Timing vs. Long-run Charitable Giving Behavior: Reconciling Divergent Approaches and Estimates

December 2004

Jon Bakija
Williams College

and

Rob McClelland
Congressional Budget Office

Abstract:

In this paper we examine the effect of the income tax on charitable giving. An important challenge in this literature has been to estimate the long-run response of giving to a persistent change in tax-price, which can be difficult to distinguish from intertemporal substitution arising from differences between current and expected future tax prices, arising for example due to transitory fluctuations in incomes, life-cycle factors, or pre-announced tax reforms. Several papers that have attempted to distinguish these effects have found that the elasticity of charitable giving with respect to a persistent price change is small, while the elasticity with respect to a transitory difference between current and expected future prices is large. Auten, Sieg, and Clotfelter (2002) advance this literature by developing an estimation procedure that incorporates a more sophisticated model of the stochastic process for income. In contrast to previous research on the topic, they find the counterintuitive result that the immediate response to a persistent price change is much *larger* than the immediate response to a one-period transitory price change. In this paper, we present a new estimation procedure that allows us to implement their assumptions about the stochastic process of income in a more conventional regression framework, and then adapt the procedure to take into account the pre-announced and phased-in nature of tax reforms that occurred during the sample period. In preliminary analysis based on a public-use panel of individual tax return data, we are able to replicate their counterintuitive pattern of price elasticities, and find that incorporating information about pre-announced and phased-in tax law changes reverses their result – the persistent price elasticity is reduced substantially, and the transitory price elasticity is now the larger of the two. We also try an instrumental variables strategy that relies exclusively on federal and state tax reforms for identification, and this yields similar results. Finally, we incorporate a dynamic adjustment process into the empirical specification, and find evidence that the long-run response to persistent price and income changes is larger than the immediate response.

Introduction

In the United States, the federal government and most state governments offer a price subsidy for charitable giving, in the form of a deduction against income tax for taxpayers who itemize. In this way, the marginal cost (or "tax price") of a charitable contribution relative to non-deductible consumption is reduced from one to one minus the marginal income tax rate. The role of this price subsidy in encouraging charitable giving has long been a topic of political debate. For instance, the creation of the standard deduction in 1944 was opposed because of a fear that it would reduce charitable donations.[1] More recently, the issue has come up in the context of tax reform plans that would limit or eliminate deduction (for example, the initial proposal that led to the Tax Reform Act of 1986, and the Hall-Rabushka "flat tax"), and the recent Bush administration proposal to extend the deduction to non-itemizers.[2]

The sensitivity of charitable contributions to the price subsidy provided by the income tax has been the subject of extensive empirical research among economists, in part because economic theory suggests that the optimal size of subsidy for charitable giving depends positively on the degree of responsiveness to the subsidy. Contributions to charitable organizations may be thought of as a form of private consumption. After all, the contribution would not be made voluntarily if doing so did not yield some utility to the donor -- for example, a "warm glow" feeling, as in Andreoni (1990). But a charitable contribution may also involve some public good aspect or positive externality, where the benefit to society is greater than the benefit to the donor at the margin. In that case, some form of price subsidy to charitable giving may be justified on efficiency grounds. It has long been recognized that in an optimal tax framework, the efficient rate of subsidy for a good with a positive externality is an increasing function of the price elasticity of demand for that good (Sandmo, 1975). Intuitively, the efficient subsidy rate is the Pigouvian subsidy rate (which depends on the size of the marginal external benefit at the efficient quantity), minus an optimal tax rate, and the latter becomes smaller when demand for the good is more responsive to price. Saez (2004) shows that this same positive relationship between the optimal subsidy and the price elasticity continues to hold in a richer model that allows for direct government provision of the same public good supported by the charity, crowding out of private contributions by public provision, and social welfare function that puts positive weight on redistribution towards the poor. Thus, a good estimate of the price elasticity of charitable giving would be a particularly useful piece of information for evaluating public policy towards charity.

Econometric challenges

Efforts to estimate the sensitivity of charitable giving to its price face a number of difficult challenges. An especially important concern is that current prices and disposable incomes may differ in a systematic way from their expected future values. Such differences may occur for three main reasons: fluctuations in pre-tax income that are known by the taxpayer to be transitory; changes in tax law that are pre-announced or gradually phased in over time; and variation over time in income arising from life-cycle factors, economy-wide productivity trends, and other influences that are predictable with some degree of error. Each of these three factors may affect price as well as income. Any predictable difference between current and future income can also produce a predictable difference between current and future price by pushing the taxpayer into a different marginal tax rate bracket. On the price side, all of these factors create opportunities for intertemporal substitution – taxpayers have a strong incentive to give to charity during those years when their marginal tax rates are temporarily high (and thus the price of giving is low). In the cross section, we observe that people with low prices give more to charity, but this may partly reflect people substituting giving from other years when they face higher prices, and thus does not necessarily indicate any increase in the total amount of lifetime giving in response to tax incentives. On the income side, Milton Friedman's permanent income hypothesis (1957) suggests that consumption should respond more weakly to a transitory fluctuation in income than to a persistent change in lifetime income. Since current income is a mixture of transitory and permanent components, income elasticities based on cross-sectional data would tend to be downwardly biased estimates of the effect of a persistent change in income.

Figure 1 offers a stylized illustration of the problem. In the top panel, the bold black line a-b represents a predictable life-cycle pattern of disposable income for an individual, and the dashed line around a-b represents fluctuations of disposable income along that life-cycle pattern that are known by the individual to be temporary at the time they occur. Ideally, we would like to have data on the expected present value of lifetime resources, which the permanent income hypothesis would suggest is an important determinant of current consumption decisions. Instead, what we typically observe in the data is current income, the height of the dashed line around a-b, which may be only weakly related to the present value of lifetime resources. On the income side, for policy purposes, we are generally most interested in identifying the elasticity of charitable giving with respect to a permanent change in disposable income. This would be helpful, for example, in an effort to estimate how a tax reform that affects disposable income in a persistent way would affect charitable giving. Conceptually, we would like to identify persistent shocks to disposable income that shift the whole life-cycle profile of disposable income, illustrated by the upward shift to line segment c-d that occurs at time *t* in the top panel of Figure 1. The response of charitable giving to such a shock would tell us something about the effects of a tax reform that changes disposable income in a persistent way. Changes in disposable income arising from the transitory fluctuations along the long-run profile, as well as changes due to predictable life-cycle patterns, may not be associated with any change in the expected present value of lifetime resources, and thus we would like to distinguish those sorts of changes from the kind of change illustrated by the shift to line c-d.

The bottom panel of Figure 1 illustrates how similar considerations apply to price. The predictable life-cycle path of prices is inversely related to the pattern of income. Because of the progressivity of the income tax, those periods of life when income is highest are also the periods when marginal tax rates are highest, and thus the price of charitable giving is lowest. The bold black line e-f represents the predictable life-cycle pattern of prices, and the dashed line around it represents fluctuations in price that are known to be transitory at the time they occur. Three types of opportunities to reduce taxes through intertemporal substitution are illustrated here. First, taxpayers may want to concentrate their giving during the portion of the life cycle when the price of giving is the lowest (coinciding with the time when income is the highest). Second, taxpayers may want to concentrate giving into years when transitory income fluctuations make the price temporarily low (represented by the dips in the dashed line). Third, suppose that a tax reform that permanently increases the price of giving (represented by the upward shift in the whole time path of prices to line g-h) is implemented at time t. If this tax reform were announced in advance at time t-1, this would create a strong incentive to accelerate charitable giving into year t-1, in anticipation of the future price increase. Conceptually, we would like to distinguish these various sorts of intertemporal substitution behavior from the long-run response of giving to the shift in the whole path of prices occurring at time t illustrated in the figure. It is this latter response that is most interesting from a policy perspective, as it is the size of this response that is relevant to determining the optimal degree of subsidy to charity.

There is well-known and striking time-series evidence of intertemporal substitution of charitable giving among very high-income individuals in anticipation of the implementation of the Tax Reform Act of 1986 (TRA86), suggesting this is an important concern. TRA86 was enacted in October of 1986, and its provisions were gradually phased-in between 1987 and 1988. The act reduced the marginal income tax rate in the top bracket from 50 percent to 28 percent, and also amended the alternative minimum tax base to include any accumulated capital gains on charitable donations of appreciated property (the latter provision was eliminated for tangible property in 1991, and for all property in 1993). Both of these created a strong incentive for high-income people to accelerate charitable donations into 1986, before the aforementioned provisions began to take effect. Indeed, towards the end of 1986, the financial press ran stories advising readers to make any large charitable donations before the end of 1986 rather than waiting (see for example, Simpson, 1986). Figure 2 illustrates how charitable giving changed as a share of 10-year average income between 1979 and 1988 for taxpayers in different 10-year average income classes (from data reported in Auten, Cilke, and Randolph, 1992). There is clear evidence of an enormous spike in giving in 1986 among taxpayers with persistently high incomes. Among millionaires (in 1991 dollars), giving surged from 12.3 percent to 20.4 percent of 10-year average income between 1985 and 1986, and then dropped back to 12.3 percent in 1987. There is no obvious alternative explanation, aside from intertemporal substitution in response to an anticipated price increase, for this 1986 spike. The figure also makes clear that only very-high-income taxpayers exhibited a large degree of intertemporal substitution. For instance, even those with incomes between $100,000 and $200,000 barely registered a blip in giving in 1986. This may be explained by the fact that the expected change in marginal tax rates was smaller for this group, and because they were less likely to face the alternative minimum tax. There may also be some difference in the size of the transitory price elasticity between this group and the highest income group. In any event, it suggests that controlling for year dummies will not remove this effect from an econometric analysis, and that timing in response to pre-announced tax changes is likely to be particularly important in data that heavily oversamples high-income people, as has been the case with much of the recent literature.

