simulations, $\epsilon^D = 0.1$. The solid red, dashed blue, and dashed-dot green line show expenditure when $\sigma = 0.2, 0.5,$ and 2, respectively.

Figure 11. Simulated Durable Expenditure Patterns for Different Values of σ When $\sigma > \epsilon^D$

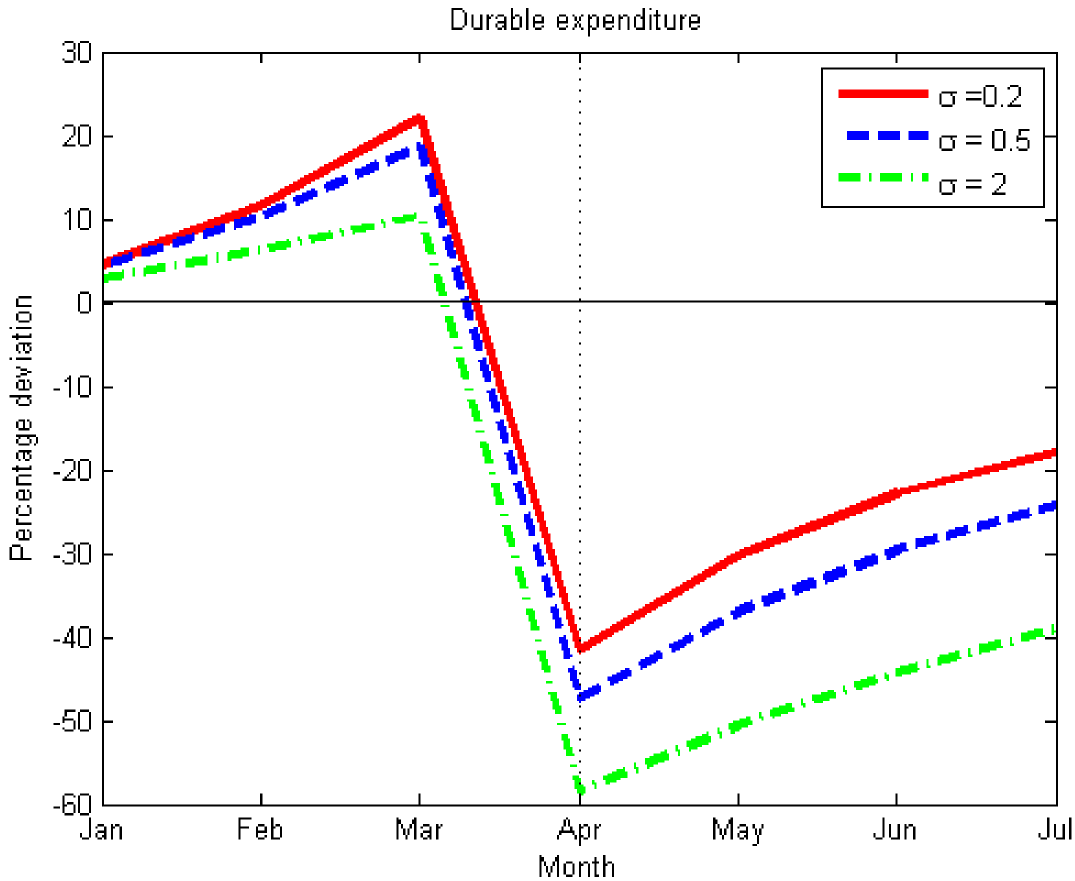

Figure 11 presents simulations of durable expenditure in the months prior to and following implementation of the VAT increase. For all three simulations, $\epsilon^D = 0.1$. The solid red, dashed blue, and dashed-dot green line show expenditure when $\sigma = 0.2, 0.5,$ and 2, respectively.

