

Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply*

Raj Chetty

Harvard University and NBER

September 2011

Abstract

How can price elasticities be identified when agents face optimization frictions such as adjustment costs or inattention? I derive bounds on structural price elasticities that are a function of the observed effect of a price change on demand, the size of the price change, and the degree of frictions. The degree of frictions is measured by the utility losses agents tolerate to deviate from the frictionless optimum. The bounds imply that frictions affect intensive margin elasticities much more than extensive margin elasticities. I apply these bounds to the literature on labor supply. The utility costs of ignoring the tax changes used to identify intensive margin labor supply elasticities are typically less than 1% of earnings. As a result, small frictions can explain the differences between micro and macro elasticities, extensive and intensive margin elasticities, and other disparate findings. Pooling estimates from existing studies, I estimate a Hicksian labor supply elasticity of 0.33 on the intensive margin and 0.25 on the extensive margin after accounting for frictions.

*E-mail: chetty@fas.harvard.edu. I would like to thank Tim Armstrong, Soren Blomquist, Richard Blundell, David Card, Daniel Feenberg, Christopher Ferrall, John Friedman, Yuriy Gorodnichenko, James Hines, Caroline Hoxby, Susumu Imai, Guido Imbens, Damon Jones, Patrick Kline, Adam Looney, Day Manoli, Ariel Pakes, Valerie Ramey, Richard Rogerson, Emmanuel Saez, Johannes Spinnewijn, Harald Uhlig, Ken Wolpin, anonymous referees, and numerous seminar participants for very helpful comments. Gregory Bruich, Jane Choi, Peter Ganong, Nate Hilger, Tord Krogh, Jessica Laird, Zachary Liscow, Keli Liu, Laszlo Sandor, Heather Sarsons, and Danny Yagan provided outstanding research assistance. Funding from NSF Grant SES 0645396 and the Sloan Foundation is gratefully acknowledged. The code used to produce the results in this paper is available on the author's website at www.economics.harvard.edu/faculty/chetty.

1 Introduction

The identification of structural parameters of stylized models is one of the central tasks of applied economics. Unfortunately, most models omit various frictions that make agents deviate systematically from their theoretical predictions. For instance, canonical models of labor supply or consumption behavior do not permit adjustment costs, inattentive agents, or status quo biases. How can structural parameters be identified when agents face such optimization frictions?

One natural solution is to estimate the structural parameters of a model that incorporates the frictions. This approach has two limitations in practice. First, it is difficult to incorporate all frictions in a tractable model. Second, estimating even simple dynamic models with frictions, such as Ss adjustment, requires strong assumptions and is computationally challenging (Attanasio 2000). Motivated by these limitations, I propose an alternative solution: bounding structural preference parameters without identifying how frictions affect behavior.

I analyze a standard dynamic lifecycle model in which the effect of income-compensated (Hicksian) price changes on demand is determined by a structural parameter of utility ε . I introduce optimization frictions into this nominal model through an error term in the demand function whose conditional expectation is unknown. These optimization errors generate differences between mean observed demand and the mean optimal demand predicted by the frictionless model. Because the optimization errors need not be orthogonal to the price, the observed Hicksian elasticity $\hat{\varepsilon}$ estimated from demand responses to a price change differs from the structural elasticity parameter ε . The observed elasticity $\hat{\varepsilon}$ confounds preferences (ε) with the effect of the frictions. For example, agents may under-react to a price increase in the short-run because of adjustment costs.

This paper seeks to identify ε from estimates of $\hat{\varepsilon}$. I focus on identifying ε because it is important for both positive and normative analysis. The impacts of prices in steady-state are determined purely by ε in many models. Moreover, the recovery of preference parameters is essential for welfare analysis.

I bound ε from observations of $\hat{\varepsilon}$ by assuming that agents choose points near the frictionless optimum. Specifically, I allow agents to deviate arbitrarily from the nominal model's prediction as long as the expected lifetime utility cost of doing so is less than δ percent of

expenditure. This property is satisfied by standard dynamic adjustment cost models, where agents remain on average within some utility threshold of their optimum. In the case of other frictions such as inattention or status quo biases, this restriction requires that agents respond to incentives that are sufficiently important.

I derive a closed-form representation for bounds on the structural Hicksian elasticity ε as a function of the observed Hicksian elasticity $\widehat{\varepsilon}$, the size of the price change used for identification $\Delta \log p$, and the degree of frictions δ .¹ The bounds shed light on what can be learned from reduced-form elasticity estimates in an environment with frictions. The bounds shrink at a quadratic rate with $\Delta \log p$. As a result, pooling several small price changes – although useful in improving statistical precision – yields less information about the structural elasticity than studying a few large price changes. If $\widehat{\varepsilon} > 0$, the lower bound on the structural elasticity ε is strictly positive, showing that frictions do not affect tests of a null hypothesis of zero response. If the observed elasticity $\widehat{\varepsilon} = 0$, the upper bound on ε can be expressed in terms of the utility cost of ignoring the price change. This permits straightforward calculations of the range of elasticities consistent with zero behavioral response, analogous to power calculations used to evaluate statistical precision.

The preceding results apply to an intensive margin model in which consumption is perfectly divisible. I also derive bounds on extensive margin elasticities by analyzing a model in which agents choose whether to buy an indivisible good. The bounds on the structural extensive margin elasticity (η) shrink linearly with δ and are therefore an order of magnitude tighter than those on the Hicksian intensive margin elasticity (ε). The bounds are tighter because the utility costs of ignoring price changes are first-order on the extensive margin, in contrast with the second-order costs on the intensive margin. Hence, frictions such as adjustment costs or inattention have smaller effects on aggregate demand when microeconomic choices are discrete rather than continuous.

One can obtain tighter bounds on ε or η by calculating the least upper bound and the largest lower bound implied by multiple observed elasticities. The sensitivity of structural elasticity estimates to frictions can be evaluated by computing these unified bounds as a function of δ . The smallest level of frictions δ_{\min} that reconciles a group of observed elasticities provides

¹The value of δ must be specified exogenously and may vary across applications. I consider the assumption that aggregate welfare would be 1% higher absent frictions ($\delta = 1\%$) to be a plausible benchmark.

measures the “economic significance” of the differences in estimates. If δ_{\min} is small, the differences are not economically significant in that they can be explained simply by allowing for small frictions. The value of ε or η when $\delta = \delta_{\min}$ converges to the true value as the number of observed elasticities grows large, providing a point estimate of the structural elasticity adjusted for frictions.

I apply these methods to investigate what can be learned about structural labor supply elasticities from the empirical literature on labor supply. The application consists of four components, each of which addresses a different strand of the labor supply literature.

First, I analyze the impact of frictions on the intensive margin elasticity – the effects of tax changes on hours of work for employed individuals. Based on a large body of microeconomic evidence, “the profession has settled on a value for this elasticity close to zero” (Saez, Slemrod, and Giertz 2011). I show that small frictions could explain why observed elasticities are often near zero by calculating the utility costs of ignoring tax reforms. For instance, the utility costs of ignoring the widely studied Tax Reform Act of 1986 (TRA86) – and instead choosing the optimal pre-reform level of work hours – are less than 2% of income per year for all except top income earners. Accordingly, empirical studies find that TRA86 induced behavioral responses in the short run only for top income earners. To assess what can be learned about ε from existing estimates of intensive margin labor supply elasticities, I calculate bounds on ε using estimates from studies of hours elasticities, taxable income elasticities, elasticities for top income earners, and macroeconomic cross-country estimates. Even though the observed elasticity estimates vary widely, all the estimates are consistent with a single structural elasticity ε if one permits frictions of 1% of post-tax earnings in choosing labor supply. Pooling the fifteen hours and taxable income elasticity estimates yields bounds on ε of (0.28, 0.54) when $\delta = 1\%$, with a 95% confidence interval of (0.23, 0.61). The minimum level of frictions required to reconcile these fifteen estimates is $\delta_{\min} = 0.5\%$ of net earnings and the corresponding point estimate of the structural Hicksian elasticity is $\varepsilon_{\delta_{\min}} = 0.33$.

Second, I analyze how frictions affect extensive margin elasticities – the effects of tax changes on employment rates. The utility costs of ignoring tax changes on the extensive margin are between 5 to 10% of income for many tax policy changes in the U.S. These large costs could explain why microeconomic studies uniformly detect significant impacts of tax changes on employment rates despite finding negligible intensive margin responses.

This result challenges the prevailing consensus that extensive margin elasticities are much larger than intensive margin elasticities. Instead, current empirical methods may simply be better suited to detecting responses on the extensive margin than the intensive margin in the presence of frictions. I calculate bounds on the structural extensive margin elasticity η using estimates from existing studies. The bounds on the extensive margin elasticities implied by each study are very tight, confirming that observed elasticities provide accurate estimates of structural elasticities on the extensive margin. The mean extensive margin elasticity among the microeconomic studies I consider is 0.25.

Third, I turn to the literature on non-linear budget set estimation, which analyzes the impacts of progressive income taxation on labor supply. One well known issue in fitting such models is that much fewer individuals bunch at kink points of the tax system than one would predict in a frictionless model (Saez 2010). I show that the utility gains from bunching at kinks are very small, typically less than 1% of consumption. Allowing for optimization frictions in NLBS models can explain the lack of bunching at most kinks and provide a more disciplined error structure for such models.

Finally, I assess whether frictions can explain the discrepancy between micro and macro estimates of Hicksian (steady-state) and Frisch (intertemporal substitution) elasticities. Macro estimates of the Hicksian elasticity, which are based on cross-country comparisons, are larger than micro estimates. Frictions and indivisible labor (Rogerson 1988, Rogerson and Wallenius 2009) fully account for this gap. On the intensive margin, the micro estimate of $\varepsilon = 0.33$ after accounting for frictions matches cross-country evidence. On the extensive margin, micro estimates match macro cross-country estimates even without accounting for frictions, as expected given the results above. Macro estimates of the Frisch elasticity, which are based on fluctuations in labor supply over the business cycle, are also larger than micro estimates of the Frisch elasticity. To assess whether this discrepancy can be explained by frictions, I bound the structural intensive margin Frisch elasticity using the estimated structural Hicksian elasticity. The bound on the intensive Frisch elasticity is consistent with macro evidence on fluctuations in hours conditional on employment over the business cycle. However, micro estimates are not consistent with macro evidence on the extensive margin intertemporal substitution elasticity, as shown in Chetty et al. (2011b). I conclude that frictions explain the gap between micro and macro estimates of steady-state elasticities relevant for cross-country comparisons, but can-

not reconcile the differences between micro and macro estimates of intertemporal substitution elasticities relevant for business cycles.

The analysis here should be viewed as one step toward characterizing how frictions affect labor supply elasticities. The results are based on a standard lifecycle model of labor supply as in MaCurdy (1981) and do not account for factors incorporated into more recent models, such as human capital accumulation, credit constraints, or uncertainty. One would have to specify a nominal model that incorporates all of these structural features to bound labor supply elasticities in such an environment. This point illustrates a more general caveat: the ability to account for frictions using bounds does not provide an excuse for failing to build an accurate model. The bounds are valid only if the nominal model is correct up to optimization frictions.