While TRA86 is the most dramatic example of an incentive to re-time charitable contributions, there have in fact been numerous other examples of this phenomenon.[3] The Economic Recovery Tax Act of 1981 reduced the top marginal rate from 70 percent to 50 percent and eventually reduced all other marginal rates by 23 percent of their former levels. This law was enacted in August of 1981, but only implemented minor rate reductions in 1981 (except for an immediate reduction in the capital gains rate) – the bulk of rate reductions were gradually phased in between 1981 and 1984. The 1990 Omnibus Reconciliation Act of 1990 (OBRA90), enacted in November 1990, increased the marginal rate in the top bracket to 31 percent, and also imposed a limitation on itemized deductions for high-income taxpayers – both provisions were implemented beginning in 1991. An increase in the top marginal rate to 39.6 percent was enacted in 1993, applying retroactively to the beginning of 1993. Although the tax increase was not officially enacted in 2002, it was likely anticipated by the end of 2002 due to Clinton's victory in the November Presidential election. There is substantial evidence that high-income taxpayers accelerated realization of income (for instance, bonuses and exercises of stock options) into 1992 in anticipation of the 1993 tax increases (Goolsbee, 2000).

Another set of challenges in estimating the responsiveness of charitable giving to taxes arises because of the potential for omitted variable bias and specification error. First, influences on charitable behavior such as education, alumni affiliation, degree of religiosity, and personal generosity are generally omitted from available data. These are likely to be associated with income in a non-linear fashion, and thus may introduce omitted variable bias into the estimates of both income and price effects. To the extent that these omitted variables are time-invariant, this problem may be addressed with panel data, through the use of fixed effects or differencing. In addition, as Feenberg (1987) has pointed out, the marginal federal income tax rate in any given year is a non-linear function of income and a few other variables (such as marital status) that may also influence giving independently. As a result, in a cross sectional analysis, the separate identification of the price and income effects depends almost entirely on the non-linear relationship between the two. The true functional form of the relationships among giving, price, and income cannot be known with certainty, and if the functional form used for estimation is incorrect, estimated price and income elasticities may be biased because each estimate may partly reflect the effects of non-linear functions of the other. Because federal and state tax reforms cause changes in tax rates that are independent of changes in income over time, panel data estimators that rely on these sorts of changes for identification, and remove the identification caused by the cross-sectional non-linear relationship between income and price, may be more robust to this potential source of bias.

A third kind of complication arises because the response of charitable giving to a change in income or price may be gradual. One reason is that charitable giving may exhibit habit formation, where giving in one period may be positively related to giving in previous periods. For example, suppose a tax reform reduces the price of charitable giving, and this causes individuals to initially increase their giving by some small amount. This may in turn lead to further increases in giving in the future, for instance because each donor begins to develop relationships additional charitable organizations, becomes the target of greater fundraising activity, and so forth. Another possibility is that when a shock to income or price occurs in one period, it may take several subsequent periods of information for an individual to determine whether the shock is going to be permanent or transitory. In either case, the long-run steady state response to a change in price or income may exceed the immediate response.

Prior research

Early empirical literature on the effects of income taxes on charitable giving typically involved regressions on single-year cross-sectional data where the dependent variable was log charitable giving, and explanatory variables included the log tax-price, log disposable income, and a vector of demographic characteristics. Using this approach, Clotfelter and Steuerle (1981) estimated a price elasticity of -1.27 and an income elasticity of 0.78, a result that Clotfelter (1985) finds to be representative of the results in a very extensive literature of this type. These early studies, however, are particularly vulnerable to the problems discussed above.

Several empirical studies have addressed at least some significant subset of the challenges discussed above using panel data sets of individual income tax returns, taking advantage of the large changes in tax law during the 1980s to help identify their models. Early examples include Broman (1989), Barrett (1991), Randolph (1995), and Barrett, McGuirk, and Steinberg (1997, henceforth BMS).[4] Estimates in these studies generally suggest that elasticity of giving with respect to a persistent price change is small, while the elasticity with respect to a transitory fluctuation in price is large, implying that people are more willing to change the timing of their giving than the long-run level of their giving in response to tax incentives. Among these studies, Randolph and BMS provide the most complete treatment of the issues discussed above.

Randolph estimates his model on a 1979-1988 panel of individual tax returns that heavily over-samples high-income returns, and uses the charity expenditure share (price times giving divided by disposable income) as the dependent variable. He treats current-year price and disposable income as measurements with error of their permanent values, and uses 10-year average pre-tax income, interacted with dummies for time periods representing different tax regimes (1979-80, 1983-85, 1987-88) as the instruments for predicting the permanent components of current income and price. Transitory fluctuations in price and income are then identified based on the difference between current and ten-year average pre-tax incomes, again interacted with tax regime dummies. Data for 1979-80 and 1986 are dropped from Randolph's analysis because of substantial differences between current and announced future tax law during these years, an issue for which Randolph's procedure offers no solution. Because the instruments for predicting permanent income and price rely so heavily on a time-invariant variable (the individual's ten-year average income), Randolph estimates a random effects rather than fixed effects model. When evaluated at the mean charity expenditure share for his sample, Randolph's estimates imply an elasticity of giving with respect to permanent price of -0.08, and an elasticity of giving with respect to a transitory deviation in price of -2.27.[5]

BMS estimate a log-linear fixed effects model on years 1979-86 of the University of Michigan public use panel of individual income tax returns. To address the possibility of intertemporal substitution, they include price from the current period, one period ahead, and one period lagged as explanatory variables. Similarly, current, one period ahead, and one period lagged values of disposable income are included in an effort to distinguish the effects of persistent and transitory changes in income. Lagged giving is included as an explanatory variable to address the possibility of gradual adjustment. They estimate a transitory price elasticity of -1.18, and a long-run permanent price elasticity of -0.47. Adjustment to price and income changes is found to occur quickly, as the estimated coefficient on lagged giving is only 0.16.[6]

The main contribution to this literature by Auten, Sieg and Clotfelter (2002) (ASC hereafter) is to develop an estimation procedure that allows income to follow a more reasonable stochastic process over time. The prior literature typically used averages of income from both past and future years, lagged averages of income, or the next year's actual income as proxies for expected values of future income. As ASC point out, these approaches involve restrictive assumptions about how the persistent and transitory components of income evolve over time, and may mischaracterize that process. For instance, in Randolph's procedure, because the instrument for permanent income does not change over time except when there is a change in tax regime, and then only in the same way for everyone who has the same average income, many unanticipated persistent shocks to income will be mischaracterized as transitory shocks, and this will lead to similar mischaracterization of persistent price shocks as transitory. In addition, many of these approaches assume that individuals have precise knowledge in advance of actual future values of income (which would be necessary, for example, to calculate a 10-year average of income that includes future years of income).

ASC begin by regressing change in log disposable income, change in log price, and change in log giving on demographic and life-cycle characteristics and year dummies, and then proceed to analyze the residuals from each of these regressions. They assume that residual log disposable income consists of a persistent component that follows a random walk, and a transitory component that is pure white noise. Persistent and transitory changes in residual log price are assumed to be linear functions of persistent and transitory income changes, respectively, plus permanent and transitory shocks that are independent of income and follow a similar stochastic process to income. Change in residual log giving is then assumed to be a linear function of persistent and transitory shocks to residual log price and residual log disposable income. The assumptions about the stochastic processes, together with the coefficients in the assumed linear relationships among the variables, imply a particular pattern of covariances among the contemporaneous and one-period lagged residual first-differenced values of log giving, log disposable income, and log price. ASC then estimate the parameters of the model using a minimum distance estimator, which chooses the parameters of the model to minimize the difference between the covariance structure implied by those parameters, and the empirical covariance structure estimated from the sample. Their model is estimated on essentially the same data set that Randolph used, extended forward to 1993, and all variables are first-differenced to control for unobserved heterogeneity.

In contrast to virtually all of the previous literature that addresses intertemporal substitution and unobserved heterogeneity in charitable giving, the estimates in ASC imply that the persistent price elasticity is large, generally larger than -1 (although this depends to some extent on the sample period chosen, and the assumptions about the stability of the variances of shocks over time). ASC consistently find, unlike previous research, that the transitory price elasticity is much smaller than the persistent price elasticity. This pattern is difficult to reconcile with economic theory. For instance, a model with habit formation or slow adjustment would suggest that the *long-run* response to a persistent price change should be larger than the response to a temporary price change – but there is no reason why the *immediate* response to the persistent price change should be larger. The pattern is also difficult to reconcile with the spike in giving in 1986 illustrated earlier in Figure 1, which appears to be strong evidence of responsiveness to a transitory difference between current and expected future prices, at least among those with very high incomes.

The counterintuitive pattern of persistent and transitory price elasticities in ASC may arise because their procedure does not deal well with pre-announced and phased-in changes in tax law. For instance, there is nothing in the model that explicitly takes into account the fact that the large price increases occurring in 1987 and again in 1988 were already fully anticipated in 1986. As a result, the ASC technique must treat the large negative first-difference of giving between 1986 and 1987 among high income taxpayers largely as a response to the persistent increase in observed *implemented* prices that begins in 1987. In fact a large part of the 1986-87 change in giving almost certainly reflects inter-temporal substitution into 1986 as contributors responded to the change in price before it actually began to be implemented. This would tend to bias the estimated persistent price elasticity upwards in absolute value, particularly in data where very high-income people are overrepresented. The transitory price elasticity in ASC may be biased towards zero essentially due to measurement error -- their measure of transitory price shock exhibits little independent variation, so that the effective noise-to-signal ratio in this variable may be large. Transitory price shocks arising from transitory income shocks are assumed to be a linear function of those income shocks, so that part of the transitory price shock exhibits no variation independent from income. Transitory price shocks that are unrelated to income shocks must come from tax reforms. But the kind of tax reform that would yield a transitory shock to price in the ASC approach would be one that changes the price for a single year, and then immediately reverses itself the following year – this sort of tax change did not occur during the sample period in question.