Figure 12. Simulated Durable Expenditure Patterns for Different Values of ζ^D

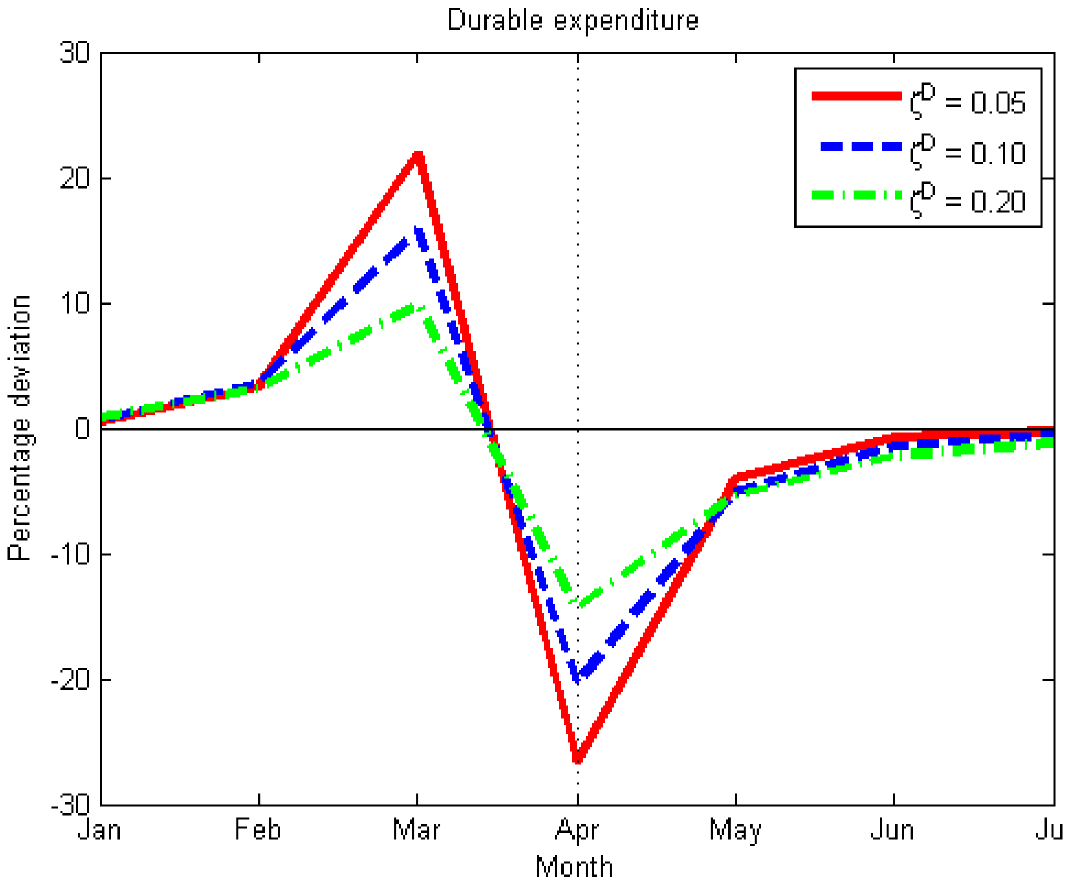

Figure 12 presents simulations of durable expenditure in the months prior to and following implementation of the VAT increase. For all three simulations, $\sigma = 0.07$, $\epsilon^D = -0.01$, and $\zeta^S = 1.18$. The solid red, dashed blue, and dashed-dot green line show expenditure when $\zeta^D = 0.05, 0.10$, and 0.20, respectively.