This paper builds upon and relates to the partial identification, near rationality, robust control, and durable goods literatures. The econometric literature on partial or set identification addresses problems such as missing data or imperfect instruments (Manski 2007, Nevo and Rosen 2008). The present paper uses set identification to estimate structural parameters with model mis-specification. I derive bounds by assuming that agents are “near rational,” as in the menu cost and near rationality literature in macroeconomics (Akerlof and Yellen 1985, Mankiw 1985, Cochrane 1989). The focus on a class of models around a pre-specified nominal model parallels the robust control literature (Hansen and Sargent 2007). The robust control literature analyzes optimal policy with a minimax criterion and model uncertainty, whereas I consider identification of the nominal model’s parameters in the same setting. Finally, the bounds provide an alternative method of estimating preferences or production functions in models with adjustment costs. This approach requires fewer assumptions than existing methods of identifying such models (e.g. Eberly 1994, Attanasio 2000) because it uses inputs that can be estimated using quasi-experimental techniques. However, it does not permit as rich an analysis of short-run counterfactuals because it only partially identifies the model’s parameters.

The paper is organized as follows. The next section sets up a dynamic model with frictions. The bounds on intensive and extensive margin price elasticities are derived in Section 3. Section 4 presents the application to labor supply and taxation. Section 5 concludes.

2 Demand Models with Frictions

Consider a dynamic model with N individuals who have heterogeneous tastes over two goods, x and y . The price of x in period t is p_t and the price of y is fixed at 1. Individual i has wealth Z_i and chooses demand by solving

$$\max_{x_t, y_t} \sum_{t=1}^T v_{i,t}(x_t, y_t) \quad \text{s.t.} \quad \sum_{t=1}^T [p_t x_t + y_t] = Z_i \quad (1)$$

To simplify exposition, I make two simplifying assumptions in the main text. First, I assume that agents face no uncertainty: prices p_t evolve deterministically. Second, I use the following specification of flow utility:

$$\begin{aligned} v_{i,t}(x_t, y_t) &= y_t + a_{i,t} \frac{x_t^{1-1/\varepsilon}}{1-1/\varepsilon} \quad \text{if } \varepsilon \neq 1 \\ v_{i,t}(x_t, y_t) &= y_t + a_{i,t} \log x_t \quad \text{if } \varepsilon = 1 \end{aligned} \quad (2)$$

This quasilinear utility specification has three convenient properties: (1) it is a money metric, (2) it makes the agent's problem static because optimal demand for x_t depends only on p_t and $a_{i,t}$, and (3) it permits heterogeneity in demand levels but generates a constant price elasticity ε . I show that the main result (Proposition 1) applies to the general case where $v_{i,t}(x_t, y_t)$ is not quasilinear and prices p_t are stochastic in Appendix A.

The utility function in (2) implies that optimal demand is $x_{i,t}^*(p_t) = (\frac{a_{i,t}}{p_t})^\varepsilon$. Let $\alpha = \sum_i \sum_t \log x_{i,t}^*(p_t = 1)/NT$ denote the mean log demand in the population when $p_t = 1$ and $\nu_{i,t} = \log x_{i,t}^*(p_t = 1) - \alpha$ denote the deviation of individual i in period t from the mean. Then we can write agent i 's demand function as

$$\log x_{i,t}^*(p_t) = \alpha - \varepsilon \log p_t + \nu_{i,t}$$

My objective is to identify ε , the structural preference parameter that controls the price elasticity of demand. More compactly, I refer to ε as the ‘‘structural elasticity.’’ When utility is quasilinear, the Hicksian (utility constant), Marshallian (wealth constant), and Frisch (marginal utility constant) elasticities are all equal to ε . The bounds derived below apply to the Hicksian elasticity when utility is not quasilinear (see Appendix A). I therefore use ε to denote the Hicksian elasticity in the general model in (1), in which the three elasticities differ.

Consider identification of ε using a price change from p_A in period A to $p_B \neq p_A$ in period B .² The standard assumption made to identify ε from such variation is the following orthogonality condition on the error term $v_{i,t}$.

A1 Tastes are orthogonal to the identifying price variation: $\mathbb{E}v_{i,A} = \mathbb{E}v_{i,B}$.

Under this assumption,

$$\varepsilon = -\frac{\mathbb{E} \log x_{i,B}^*(p_B) - \mathbb{E} \log x_{i,A}^*(p_A)}{\log p_B - \log p_A} \quad (3)$$

Equation (3) shows that the observed response to a price change point identifies ε in the frictionless model in (1). I refer to (1) as the “nominal” model, following the robust control literature. I now show how optimization frictions affect the link between ε and the observed response using two examples.

Example 1: Adjustment Costs. Suppose the agent must pay an adjustment cost of $k_{i,t}$ to change his consumption of x in period t . In this model, agent i chooses consumption $x_{i,t}$ in period t by solving:

$$\max_{x_t} \sum_{t=1}^T [a_{i,t} \frac{x_t^{1-1/\varepsilon}}{1-1/\varepsilon} - p_t x_t - k_{i,t} \cdot (x_t \neq x_{t-1})]. \quad (4)$$

Observed demand in this model, $x_{i,t}$, differs from the frictionless optimum $x_{i,t}^*$. Let the observed elasticity estimated from a price change between periods A and B be denoted by

$$\hat{\varepsilon} = -\frac{\mathbb{E} \log x_{i,B}(p_B) - \mathbb{E} \log x_{i,A}(p_A)}{\log p_B - \log p_A}$$

In this model, $\hat{\varepsilon}$ no longer identifies the structural elasticity ε . The observed elasticity $\hat{\varepsilon}$ may be smaller or larger than ε depending upon the evolution of prices, adjustment costs, and tastes. Nevertheless, the structural elasticity ε still plays a central role in determining behavior in steady-state. For example, the effects of permanent price variation across economies starting in period 1 (e.g., countries with different tax regimes) is determined purely by ε . Intuitively, adjustment costs affect observed elasticities primarily in the short-run, as agents delay adjustment until periods with low switching costs.

²The analysis is unaffected if the identifying price variation comes from comparing two different individuals facing different prices in the same period A , provided that the variation in p_A is orthogonal to the variation in tastes across individuals $\nu_{i,A}$.

Example 2: Price Misperceptions. A growing body of evidence indicates that individuals misperceive prices, e.g. because of inattention to tax rates (DellaVigna 2009). To model this class of deviations from (1), let $\tilde{p}_{i,t}(p_t)$ denote agent i 's perceived price as a function of the true price in period t . The agent chooses $x_{i,t}$ by solving

$$\max_{x_t} \sum_{t=1}^T [a_{i,t} \frac{x_t^{1-1/\varepsilon}}{1-1/\varepsilon} - \tilde{p}_{i,t}(p_t) \cdot x_t] \quad (5)$$

The resulting observed elasticity is

$$\hat{\varepsilon} = \varepsilon \frac{\mathbb{E} \log \tilde{p}_{i,B}(p_B) - \mathbb{E} \log \tilde{p}_{i,A}(p_A)}{\log p_B - \log p_A}$$

Again, the observed elasticity $\hat{\varepsilon}$ confounds the structural elasticity of interest ε with other parameters, in this case the effect of the price change on mean perceived prices. But if perceptions converge to the truth over time, steady-state behavior is determined solely by ε .

Optimization Frictions and Partial Identification. The two examples above illustrate why it is challenging to accurately model and fully identify structural models with frictions. In the first example, full identification requires estimation of many primitives. The second example is more challenging because it requires specification of a theory of perceptions $\tilde{p}_{i,t}(p_t)$. This problem motivates a less ambitious strategy: identifying ε without fully identifying the primitive sources of optimization frictions. Identifying ε is useful (though not always sufficient) for both positive and normative analysis. As discussed in the examples above, ε is sufficient to predict steady-state responses under plausible conditions. The structural elasticity ε and the observed elasticity $\hat{\varepsilon}$ are together sufficient for welfare calculations in many applications (Chetty, Looney, and Kroft 2009).

It is useful to recast the problem of identifying ε with unknown frictions as a partial identification problem. Define agent i 's "optimization error" as the log difference between his optimal demand under the nominal model and his observed demand: $\phi_{i,t} = \log x_{i,t} - \log x_{i,t}^*$.³ Then observed demand for agent i can be written as

$$\log x_{i,t} = \alpha - \varepsilon \log p_t + \nu_{i,t} + \phi_{i,t} \quad (6)$$

Define $x_t(p_t) = [\prod_{i=1}^N x_{i,t}(p_t)]^{1/N}$ and $x_t^*(p_t) = [\prod_{i=1}^N x_{i,t}^*(p_t)]^{1/N}$ as the geometric means of ob-

³The optimization error is an error from the econometrician's perspective but not necessarily from the agent's perspective. In the adjustment cost model, the agent optimizes by choosing $x_{i,t}$ according to (4).

served and optimal demands.⁴ Mean observed (log) demand is

$$\log x_t = \mathbb{E} \log x_{i,t} = \log x_t^*(p_t) + \mathbb{E} \phi_{i,t}$$

Unlike the preference heterogeneity error $\nu_{i,t}$, the optimization errors $\phi_{i,t}$ generated by frictions are not orthogonal to changes in prices. For example, in the adjustment cost model, mean observed demand may be at the optimum in period A ($\mathbb{E} \phi_{i,A} = 0$), but above the new optimum following a price increase in period B ($\mathbb{E} \phi_{i,B} > 0$). Without assumptions on $\phi_{i,t}$, ε is unidentified by the observed response $\mathbb{E} \log x_{i,B} - \mathbb{E} \log x_{i,A}$. Intuitively, if one places no restrictions on perceptions or adjustment costs, an observed response to a price change can be reconciled with any structural price elasticity.

Restricting the Degree of Frictions. One can obtain bounds on ε by restricting the support of $\phi_{i,t}$ without making additional assumptions about $\mathbb{E} \phi_{i,t}$. I restrict the support of $\phi_{i,t}$ by requiring that agents make choices near the optimal choice under the nominal model. I obtain a money-metric measure of the utility cost of setting x suboptimally for the general nominal model in (1) using an expenditure function. Let $U_{i,t}^*$ denote agent i 's total utility from periods t to T under his optimal consumption plan. The minimum expenditure needed to attain $U_{i,t}^*$ when the agent sets x_t at \tilde{x}_t is

$$e_{i,t}(\tilde{x}_t) = \min_{x_s, y_s} \sum_{s=t}^T (p_s x_s + y_s) \text{ s.t. } \sum_{s=t}^T v_{i,t}(x_s, y_s) \geq U_{i,t}^* \text{ and } x_t = \tilde{x}_t.$$

The agent's utility cost (measured in dollars) from setting $x_{i,t}$ suboptimally is $e_{i,t}(x_{i,t}) - e_{i,t}(x_{i,t}^*)$. I restrict the size of optimization errors by requiring that the mean utility cost as a fraction of optimal expenditure on good x is less than an exogenously specified threshold δ :

$$\frac{1}{N} \sum_i [e_{i,t}(x_{i,t}) - e_{i,t}(x_{i,t}^*)] / p_t x_{i,t}^* \leq \delta \quad (7)$$

The threshold δ measures the degree of optimization frictions, scaled as a percentage of expenditure on good x . For instance, $\delta = 1\%$ permits deviations from optimal demand with an average utility cost of up to 1% of expenditure on x_t .⁵ I measure utility costs under the

⁴The geometric mean is analytically convenient because individuals with different levels of expenditure are weighted equally in calculations of aggregate demand elasticities. If one defines mean demand as an arithmetic mean, the results below hold if the δ class of models in (7) is defined as requiring that the expenditure-weighted mean of utility costs is less than δ .