In what follows, we develop an estimation strategy that takes advantage of the improved specification of the stochastic process for income presented in ASC, but that assumes that taxpayers take advantage of publicly available information about future changes in the nature of the tax system that have already been enacted. Our procedure makes efficient use of the highly non-linear ways in which the relationship between income and price changed over time. In addition, we strengthen the identification of our empirical models by incorporating both federal and state income tax laws – with the exception of Feenberg (1987), state income tax variation has not been used in empirical analysis of the effects of income taxes on charitable giving. As a first step we show how the ASC approach can be replicated in a multi-step regression framework. Breaking the procedure down into an approximate regression equivalent will help clarify how the approach used in ASC compares to the previous literature, which has been exclusively regression based. It will also permit us to adapt the procedure to allow for a richer specification of how tax laws changed current and expected future prices and incomes over time, which we later find to be crucial to reconciling the results of ASC with the rest of the literature and to obtaining consistent estimates.

Basic empirical model

The most basic form of model we wish to estimate is:

$$(1)\ \Delta g_{it} = \Delta\alpha_t + \beta_1 \Delta p^*_{it} + \beta_2 \Delta p^T_{it} + \beta_3 \Delta y^*_{it} + \beta_4 \Delta y^T_{it} + \boldsymbol{\beta_5} \Delta \mathbf{X}_{it} + \Delta \nu_{it}\,,$$

where Δg_{it} is the change in the natural log of giving by taxpayer *i* between time *t*-1 and time *t*, $\Delta\alpha_t$ is the change in time specific intercept between *t*-1 and *t*, Δp^*_{it} and Δy^*_{it} are persistent changes in the natural logs of price and disposable income, Δp^T_{it} and Δy^T_{it} are transitory changes in the natural logs of price and disposable income, $\Delta \mathbf{X}_{\mathbf{it}}$ is the change in a vector of demographic characteristics, and $\Delta \nu_{it}$ is an error term. Throughout the paper, all variables are measured in real constant dollars. The equation is first-differenced to remove an individual-specific fixed effect that may be correlated with the levels of the other explanatory variables. Conceptually, we want Δp^*_{it} and Δy^*_{it} to capture vertical shifts in an individual's whole expected future time path of prices and disposable income that were unanticipated before time *t*, like those illustrated by the shift from line a-b to line c-d and the shift from line e-f to line g-h in Figure 1. Ideally, if information becomes available at time *t* that such a shift in the lifetime profile of *p* and *y* will occur in the future (but will not yet be fully implemented time *t*), we want to incorporate the expected long-run vertical shift illustrated in Figure 1 into Δp^*_{it} and Δy^*_{it} immediately at time *t*, to allow for the possibility that taxpayers begin to respond immediately in anticipation of the future change. Δp^T_{it} and Δy^T_{it} then represent the change in actual log price and log disposable income between t-1 and t, minus the persistent change in expected log price and income that had been unanticipated before time t. So, for example, suppose that there is a change in price between time t-1 and time t, but there is no change in the expected future path of prices that had been unanticipated before time t. In that case, the full change in p represents a "transitory" change in price, in the sense that it arose either due to a fluctuation in income that is known by the taxpayer to be transitory, or in the sense that the change was already anticipated prior to time t. Either sort of price change represents an opportunity for intertemporal substitution. Similarly, if there is no change in current price between t-1 and time t, but there is a positive shock to the expected future path of prices at time t, there should be an equal negative shock to Δp^T_{it}, reflecting the fact that the current price has been made temporarily low relative to the now higher expectation of future prices. This would be the case, for example, in 1986, when the announcement of TRA86 caused the current price to be temporarily below the expected future price, creating an opportunity for intertemporal substitution. If some of the responsiveness to prices represents intertemporal substitution, then we should find that the coefficient on Δp^T_{it} is larger than the coefficient on Δp^*_{it}. Later in the paper, we will also consider a more general form of specification that allows lagged changes in giving, prices, and incomes to affect subsequent realizations of g.

Replicating the ASC approach in a regression framework

The ASC approach to estimating equation (1) can be expressed as a system of three equations:

$$(2)\ \Delta y_{it}^R = \xi_{it} + (\eta_{it} - \eta_{it-1}) = \Delta y_{it}^* + \Delta y_{it}^T$$

$$(3)\ \Delta p_{it}^R = (\omega_{it} + a_1\xi_{it}) + [\zeta_{it} - \zeta_{it-1} + a_2(\eta_{it} - \eta_{it-1})] = \Delta p_{it}^* + \Delta p_{it}^T$$

$$\begin{aligned}(4)\ \Delta g_{it}^R &= \beta_1(\omega_{it} + a_1\xi_{it}) + \beta_2[\zeta_{it} - \zeta_{it-1} + a_2(\eta_{it} - \eta_{it-1})] + \beta_3\xi_{it} + \beta_4(\eta_{it}-\eta_{it-1}) + \psi_{it} + \varepsilon_{it}-\varepsilon_{it-1} \\ &= \beta_1\Delta p_{it}^* + \beta_2\,\Delta p_{it}^T + \beta_3\,\Delta y_{it}^* + \beta_4\Delta y_{it}^T + \Delta\nu_{it}\end{aligned}$$

Δy_{it}^R, Δp_{it}^R, Δg_{it}^R are the residuals from regressing Δy_{it}, Δp_{it}, and Δg_{it} on time dummies and the demographic characteristics $\Delta \mathbf{X}_{it}$.[7] ξ_{it}, η_{it}, ω_{it}, ζ_{it}, ψ_{it}, ε_{it} are all zero-mean random shocks that are assumed to be independently, identically, and normally distributed. Each shock has its own variance, and each variance is assumed to be constant across individuals and time (although ASC also test the effects of allowing each variance to change over time). In equation (2), ξ_{it} is a persistent shock to log disposable income, and η_{it} is a transitory shock to log disposable income that dies off after one period. Thus, Δy_{it}^R, is decomposed into a persistent component Δy_{it}^*, plus a transitory component Δy_{it}^T, where $\Delta y_{it}^* = \xi_{it,}$ and $\Delta y_{it}^T = \eta_{it} - \eta_{it-1}$. In other words, the portion log disposable income that cannot not explained by year dummies or demographic / life cycle factors is assumed to follow a random walk, with a white noise transitory fluctuation around that random walk.

In equation (3), ω_{it} is a persistent shock to *p* that is independent of income, ζ_{it} is a transitory shock to *p* that that is independent of income and dies off after one period, and a_1 and a_2 are coefficients that represent the marginal effect on price of each shock to income. Thus, ASC assume that the total change in residual log price Δp_{it}^R consists of a persistent component Δp_{it}^*, which includes an independent persistent price shock ω_{it} and a linear function of the persistent income shock, $a_1\xi_{it}$, plus a transitory component Δp_{it}^T, which depends on independent transitory price shocks $\zeta_{it} - \zeta_{it-1}$ and a linear function of the transitory income shock, $a_2(\eta_{it} - \eta_{it-1})$. Equation (4) is just another representation of equation (1), where the effects of time dummies and the variables in $\Delta \mathbf{X}_{it}$ have already been partialed out of the remaining variables. The top row of equation (4) shows how the shocks specified in equations (2) and (3) can be substituted in for Δp_{it}^*, Δp_{it}^T, Δy_{it}^*, and Δy_{it}^T. Finally, ASC assume that the error term in the giving equation, $\Delta \nu_{it}$, can further be decomposed into a persistent shock ψ_{it}, and the first difference of a transitory shocks ε_{it}-ε_{it-1}.

The parameters a_1, a_2, β_1, β_2, β_3, and β_4 are estimated by ASC in a minimum-distance estimator framework. The values of these parameters, together with equations (2), (3), and (4), imply particular values for the variances and covariances of Δy_{it}^R, Δp_{it}^R,

$\Delta g_{it}{}^{R}$, $\Delta y_{it-1}{}^{R}$, $\Delta p_{it-1}{}^{R}$, and $\Delta g_{it-1}{}^{R}$. The minimum distance estimator chooses values of the parameters to minimize the difference between the variances and covariances predicted by the model, and those observed in the sample.

Instead of relying on the covariance structure estimation approach used in ASC, their procedure can alternatively be broken down in to a series of simple ordinary least squares regressions if one is able to decompose income and price into estimated transitory and permanent components at the individual level. It is possible to do so in a way that is consistent with the assumptions of the ASC model by applying a "filtering" technique from the macroeconometric literature. The filtering method we use is closely related to the Hodrick-Prescott (HP) filter, which is commonly used in the real business cycle literature, for example to distinguish the cyclical component of GDP from its long-run trend component. The HP filter assumes that a variable y_t follows an I(2) persistent trend (in other words, the trend follows a random walk), with a transitory component made up of pure white noise around that trend. In the HP filter, the permanent trend component y_t* is chosen to minimize

(5) $\sum (y_t - y_t^*)^2 + \lambda \sum (\Delta y_t^* - \Delta y_{t-1}^*)^2$,

where λ is a parameter that determines the smoothness of estimated path of Δy_t*. As the smoothing parameter λ gets larger, permanent income becomes smoother, approaching a constant linear trend. Hodrick and Prescott (1997, p. 4) note that if the transitory and permanent shocks are independently, identically, and normally distributed (as is assumed in ASC), then the optimal value for λ is the ratio of the variances of the transitory shock to the variance of the persistent shock. Intuitively, when the variance of transitory shocks is large and the variance of persistent shocks is small, then the true path of the persistent component will be fairly smooth, rarely diverging from its existing path to any great extent. In that case, if λ is appropriately set to equal the ratio of transitory to permanent variances, more weight will be put on minimizing the estimated squared permanent shocks, and the filter will produce a smoother time-series for the permanent component.