Figure 13. The costs and benefits of adjusting the durable stock at the margin

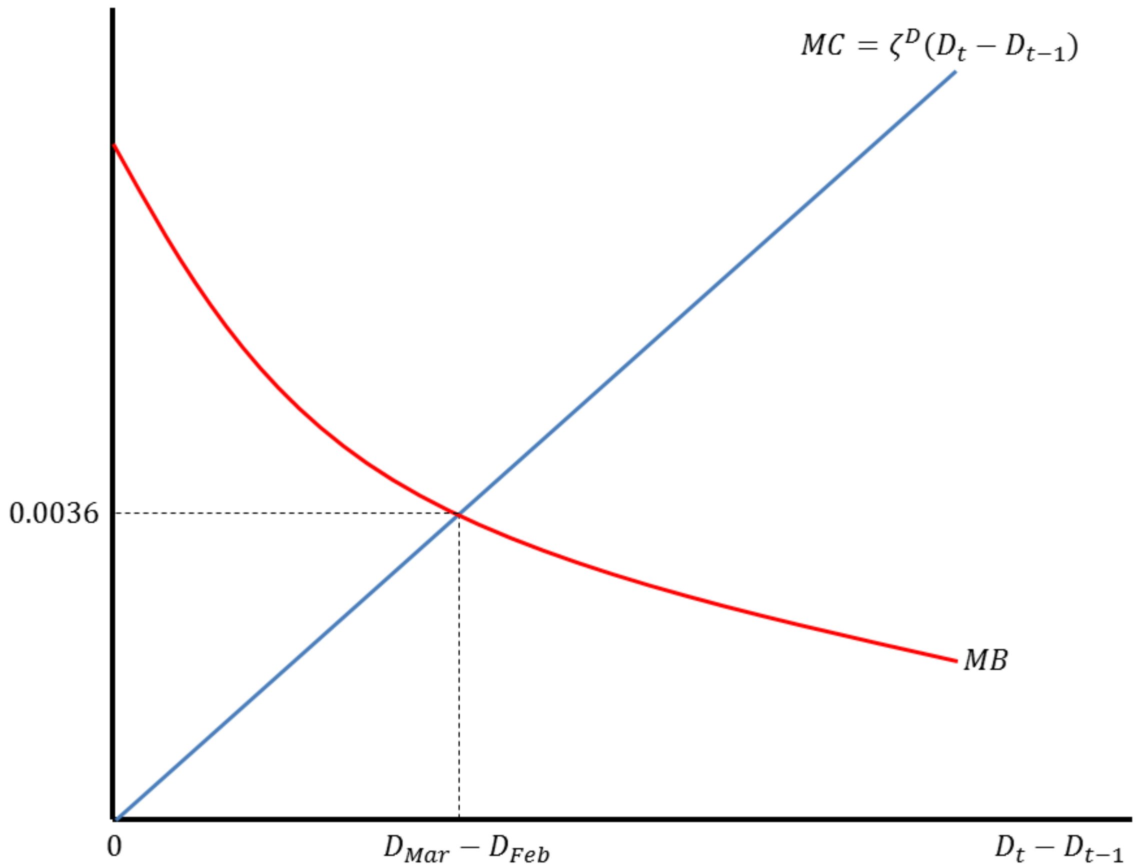

Figure 14. The costs and benefits of stockpiling at the margin

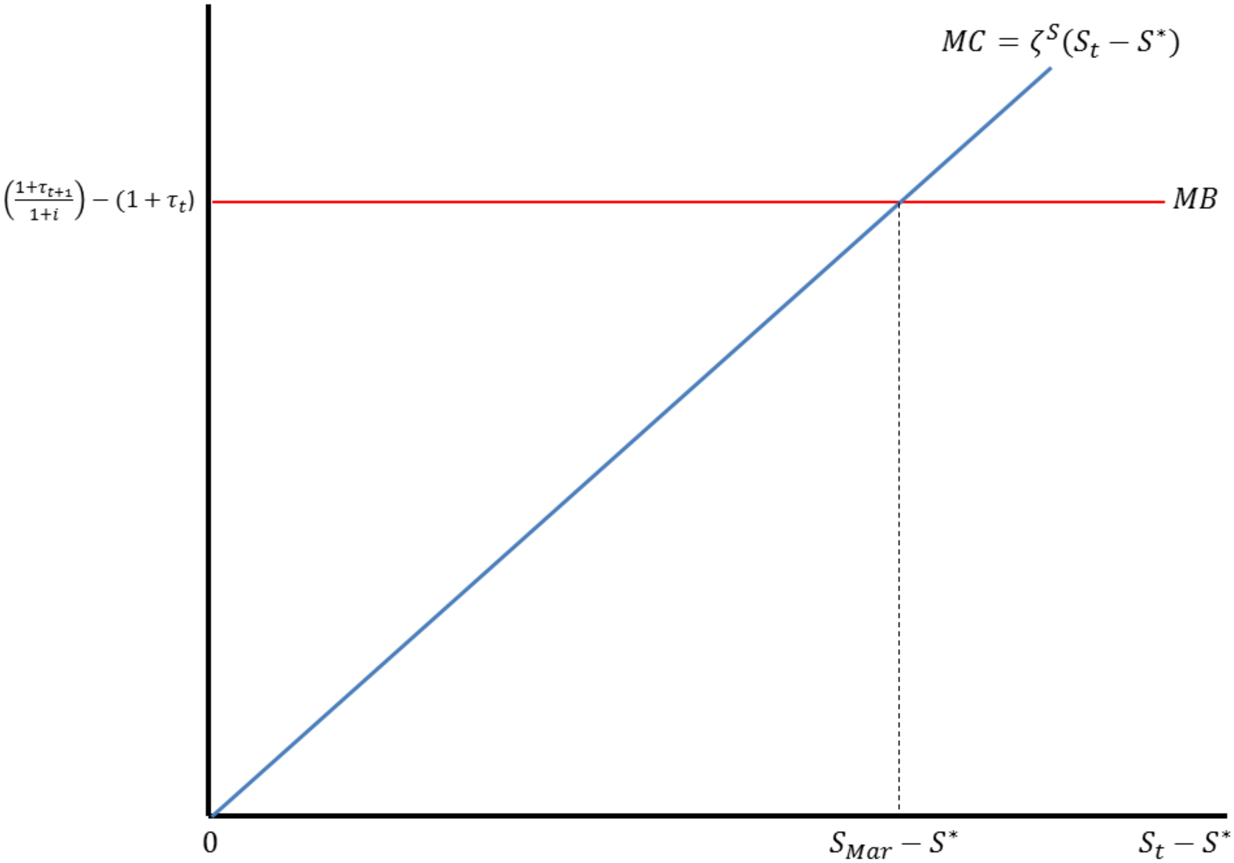

Figure 15. Comparison of the Time Path of Expenditures Generated by the Model to the Empirical Estimates

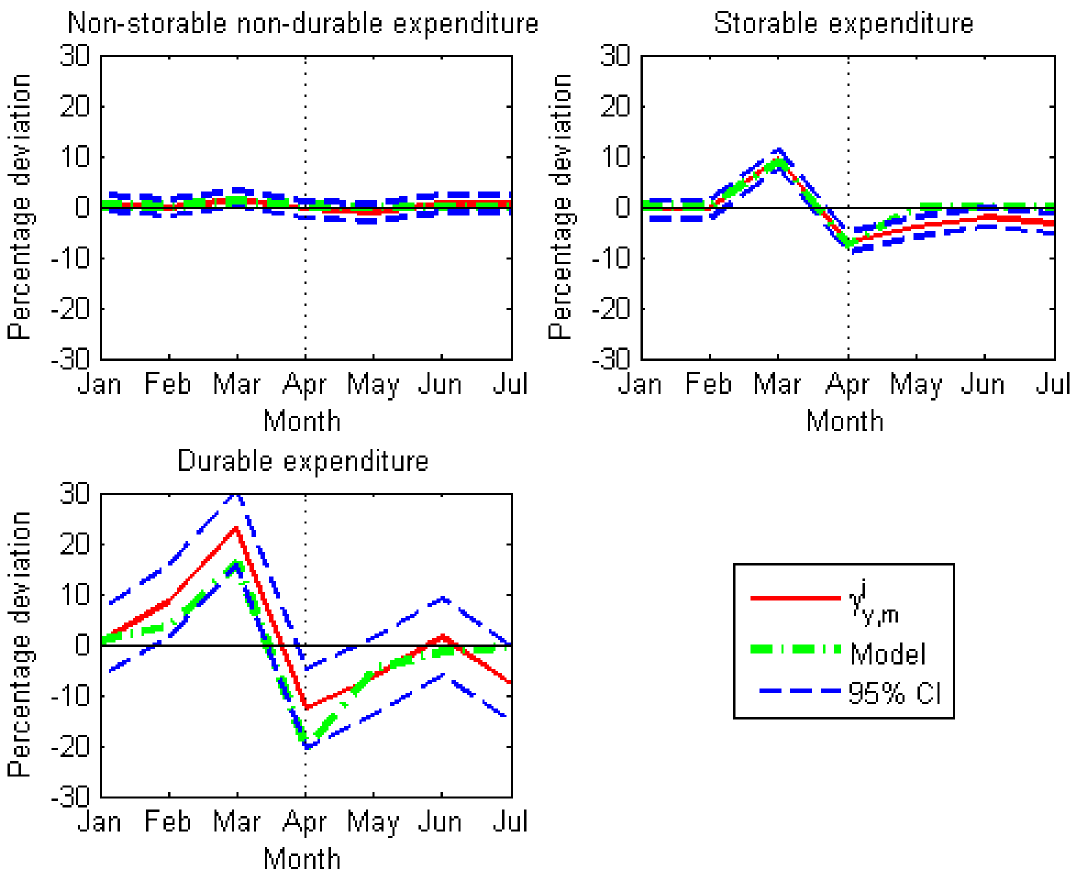

Figure 15 compares the time path of expenditures generated by the model to the empirical estimates shown in Figure 6. The dashed green line shows the time path generated by the model, while the solid red line shows the empirical estimates of the expenditure response to the VAT rate increase based on the JFIES survey data and the specification in Equation (4). The dashed blue lines are 95 percent confidence intervals for the empirical estimates. The dashed vertical line represents April 1997, the month the VAT rate increase was implemented.