⁵The appropriate choice of δ depends on the length of time that a period represents because δ is scaled by expenditure per period $p_t x_{i,t}^*$. For instance, in a fixed adjustment cost model, one should set δ to be 12 times larger when periods correspond to months rather than years.

nominal model because in standard models with frictions (e.g. Example 1 above), agents' choices depend upon whether the gains from reoptimization – as calculated under the frictionless model – exceed the size of the frictions. Because utility costs are calculated under the nominal model, the results that follow require that the nominal model is correct in a frictionless environment.

I refer to the models that generate observed demand levels $x_{i,t}$ that satisfy (7) as a “ δ class of models” around the nominal model.⁶ The adjustment cost model in (4) lies in the δ class of models around (1) if the average adjustment cost as a percentage of expenditure $\frac{1}{N} \sum_i k_{i,t}/p_t x_{i,t}^* \leq \delta/2$ in all periods t . Intuitively, if agents face adjustment costs of less than $\delta/2$, they will never tolerate a utility loss of more than δ by setting x_t suboptimally because they could always switch to x_t^* and then back to x_t in period $t + 1$. Similarly, the model of price misperceptions in (5) lies in the δ class of models around (1) if the expected utility losses due to misperceptions are less than δ – that is, if perceptions are not too inaccurate on average.

Although (7) is defined based on the utility cost of setting demand suboptimally in a single period, the δ class of models includes dynamic models in which agents make choices based upon the present value of utility gains over their lifetimes. The reason is that with a suitable choice of δ , (7) provides a worst-case scenario for the choice of x_t . For example, in the adjustment cost model, forward-looking agents might switch x_t to x_t^* even if the flow utility gains from doing so are smaller than $k_{i,t}/2$, because they can reap utility gains over their lifetimes by paying the switching cost once. However, irrespective of the path of prices and tastes, these forward-looking agents' behavior will always satisfy (7) if δ is specified as twice the mean adjustment cost. The choices of myopes who consider only flow utility gains will also satisfy (7). The δ class of models thus encompasses a rich set of dynamic models of behavior around the nominal model.

A δ class of models maps prices and primitives to a set of mean demand levels, which I denote by

$$X_t(p_t, \delta) = \{x_t : \frac{1}{N} \sum_i [e_{i,t}(x_{i,t}) - e_{i,t}(x_{i,t}^*)]/p_t x_{i,t}^* \leq \delta\} \quad (8)$$

When utility is quasilinear, the choice set $X_t(p_t, \delta)$ takes a particularly simple form. In the

⁶The restriction on $x_{i,t}$ in (7) is effectively a restriction on the support of the optimization error $\phi_{i,t}$ because $x_{i,t} = x_{i,t}^* e^{\phi_{i,t}}$

quasilinear case, we can assume without loss of generality that the agent splits his wealth equally across periods because the consumption path of y_t does not affect utility. Then flow utility as a function of x_t is given by

$$u_{i,t}(x_t) = Z_i/T - p_t x_t + a_{i,t} \frac{x_t^{1-1/\varepsilon}}{1-1/\varepsilon}$$

In this case, (7) can be written as the set of demands that yield flow utility within δ units of the optimum on average:

$$X_t(p_t, \delta) = \{x_t : \frac{1}{N} \sum_i [u_{i,t}(x_{i,t}^*) - u_{i,t}(x_{i,t})] / p_t x_{i,t}^* \leq \delta\}$$

Because the demand problem under the nominal model is effectively static with quasilinear utility, the lifetime utility cost of setting x_t suboptimally in period t is just the flow utility cost of the error. Figure 1 illustrates the construction of the choice set $X(p_t, \delta)$ with quasilinear utility when there is no heterogeneity across agents and $\delta = 1\%$. The figure plots flow utility $u(x_t)$ when $a_{i,t} = e^{3.5}$, $\varepsilon = 1$, $\log p_t = 1$, and $Z/T = 100$. The set of choices that yield utility within $\delta = 1\%$ of the optimum, $X(p_t, \delta) = [10.2, 14]$, is depicted by the interval on the x axis.

Now consider how a price increase from p_A to p_B affects mean observed demand in a δ class of models. Figure 2a illustrates the choice sets at the two prices, $X(p_A, \delta)$ and $X(p_B, \delta)$, with the same parameters as in Figure 1. The structural elasticity ε controls the movement of the choice sets with the price p , as illustrated by the dashed line. The solid lines illustrate that various mean demand changes $[\log x_B(p_B) - \log x_A(p_A)]$ may be observed for a given value of ε . Each solid line is generated by a different model. For instance, the flat line could be generated by a model with status quo bias or satisficing consumers. Over-reaction could be observed in a model with adjustment costs, e.g. if there has been a history of price increases in the past. One may even observe an *increase* in demand, for instance if the price increase reflects a change in tax policy that raises tax rates but makes taxes less salient.

These examples show that optimization frictions destroy the 1-1 map between the observed response and the structural elasticity in (3). Let the range of structural elasticities consistent with a given observed elasticity $\hat{\varepsilon}$ in a δ class of models be denoted by $(\varepsilon_L(\hat{\varepsilon}, \delta), \varepsilon_U(\hat{\varepsilon}, \delta))$. The objective of this paper is to characterize ε_L and ε_U in terms of empirically estimable parameters. The bounds $(\varepsilon_L, \varepsilon_U)$ measure the uncertainty in the structural elasticity due to potential mis-specification of the behavioral model, much as a statistical confidence interval

measures the uncertainty in the parameter estimate due to sampling error.⁷

3 Bounds on Price Elasticities

I derive bounds on intensive margin elasticities in two steps. First, I characterize $X_t(p_t, \delta)$, the set of mean observed demands at a price p_t for a given value of ε . Second, I identify the set of structural elasticities ε consistent with an observed elasticity $\widehat{\varepsilon}$. After establishing these results for the intensive margin case, I replicate the analysis for an extensive margin model in which x is an indivisible good. Finally, I show how multiple observed elasticities can be combined to obtain more informative bounds on the structural elasticity.

Throughout, I focus on identification of bounds on ε , taking $\widehat{\varepsilon}$ as an estimate from an infinite sample. Inference about the bounds in finite samples, where there is statistical imprecision in the estimate of $\widehat{\varepsilon}$, can be handled using the techniques proposed by Imbens and Manski (2004) or Chernozhukov, Hong, and Tamer (2007).

3.1 Bounds on the Choice Set

The following lemma analytically characterizes $X_t(p, \delta)$ for small δ using a quadratic approximation to flow utility $u_{i,t}(x)$ in the quasilinear case.

Lemma 1. For small δ , the set of mean observed demands is approximately

$$X_t(p_t, \delta) = \{x_t : |\log x_t - \log x_t^*| \leq [2\varepsilon\delta]^{1/2}\}$$

Proof. It is convenient to rewrite the definition of the choice set in (8) as requiring that $u_{i,t}(x_{i,t}^*) - u_{i,t}(x_{i,t}) \leq \delta_{i,t} p_t x_{i,t}^*$ and $\frac{1}{N} \sum_i \delta_{i,t} \leq \delta$. Here $\delta_{i,t}$ can be interpreted as the degree of frictions faced by agent i in period t . Taking a quadratic approximation to $u_{i,t}(x) = u_{i,t}(e^{\log x})$ around $\log x_{i,t}^*$ and exploiting the first-order condition under the nominal model $u'_{i,t}(x_{i,t}^*) = 0$ yields

$$u_{i,t}(x_{i,t}^*) - u_{i,t}(x) \simeq -\frac{1}{2}(x_{i,t}^*)^2 u''_{i,t}(x_{i,t}^*)(\log x - \log x_{i,t}^*)^2 \quad (9)$$

Therefore, agent i 's observed demand in period t must satisfy

$$|\log x_{i,t} - \log x_{i,t}^*| \leq [-2\delta_{i,t} \frac{p_t}{x_{i,t}^*} \frac{1}{u''_{i,t}(x_{i,t}^*)}]^{1/2} \quad (10)$$

⁷I characterize the range of ε rather than other measures of dispersion because we typically lack a prior distribution over the models within the δ class. A natural approach in such cases is to consider worst-case scenarios (Hansen and Sargent 2007).

With the quasilinear utility specification in (2), $u''_{i,t}(x_t) = \frac{\partial^2 v_{i,t}(x_t)}{\partial x_t^2}$ and the first order condition in the nominal model for $x_{i,t}$ is $\frac{\partial v_{i,t}}{\partial x}(x_{i,t}^*(p_t)) = p_t$. Implicitly differentiating this first order condition yields

$$u''_{i,t}(x_{i,t}^*) \frac{dx_{i,t}^*}{dp_t} = 1 \quad (11)$$

Substituting (11) into (10) gives the following restriction on demand for each agent:

$$|\log x_{i,t} - \log x_{i,t}^*| \leq [2\varepsilon\delta_{i,t}]^{1/2}$$

To derive bounds on mean observed demand x_t , use Jensen's inequality to obtain:

$$|\log x_t - \log x_t^*| = |\mathbb{E} \log x_{i,t} - \mathbb{E} \log x_{i,t}^*| \leq \mathbb{E}[2\varepsilon\delta_{i,t}]^{1/2} \leq [2\varepsilon\delta]^{1/2}$$

It follows that mean observed demand x_t in a δ class of models satisfies

$$|\log x_t - \log x_t^*| \leq [2\varepsilon\delta]^{1/2}$$

Note that the approximation error in this equation vanishes as $\delta \rightarrow 0$ because the remainder of the Taylor approximation in (9) involves higher-order terms.

Lemma 1 captures three intuitions. First, the width of the choice set, which is $2[2\varepsilon\delta]^{1/2}$ log units, shrinks at a square-root rate as δ goes to zero. This result implies that even small optimization frictions δ can generate substantial variation in observed behavior. With a price elasticity of $\varepsilon = 1$ and $\delta = 1\%$, the choice set extends approximately +/-14% around $x^*(p_t)$, as illustrated in Figure 1. The root- δ shrinkage of the choice set is driven by the second-order losses of deviating from the maximum of a smooth function (Akerlof and Yellen 1985, Mankiw 1985).