In order to be consistent with the ASC assumptions about the stochastic process for income, it is straightforward to modify the HP filter to suit an I(1) process for residual permanent income (in other words, the part of log permanent income that cannot be explained by other factors follows a random walk in levels), again with a white noise component for transitory fluctuations in income. To decompose the residual component of y_{it}^R into transitory and permanent components, the filter is applied to each individual time series separately, and y_{it}^* is chosen to minimize:

(6) $$\sum_t \left(y_{it}^R - y_{it}^*\right)^2 + \lambda\left(y_{it}^* - y_{it-1}^*\right)^2,$$

where λ is set equal to the ratio of the estimated variance of the transitory shock to the estimated variance of permanent shock, that is, $\sigma_\eta^2 / \sigma_\xi^2$. Since we do not know the true variances, we use estimates from the sample as a whole, which is appropriate for replicating the versions of the ASC estimator where variances are assumed to be constant across both individuals and time. ASC show that under the assumptions of their model, $Var(\Delta y_{it}^R) = \sigma_\xi^2 + 2\sigma_\eta^2$ and $Cov(\Delta y_{it}^R, \Delta y_{it-1}^R) = -\sigma_\eta^2$. Under those assumptions, we can obtain sample estimates of the variances needed to construct λ by calculating:

$\hat{\sigma}_\eta^2 = -Cov(\Delta y_{it}^R, \Delta y_{it-1}^R)$,

and

$\hat{\sigma}_\xi^2 = Var(\Delta y_{it}^R) - 2\hat{\sigma}_\eta^2$.

In all of our models, λ is estimated to be in the vicinity of 0.7. [8] Since this procedure involves precisely the same assumptions about the stochastic process for income and the structure of variances and covariances as in ASC, it achieves the same decomposition of income into transitory and permanent components as is implicitly occurring in the ASC approach.[9]

A four step procedure for replicating ASC in a regression framework is as follows:

STEP 1:

First, regress the levels of y_{it}, p_{it}, and g_{it} on year dummies, $\mathbf{X}_{it}$, and an individual-specific fixed effect.[10] The residuals from each of these regressions are y_{it}^R, p_{it}^R, and g_{it}^R, respectively.

STEP 2:

Next, we use our filter on each individual time series of y_{it}^{R} to estimate $\hat{\xi}_{it} = \hat{y}_{it}^{*} - \hat{y}_{it-1}^{*}$.

Once we have $\hat{\xi}_{it}$, we can construct an estimate of the first-difference of transitory shocks, $(\hat{\eta}_{it} - \hat{\eta}_{it-1}) = \Delta y_{it}^{R} - \hat{\xi}_{it}$.

STEP 3:

Again pool the data and use OLS to estimate:

(7) $\Delta p_{it}^{R} = a_1 \hat{\xi}_{it} + a_2(\hat{\eta}_{it} - \hat{\eta}_{it-1}) + u_{it}$

This is just equation (3) above with the estimated permanent and transitory income shocks substituted for their true values. The residual from this regression, u_{it}, equals $\omega_t + \zeta_t - \zeta_{t-1}$. We then estimate ω_t and $(\zeta_t - \zeta_{t-1})$ with the filter, using the same approach as for income above (it is necessary, but straightforward, to integrate the u_{it}'s for an individual into a series of levels in order to apply the filter).

STEP 4:

Finally, estimate equation (4) by linear regression. Because of step 2 and step 3, we now know everything on the right hand side of equation (4) except for β_1, β_2, β_3, and β_4, which are the coefficients (also elasticities) that we are estimating now via linear regression, and the error term in the giving equation. Our replication procedure thus far does not address the structure of the error term in the giving equation assumed by ACS, but does compute standard errors that are robust to arbitrary forms of autocorrelation (and heteroskedasticity).

Incorporating a richer specification of tax law

Our new estimation procedure builds upon the ASC specification of a stochastic process for income, but differs from their approach in several key ways. We assume that it is Y_{it}^{R}, the natural log of *pre-tax* income (after partialing out the effects of variables in $\mathbf{X}_{it}$, a time-trend, and an individual specific fixed effect) that follows the stochastic process specified in ASC for log *after-tax* income y_{it}.[11] In other words, we now specify that $\Delta Y_{it}^{R} = \xi_{it} + (\eta_{it} - \eta_{it-1})$. We then construct estimates of transitory and persistent changes in y_{it} and p_{it} by combining the estimated information about the transitory and permanent components of pre-tax income with detailed information on actual tax laws known at time *t* and time *t*-1. A particularly important advantage of our approach here is that we can address the issue of pre-announced and phased-in changes in tax law by using the fully-phased in version of currently announced steady-state future tax law to construct estimates of expected future price and disposable income.

Another advantage of our approach is that it allows shocks to pre-tax income to affect tax liability and tax rates in different ways depending on the specifics of the actual tax law applying to that individual at that time. By estimating a_1 and a_2 as fixed coefficients, ASC implicitly impose the restriction that a shock to income of a given size always has the same marginal effect on price across all individuals in all time periods. In fact, the marginal effect of an income shock on price can vary greatly across individuals and time, as the relationship is determined by the complicated non-linear tax bracket structure, which varies over time (due to tax reforms) and across states and filing statuses, and also depends on where one's taxable income is located in the tax bracket structure, which varies across individuals.

An additional important consideration arises because ASC find that in models where the variance of shocks is assumed to be constant over time, the estimated persistent price elasticity is quite sensitive to the sample period chosen. They suggest that the instability of the parameter estimates arises because of the assumption of constant variances over time – in particular, they argue that tax reforms during the 1980s greatly reduced the variance of shocks to price, by compressing the marginal rate distribution. When they relax this assumption and allow variances to change over time, the price elasticity estimates become more stable over different sample periods, and uniformly large, while other elasticities are not much affected. Although we do not yet allow non-stationary variances in our procedure, we address the concern motivating the non-stationarity assumption more directly, because in our approach shocks to income are fed through the actual tax rate schedules themselves in order to determine their effects on price. Thus, the effects of compressing the rate distribution are accounted for directly and precisely, rather than assuming that this effect can be approximated through a changing variance structure.

Our specific procedure for estimating equation (1) above can be broken down into four steps.

STEP 1

Regress log pre-tax income on an individual-specific fixed effect (α_i), a time trend (t), and the vector of demographic and life cycle characteristics ($\mathbf{X}_{it}$). The first stage regression equation is $Y_{it} = \gamma \mathbf{A}_{it}$, where $\mathbf{A}_{it} = [\alpha_i \mid t \mid \mathbf{X}_{it}]$. The residual from this regression is $\hat{Y}_{it}^R$.

STEP 2

Run each individual time-series of $\hat{Y}_{it}^R$ through the filter in order to estimate $\hat{Y}_{it}^*$ and $\hat{\xi}_{it}$.

STEP 3

Use the information above, in conjunction with information about tax law known at time t and time t-1, to construct estimates of the transitory and persistent changes in price and disposable income at time t. We will label these constructed estimates $\Delta\hat{p}_{it}^*$, $\Delta\hat{p}_{it}^T$, $\Delta\hat{y}_{it}^*$ and $\Delta\hat{y}_{it}^T$. It will be useful to define $y_t(Y_{it},\mathbf{Z}_{it})$ and $p_t(Y_{it},\mathbf{Z}_{it})$ as general functions that map the log of pre-tax income Y_{it}, and a vector of other individual characteristics $\mathbf{Z}_{it}$, into log after-tax income and log price. The vector of other characteristics entering tax calculations, $\mathbf{Z}_{it,}$ is defined to include tax-relevant demographic characteristics such as age, marital status, number of dependents, and state of residence, as well as information on the allocation of pre-tax income across various components of income and deductions, expressed as shares of pre-tax income. An "*S*" superscript indicates a "steady state" function, that is, the function that is expected to apply after currently announced tax law is fully phased in, based on information available at the relevant date. So, for example, $y_t^S(.)$ and $p_t^S(.)$ represent the functions that would map Y_{it} and $\mathbf{Z}_{it}$ into log after-tax income and log price based on steady state tax law that is known as of time t. Similarly, $y_{t-1}^S(.)$ and $p_{t-1}^S(.)$ represent the equivalent functions based on steady state tax law known as of time *t*-1. For federal taxes, we assume that in years when a new tax law is enacted or when a recently enacted law is being gradually phased-in, known steady-state tax law is the fully-phased in version of that newly or recently enacted law. In other years, steady state federal law is assumed to be currently applicable law. In the absence of detailed information on the legislative histories of various state tax reforms, we assume that for state taxes, steady state law known at time *t* is the law that applies in year t+1.

Our goal in constructing the estimates of persistent changes in price and disposable income $\Delta\hat{p}_{it}^{*}$ and $\Delta\hat{y}_{it}^{*}$, is to isolate long-run vertical shifts in the whole expected future profile of prices and disposable income that become known at time *t* and had not been anticipated before then, as illustrated in Figure 1. In our model this sort of shift arises for two reasons: an unanticipated persistent shock to income ξ_{it}; and a tax reform that becomes known at time *t* and which changes the expected steady-state future tax function in a persistent manner. Note that ξ_{it} affects disposable income directly, and also indirectly through its influence on tax liability. It also affects price to the extent that it pushes the taxpayer into a different marginal rate bracket. This is all subsumed in the functions that map *Y* and **Z** into *p* and *y*. In order to isolate these two sources of previously unanticipated persistent shocks, we construct estimates of $\Delta\hat{p}_{it}^{*}$ and $\Delta\hat{y}_{it}^{*}$ that change only because of $\hat{\xi}_{it}$ and any changes in the p_t^S and y_t^S functions that become known between time *t*-1 and time *t*, holding all other characteristics constant at their *t*-1 levels.[12] Thus we construct:

(8) $$\Delta\hat{y}_{it}^{*} = y_t^S(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1} + \xi_{it}, \mathbf{Z}_{it-1}) - y_{t-1}^S(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})$$

(9) $$\Delta\hat{p}_{it}^{*} = p_t^S(\hat{Y}_{it}^{*} + \hat{\gamma}\mathbf{A}_{it-1} + \xi_{it}, \mathbf{Z}_{it-1}) - p_{t-1}^S(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})$$

Transitory changes in price and income between time *t* and time *t*-1 are just the actual changes minus the persistent changes.