Figure 16. Simulating the Response to New Zealand's July 1989 GST Rate Increase

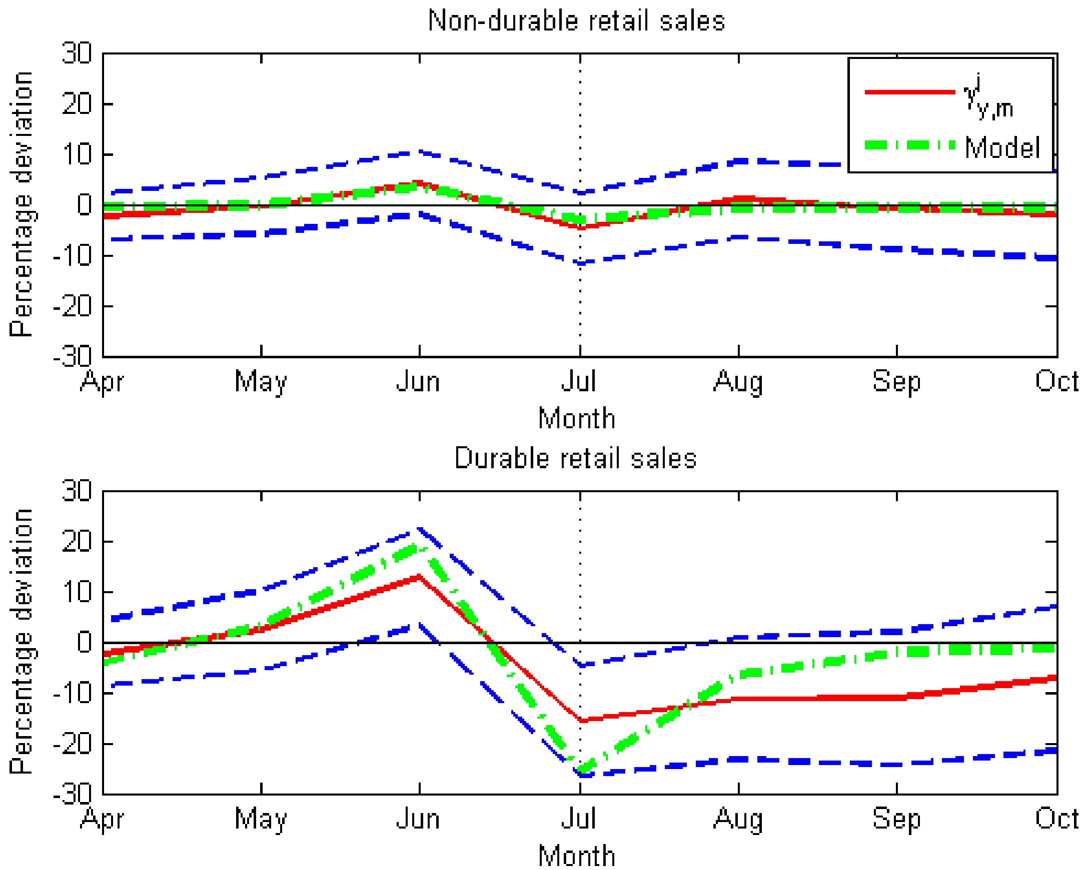

Figure 16 compares the time path of expenditures generated by the model to empirical estimates of the retail sales response to New Zealand's July 1989 Goods and Services Tax (GST) rate increase from ten to 12.5 percent, which is documented in Cashin (2011). The rate increase was announced in March 1989, four months prior to implementation, and was uncompensated. Because the tax rate increase was uncompensated, I allow for an income effect, while the structural parameter estimates characterizing intertemporal substitution remain unchanged.

Figure 17. Simulating the Expenditure Response to Japan's Phased-in VAT Rate Increase

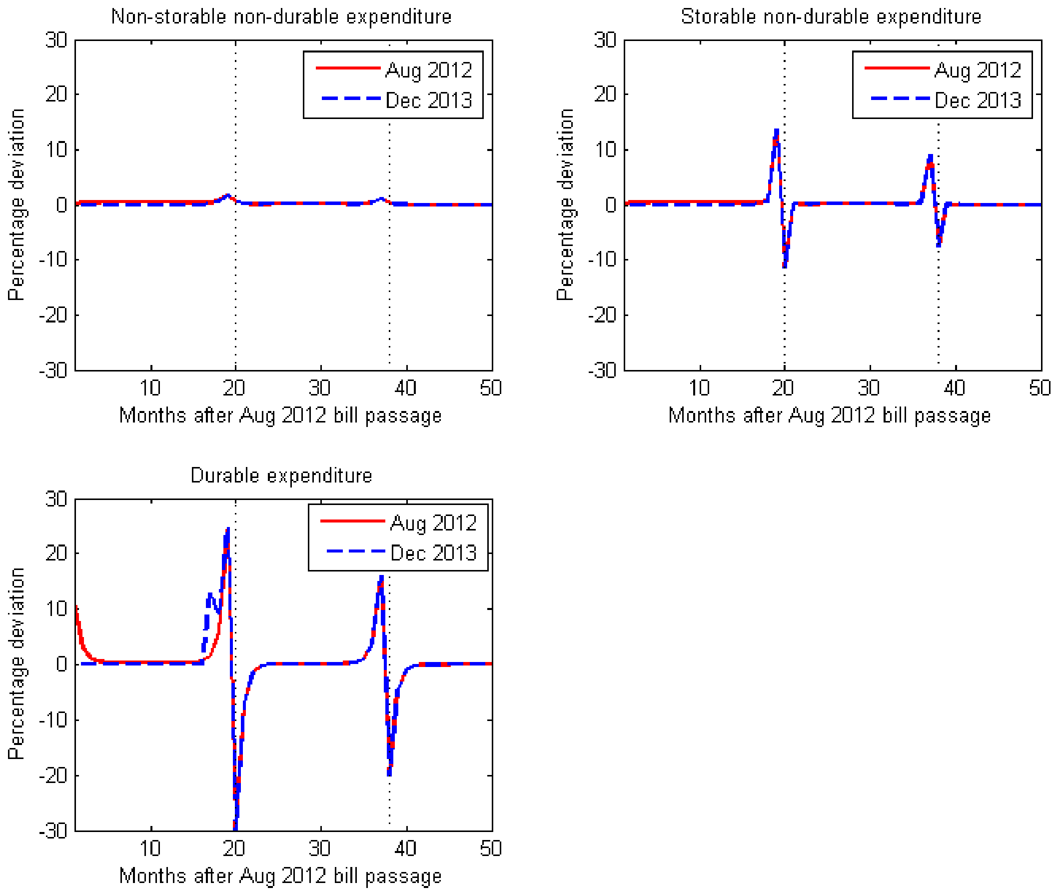

Figure 16 simulates the (Hicks) compensated expenditure response to the proposed VAT rate increase in Japan under two scenarios. The first, labeled 'Aug 2012' and represented by a solid red line, assumes that the August 2012 passage of a bill to increase the VAT rate constitutes announcement, with the rate increased from five to eight percent in April 2014, and from eight to ten percent in October 2015. The second scenario, labeled 'Dec 2013' and represented by a dashed blue line, assumes that passage of the fiscal year 2014 budget constitutes announcement.