Second, equation (10) shows that the width of the choice set is inversely related to the curvature of utility around the optimum, $u''_{i,t}(x_{i,t}^*)$. A useful property of the model is that $u''_{i,t}(x_{i,t}^*)$ is pinned down by ε , the structural parameter of interest. Highly curved utilities generate small structural elasticities because the agent has a strong preference to locate near $x_{i,t}^*$. For example, suppose the demand for an essential medicine is perfectly price inelastic at a level $x_{i,t}^*$. The price elasticity of demand approaches zero as the curvature of the utility function approaches infinity – agents demand the medicine at any price only if they lose infinite utility by not having it. Because the utility costs of deviating from $x_{i,t}^*$ are infinitely large,

the choice set $X_i(p_t, \delta)$ collapses to the singleton x_t^* for any δ when $\varepsilon = 0$, as illustrated in Figure 2b. The choice set expands as ε rises. This connection between ε and the curvature of utility is critical because it eliminates the need to estimate the additional parameter $u''_{i,t}(x_{i,t}^*)$ when bounding ε .

Finally, the set of mean observed demands depends only upon the mean level of frictions δ , and not the distribution of frictions at the individual level $\delta_{i,t}$. Because each individual's choice set is proportional to $[\delta_{i,t}]^{1/2}$, the potential difference between mean observed and optimal demand is largest (the worst-case scenario) when $\delta_{i,t} = \delta$ for all i, t .

3.2 Bounds on the Structural Elasticity

Figure 3a depicts the largest structural elasticity ε that could have generated an observed elasticity $\hat{\varepsilon}$ for a price increase from p_A to p_B . When $\varepsilon = \varepsilon_U$, mean observed demand lies at the bottom of the choice set at price p_A ($\log x_A(p_A) = \log x_A^*(p_A) - (2\varepsilon\delta)^{1/2}$) and the top of the choice set at price p_B ($\log x_B(p_B) = \log x_B^*(p_B) + (2\varepsilon\delta)^{1/2}$). The upper bound ε_U therefore satisfies the condition

$$\hat{\varepsilon} = -\frac{\log x_B(p_B) - \log x_A(p_A)}{\log(p_B) - \log(p_A)} = -\frac{\log x_B^*(p_B) - \log x_A^*(p_A) + 2(2\varepsilon\delta)^{1/2}}{\log(p_B) - \log(p_A)} = \varepsilon_U - 2\frac{(2\varepsilon_U\delta)^{1/2}}{\Delta \log p} \quad (12)$$

where $\Delta \log p = |\log(p_B) - \log(p_A)|$.⁸ Similarly, the lower bound structural elasticity ε_L consistent with $\hat{\varepsilon}$, illustrated in Figure 3b, is defined by the equation

$$\hat{\varepsilon} = \varepsilon_L + 2\frac{(2\varepsilon_L\delta)^{1/2}}{\Delta \log p} \quad (13)$$

The following proposition characterizes the solutions to (12) and (13).

Proposition 1. Under assumption A1, for small δ , the range of structural elasticities consistent with an observed elasticity $\hat{\varepsilon}$ is approximately $(\varepsilon_L, \varepsilon_U)$ where

$$\begin{aligned} \varepsilon_L &= \hat{\varepsilon} + \frac{4\delta}{(\Delta \log p)^2}(1 - \rho) \text{ and } \varepsilon_U = \hat{\varepsilon} + \frac{4\delta}{(\Delta \log p)^2}(1 + \rho) \\ \text{with } \rho &= \left(1 + \frac{1}{2}\frac{\hat{\varepsilon}}{\delta}(\Delta \log p)^2\right)^{1/2} \end{aligned} \quad (14)$$

Proof. Equations (12) and (13) both reduce to the quadratic equation $(\hat{\varepsilon} - \varepsilon)^2 = \frac{8\varepsilon\delta}{(\Delta \log p)^2}$. The upper and lower roots of this quadratic equation are the bounds.

⁸With $\Delta \log p$ defined as the absolute value of the log price change, the results below also apply to price reductions.

Equation (14) maps the magnitude of the price change ($\Delta \log p$), the observed elasticity $\hat{\varepsilon}$, and the degree of frictions δ to bounds on the structural elasticity ε when flow utility is quasilinear.⁹ In Appendix A, I show that when utility is not quasilinear, Proposition 1 applies to the Hicksian elasticity. In particular, if the demand function is isoelastic between p_A and p_B , an observed Hicksian elasticity $\hat{\varepsilon}$ generates bounds on the structural Hicksian elasticity ε given by exactly the same formula as (14). The discussion that follows therefore applies to Hicksian elasticities in the general model in (1).

The dashed lines in Figure 4 show the bounds $(\varepsilon_L, \varepsilon_U)$ as a function of $\hat{\varepsilon}$ with $\delta = 1\%$ of expenditure.¹⁰ Panel A considers a price change of $\Delta \log p = 20\%$, while Panel B considers $\Delta \log p = 40\%$. The bounds offer several insights into what can be learned about structural elasticities from reduced-form estimates of observed elasticities. First, larger price changes are much more informative about ε because the bounds shrink at a quadratic rate with $\Delta \log p$. With a price change of 20%, an observed elasticity of $\hat{\varepsilon} = 0.2$ is consistent with structural elasticities up to $\varepsilon_U = 2.3$. With $\Delta \log p = 40\%$ and $\hat{\varepsilon} = 0.2$, $\varepsilon_U = 0.85$. The reason for this rapid shrinkage is that the movement in the choice sets for a given value of ε is larger when $\Delta \log p$ is larger, resulting in a narrower set of observed responses $\hat{\varepsilon}$ consistent with any given ε .

Second, the bounds are asymmetric around the observed elasticity: $\varepsilon_U - \hat{\varepsilon} > \hat{\varepsilon} - \varepsilon_L$. This asymmetry is driven by the proportional relationship between the width of the choice sets and ε , as shown in Lemma 1. Large structural elasticities generate wide choice sets and are therefore consistent with a broader range of $\hat{\varepsilon}$ than small structural elasticities. A related implication is that if ε is small, there will be little dispersion in observed elasticities across studies, whereas a large ε may lead to substantial variation in observed elasticities.

Third, the lower bound is strictly positive ($\varepsilon_L > 0$) whenever $\hat{\varepsilon} > 0$ regardless of δ . If $\varepsilon = 0$, the choice sets collapse to a single point $x_t^*(p_A) = x_t^*(p_B)$ as shown in Lemma 1, and one will therefore never observe positive values of $\hat{\varepsilon}$. Agents intent on maintaining a fixed

⁹When $\hat{\varepsilon}$ is a finite-sample estimate, a 95% confidence set for ε can be obtained by computing ε_L using the lower limit of the 90% confidence interval for $\hat{\varepsilon}$ and ε_U using the upper limit of the 90% confidence interval under certain regularity conditions (Imbens and Manski 2004).

¹⁰These bounds are computed using (14), which relies on a quadratic approximation to utility. To evaluate the quality of the approximation, I calculated the exact bounds with the utility in (2) numerically for a range of values of $\hat{\varepsilon} < 1$, $\Delta \log p < 100\%$, and $\delta = 1\%$. In all cases, the exact and approximate bounds differ by less than 0.001, showing that (14) is sufficiently accurate for most applications.

value of x must face very large costs of deviating from the optimum and therefore will never do so.¹¹ This result is useful for hypothesis testing: finding $\widehat{\varepsilon} > 0$ is adequate to reject the null of a zero structural elasticity regardless of frictions.

Finally, consider the converse case of a study that detects zero observed behavioral response ($\widehat{\varepsilon} = 0$).¹² When $\widehat{\varepsilon} = 0$, the bounds take a particularly simple form. The lower bound is $\varepsilon_L = 0$. The upper bound can be conveniently expressed in terms of the utility cost of ignoring the price change for an optimizing agent with time-invariant preferences. Consider a hypothetical agent who has fixed tastes $a_{i,t} = a_i$ across periods A and B and is initially at his nominal optimum $x_i^*(p_A)$. Using a quadratic approximation analogous to that in Lemma 1, this agent's utility loss from failing to change demand to $x_i^*(p_B)$ in period B is

$$\Delta u_i \equiv u_{i,B}(x_i^*(p_B)) - u_{i,B}(x_i^*(p_A)) \simeq -\frac{1}{2}u''_{i,B}(x_{i,B}^*)(\log x_{i,B}^* - \log x_{i,A}^*)^2(x_{i,B}^*)^2.$$

Using equation (11), the utility loss from failing to reoptimize in response to a price change as a percentage of the optimal expenditure level at price p_B is

$$\Delta u_{\%}(\varepsilon) = \frac{\Delta u_i}{p_B x_i^*(p_B)} = \frac{1}{2}(\Delta \log p)^2 \varepsilon \quad (15)$$

The utility loss $\Delta u_{\%}(\varepsilon)$ is an increasing function of the structural elasticity ε . The following result shows that the upper bound on ε when $\widehat{\varepsilon} = 0$ can be expressed in terms of $\Delta u_{\%}(\varepsilon_U)$.

Corollary 1. Under assumption A1, for a given value of ε , the observed elasticity $\widehat{\varepsilon}$ can be 0 only if $\Delta u_{\%}(\varepsilon) \leq 4\delta$.

Proof. When $\widehat{\varepsilon} = 0$, (14) implies $\varepsilon_U = 8\delta/(\Delta \log p)^2$. Combining this equation with (15) yields the result.

Corollary 1 provides a simple method of determining the range of structural elasticities for which one can be sure to detect a behavioral response. Starting from the optimum, the percentage utility cost of ignoring a price change given an elasticity of ε must exceed 4δ to guarantee an observed elasticity $\widehat{\varepsilon} > 0$. The 4δ condition is obtained because the cost of deviating from the optimum rises at a quadratic rate (see Appendix A for details). When $\widehat{\varepsilon} = 0$, ε_U shrinks at a quadratic rate with $\Delta \log p$ but only a linear rate with δ . Studying

¹¹By the same reasoning, $\widehat{\varepsilon} < 0$ implies $\varepsilon > 0$, as one could never observe a negative response if $\varepsilon = 0$. Note that negative structural elasticities ($\varepsilon < 0$) are ruled out by agent optimization in the nominal model.

¹²Among the feasible responses in a δ class of models, a zero response is perhaps the most likely outcome, as it requires no adjustments or attention.

a price change that is twice as large yields more information about ε even if frictions are also twice as large, underscoring the value of placing greater weight on large treatments for identification.

3.3 Extensive Margin Elasticities

I now replicate the analysis above for the case where x is an indivisible good and agents make extensive margin choices about whether to buy x . To analyze extensive margin responses, consider the model in (1) with the quasilinear flow utility in (2), but assume that $x \in \{0, 1\}$, so that agents make a discrete choice.