(10) $$\Delta\hat{y}_{it}^{T} = \Delta y_{it} - \Delta\hat{y}_{it}^{*}$$

(11) $\Delta\hat{p}_{it}^{T} = \Delta p_{it} - \Delta\hat{p}_{it}^{*}$

These transitory components include changes in price caused by transitory fluctuations in income, predictable changes caused by life cycle and demographic factors, and changes in the difference between current and expected future prices and incomes caused by pre-announced changes in tax law. As an example of the latter, suppose that a persistent 10 percent increase in price that is implemented beginning at time *t*+1 is announced at time *t*, and there is no change in current price between *t*-1 and t. In that case, there will be a positive shock of approximately 0.1 to $\Delta\hat{p}_{it}^{*}$, and a corresponding shock of -0.1 to $\Delta\hat{p}_{it}^{T}$. The effects of the two shocks can be distinguished because in the next period, $\Delta\hat{p}_{it}^{*}$ will be zero and $\Delta\hat{p}_{it}^{T}$ will be +0.1. Intuitively, the portion of the change in giving at time t that reverses itself at time t+1 contributes to the estimate of the effect of the transitory price change, while the portion of the change at time *t* that does not reverse itself at *t*+1 contributes to the estimate of the effect of the persistent price change.

Changes in price or income that are driven by predictable life cycle and demographic factors are largely controlled for in any event through $\Delta\mathbf{X}_{it}$, but these factors may induce some residual independent variation in *p* and *y* because the latter are complicated non-linear functions of many of the components of $\Delta\mathbf{X}_{it}$. To the extent that the influences of these factors are not absorbed by controlling for $\Delta\mathbf{X}_{it}$, we want them to be relegated to the transitory components of price and income, because they are likely to be a source of predictable differences in price across time, and thus the response to them may reflect intertemporal substitution. This procedure achieves that.

STEP 4:

Substitute (8) through (11) into equation (1) in order to estimate:

(12) $\Delta g_{it} = \alpha_t + \beta_1\Delta\hat{p}_{it}^{*} + \beta_2\Delta\hat{p}_{it}^{T} + \beta_3\Delta\hat{y}_{it}^{*} + \beta_4\Delta\hat{y}_{it}^{T} + \boldsymbol{\beta}_\mathbf{5}\Delta\mathbf{X}_{it} + \Delta\nu_{it}$

Instrumental variables

While the estimation procedure outlined above has many advantages, one potential concern is that, although our estimated decomposition of persistent and transitory components should be correct on average if our assumptions about the stochastic process for income are correct, there will inevitably be errors in the decomposition for particular individuals and years. This is difficult to avoid when the variables we are interested in are unobservable. One method to address any potential biases that might occur due to this measurement error problem is to use instrumental variables that rely entirely on tax reforms for identification. Tax reforms are likely to be exogenous from the perspective of the individual, and their effects on price and income can be decomposed into transitory and permanent components with greater certainty than can pre-tax income. That is because the changes in current and future tax law caused by tax reforms were published and known with relative certainty, while the portion of any change in pre-tax income that is transitory or permanent cannot be known with certainty. In this approach, we treat $\Delta\hat{p}_{it}^{*}$, $\Delta\hat{p}_{it}^{T}$, $\Delta\hat{y}_{it}^{*}$ and $\Delta\hat{y}_{it}^{T}$ as measurements with error of their true values. The excluded instruments we use here are:

$$(13)\quad \Delta pIV_{it}^{*} = p_{t}^{S}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1}) - p_{t-1}^{S}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})$$

$$(14)\quad \Delta pIV_{it}^{T} = \left[p_{t}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1}) - p_{t-1}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})\right] - \Delta pIV_{it}^{*}$$

$$(15)\quad \Delta yIV_{it}^{*} = y_{t}^{S}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1}) - y_{t-1}^{S}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})$$

$$(16)\quad \Delta yIV_{it}^{T} = \left[y_{t}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1}) - y_{t-1}(\hat{Y}_{it-1}^{*} + \hat{\gamma}\mathbf{A}_{it-1}, \mathbf{Z}_{it-1})\right] - \Delta yIV_{it}^{*}$$

Here the functions $p_t(\cdot,\cdot)$ and $y_t(\cdot,\cdot)$ use current tax law, rather than the fully-phased in tax law. These instruments isolate the portion of changes in transitory and permanent price and income that are caused by tax reforms, holding life-cycle and persistent components of income and the characteristics in **Z** constant at their t-1 levels. Note that equations

(13) and (15) will only be non-zero in years when tax reforms are announced, and equations (14) and (16) will only be non-zero in years when current tax law differs in a systematic way from known future tax law due to a pre-announced or gradually phased-in tax reform. As an example, in 1986, the instrument for transitory price change will be the price calculated using the 1988 tax function minus the price calculated using the 1986 tax function, both evaluated at time *t*-1 levels of persistent and life-cycle income and characteristics in $\mathbf{Z}_{it}$.

Data for preliminary empirical analysis

Our preliminary empirical work on this subject utilizes the 1979-90 University of Michigan public-use panel of individual tax returns. The panel consists of a random sample of tax returns selected based on the last four digits of the social security number. We follow ASC closely in sample selection and variable definition. The sample includes those who were exogenous itemizers (would have itemized even with zero charitable giving), and had income above the filing threshold, for at least five consecutive years without a change in marital status. We exclude members of married couples who file their federal income taxes separately (a very rare occurrence) and returns filed for a prior tax year. Our definition of income, which follows Auten Cilke and Randolph (1992), begins with adjusted gross income, and makes a number of modifications to keep the measure consistent across time, most notably by adding back in any excluded capital gains realizations. The resulting sample has 26,012 observations, or 22,844 after dropping the first year of data for each individual due to first-differencing.

Variables in $\mathbf{X}_{it}$ in the public use data include number of dependents, and whether the taxpayer or spouse are aged 65 or over, whether the return reported wage and salary income, whether the return reported self-employment income, whether the taxpayer has any children.[13] In addition, each of these dummies, as well as a dummy for marital status, are interacted with time trends. Thus, most demographic variables are allowed to cause a discrete shift in the dependent variable when they change, and are allowed to affect the trend in the dependent variable as well. The dependent variable is the natural logarithm of charitable donations plus $10 (to avoid taking the log of zero).

All tax liabilities and tax rates are computed using a very detailed federal and state income tax calculator which is described in Bakija (2004).[14] Following ASC, tax variables are calculated setting charitable giving equal to 1 percent of pre-tax income, to avoid endogeneity (we also tried using this as an instrument for the values calculated using actual charitable giving, and found that it made little difference to the results). Marginal tax rates are then calculated based on a $0.10 increase in charitable donation above that level.[15] As is the convention in the literature, we define the price so that it includes the extra tax savings that occur through the avoidance of capital gains tax on donations of appreciated property. We also take into account that between 1987 and 1992 such capital gains would be taxed by the alternative minimum tax should the taxpayer become subject to it. We define:

$p_{it} = ln\{1\text{-}mtr_t\text{-}0.5*(ncshare)*[mtrcg_{it}^S - mtra_{it}]\}$,

where mtr_t is the marginal income tax rate at which charity is deducted, $mtrcg_{it}^S$ is the steady-state marginal income tax rate on capital gains realizations, $mtra_{it}$ is marginal rate of alternative minimum tax on capital gains on gifts of appreciated assets (equal zero if not subject to AMT), and *ncshare* is the average ratio of non-cash charitable gifts to total charitable gifts from 1979-88 for all taxpayers in the real income class of which this particular taxpayer is a member, from Auten, Cilke, and Randoph (1992, pp. 282-283). The 0.5 is the rule-of thumb ratio of the present value of capital gains realizations to accrued capital gains, times the gain-to-value ratio for gifts of appreciated assets, estimated by Feldstein (1975) and Feldstein and Clotfelter (1976), which has been used throughout the literature. Note that the expected future steady marginal capital gains tax rate ($mtrcg_{it}^S$) is the relevant capital gains tax rate here, as the choice is between giving an appreciated asset to charity today, or realizing the gain on the asset at some unspecified point in the future.

Preliminary results

Table 1 shows preliminary results from estimating each of the models discussed above, as well as some more traditional models for purposes of comparison. Row (1) depicts estimates from a conventional (pooled) cross-sectional regression, where no attempt is made to distinguish transitory from permanent variation in price or income nor to deal with unobserved heterogeneity. The price elasticity is -1.22 and the income elasticity is 0.84 – both are quite typical of results found in the traditional cross-sectional literature. Row (2) takes the same approach, but now first differences the data to control for unobserved heterogeneity. Both elasticities are substantially reduced, to -0.45 for price and 0.28 for income. This is also broadly consistent with the findings of previous studies that analyzed early years of this same data set, e.g., Broman (1989) and Barrett, McGuirk, and Steinberg (1997).

Rows (3) and (4) show elasticities reported in ASC for the two specifications and sample periods that are closest to our replication. As noted above, their persistent price elasticity estimates are sensitive to the sample period when the variances of the shocks are constrained to be constant over time. Their transitory price elasticity estimates are consistently low, and substantially lower than their persistent price elasticity estimates. Row (5) shows the results of our replication of their approach using a multi-step method on the public-use data. The basic pattern of results in the replication is reasonably similar to what ASC find, considering the differences in sample period and composition of sample. In particular, the replication reproduces the odd result that the persistent price elasticity (-0.64) is larger than the transitory price elasticity (-0.34).