TABLE 1. SUMMARY STATISTICS

Variable	Mean	Std.	Min	Max
Age of head	51.5	13.7	17	99
Number of household members	3.38	1.24	2	11
Number of household members under age 15	0.68	0.98	0	7
Number of household members aged 65+	0.47	0.75	0	4
Number of working members	1.52	0.95	0	7
Yearly income (1,000 yen)	7,113	4,652	0	97,043
Total expenditure (1,000 yen)	317	266	20	14,346
Excluding Tax Exempted items (1,000 yen)	221	195	15	9,255
Non-storable non-durables (N) (1,000 yen)	120	78	7	5,523
Storable non-durables (S) (1,000 yen)	52	32	.58	3,790
Durables (D) (1,000 yen)	47	138	0	7,678
Number of Observations	646,900			
Number of Households	129,380			

Note: Yearly household income and monthly household expenditures are listed in thousands of yen, with 2005 serving as the base year.

Table 2. Baseline Structural Parameter Estimates

Parameter	Estimate
σ	0.07 [-0.03, 0.17]
ϵ^D	-0.01 [-0.04, 0.02]
ζ^D	0.10** [0.01, 0.18]
ζ^S	1.18*** [1.02, 1.34]

95 percent confidence intervals for the structural parameter estimates above are listed in brackets, and are computed using the delta method (see Appendix B for a full explanation). ***, **, and * represent significance at the one, five, and ten percent levels, respectively.

Table 3. Sensitivity of Structural Parameter Estimates to the Fixed Parameter Values					
Calibrated parameter	**Value**	$\boldsymbol{\sigma}$	$\boldsymbol{\epsilon^D}$	$\boldsymbol{\zeta^D}$	$\boldsymbol{\zeta^S}$
Baseline	See Section 4.2	0.07	-0.01	0.10	1.18
r	0.0008 (0.01)	0.07	-0.01	0.09	1.18
	0.0025 (0.03)	0.07	-0.01	0.09	1.18
δ	0.01 (0.15)	0.05	-0.01	0.11	1.17
	0.03 (0.30)	0.11	-0.01	0.10	1.18
ψ^S	0.25	0.07	-0.01	0.09	1.39
	0.35	0.08	-0.01	0.11	0.96
X^D	0.35	0.09	-0.01	0.12	1.13
$\overline{C^N}$	0.50	0.06	-0.01	0.08	1.23

Numbers listed in parentheses represent annual values for the fixed parameters.

Appendix A

Figure A.1. The Impact of an Increase in the Real Interest Rate when Households are Net Savers

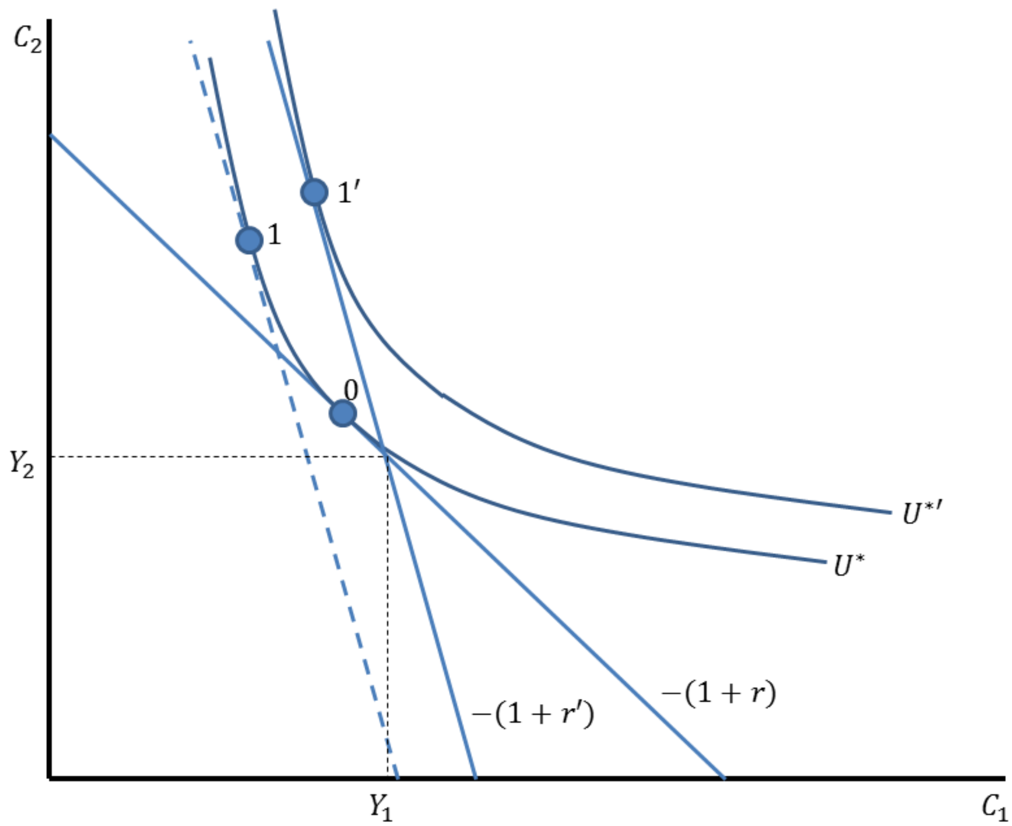

In the figure above, the real interest rate increases from r to r'. Prior to the change, the optimizing bundle for the representative consumer is given by bundle 0. Note that the consumer is a net saver since $C_1 < Y_1$, where C_1 is first period consumption and Y_1 is first period income. Given the increase in the real interest rate, the true intertemporal substitution effect is identified by holding utility constant at U^* while allowing for the increase in the interest rate. The new optimizing bundle would be given by bundle 1. However, the increase in the interest rate also induces an income effect, so the optimizing consumption bundle is given by bundle $1'$. The ratio of C_2 to C_1 is smaller at bundle $1'$ than it is at bundle 1, and thus a simple regression of the first difference of the log of the consumption ratio on the real interest rate will lead to a downward biased estimate of the IES.