It is optimal for an agent to buy the good if its utility exceeds its price, i.e. if $b_{i,t} \equiv \frac{\alpha_{i,t}}{1-\frac{1}{\varepsilon}} > p_t$. Let the distribution of the rescaled taste parameter $b_{i,t}$ in the population be given by a smooth cdf $F_t(b_{i,t})$ with positive support for all $b_{i,t} > 0$. I make an identification assumption analogous to A1 to ensure that elasticity estimates are unbiased without frictions:

A1' Tastes are orthogonal to the identifying price variation: $F_A = F_B$.

Let $\theta_t^* = 1 - F_t(p_t)$ denote the optimal fraction of agents who buy x and θ_t denote the observed fraction who buy x in period t . The structural extensive margin demand elasticity for a price change from p_A to p_B is $\eta(p_A, p_B) \equiv \frac{\log \theta_B^*(p_B) - \log \theta_A^*(p_A)}{\log(p_B) - \log(p_A)}$. The corresponding observed extensive margin elasticity is $\hat{\eta}(p_A, p_B) \equiv \frac{\log \theta_B(p_B) - \log \theta_A(p_A)}{\log(p_B) - \log(p_A)}$. Because the density $f(p)$ varies with the price, $\eta(p_A, p_B)$ varies with the price. To bound $\eta(p_A, p_B)$, I assume that the aggregate demand function is locally iso-elastic.

A2 The extensive margin elasticity is constant between p_A and p_B : $\eta(p_t) = -\frac{\partial \theta_t^*}{\partial p_t} \frac{p_t}{\theta_t^*} = \eta(p_A, p_B)$ for $p_t \in [p_A, p_B]$.

Let δ denote the degree of frictions permitted as a fraction of expenditure when buying the good, p_t . Then a δ class of models around the nominal extensive margin model can be defined by requiring that average utility losses are less than δp_t , as shown in Appendix A. I now establish a set of results analogous to those in the intensive margin case. The proofs of these results, which are given in the appendix, use first-order Taylor approximations and parallel those for the intensive margin.

Lemma 2. For small δ , the set of participation rates is approximately

$$\Theta_t(p_t, \delta) = \{\theta_t : |\log \theta_t - \log \theta_t^*| \leq \eta(p_t) \delta\}$$

The width of the choice set is directly proportional to the structural elasticity η and the degree of frictions δ . The structural elasticity matters because it is proportional to the density of the taste distribution $f(p_t)$. If $\eta(p_t)$ is large, many agents are near indifferent between buying x and not buying x at price p_t and experience small utility costs by choosing x suboptimally. This leads to greater variation in θ_t .

The critical difference between Lemma 2 and its intensive margin analog, Lemma 1, is that the width of the choice set for participation rates is proportional to δ rather than $\delta^{1/2}$. This makes the choice set much narrower on the extensive margin than the intensive margin. With $\delta = 1\%$ and a structural elasticity of 1, the choice set spans $\pm 1\%$ of optimal aggregate demand on the extensive margin, compared with $\pm 14\%$ on the intensive margin. The choice set is much narrower because individuals incur first-order utility losses from choosing x suboptimally on the extensive margin since they are not near interior optima.

The lower and upper bounds on η given an observed elasticity $\hat{\eta}$ can be characterized as in Figure 3b, leading to the following analog of Proposition 1.

Proposition 2. Under assumptions A1' and A2, for small δ , the range of structural elasticities consistent with an observed elasticity $\hat{\eta}$ is approximately (η_L, η_U) where

$$\eta_L = \hat{\eta}/(1 + \rho_\eta) \text{ and } \eta_U = \begin{cases} \hat{\eta}/(1 - \rho_\eta) & \text{if } \rho_\eta < 1 \\ \infty & \text{if } \rho_\eta \geq 1 \end{cases} \quad (16)$$

where $\rho_\eta = \frac{2\delta}{\Delta \log p}$.

Because the choice set grows linearly with η on the extensive margin, the bounds on η are more sensitive to the level of frictions. If the level of frictions is sufficiently large relative to the size of the identifying variation ($2\delta > \Delta \log p$), the η is unbounded above because the choice sets widen more rapidly than they shift as η rises. Intuitively, even if no one responds to a small price change on the extensive margin, there could nevertheless be a large lurking density of agents who are very close to indifferent between buying and not buying x , generating arbitrarily large η . In contrast, on the intensive margin, we obtain a finite upper bound on ε for any price change because the choice set grows more slowly (in proportion to $(\varepsilon)^{1/2}$) with ε .

Conversely, when frictions are relatively small, the bounds on η are much tighter than those on ε for a given δ because the choice set is much narrower on the extensive margin. This is illustrated by the dotted lines in Figure 4, which show the extensive margin bounds (η_L, η_U)

as a function of $\hat{\eta}$. With $\delta = 1\%$ and a price change of 20%, an observed elasticity of $\hat{\eta} = 0.2$ is consistent with extensive margin structural elasticities up to $\eta_U = 0.22$, in contrast with the upper bound of $\varepsilon_U = 2.3$ for the same parameters on the intensive margin. In practice, most empirical studies generate tight bounds on η for plausible levels of δ , as shown in the application below. For instance, with $\delta = 1\%$ frictions, one needs a price change of just 2% to obtain a finite upper bound on η .

One can also establish an analog to Corollary 1 by considering the utility cost of not responding to a price change for the agent who is just indifferent between buying and not buying at price p_A , i.e. the agent with $b_i = p_A$. Let

$$\Delta u_{\text{ext},\%} = \frac{|p_B - p_A|}{p_B} \simeq \Delta \log p \quad (17)$$

denote the utility cost to this agent (as a percentage of expenditure on x when participating) of choosing x suboptimally when the price is changed to p_B .¹³ The utility cost of ignoring a price change is a first-order function of $\Delta \log p$ on the extensive margin, in contrast with the second-order cost on the intensive margin in (15). Intuitively, non-participants enjoy the benefits of a price cut only if they reoptimize their behavior and enter the market. In contrast, on the intensive margin, the first-order increase in wealth from the price cut is automatically obtained; the benefit of reoptimization is only the second-order gain of choosing a better level of consumption.

Corollary 2. Under assumptions A1' and A2, if $\eta > 0$ then $\hat{\eta}$ can be 0 only if $\Delta u_{\text{ext},\%} \leq 2\delta$.

If the utility cost of ignoring the price change for the marginal agent exceeds 2δ , we must observe $\hat{\eta} > 0$ if $\eta > 0$. Because the utility losses from ignoring price changes are first-order on the extensive margin, price changes induce behavioral responses even with substantial frictions. A 20% change in the price could produce $\hat{\eta} = 0$ only with frictions of $\delta > 10\%$ when $\eta > 0$. In contrast, the same 20% change could produce $\hat{\varepsilon} = 0$ on the intensive margin with a structural elasticity of $\varepsilon = 0.5$ even with $\delta = 0.25\%$. Frictions have smaller effects on aggregate demand when microeconomic choices are discrete rather than continuous because the costs of suboptimal choice are concentrated among the marginal agents with $b_i \simeq p_A$.¹⁴

¹³For price cuts, the relevant utility cost is for a marginal agent who was not buying x at price p_A ; for price increases, the relevant utility cost is for an agent who was buying x at p_A . Intuitively, the agent who experiences the largest utility cost $\Delta u_{\text{ext},\%}$ determines the lower bound on $\hat{\eta}$.

¹⁴If the aggregate costs of suboptimal choice were shared across all agents, they would become a second-order function of $\Delta \log p$ because the fraction of agents who lose utility by not reoptimizing is proportional to $\Delta \log p$.

3.4 Combining Multiple Observed Elasticities

One can obtain more information about the structural elasticity by combining multiple observed elasticities. I demonstrate this for the intensive margin, but the results that follow apply identically to extensive margin elasticities. Suppose we have a set of observed elasticities $\{\widehat{\varepsilon}_1, \dots, \widehat{\varepsilon}_J\}$ from J empirical studies. Let $\Delta \log p_j$ denote the size of the price change used to identify observed elasticity j . Let ε_L^j and ε_U^j denote the lower and upper bounds implied by study j , derived using Proposition 1. Let $\varepsilon_L^{\max} = \max(\varepsilon_L^j)$ denote the largest lower bound and $\varepsilon_U^{\min} = \min(\varepsilon_U^j)$ denote the least upper bound. Then it follows that $\varepsilon \in (\varepsilon_L^{\max}, \varepsilon_U^{\min})$.

By calculating $(\varepsilon_L^{\max}, \varepsilon_U^{\min})$ as a function of δ , one can assess how sensitive estimates of ε are to frictions. One value of special interest is the smallest δ that reconciles the observed elasticities, δ_{\min} . When $\delta = \delta_{\min}$, the structural elasticity ε is point identified. To characterize this minimum- δ value of ε , let $\widehat{\varepsilon}_1$ denote the observed elasticity that produces the least upper bound and $\widehat{\varepsilon}_2$ the observed elasticity that produces the highest lower bound when $\delta = \delta_{\min}$. The minimum- δ estimate of ε satisfies $\varepsilon_{\delta_{\min}} = \varepsilon_U(\widehat{\varepsilon}_1, \delta_{\min}) = \varepsilon_L(\widehat{\varepsilon}_2, \delta_{\min})$. Solving these two equations using the definitions of ε_U and ε_L in (12) and (13) yields the following estimator:

$$\varepsilon_{\delta_{\min}} = \frac{\Delta \log p_1 \widehat{\varepsilon}_1 + \Delta \log p_2 \widehat{\varepsilon}_2}{\Delta \log p_1 + \Delta \log p_2} \quad (18)$$

Equation (18) also applies to extensive margin elasticities: $\eta_{\delta_{\min}}$ is the same weighted average of the pivotal observed elasticities.

The $\varepsilon_{\delta_{\min}}$ estimator for the structural elasticity has two attractive features. First, it does not require exogenous specification of δ . Second, if one were to observe all possible elasticities $\widehat{\varepsilon}$ generated by a δ class of models, the smallest level of frictions that could reconcile the observed values of $\widehat{\varepsilon}$ would be $\delta_{\min} = \delta$, resulting in $\varepsilon_{\delta_{\min}} = \varepsilon$. In this sense, $\varepsilon_{\delta_{\min}}$ converges to ε if observed elasticities are estimated in a sufficiently rich set of environments.

The value of δ_{\min} can be used to formally define “economically significant” differences. If δ_{\min} is small, the differences in estimates are not economically significant in that they can be reconciled simply by allowing for small frictions rather than fundamentally changing the economic model. In analogy with reporting the statistical significance of differences between estimates, the economic significance of a new estimate can be quantified by reporting the δ_{\min} required to reconcile it with prior evidence.

4 Application: Labor Supply

The wage elasticity of labor supply is a parameter of central interest for tax policy analysis and macroeconomic models. A large literature in labor economics, macroeconomics, and public finance estimates this elasticity using various methods. There are many frictions that may make observed labor supply differ from optimal labor supply, such as costs of switching jobs (Altonji and Paxson 1992), inertia (Jones 2008), and inattention (Chetty and Saez 2009). But few studies that estimate labor supply elasticities account for such frictions. The methods developed above are therefore well suited to extracting the information these studies contain about the structural labor supply elasticity.