Row (6) shows the results of implementing our new estimation procedure (from equation 12 above), which combines the ASC stochastic process with income with a more precise modeling of tax law information, particularly with regards to pre-announced and phased-in changes. We find that our approach undoes the counterintuitive pattern of price elasticities found by ASC, as the persistent price elasticity falls from -0.64 to -0.24, and the transitory price elasticity rises slightly, to -0.4. This is consistent with our prior expectation that failure to model pre-announced changes in tax law, particularly with respect to TRA86, would bias the persistent price elasticity upwards. Because the transitiory elasticity is now larger than the persistent price elasticity, our results suggest that there is some degree of intertemporal substitution of charitable giving, although the responsiveness still appears to be rather small. We suspect that intertemporal substitution may not be an especially large issue in the public use data set, which largely consists of middle class taxpayers, but is likely to be a bigger issue among the high income taxpayers in the ASC data set, as evidenced by the large 1986 spike in Figure 1. We also find that our new procedure increases the estimated elasticity of giving with respect to persistent changes in income to 0.72, and that the transitory income elasticity is essentially zero. If people behave according to the permanent income hypothesis, this would suggest that our procedure is doing a relatively good job of separating out persistent and transitory components of income.

Row (7) shows the results of estimating the instrumental variables approach described above on the public-use data. The main changes are that the transitory price elasticity becomes a bit larger at -0.56, and the persistent income elasticity falls to 0.27. The public use data set's small sample size, particularly with regards to high income people who experienced the most interesting changes in tax law during this period, limits what we can learn from this exercise, as the standard errors are quite large. Nonetheless, the fact that the point estimate the persistent price elasticity barely changes under the IV approach is at least somewhat reassuring that measurement error is not driving our results.

Dynamic model

A next step is to make the model more dynamic, allowing lagged values of giving, price, and income variables to enter the equation. That allows for slow adjustment by taxpayers, which would mean that the long-run elasticities with respect to persistent price and income changes might be larger than the immediate effects we estimated above. In addition, if intertemporal substitution is an important part of the story, then including lagged dependent and explanatory variables would be appropriate. For instance, if there is intertemporal substitution of giving into 1986 in response to the expected future increase in price caused by TRA86, one would expect this to reduce giving in future years relative to what it otherwise would have been, ceteris paribus. It makes sense to include lags of the price and income variables as well as lagged giving because the effect of a change in lagged giving should depend on the reason it occurred. For instance, if giving was large last year because of intertemporal substitution from this year to last in response to a temporarily low price, then we would expect that surge in giving to increase giving this year. On the other hand, if the increase in giving last period was because it was the first year of a persistent reduction in price, and there is slow adjustment to the new permanent price, then we might expect the increase last period to be associated with an increase this period.

Our more dynamic estimation equation is:

$$(17)\quad \begin{aligned}\Delta g_{it} &= \delta_{0t} + \delta_1 \Delta \hat{p}^*_{it} + \delta_2 \Delta \hat{p}^T_{it} + \delta_3 \Delta \hat{y}^*_{it} + \delta_4 \Delta \hat{y}^T_{it} \\ &+ \delta_5 \Delta \hat{p}^*_{it-1} + \delta_6 \Delta \hat{p}^T_{it-1} + \delta_7 \Delta \hat{y}^*_{it-1} + \delta_8 \Delta \hat{y}^T_{it-1} + \delta_9 \Delta g_{it-1} + \boldsymbol{\delta}_{10} \Delta \mathbf{X}_{it} + \Delta \nu_{it}\end{aligned}$$

In this specification, the long-run elasticity of giving with respect to a permanent change in price is $(\delta_1 + \delta_5)/(1-\delta_9)$, and the long-run elasticity of giving with respect to a permanent change in disposable income is $(\delta_3 + \delta_7)/(1-\delta_9)$. The effects of one-period shocks to the transitory variables converge to zero in the long-run.[16]

It is well known that in the first-differenced equation (17) will suffer from endogeneity bias, because Δv_{it} depends in part on v_{it}, which is correlated with g_{it-1}. This would tend to bias our estimate of δ_9 downward, Additional endogeneity problems arise with a lagged dependent variable if there is serial correlation in v_{it}. In order to address these problems in an efficient manner, we use the instrumental variables approach developed by Arellano and Bond (1991), which is sometimes referred to as "difference GMM."[17] The Arellano Bond test for autocorrelation finds evidence of first- and second-order autocorrelation, but rejects higher orders of autocorrelation. Thus all available lagged levels g_{it} from periods t-3 and earlier are used as instruments for Δg_{it-1}. In addition, based on equation (17) above, lagged values of $\Delta \mathbf{X}_{it}$ from periods t-1 to t-3, as well as lagged values of $\Delta \hat{p}^*_{it}$, $\Delta \hat{p}^T_{it}$, $\Delta \hat{y}^*_{it}$, $\Delta \hat{y}^T_{it}$ from t-2 and t-3 should be correlated with Δg_{it-1}, either directly or through their correlation with Δg_{it-2}, so these are used as instruments for

Δg_{it-1} in any periods when they are available. We use the two-step version of this estimator, which uses initial consistent estimates of the error term from equation (17) above ($\hat{v}_{it}$) to construct an optimal weighting matrix for GMM, which improves efficiency, and we compute standard errors that are robust to arbitrary forms of heteroskedasticity and autocorrelation. The finite-sample bias in the two-step standard errors is corrected using the procedure outlined in Windmeijer (2000). We also try an instrumental variables approach, where instruments depending entirely on tax reforms for identification, as outlined in equations (13) through (16) above, are used for current and lagged values of the price and disposable income variables.

Results from the dynamic Arellano Bond estimator are shown in Table 2. The top panel depicts the conventional estimates (which do not instrument for the price and income variables). The estimates support the notion that there is some degree of slow adjustment and/or habit formation in the data, as the elasticity of current giving with respect to one-period lagged giving is found to be 0.36 and highly statistically significant. The long-run elasticity of giving with respect to a persistent shock to price is estimated to be -0.75, and the long run persistent income elasticity is 0.56. The transitory price elasticity continues to be somewhat higher than the immediate elasticity of giving with respect to a persistent price change, but is now lower than the long-run effect. Results from the version which instruments for price and disposable income changes are quite similar in the long run effects and in the short run price elasticities, except for a somewhat higher transitory price elasticity.

Conclusion

In this paper we develop a new estimation procedure for analyzing the effect of the income tax on charitable giving. It combines a comparatively unrestrictive stochastic process for income with a method for addressing anticipated changes in tax law arising from pre-announcement. We find that taking into account the pre-announced and phased-in nature of tax reforms has an important effect on the results, reducing the persistent price elasticity substantially and raising the transitory price elasticity slightly, so that the latter is again larger than the former. At the same time, we do not find particularly strong evidence of intertemporal substitution of charitable giving in response to anticipated price variation, at least in a small public use data set consistng largely of middle-class households. We then estimate an Arellano-Bond dynamic panel data model that allows for gradual adjustment of giving and habit formation, and we find that the long-run responses to persistent price changes are larger than the immediate responses. The net result is a long-run persistent price elasticity in the vicinity of -0.75.

Table 1 – Estimated price and income elasticities of charitable giving

	Persistent price elasticity	Transitory price elasticity	Persistent income elasticity	Transitory income elasticity
(1) Pooled cross-section, no separation of transitory and permanent	-1.22 (0.16)		0.84 (0.06)	
(2) First differenced, no separation of transitory and permanent	-0.45 (0.10)		0.28 (0.04)	
(3) Results reported in ASC, 1980-87, assuming time-invariant variances	*-2.13* *(1.17)*	*-0.14* *(0.20)*	*0.74* *(0.11)*	*0.29* *(0.02)*
(4) Results reported in ASC, 1980-92, assuming time-invariant variances	*-0.31* *(0.15)*	*-0.02* *(0.10)*	*0.91* *(0.02)*	*0.30* *(0.01)*
(5) Our replication of ASC on public-use data	-0.64 (0.30)	-0.34 (0.17)	0.55 (0.10)	0.12 (0.06)
(6) Estimation of equation (12) on public use data	-0.24 (0.13)	-0.40 (0.11)	0.72 (0.08)	0.02 (0.07)
(7) Instrumental variables relying solely on tax reforms for identification	-0.28 (0.29)	-0.56 (0.32)	0.27 (0.24)	-0.06 (0.76)

Standard errors robust to arbitrary forms of autocorrelation and heteroskedasticity reported in parentheses. Results in italics are directly from Auten, Sieg, and Clotfelter (2002), presented here for purposes of comparison.

Table 2 – Elasticity estimates from Arellano Bond dynamic specification (equation 17)

	Persistent price $\Delta\hat{p}_i^*$	Transitory price $\Delta\hat{p}_i^T$	Persistent income $\Delta\hat{y}_i^*$	Transitory income $\Delta\hat{y}_i^T$	Lagged giving Δg_i
Current	-0.41 (0.14)	-0.50 (0.11)	0.40 (0.11)	0.14 (0.08)	
Lagged	-0.07 (0.12)	-0.08 (0.09)	-0.07 (0.11)	-0.09 (0.08)	0.36 (0.05)
Long-run elasticity	-0.75 (0.21)		0.52 (0.10)		
Instrumental variables relying solely on tax reforms for identification					
Current	-0.48 (0.25)	-0.66 (0.27)	0.04 (0.26)	0.11 (0.25)	0.35 (0.05)
Lagged	-0.00 (0.20)	-0.10 (0.17)	0.32 (0.21)	-0.14 (0.16)	
Long-run elasticity	-0.74 (0.32)		.56 (0.31)		

Standard errors robust to arbitrary forms of autocorrelation and heteroskedasticity, with Windmeijer finite sample bias correction, reported in parentheses. Standard errors for the long-run elasticities, which are non-linear functions of coefficients, are approximated via the delta method.