TABLE A.1. CATEGORIZATION OF GOODS AND SERVICES SUBJECT TO THE VAT

Durables	Storables	Non-Storable Non-Durables
Tools	Grains (e.g. noodles)	Bread
Cooking appliance	Fish (dried, fish paste)	Fish (fresh)
Refrigerator	Meat (processed)	Meat (raw)
Vacuum	Dairy (e.g. butter)	Dairy (e.g. milk)
Washing machine/dryer	Vegetable (e.g. beans)	Vegetable (fresh)
Other household durables (e.g. microwave)	Fruit (canned)	Fruit (fresh)
Air conditioner	Oils, spices, and seasonings	Cake
Fan heaters	Sugar	Cooked food (e.g. sushi)
Stove	Sweets (e.g. chocolate)	Electricity
Other heating and cooling appliances	Cooked food	Natural gas
General furniture	Beverages (e.g. tea)	Water
Clock	Alcoholic beverages	Flowers
Lighting	Light bulbs	Newspaper
Floor coverings and curtains	Domestic goods (e.g. laundry detergent)	Eating out
Other interior furnishings	Cloth	Domestic services
Bedding	Medicine	Bus fare
Utensils	Medical supplies (e.g. bandages)	Taxi fare
Japanese clothing	Stationery	Airfare
Western clothing	Film	Other public transit
Women’s coats	Recording media (e.g. CD)	Automotive fees
Shirts	Pet food	Automotive insurance
Underwear	Personal care items (e.g. shaving cream)	Telephone service
Other clothing	Tobacco	Recreational good repair
Footwear	Rail service	Recreational durable good repair
Automobile	Gasoline	Lodging
Other vehicle		Package tour
Bicycle		Lesson fees
Auto parts		Television service
Telephone		Movie or play admission
Textbook		Other admissions
Television		Other recreational services
Stereo		Other insurance
Portable audio equipment		Social expenses (e.g. money gifts)
Video recorder		
Camera		
Computer	**(Durables Cont.)**	
Musical instrument	Personal effects (e.g. umbrella)	
Desk	Handbag	
Other recreational durable goods	Accessories (e.g. watch)	
Golf equipment	Other personal effects (e.g. cane)	
Other sporting goods	Home repair (e.g. plumbing)	
Sport outfits	Clothing services (e.g. tailoring)	
Toys	Auto repair	
Other recreational goods	Personal care services (e.g. haircut)	
Books	Personal effect services (e.g watch repair)	

Appendix B: Computing the standard errors

Denote the mapping in (5) as

$$\hat{\theta} = f(\hat{\gamma}).$$

By the delta method, it can be shown that

$$\sqrt{n}\big(\hat{\theta} - \theta_0\big) \xrightarrow{d} N(0, f'(\gamma_0)W\, f'(\gamma_0)^T),$$

where n is the number of observations used in the regressions that yield the $\hat{\gamma}'s$, $f'(\gamma_0)$ is a P x M matrix of derivatives, and W is the asymptotic variance-covariance matrix of $\sqrt{n}(\hat{\gamma} - \gamma_0)$. In practice, W is replaced by its sample estimate.

To compute $f'(\gamma_0)$, let

$$L\big(\hat{\gamma}, \gamma(\theta)\big) = \big(\hat{\gamma} - \gamma(\theta)\big)^T V^{-1}\big(\hat{\gamma} - \gamma(\theta)\big).$$

Then

$$\frac{\partial L(\hat{\gamma}, \theta)}{\partial \theta} = L_\theta(\hat{\gamma}, \theta) = 2\big(\hat{\gamma} - \gamma(\theta)\big)^T V^{-1} \frac{\partial\big(\hat{\gamma} - \gamma(\theta)\big)}{\partial \theta} = -2\big(\hat{\gamma} - \gamma(\theta)\big)^T V^{-1} \frac{\partial \gamma(\theta)}{\partial \theta}$$

is a 1 x P vector, where $\frac{\partial \gamma(\theta)}{\partial \theta}$ is an M x P matrix, and $L_\theta\big(\hat{\gamma}, \hat{\theta}\big) = \mathbf{0}$. By the implicit function theorem,

$$L_{\theta,\hat{\gamma}}(\hat{\gamma}, \theta) + L_{\theta,\theta}(\hat{\gamma}, \theta)\, f'(\gamma) = 0,$$

and it follows that $f'(\gamma)$ can be approximated as

$$f'(\gamma) = -L_{\theta,\theta}\big(\hat{\gamma}, \hat{\theta}\big)^{-1} L_{\theta,\hat{\gamma}}\big(\hat{\gamma}, \hat{\theta}\big),$$

where $f'(\gamma)$ is a P x M matrix, $L_{\theta,\theta}(\hat{\gamma},\theta)$ is a P x P matrix, and $L_{\theta,\hat{\gamma}}(\hat{\gamma},\theta)$ is a P x M matrix. Each element of $L_{\theta,\theta}(\hat{\gamma},\theta)$ can be expressed as

$$L_{\theta_p,\theta_q} = 2\sum_{m=1}^{M}\frac{1}{\widehat{\sigma_m^2}}\frac{\partial\gamma_m(\theta)}{\partial\theta_p}\frac{\partial\gamma_m(\theta)}{\partial\theta_q}\ \forall\ p = 1, \dots, P; q = 1, \dots, P$$

where $\widehat{\sigma_m^2}$ is the sample variance of $\hat{\gamma}_m$. Each element of $L_{\theta,\hat{\gamma}}(\hat{\gamma},\theta)$ can be expressed as

$$L_{\theta_p,\hat{\gamma}_m} = -2\frac{1}{\widehat{\sigma_m^2}}\frac{\partial\gamma_m(\theta)}{\partial\theta_p}\ \forall\ m = 1, \dots, M;\ p = 1, \dots, P.$$

$\frac{\partial\gamma_m(\theta)}{\partial\theta_p}$ is computed numerically as

$$\frac{\partial\gamma_m(\theta)}{\partial\theta_p} = \frac{\gamma_m\left(\hat{\theta} + he_p\right) - \gamma_m\left(\hat{\theta} - he_p\right)}{2h}\ \ \forall\ m = 1, \dots, M;\ p = 1, \dots, P,$$

where h is small and e_p is a P x 1 vector with a one in the p^{th} row and a zero in all others.[44]

[44] In practice, $h = 1 \text{ x } 10^{-6}$. The standard error estimates are robust to larger choices of h.