I analyze the effects of frictions on four strands of the labor supply literature: (1) intensive margin elasticities, (2) extensive margin elasticities, (3) non-linear budget set estimation, and (4) macroeconomic elasticity estimates. Throughout, I focus on identifying Hicksian elasticities relevant for steady-state comparisons. I discuss the implications of the analysis for the Frisch (intertemporal substitution) elasticity relevant for understanding business cycle fluctuations in the context of the fourth application.

4.1 Intensive Margin Elasticities

Following MaCurdy (1981), I characterize structural labor supply elasticities in a lifecycle model in which agents choose consumption (c_t) and hours of work (l_t) to solve

$$\max_{c_t, l_t} \sum_{t=1}^T v_{i,t}(c_t, l_t) \quad \text{s.t.} \quad \sum_{t=1}^T [Y_{i,t} + (1 - \tau_t)wl_t - c_t] = 0 \quad (19)$$

where τ_t denotes the tax rate in period t , w denotes the wage rate, and $Y_{i,t}$ denotes unearned (non-wage) income. Let $l_t^{c,*}(\tau_t)$ denote the structural Hicksian labor supply function generated by (19). Note that (19) is equivalent to the demand model in (1) with leisure as one of the consumption goods. Because the Hicksian wage elasticity of leisure coincides with the Hicksian wage elasticity of labor supply, Proposition A1 can be used to bound the structural labor supply elasticity $\varepsilon = \frac{\log l_B^{c,*}(\tau_B) - \log l_A^{c,*}(\tau_A)}{\log(1 - \tau_B) - \log(1 - \tau_A)}$. In this application, $\Delta \log p = \Delta \log(1 - \tau)$ and δ measures frictions in choosing labor supply as a percentage of net-of-tax earnings $(1 - \tau_t)wl_t^{c,*}$.

I evaluate the impact of frictions on intensive margin elasticities in two steps. I begin by simulating the utility costs of ignoring the tax changes used for identification in the micro-

econometric literature. I find that the costs are typically quite small, suggesting that frictions might substantially attenuate observed elasticities (Corollary 1). I then calculate bounds on the structural Hicksian labor supply elasticity using existing estimates of observed elasticities.

4.1.1 Utility Costs of Ignoring Tax Changes

I calculate the costs of ignoring tax changes with quasilinear, iso-elastic flow utility:

$$v_{i,t}(c_t, l_t) = c_t - a_i \frac{l_t^{1+1/\varepsilon}}{1 + 1/\varepsilon}. \quad (20)$$

Let $T_t(wl)$ denote an agent's tax liability as a function of his taxable income in year t . Since the path of consumption has no impact on the utility costs of choosing l suboptimally when utility is quasilinear, I assume without loss of generality that the agent sets consumption equal to net-of-tax income. Then flow utility as a function of the labor supply choice and tax regime is

$$u_i(l; T_t) = wl - T_t(wl) - a_i \frac{l^{1+1/\varepsilon}}{1 + 1/\varepsilon} \quad (21)$$

I consider tax changes over a three year interval, following the convention in the literature (Gruber and Saez 2002). Let $l_{i,t}^*$ denote optimal labor supply in period t under the nominal model in (19). The utility loss in dollars from ignoring the tax changes that occur between years $t - 3$ and t for an individual who sets labor supply at the optimum in year $t - 3$ is:¹⁵

$$\Delta u_{i,t} = u_i(l_{i,t}^*; T_t) - u_i(l_{i,t-3}^*; T_t)$$

I calculate $l_{i,t}^*$ and $\Delta u_{i,t}$ numerically for various values of a_i and years t .¹⁶ I use a structural elasticity of $\varepsilon = 0.5$, the upper bound on ε estimated below, to obtain upper bounds on utility losses. The tax rates $T_t(wl)$ are obtained from the NBER TAXSIM calculator, including both employer and employee payroll taxes but ignoring state taxes. I consider a single tax filer with two children who has only labor income and no deductions other than those for children.

I adjust for inflation in the wage w using the CPI over the relevant three-year period.¹⁷

¹⁵The results below do *not* assume that all agents start at the optimum in the base year; they only require that choices in the base year lie within a δ class of models. I calculate utility costs for agents who start at the optimum because this calculation tells us whether $\hat{\varepsilon} = 0$ is consistent with a given structural elasticity (Corollary 1).

¹⁶The only heterogeneity across agents in these calculations is the disutility of labor a_i , which generates differences in pre-tax earnings. Heterogeneity in w_i is isomorphic to heterogeneity in a_i in (20).

¹⁷A STATA program TAXCOST.ado that calculates the utility cost of ignoring tax reforms has been posted on the NBER server. TAXCOST takes exactly the same inputs as TAXSIM. By running TAXCOST instead of TAXSIM, researchers can calculate the utility costs of ignoring the tax changes they are using for identification. See <http://obs.rc.fas.harvard.edu/chetty/taxcost.html> for further information.

Tax Reform Act of 1986. The Tax Reform Act of 1986 (TRA86) is one of the largest reforms in the U.S. tax code and the focus of many empirical studies. Figure 5 evaluates the costs of ignoring this tax reform. Panel A shows the marginal tax rate schedules in 1985 (thick red line) and 1988 (thin blue line). The dashed blue line, which is replicated in all the panels as a reference, shows the log change in the marginal net-of-tax rate (NTR), $\Delta \log(1 - \text{MTR})$. TRA86 increased the NTR by 15-20% for those with incomes below \$100,000 and by nearly 40% for those with incomes close to \$200,000.

Panel B plots the utility cost (measured in dollars) of ignoring the tax change ($\Delta u_{i,1988}$) vs. gross taxable income in 1985.¹⁸ For instance, an individual whose taste parameter a_i placed him at an optimal taxable income of \$100,000 prior to TRA86 would lose \$1,000 by failing to reoptimize labor supply in response to the change in the tax code. Panel C plots the cost of ignoring the tax reform as a percentage of consumption, $\Delta u_{i,1988,\%} = \Delta u_{i,1988} / (wl_{i,1988}^* - T_{1988}(wl_{i,1988}^*))$. Most individuals earning less than \$100,000 lose less than 1% of net earnings by ignoring TRA86 when choosing labor supply in 1988. Using Corollary 1, this result implies that frictions of $\delta = 1\%$ could lead to an observed elasticity of $\hat{\varepsilon} = 0$ from TRA86 for an individual earning less than \$100,000 even if his underlying structural elasticity were $\varepsilon = 0.5$.¹⁹

Finally, Panel D plots the change in taxable income ($wl_{i,1988}^* - wl_{i,1985}^*$) required to reoptimize relative to TRA86. With $\varepsilon = 0.5$, a taxpayer earning \$100,000 prior to the reform would have to increase his pre-tax earnings by \$13,000 in order to reach his new optimum. This substantial change would yield a utility gain (net of the disutility of added labor) of only \$1,000. Given that the search costs of immediately finding additional work that pays an extra \$13,000 could well exceed \$1,000, it is plausible that many individuals would not respond to TRA86 within a three-year horizon.²⁰

¹⁸Values at non-convex kinks in the base year are interpolated to obtain a continuous curve. Since no individual would optimally locate at a non-convex kink, the utility cost is undefined at these points.

¹⁹Corollary 1 applies to individuals who are at an interior optimum both before and after the tax change. In particular, a tax change could produce an observed elasticity $\hat{\varepsilon} = 0$ if the level of frictions $\delta > \Delta u_{i,t,\%}(\varepsilon)/4$ for such individuals. For individuals who optimally locate at kinks between tax brackets, the tangency conditions used to derive Corollary 1 do not hold. However, even for these agents, it is clear that a tax change could produce $\hat{\varepsilon} = 0$ if $\delta > \Delta u_{i,t,\%}(\varepsilon)$.

²⁰The total lifetime gain from reoptimizing labor supply is much larger because the agent gains \$1,000 every year. However, because the flow utility gains are relatively small, many agents may delay adjustment until a period where frictions (e.g. job switching costs) are lower. Thus, micro studies might not detect much change in labor supply between 1985 and 1988 even if TRA86 induced individuals to reoptimize in the long run.

The costs of ignoring TRA86 are considerably larger for high income earners. An individual earning \$200,000 in 1985 would lose \$4,500 per year (nearly 3% of net earnings) by ignoring the tax reform. High income individuals gain a lot more from reoptimizing both because the dollars at stake rise with income and because the change in tax rates was larger for high incomes.

Figure 6a extends the analysis of tax reforms to cover all tax changes from 1970-2006. I compute the percentage utility loss ($\Delta u_{i,t,\%}$) from ignoring tax changes at the 20th, 50th, and 99.5th percentile of the household income distribution. The value plotted for year t is the percentage utility cost of choosing $l_{i,t-3}^*$ instead of $l_{i,t}^*$ in year t . There is no tax change since 1970 for which the utility cost of failing to reoptimize on the intensive margin exceeds 1% of net earnings for the median taxpayer. The utility costs of ignoring tax reforms are substantial only for the top 1% of income earners around TRA86. Correspondingly, the largest observed elasticities in historical time series are for top income earners around TRA86; for lower income groups and other time periods, observed intensive margin elasticities are near zero (Saez 2004).

While there is little gain from adjusting behavior to optimally react to tax *changes* on the intensive margin over any three year interval, it is not the case that ignoring taxes completely imposes little cost. For example, using equation (15), the utility cost of ignoring a tax rate of $\tau = 40\%$ and working $l^*(\tau = 0)$ hours is $\frac{1}{2} \cdot \frac{1}{2} \cdot (0.4)^2 = 4\%$ of net earnings per year when $\varepsilon = 0.5$. This is why short-run responses to tax reforms may not be very informative about how the tax system affects labor supply on the intensive margin in steady state.

4.1.2 Bounds on the Intensive Margin Hicksian Elasticity

How much can be learned about the structural Hicksian labor supply elasticity (ε) from existing elasticity estimates? To answer this question, I apply Proposition A1 to calculate the bounds on ε implied by a set of well-known studies of intensive margin labor supply. One should keep two caveats in mind when interpreting the results of the exercise. First, I assume a constant structural elasticity ε across all the studies, ignoring potential variation in local preferences across tax regimes, income levels, demographic groups, or countries. Second, I assume that each study provides an unbiased estimate of the observed elasticity $\hat{\varepsilon}$. Econometric issues such as omitted variables and mean reversion may bias some of the estimates (Saez, Slemrod, and Giertz 2011). Any such biases would pass through to the bounds.