Figure 1 – Stylized illustration of transitory, permanent, and life-cycle variation in income and prices

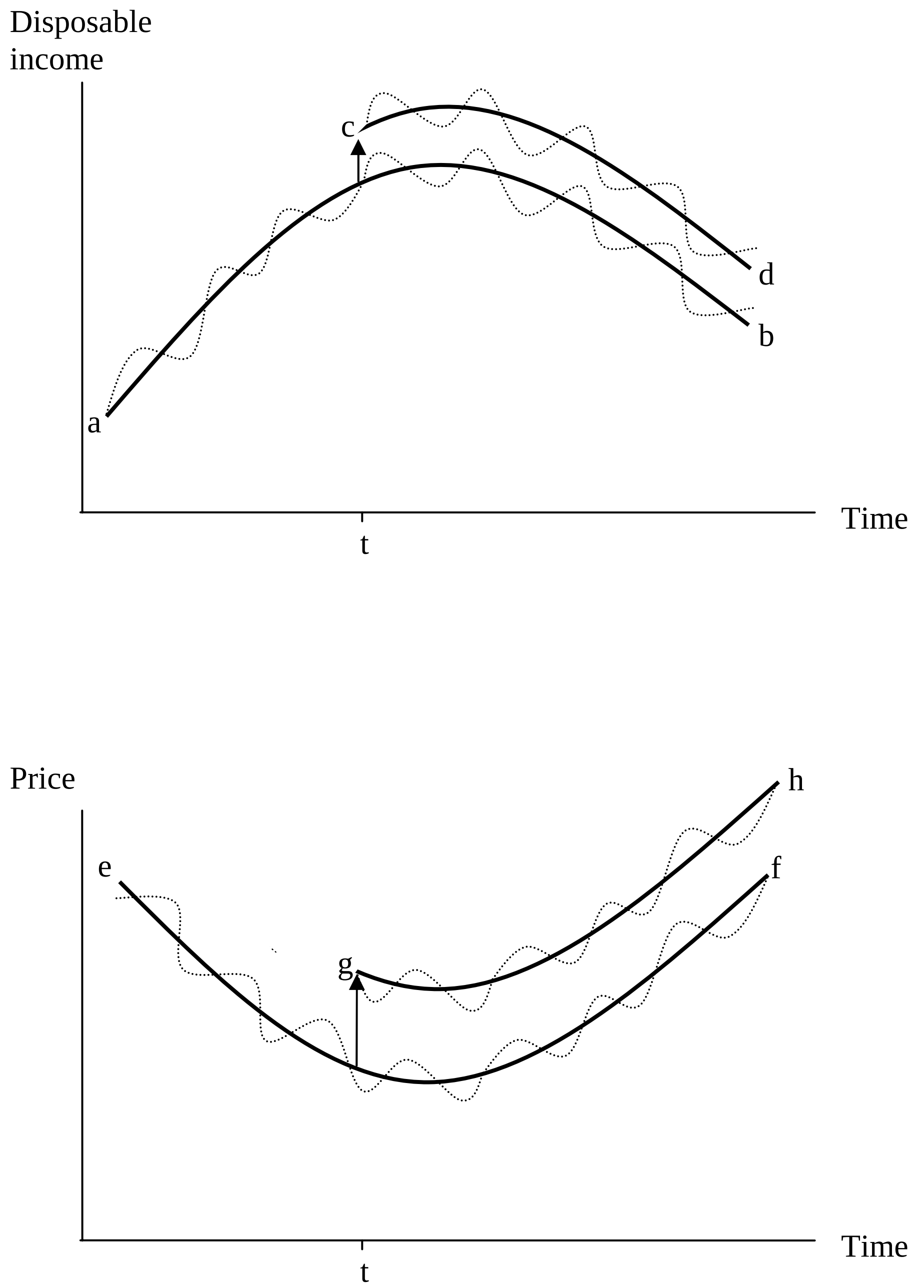

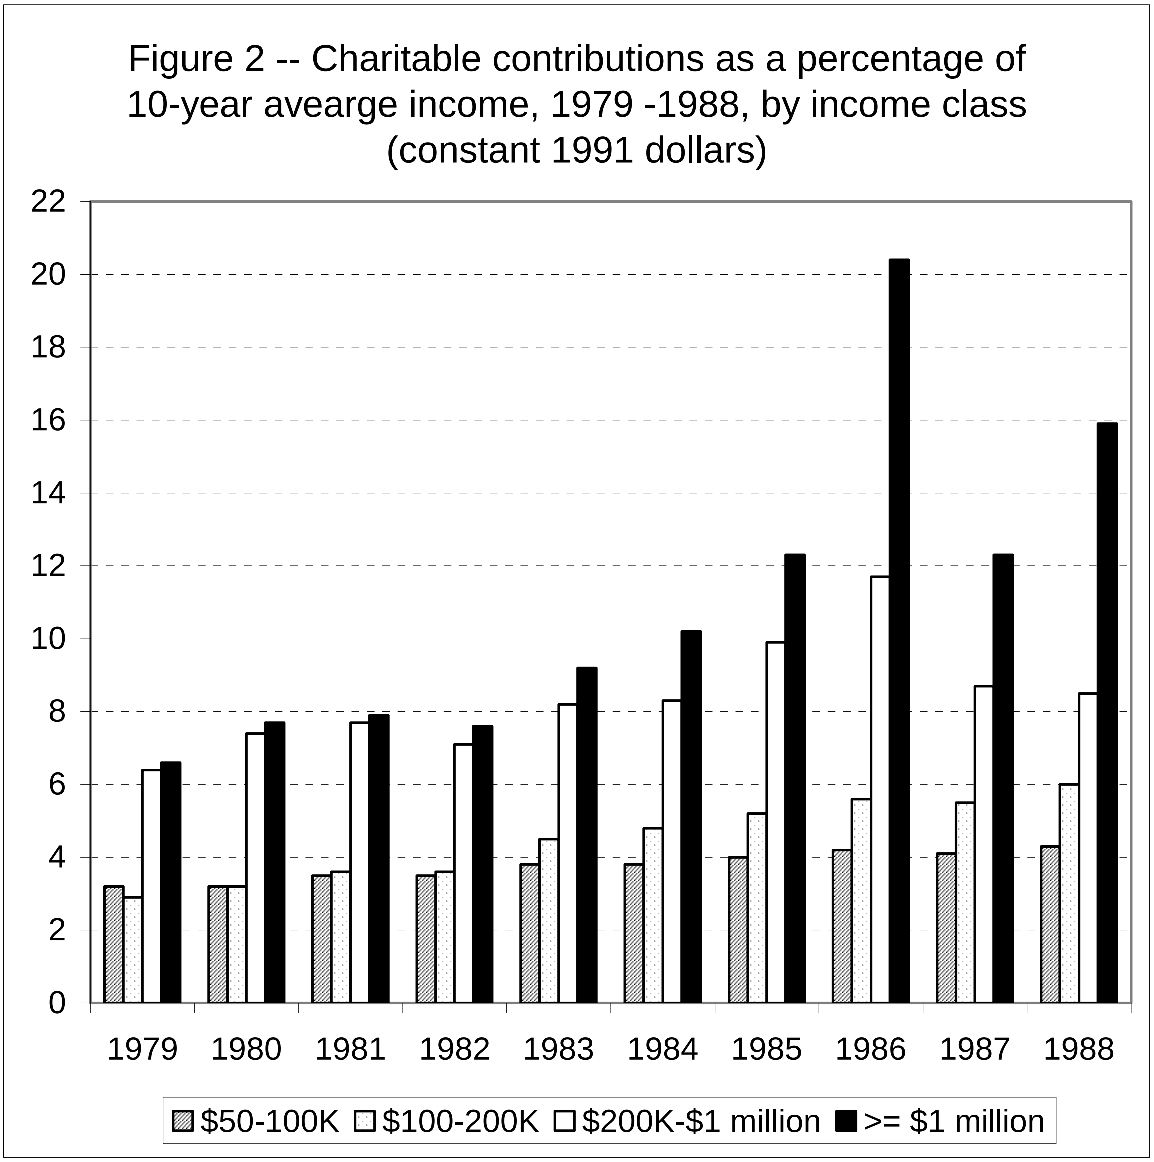

Source: Auten, Cilke, and Randolph (1992)

References

Andreoni, J., 1990. "Impure Altruism and Donation to Public Goods: A Theory Of Warm Glow Giving." *Economic Journal* 100, 464–477.

Aprill, Ellen P. 2001. "Churches, Politics, and the Charitable Contribution Deduction." *Boston College Law Review*. Vol. 42: 843-874.

Arellano, M and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economics Studies*. Vol. 58: 277-297.

Auten, G.E., Cilke, J.M., Randolph, W.C. 1992. "The Effects of Tax Reform on Charitable Contributions." *National Tax Journal*. 45 (3), 267-290.

Auten, G.E., Sieg, H., Clotfelter, C.T. 2002. "The Distribution of Charitable Giving, Income and Taxes: An Analysis of Panel Data." *American Economic Review*, Vol. 92, No. 1 (March), 371-382.

Auten, G.E., Sieg, H., Clotfelter, C.T. 2000. "Apppendix of the Paper: The Distribution of Charitable Giving, Income and Taxes: An Analysis of Panel Data." February 16.

Bakija, Jon. 2004. "Documentation for IncTaxCalc: A Federal-State Personal Income Tax Calculator Covering the Years 1900-2002 " Working paper, Williams College. <http://wso.williams.edu/~jbakija/Documentation_IncTaxCalc.pdf>

Bakija, Jon. 2002. "Distinguishing Transitory and Permanent Price Elasticities of Charitable Giving with Pre-Announced Changes in Tax Law." Working Paper, Williams College. <http://wso.williams.edu/~jbakija/charity_jb2.pdf>

Barrett, Kevin S. "Panel Data Estimates for Charitable Giving: A Synthesis of Techniques." *National Tax Journal*. Vol. 44, No. 3 (September): 365-81.

Barrett, K.S., McGuirk, A.M., Steinberg, R. 1997. "Further Evidence on the Dynamic Impact of Taxes on Charitable Giving." *National Tax Journal*. 50 (2), 321-334.

Birnbaum, Jeffrey H. and Alan S. Murray. 1987. *Showdown at Gucci Gulch: Lawmakers, Lobbyists, and the Unlikely Triumph of Tax Reform*. New York: Random House.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics*. Vol. 87, 115-143.

Blundell, Richard, Stephen Bond, and F Windmeijer. 2000. "Estimation in Dynamic Panel Data Models: Improving on the Performance of the Standard GMM Estimator." in B Baltagi, ed., *Advances in Econometrics, Vol. 15: Nonstationary Panels, Panel Cointegration, and Dynamic Panels*. Amsterdam: JAI Elsevier.