Appendix C: Computing deadweight loss

This section describes the methodology used to compute the deadweight loss measures presented in Section 4.7. The initial tax rate on expenditure is $\tau_0 = 0.05$. In period $t = t_A^*$, it is announced that the tax rate on expenditure will be increased from 0.05 to 0.08 in period $t = t_I^*$, and from 0.08 to 0.10 in period $t = t_I^{**}$. Producer prices are assumed to be fixed, so the entire burden of the tax rate increase is borne by consumers in the form of higher prices. Let p_0 represent the vector of intertemporal prices under the old tax regime (i.e. $p_0 = 1.05 \;\; \forall\, t$), and p_1 represent a vector of prices under the new tax regime, where

$$p_1 = [1.05, t = t_A^*, \ldots, t_{I-1}^*; 1.08, t = t_I^*, \ldots, t_{I-1}^{**}; 1.10, t = t_I^{**}, \ldots, \infty]$$

To calculate the deadweight loss associated with the pre-announced and phased-in tax rate increase, begin with the compensating variation measure of deadweight loss (in present value), given by

$$DWL = \sum_{t=1}^{\infty} \left(\frac{1}{1+i}\right)^{t-1} \left[\left[E(p_{1,t}, U_0) - E(p_{0,t}, U_0)\right] - R(p_{0,t}, p_{1,t}, U_0)\right],$$

where i is the nominal interest rate, $\left[E(p_{1,t}, U_0) - E(p_{0,t}, U_0)\right]$ is the amount required in period t to leave a household as well off (in terms of the present value of liftetime utility, U_0) after the tax change as it was beforehand, and $R(p_{0,t}, p_{1,t}, U_0)$ is the change in revenue in period t between the two tax regimes, holding U_0 constant.

Using a second-order Taylor series expansion around $p_{0,t}$, $\left[E(p_{1,t}, U_0) - E(p_{0,t}, U_0)\right]$ can be approximated as

$$\begin{aligned}\left[E(p_{1,t}, U_0) - E(p_{0,t}, U_0)\right] &\approx \left.\frac{\partial E(p, U_0)}{\partial p}\right|_{p_{0,t}} (p_{1,t} - p_{0,t}) + \frac{1}{2}\frac{\partial}{\partial p}\left.\frac{\partial E(p, U_0)}{\partial p}\right|_{p_{0,t}} (p_{1,t} - p_{0,t})^2 \\ &= \Delta\tau_t x^c(p_{0,t}, U_0) - \frac{1}{2}\Delta\tau_t \Delta x^c(p_{0,t}, U_0), \qquad (1)\end{aligned}$$

where $\Delta\tau_t$ is the difference in the tax rate in period t between the new and old tax regime, $x^c(p_{0,t}, U_0)$ is net-of-tax expenditure in period t under the original tax regime, and $\Delta x^c(p_{0,t}, U_0)$ is the difference in net-of-tax expenditure between the old and new tax regime, holding utility constant at its original level.

$R(p_{0,t}, p_{1,t}, U_0)$ can be expressed as

$$\begin{aligned} R(p_{0,t}, p_{1,t}, U_0) &= (p_{1,t} - p_{0,t})x^c(p_{1,t}, U_0) \\ &= \Delta\tau_t x^c(p_{1,t}, U_0), \qquad (2) \end{aligned}$$

where $x^c(p_{1,t}, U_0)$ is net-of-tax expenditure in period t under the new tax regime, holding utility constant.

Combining (1) and (2), rearranging, and simplifying, one is left with

$$DWL = \frac{1}{2}\sum_{t=t_I^*}^{\infty}\left(\frac{1}{1+i}\right)^{t-1}\Delta\tau_t\Delta x^c(p_{0,t}, U_0) \qquad (3)$$

Computing the present value of the deadweight loss for periods $t = t_I^{**}, \ldots, \infty$ is relatively straightforward. I simulate the compensated expenditure response to the phased-in tax rate increase, determine the difference between net-of-tax expenditures under the two tax regimes in each period t, and substitute these values into (3). Figure 3 provides an illustration for any period $t > t_I^{**}$.

Computing the present discounted value of the deadweight loss for periods $t = t_I^*, \ldots, t_{I-1}^{**}$ requires an additional simulation. Why? Note that the price of consumption in periods $t = t_I^{**}, \ldots, \infty$ increases relative to the price of consumption in periods $t = t_I^*, \ldots, t_{I-1}^{**}$. As a result, the compensated demand curve for goods purchased in periods $t = t_I^*, \ldots, t_{I-1}^{**}$ may shift to the right, as households substitute away from consumption during the relatively high price later periods. If one does not control for this rightward shift in the demand curve, $\Delta x^c(p_{0,t}, U_0)$ would be biased downwards in the periods $t = t_I^*, \ldots, t_{I-1}^{**}$, as illustrated in Figure C.2. To avoid this problem, it is necessary to simulate the compensated expenditure response to a permanent increase in the tax rate on expenditure from 0.05 to 0.08 in period $t = t_I^*$, which is announced in period $t = t_A^*$. Doing so ensures that there will only be movement along the compensated demand curve, as opposed to the movement along and shift in the demand curve that are induced under the first simulation. Figure C.3 illustrates. After running this simulation, I determine the difference between net-of-tax expenditures under the two tax regimes in each period t, and substitute these values into (3) as before.