Table I divides the studies of intensive margin labor supply into four groups: (A) studies that measure labor supply using hours of work; (B) studies that measure labor supply using taxable income; (C) studies that use taxable income but focus exclusively on top income earners; and (D) studies that rely on cross-sectional comparisons (across countries with different tax regimes or individuals with different wage rates) to estimate steady-state hours elasticities.²¹ The table lists the point estimate and standard error of the observed Hicksian elasticity and the change in the net-of-tax rate used for identification. Details on the calculations and sources for each study are given in Appendix B. For quasi-experimental studies that analyze a single tax change, I define $\Delta \log(1-MTR)$ as the change in the mean MTR for the treatment group (e.g. top income earners in Feldstein (1995)). For studies that pool tax or wage changes of different sizes (e.g. Gruber and Saez 2002), I define $\Delta \log(1-MTR)$ as twice the standard deviation of $\Delta \log(1-MTR)$ in the sample. This is the size of the single price change that would generate the same statistical precision as the variation in $\Delta \log(1-MTR)$ used for identification, as shown in Appendix B.

The observed elasticity estimates vary substantially across studies. Microeconomic studies of the full population find the smallest elasticities: the mean observed hours and taxable income elasticities among the studies considered in Panels A and B is 0.15. Studies of top income earners find much larger elasticities, with a mean of 0.84. The mean elasticity among macroeconomic studies of steady-state responses is 0.32.

The largest observed elasticities in Panels A and B are obtained from the studies that focus on the largest changes in tax policy: the abolition of the income tax for a year in Iceland (Bianchi, Gudmundsson, and Zoega 2001) and a Swedish tax reform in 1991 termed the “tax reform of the century” (Gelber 2010). This pattern is consistent with the view that frictions are less likely to attenuate short-run responses to very large price changes. Excluding the Bianchi, Gudmundsson, and Zoega and Gelber studies, every point estimate in Panels A and B is below all of the point estimates in Panels C and D. Moreover, many of the confidence intervals for $\hat{\epsilon}$ in Panels A and B do not overlap with the confidence intervals for $\hat{\epsilon}$ in Panels C and D. Hence, the systematic differences in point estimates of observed elasticities across

²¹In the model in (19), the hours and taxable income elasticities are the same. I therefore pool estimates from both types of studies to bound the structural labor supply elasticity in this model. In more general models, taxable income elasticities may be larger than hours elasticities because they incorporate changes in reporting and avoidance behavior as well as changes in work effort (Slemrod 1995).

the studies in the different groups cannot be explained by statistical imprecision.

Can frictions explain the differences in the estimates? Columns 6-7 of Table I show the bounds $(\varepsilon_L, \varepsilon_U)$ implied by each point estimate with frictions of $\delta = 1\%$ of net earnings. Many studies that use small tax changes are consistent with structural elasticities above 1 despite obtaining small estimates. Figure 7 gives a visual representation of the bounds in columns 6-7. For scaling purposes, I exclude studies that use variation in net-of-tax rates of less than 20% for identification. None of the intervals plotted in the figure are disjoint – that is, all the estimates are consistent with a single structural Hicksian elasticity ε if one permits 1% frictions. Hence, the differences in estimates across high and low income earners as well as the differences in estimates across macroeconomic and microeconomic studies can be fully explained by small frictions.

Although any one study by itself produces wide bounds, the studies in Table I yield informative bounds on the Hicksian elasticity when combined. Intuitively, by estimating elasticities in many environments, one can obtain much sharper bounds on ε . The unified lower bound across the studies when $\delta = 1\%$ is $\varepsilon_L = 0.47$, obtained from Goolsbee’s (1999) analysis of TRA86. The unified upper bound is $\varepsilon_U = 0.51$, obtained from Blau and Kahn’s (2007) estimate using cross-sectional wage variation in the U.S. These bounds are robust in the sense that even if one excludes these two pivotal studies, the unified bounds expand only to $(0.44, 0.54)$, with the pivotal estimates now coming from Kopczuk (2010) and Gelber (2010).

While it is instructive to demonstrate that frictions can explain the differences in estimates between Panels B and C, the large elasticities for top income earners most likely reflect manipulation of reported taxable income rather than changes in labor supply (Slemrod 1995). One may also question the validity of the estimates in Panel D because of the many omitted variables and other factors that could bias cross-sectional comparisons (Alesina, Glaeser, and Sacerdote 2005). If we only include the studies in Panels A and B, the unified bounds are $(0.28, 0.54)$. These more conservative bounds are my preferred range of estimates for the structural labor supply elasticity with $\delta = 1\%$ frictions.

Figure 8 shows how the unified bounds vary with the degree of frictions. The dark shaded region shows the values of ε consistent with the observed elasticities in Panels A and B of Table I for $\delta \in (0, 5\%)$. The bounds widen as δ rises, but remain somewhat informative even with $\delta = 5\%$, where $\varepsilon_L = 0.15$ and $\varepsilon_U = 1.23$. Given that individuals are unlikely to tolerate

utility losses equivalent to 5% of net earnings per year on average, we can rule out $\varepsilon < 0.15$ (as suggested by some microeconomic studies) or $\varepsilon > 1.23$ (as used in some macro calibrations) based on existing evidence.

The smallest value of δ that can reconcile the observed elasticity estimates in Panels A and B is $\delta_{\min} = 0.5\%$. That is, the differences in these 15 observed elasticity estimates are “economically significant” only if frictions in choosing labor supply are less than 0.5% of net earnings on average. The corresponding minimum- δ point estimate of the structural elasticity is $\varepsilon_{\delta-\min} = 0.33$. This value of 0.33 is my preferred point estimate of the structural intensive margin Hicksian elasticity adjusted for frictions. Interestingly, this value is similar to the point estimates obtained from studies that are less susceptible to frictions to begin with – the steady-state cross-sectional comparisons in Panel D and the micro studies of large tax changes discussed above.

Columns 8-9 of Table I show a 95% confidence set for the ε implied by each study. These columns use the lower endpoint of the 90% confidence interval (CI) for $\hat{\varepsilon}$ to calculate ε_L and the upper endpoint of the 90% CI to calculate ε_U (Imbens and Manski 2004), assuming that $\hat{\varepsilon}$ is normally distributed. In many cases, the 95% confidence sets are only slightly wider than the bounds obtained when ignoring sampling error. For instance, ε_U for Gelber’s estimate for men rises from 0.54 to 0.59. A 95% confidence set for the unified bounds can be constructed by using a simple Bonferroni bound.²² The 95% confidence set for the unified bounds is (0.23, 0.61) when using the studies in Panels A and B. These calculations indicate that the greater source of imprecision in labor supply elasticities is uncertainty about the economic model of behavior due to frictions rather than noise due to sampling error.

4.2 Extensive Margin Elasticities

I now apply the results in Section 3.3 to explain why microeconomic estimates of observed elasticities on the extensive margin are larger than those on the intensive margin (Heckman 1993). As above, I first calculate the utility costs of ignoring tax changes on the extensive margin and then apply Proposition 2 to bound the extensive margin Hicksian elasticity.

²² Given J estimates $\{\hat{\varepsilon}_1, \dots, \hat{\varepsilon}_J\}$, let $(\varepsilon_L^{j,CI}, \varepsilon_U^{j,CI})$ denote a $1 - .05/J$ percent confidence interval for ε for study j , calculated using the method in Imbens and Manski (2004) as above. The intersection of these J regions is a (conservative) 95% CI for the unified bounds: $P[\varepsilon \in (\varepsilon_L^{j,CI}, \varepsilon_U^{j,CI}) \text{ for all } j] = 1 - P[\varepsilon \notin (\varepsilon_L^{j,CI}, \varepsilon_U^{j,CI}) \text{ for some } j] \geq 1 - \sum_{j=1}^J P[\varepsilon \notin (\varepsilon_L^{j,CI}, \varepsilon_U^{j,CI})] \geq 1 - J \times 0.05/J = 0.95$. Thanks to Tim Armstrong for suggesting this approach.

4.2.1 Utility Costs of Ignoring Tax Changes

I calculate the utility costs of suboptimal choice on the extensive margin using the model in (21) with $l \in \{0, 1\}$, so that agents can only choose whether to work or not. I follow the same methodology as in Section 4.1.1 to calculate the utility cost of ignoring a tax change for the marginal agent in year $t - 3$ at each gross earnings level w_i . The marginal agent at w_i has $b_i = w_i - T_{t-3}(w_i)$. The utility cost (measured as a percentage of net-of-tax earnings when working) of choosing l_t suboptimally for this agent is

$$\Delta u_{i,t,\text{ext},\%} = |\log(w_i - T_t(w_i)) - \log(w_i - T_{t-3}(w_i))|.$$

Earned Income Tax Credit Expansions. Figure 9 replicates Figure 5 for another important episode in U.S. tax policy – the expansion of the Earned Income Tax Credit under the Clinton administration. Most studies find virtually no changes in labor supply in response to EITC expansions for individuals on the intensive margin, but find a substantial response on the extensive margin (Meyer and Rosenbaum 2001, Eissa and Hoynes 2006). Figure 9 shows that this pattern could be driven by frictions.

Panel A shows tax changes and utility costs on the intensive margin. The dashed blue line shows that between 1993 and 1996, net-of-tax wage rates rose by 20% for single tax filers with two children earning below \$10,000 as the phase-in subsidy was increased. Meanwhile, net-of-tax wages fell by roughly 15% for those with incomes between \$15,000 and \$30,000 because of the increase in the phase-out tax rate. The solid red curve, constructed as in Figure 5c, shows that most individuals lose less than 1% of net earnings per year by ignoring these changes on the intensive margin. Corollary 1 implies that an observed response of $\hat{\varepsilon} = 0$ would be consistent with $\varepsilon = 0.5$ if one permits $\delta = 1\%$ frictions in reoptimizing labor supply.

Panel B of Figure 9 replicates Panel A for the extensive margin. The x axis of these figures is the income that the individual would earn (w_i) were he to work prior to the EITC expansion. On the extensive margin, the relevant tax rates are average rather than marginal. The dashed blue curve shows the change in net-of-average-tax rates (i.e., the return to working) as a result of this reform. The solid red curve shows the utility cost of ignoring the EITC expansion for individuals on the margin of entering the labor force at various income levels in 1993, which coincides with the log change in the net of tax rate as shown in (17). Consider an individual

who would earn \$5,000 when working and is indifferent between working and not working in 1993, i.e. has disutility of work $b_i = 5000 - T_{1993}(5000)$. Figure 9b shows that for this marginal individual, the gain from entering the labor force in response to the Clinton EITC expansion is 18% of net income when working, roughly \$1,000. In contrast, the gain from reoptimizing hours on the intensive margin for a worker earning \$5,000 prior to the reform is 0.7% of income, roughly \$50. On the extensive margin, the agent would have lost the extra \$1,000 EITC refund if he had ignored the tax reform and stayed out of the labor force. But on the intensive margin, a worker gets the \$1,000 tax reduction even if he does not change his hours. This could explain why individuals respond to the EITC expansion in the short run on the extensive margin despite frictions. Indeed, Corollary 2 implies that one could observe an elasticity of $\hat{\eta} = 0$ on the extensive margin only if frictions in adjusting labor supply exceed $\delta = 9\%$ of net-of-tax earnings when working.