Boone, Laurence. 2000. "Comparing Semi-Structural Methods to Estimate Unobserved Variables: The HPMV and Kalman Filter Approaches." OECD Working Paper No. ECO/WKP(2000)13.

Bond, Stephen R. 2002. "Dynamic Panel Data Models: A Guide to Micro Data Methods and Practice." *Portugese Economic Journal*. Vol. 1, 141-162.

Broman, A.J. 1989. "Statutory Tax Rate Reform and Charitable Contributions: Evidence from a Recent Period of Reform." *Journal of the American Taxation Association*. (Fall), 7-21.

Brown, Eleanor. 1997. "Taxes and Charitable Giving: Is There a New Conventional Wisdom?" *Proceedings of the Eighty-Ninth Annual Conference on Taxation held under the auspices of the National Tax Association, November 10-12*, 1996. pp. 153-59

Clotfelter, Charles. *Federal Tax Policy and Charitable Giving*. Chicago: University of Chicago Press, 1985

Deaton, A., Muellbauer, J. 1980. "An Almost Ideal Demand System." *American Economic Review*. 70 (3), 323-326.

Feenberg, D. 1987. "Are Tax Price Models Really Identified: The Case of Charitable Giving." *National Tax Journal*. 40 (4), 629-633.

Feldstein, M. 1975. "The Income Tax and Charitable Contributions: Part I -- Aggregate and Distributional Effects." *National Tax Journal*. 28 (1), 81-100.

Feldstein, M., Clotfelter, C.T. 1976. "Tax Incentives and Charitable Contributions in the United States: A Microeconometric Analysis." *Journal of Public Economics*. 5 (1), 1-26.

Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press (for NBER).

Goolsbee, Austan. 2000. "What Happens When You Tax the Rich? Evidence from Executive Compensation." *Journal of Political Economy*. Vol. 108, No. 2.

Hall, Robert E., and Alvin Rabushka. 1995. *The Flat Tax*. 2nd ed. Stanford, CA: Hoover Institution Press.

Hodrick, R.J., and E.C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit, and Banking*, Vol. 29, No. 1, pp. 1-16.

Randolph, W.C. 1995. "Dynamic Income, Progressive Taxes, and the Timing of Charitable Contributions." *Journal of Political Economy*. 103 (4), 709-738.

Roodman, Malin. 2004. "XTABOND2: Stata Module to extend XTABOND Dynamic Panel Data Estimator." <http://econpapers.hhs.se/software/bocbocode/s435901.htm>

Saez, Emmanuel. 2004. "The Optimal Treatment of Tax Expenditures." *Journal of Public Economics*, Vol. 88, pp. 2657– 2684.

Sandmo, A., 1975. "Optimal Taxation in the Presence of Externalities." *Swedish Journal of Economics*. 77, 86– 98.

Simpson, Janice. 1986. "Tactics for Giving: Tax Law Spurs Interest in Charitable Contributions." *Wall Street Journal*, Dec. 4, p. 1.

Windmeijer F. 2000. "A Finite Sample Corrections for the Variance of Linear Two-Step GMM Estimators." Institute for Fiscal Studies Working Paper Series No. W00/19, London.

Endnotes

[1] See Aprill (2001)

[2] The original "Treasury I" plan that led to the Tax Reform Act of 1986 would have disallowed deductions for charitable contributions that were below 2 percent of income (Birnbaum and Murray, 1987, p. 90). The Hall and Rabushka (1995) call for the elimination of charitable deductions in their "flat tax" plan.

[3] See Auten, Cilke, and Randolph (1992) for a thorough description of the changes in tax law that occurred between 1981 and 1991, and how they affected incentives to give to charity.

[4] See Brown (1997) for a recent literature review.

[5] In the expenditure share equation approach used by Randolph, the elasticity of giving varies depending on the size of the expenditure share. The most commonly cited results for Randolph evaluate the elasticity at the giving-weighted mean expenditure share of 0.085, which produces a permanent price elasticity of -0.58 and a transitory price elasticity of -1.55. However, for reasons discussed in detail in Bakija (2002), the elasticity evaluated at the mean expenditure share (0.039) is likely to be more representative of the elasticity of aggregate giving with respect to price in the sample.

[6] The long-run effect is the immediate effect divided by $(1-\gamma)$, where γ is the coefficient on lagged giving. A small value of γ thus implies that the long-run effect is not much larger than the immediate effect. Also note that when the number of time periods is relatively small, a fixed effects estimator is inconsistent when a lagged dependent variable enters the specification as an explanatory variable, because differencing the current error term from its time-mean introduces a correlation between that error term and the time-mean differenced lagged dependent variable. As a result, when we later estimate a dynamic version of our model which includes a lagged dependent variable, we first difference the data. See Bond (2002) for an overview of the issues involved.

[7] The demographic characteristics in ASC's $\Delta\mathbf{X}_{it}$ vector include a third-order polynomial in age, as well as dummy variables for each of: receipt of wage or salary income, receipt of self-employment income, married, head of household, retired, and whether the taxpayer has children. Thus, the underlying level variables in $\mathbf{X}_{it}$ would include, for example, each of the dummy variables interacted with a time trend, and a an appropriate set of level variables for age that produce a third-order polynomial when first-differenced.

[8] If $Cov(\Delta y_{it}^R, \Delta y_{it-1}^R)$ is positive, or if it is negative but is larger than $.5*Var(\Delta y_t^R)$ in absolute value, then the estimate of either σ_ξ^2 or σ_η^2 would be negative. This would suggest an obvious violation of the assumptions of the model. This problem did not arise in our estimation.

[9] Permanent income is found by calculating $\mathbf{y^*} = \mathbf{B}^{-1}\mathbf{y^R}$, where $\mathbf{y^*}$ is a vector including y_{it}^* for each period and $\mathbf{y^R}$ is a vector including y_{it}^R for each period. If, for example, there are four periods, then:

$$\mathbf{B} = \begin{bmatrix} \lambda+1 & -\lambda & 0 & 0 \\ -\lambda & 2\lambda+1 & -\lambda & 0 \\ 0 & -\lambda & 2\lambda+1 & -\lambda \\ 0 & 0 & -\lambda & \lambda+1 \end{bmatrix}$$

[10] ASC remove the individual-specific fixed effect in this first stage by differencing. We estimate this first stage by fixed effects because it will become useful later in our own estimation procedure to be able to reconstruct the level of income predicted by this regression. In our replication of ASC, we tried doing the first stage regressions both ways and found the results to be nearly identical.

[11] We use a time trend rather than year dummies here because the coefficients on the year dummies will partly reflect aggregate transitory fluctuations in income, and we want our predicted values from the first stage regressions to capture only the smooth component of income that is largely predictable in advance based on life cycle and demographic considerations, the individual-specific intercept, and economy-wide productivity trends. In the final stage regressions, we control for year dummies separately, so this should not make much difference.

[12] A case could be made that we actually want to evaluate the vertical shift in the expected future profile of prices and disposable incomes holding other characteristics constant at their expected *future* levels. A difficulty is that many of those future characteristics are unknown at time *t* (although some, such as age, are clearly known in advance). The idea behind holding characteristics constant at their t-1 levels is that the size of the vertical shift in the profile will be approximately the same whether we use last year's values of other characteristics, or values from a few years hence. Since we are trying to measure the vertical distance of an approximately parallel shift in the profile, holding these characteristics constant matters more than which particular year they are chosen from.

[13] Marital status is not included because, given our sample selection method, it does not change over time for an individual, and thus is collinear with the fixed effect (or always zero after first differencing).

[14] Although state of residence is omitted for returns with AGI above $200,000 in the public use panel, in the vast majority of cases we were able to impute state of residence based on the observed state in years when the taxpayer had an AGI below $200,000. In cases where the taxpayer's AGI is never below $200,000, we treat the taxpayer as if he or she lived in a state without an income tax. Less than 0.4 percent of observations fall into this category. One complication is that in some states, married filing separately is advantageous, in which case tax liability and tax rates can depend on the division of income between spouses. Since our data do not provide information on the division of income between spouses (except, in some cases, for labor income in years when the two-earner credit was in effect), we assume that in married couples, 80 percent of labor and business income is earned by the primary taxpayer and 20 percent is earned by the spouse, while other income is assumed to be divided 50-50.

[15] In the event that a "notch" producing an unusually large marginal tax rate is encountered, the marginal rate is recalculated by subtracting $0.10 from the initial value.

[16] The long-run change in the level of log giving in response to a one unit unanticipated persistent shock to log price is $\delta_1 + \delta_1\delta_9 + \delta_5 + (\delta_1\delta_9 + \delta_5)\,\delta_9 + (\delta_1\delta_9 + \delta_5)\,\delta_9^2 + \ldots = (\delta_1 + \delta_5)\sum_{j=0}^{\infty}\delta_9^j = (\delta_1 + \delta_5)/(1-\delta_9)$.

The long-run effect on the level of giving of a one unit shock to log price occurring at time t=1 that is recognized as transitory and which lasts only one period is: $\delta_2 + \delta_2\delta_9 - \delta_2 + \delta_6 + (\delta_2\delta_9 - \delta_2 + \delta_6)\,\delta_9 - \delta_6 + ((\delta_6\delta_9 - \delta_2 + \delta_6)\delta_9 - \delta_6)\,\delta_9 + \ldots$, which gradually converges to zero.

[17] We implement this estimator using the "xtabond2" program developed for Stata by Roodman (2004). The Arellano Bond estimator essentially treats the estimation problem as a system of separate equations for each time period, with coefficients in each final stage equation constrained to be common across the equations. The system of equations approach allows differing numbers of instruments to be used to predict the lagged dependent variable in each equation. This increases efficiency because it allows larger numbers of lags of the instruments to be included in the later periods, information that would otherwise be thrown away in conventional two stage least squares. The system is estimated by generalized method of moments.