Figure C.1. Deadweight Loss in Periods $t > t_I^{**}$

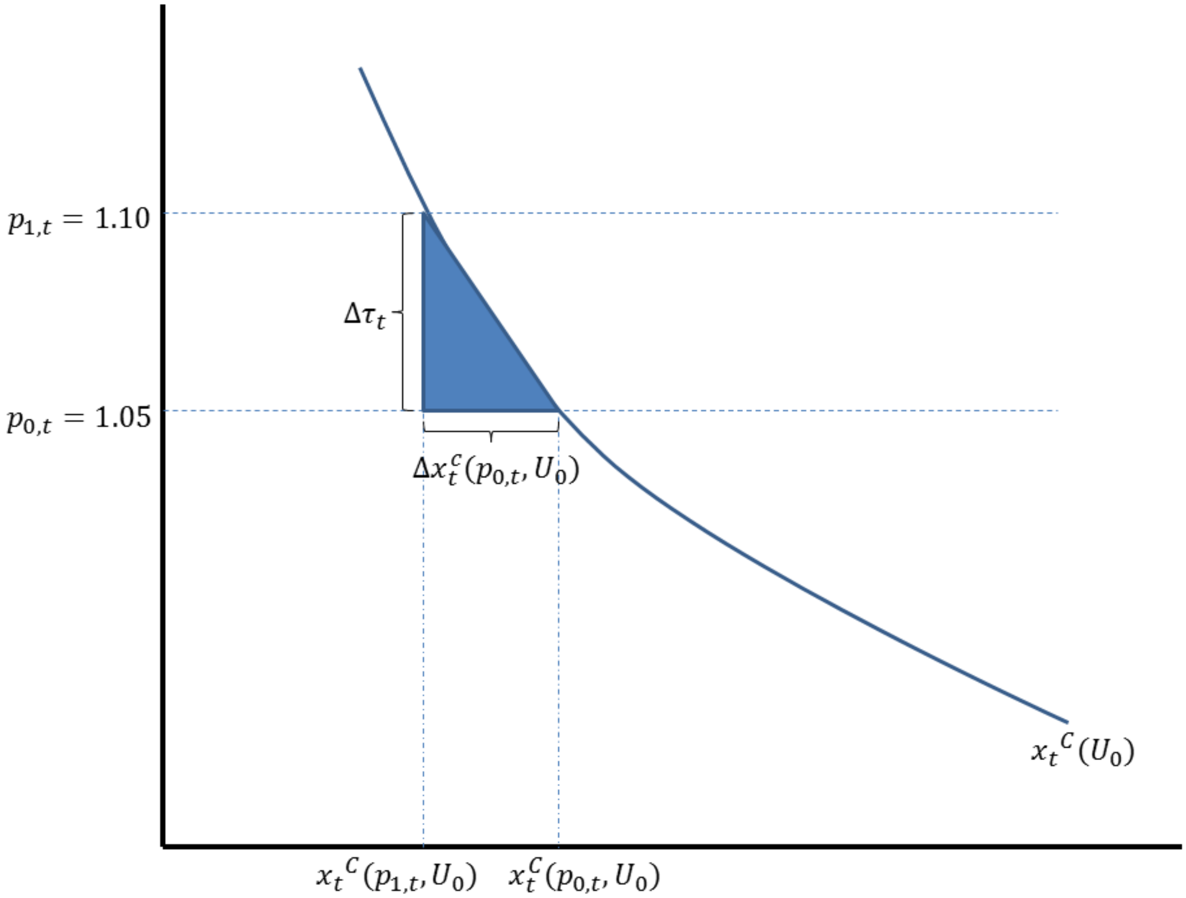

Figure C.2. Biased Deadweight Loss in Periods $t = t_I^*, \ldots, t_{I-1}^{**}$

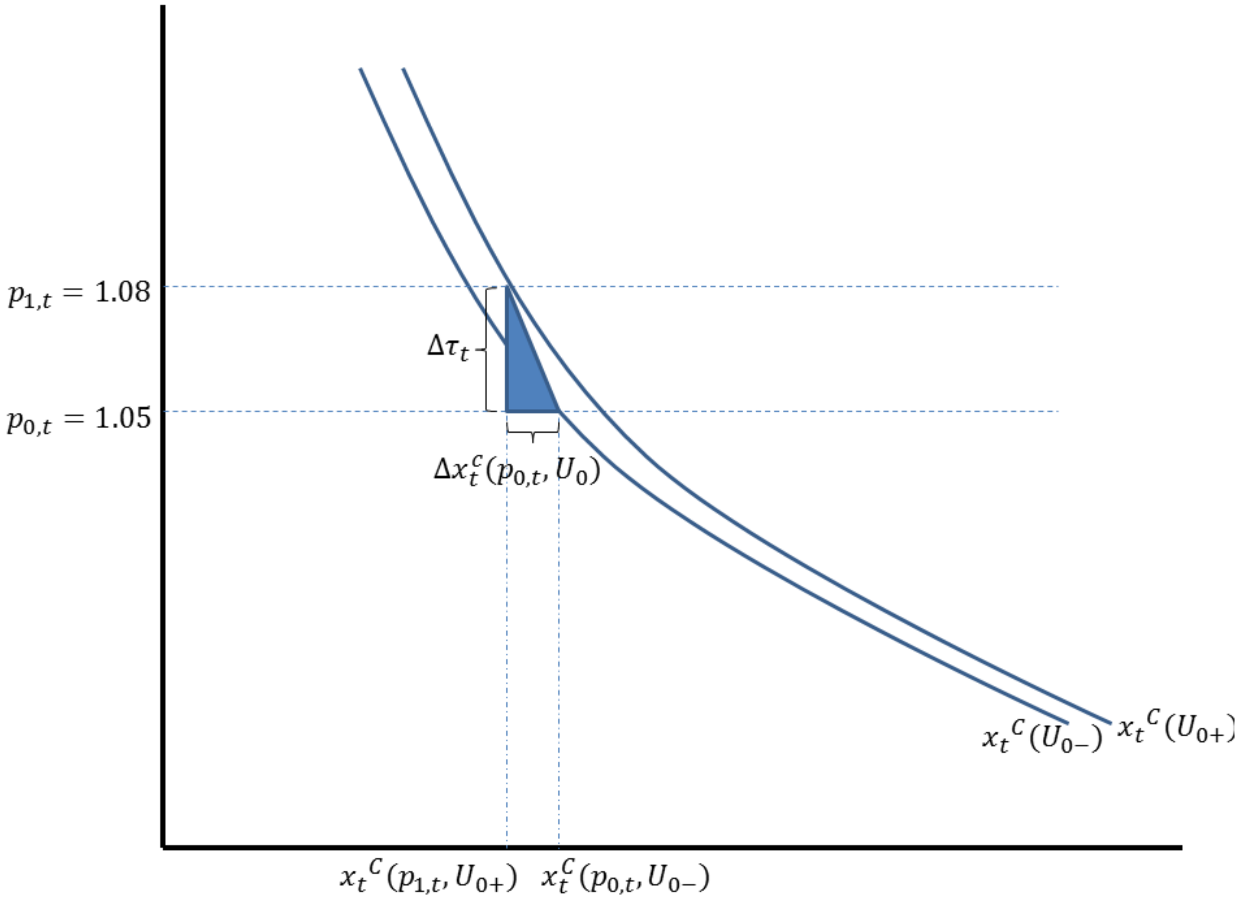

Figure C.2. Corrected Deadweight Loss in Periods $t = t_I^*, \ldots, t_{I-1}^{**}$