Figure 6b extends this analysis to cover all tax changes from 1970-2006. In contrast to the intensive margin results shown in Figure 6a, there are several tax changes that would generate large utility losses (5-10% of net earnings) if ignored on the extensive margin.²³ The utility costs are particularly large for individuals who earn low incomes when working, consistent with the literature finding the largest extensive margin responses for this group.

4.2.2 Bounds on the Extensive Margin Hicksian Elasticity

Chetty et al. (2011b, Table 1) present a meta analysis of extensive margin elasticity estimates. In Table II, I apply Proposition 2 to calculate the bounds implied by the studies they consider with $\delta = 1\%$ frictions.²⁴ Panel A considers estimates from quasi-experimental studies, while Panel B considers steady-state estimates from studies that exploit cross-sectional variation across countries or individuals. Two results emerge from this analysis.

First, the bounds on extensive margin elasticities are much tighter than those on the intensive margin, as shown in Figure 10. For instance, Eissa and Liebman's (1996) analysis

²³In these calculations, I assume that the marginal worker is in the labor force in cases where the average tax rates rises over the three years and out of the labor force for cases where it falls. This is the relevant calculation to determine when one would observe zero response on the extensive margin, as shown above. I exclude the 99.5 percentile from Figure 6b for scaling reasons and because few individuals enter the labor force at the 99.5 percentile of the income distribution.

²⁴Among the studies considered by Chetty et al. (2011), I include only those that estimate steady-state elasticities and for which I was able to compute the size of the tax change used for identification.

of EITC expansions yields $\hat{\eta} = 0.30$ and bounds on η of $(0.26, 0.36)$ with $\delta = 1\%$ frictions.²⁵ Observed labor supply elasticities appear to provide reasonably accurate estimates of structural elasticities on the extensive margin.

Second, the heterogeneity in extensive margin elasticities across groups cannot be attributed purely to frictions. The minimum level of frictions required to reconcile the extensive margin elasticities in Table II is $\delta_{\min} = 18\%$. Hence, there are economically significant differences in structural extensive margin elasticities across groups. For instance, it is plausible that participation elasticities are especially large for low-income single mothers (Meyer and Rosenbaum 2001). There may be similar heterogeneity in structural elasticities on the intensive margin, but existing evidence is inadequate to detect such heterogeneity in the presence of small frictions.

The results in Tables I and II challenge the commonly held view that extensive margin elasticities are larger than intensive margin elasticities. This view underpins some important results in modern optimal tax theory, such as providing a rationale for programs such as the Earned Income Tax Credit (Saez 2002). The analysis here suggests that estimated intensive margin elasticities may be smaller than extensive elasticities simply because of frictions. In steady state, the intensive elasticity may actually be larger than extensive elasticities, reducing the welfare gains from programs such as the EITC.

4.3 Bunching at Kinks and Non-Linear Budget Set Models

The preceding two sections have considered studies that analyze local changes in marginal tax rates without fully modelling each agent’s budget set in a progressive tax system. Another important strand of the literature on labor supply accounts for the entire tax system by estimating non-linear budget set (NLBS) models of labor supply. Frictions can also explain various patterns in the non-linear budget set literature.

(i) Bunching at Kinks. Frictionless NLBS models are rejected by the data because they predict much more bunching at kinks than observed in practice. This is illustrated in Figure 11, which plots the income tax schedule in 2006 (dashed blue line) for a single filer with two children. The solid grey curve shows the income distribution predicted by the frictionless

²⁵The level of frictions may differ on the extensive and intensive margins. However, frictions would have to be 10 times larger on the extensive margin in order to generate the same impacts as on the intensive margin.

model in (21) when $\varepsilon = 0.5$ and tastes a_i are uniformly distributed. The frictionless model predicts sharp spikes (mass points) at each kink in a kernel density plot of the income distribution. However, empirical income distributions for wage earners exhibit no such bunching at kinks (Saez 2010). Small frictions in choosing labor supply can explain why bunching is not more prevalent. The number next to each convex kink in Figure 11 shows the utility gain as a percentage of consumption (calculated using the utility in (21) from locating at that kink point relative to optimizing under the incorrect assumption that the rate in the previous bracket continues into the next bracket.²⁶ The utility losses are less than 1% of net earnings at most of the kinks.

The traditional solutions used to deal with the lack of bunching at kinks when fitting NLBS models are to introduce optimization errors that smooth the income distribution around the kink (e.g. Hausman 1981, Blomquist and Hansson-Brusewitz 1990) or to smooth the budget set itself (MaCurdy, Green, and Paarsche 1990). The approach proposed here – permitting agents to deviate systematically from their frictionless optima provided that the utility losses fall below some threshold – places more structure on the nature of these optimization errors and could thereby improve identification.

(ii) Bunching among the Self-Employed. Saez (2010) documents that unlike wage earners, self-employed individuals bunch at the first kink of the EITC schedule, where tax refunds are maximized. Audit studies show that self employment income is frequently misreported on tax returns because of the lack of double reporting. Unlike changing actual hours of work, misreporting generates a first-order utility gain because it transfers resources from the government to the taxpayer. The large utility gains from misreporting taxable income could explain why the self-employed overcome frictions and bunch at this kink.²⁷

(iii) Notches. Unlike kinks, notches in budget sets, where a \$1 change in earnings leads to a discontinuous jump in consumption, generate substantial behavioral responses. For example, income cutoffs to qualify for Medicaid (Yelowitz 1995) and social security benefits in some pension systems (Gruber and Wise 1999) induce sharp reductions in labor supply. To calculate

²⁶There are many values of a_i that can induce individuals to locate at each kink. The numbers in the figure are (unweighted) mean percentage losses for agents who would optimally locate at the kink.

²⁷Even the self-employed do not bunch at the second kink of the EITC schedule (where the phase-out region begins). The first kink in the EITC schedule maximizes the size of the EITC refund while minimizing payroll tax liabilities. There is no reason to locate at the second kink if one's goal is to reap first-order gains from income manipulation.

the utility cost of ignoring a notch, suppose that earning $wl_t > K$ triggers a penalty of P . Then the utility cost of setting $l_t > K/w$ for an individual with $l_t^* \leq K/w$ exceeds P . Because the utility cost of ignoring a notch increases at a first-order rate with the size of the penalty P , notches affect observed behavior substantially even with frictions. Notches are therefore a promising source of variation for identification of structural elasticities.

4.4 Micro vs. Macro Elasticities

The final strand of the literature I consider is the debate on micro vs. macro labor supply elasticities. Macroeconomic models calibrate labor supply elasticities to match the variation in aggregate hours of work across countries with different tax systems or over the business cycle. In both cases, macro calibrations of representative agent models imply larger elasticities than microeconomic estimates of intensive margin elasticities. Can frictions explain this gap?

The macro literature uses the term “macro elasticity” to refer to the Frisch elasticity of aggregate hours and “micro elasticity” to refer to the intensive-margin elasticity of hours conditional on employment (e.g. Prescott 2004, Rogerson and Wallenius 2009). I instead use the terms “micro” and “macro” to refer to the source of variation used to estimate the elasticity, for two reasons. First, both intensive and extensive margin responses are determined by microeconomic household-level choices. Second, the Frisch (marginal utility constant) elasticity is important for understanding business cycle fluctuations, but does not control the steady-state impacts of differences in taxes across countries. The Frisch elasticity determines intertemporal substitution responses to temporary wage fluctuations, while the Hicksian (wealth constant) elasticity controls steady-state responses and the efficiency costs of taxes (MaCurdy 1981, Auerbach 1985).²⁸ I first compare micro and macro estimates of Hicksian elasticities and then turn to Frisch elasticities.

Cross-Country Evidence and Hicksian Elasticities. The mean estimate of the intensive margin Hicksian elasticity from the two macroeconomic studies in Table I (Prescott 2004, Davis and Henrekson 2005) is 0.33.²⁹ The mean estimate of the extensive margin Hicksian

²⁸Chetty et al. (2011b) discuss these elasticity concepts in greater detail and show that some discrepancies across studies arise simply from differences in terminology.

²⁹The well known elasticity of 3 reported by Prescott (2004) is a Frisch elasticity. Regressing log hours on log net-of-tax rates using Prescott’s data yields a Hicksian aggregate hours elasticity of 0.7 and an intensive elasticity of 0.46. Prescott translates the Hicksian elasticity of 0.7 into a Frisch elasticity of 3 based on specific functional form assumptions about utility.

elasticity from the three macroeconomic studies in Table II (Nickell 2003, Prescott 2004, Davis and Henrekson 2005) is 0.17. Hence, macro cross-country evidence implies an aggregate hours elasticity of $0.33 + 0.17 = 0.5$.

These macro elasticity estimates are consistent with micro estimates once one accounts for optimization frictions. On the intensive margin, even the smallest estimates in Table I are consistent with a structural elasticity of 0.33 with $\delta = 1\%$ frictions. The minimum- δ micro estimate of $\varepsilon = 0.33$ coincides exactly with the macro intensive elasticity. Intuitively, macroeconomic comparisons are more likely to overcome frictions because they analyze steady-state behavior and because they induce coordinated changes in work patterns (Altonji and Oldham 2003, Chetty et al. 2011a).

On the extensive margin, the observed micro estimates in Panel A of Table II are similar to the macro and cross-sectional estimates in Panel B even without accounting for frictions. The mean micro estimate of η is 0.25. The similarity between micro and macro estimates of extensive margin elasticities is consistent with the prediction that frictions have little impact on extensive margin responses.

I conclude that both micro and macro evidence imply steady-state aggregate hours elasticities of approximately 0.5 once one accounts for frictions and indivisible labor. Indivisible labor models show that both intensive and extensive margins are important in accounting for aggregate hours differences (Rogerson 1988, Ljungqvist and Sargent 2006, Rogerson and Wallenius 2009). Frictions explain why micro estimates of steady-state elasticities are smaller than macro estimates on the intensive margin but are similar on the extensive margin.

Intertemporal Substitution and Frisch Elasticities. Equilibrium macro models – in which fluctuations in labor supply are driven by preferences – require intensive margin Frisch elasticities of about 0.5 and extensive margin Frisch elasticities above 2 to fit observed fluctuations in employment and hours over the business cycle (Chetty et al. 2011b). The analysis in the present paper does not directly tell us whether micro evidence is consistent with these values because it bounds the Hicksian rather than the Frisch elasticity. Chetty et al. (2011b) summarize micro estimates of the Frisch elasticity. Here, I instead show that one can obtain tight bounds on the structural intensive margin Frisch elasticity from the estimated structural Hicksian elasticity of $\varepsilon = 0.33$.

In the lifecycle labor supply model in (19), the intensive margin Frisch elasticity ε^F is re-

lated to the intensive margin Hicksian elasticity by the following equation (Ziliak and Kniesner 1999, Browning 2005):

$$\varepsilon^F = \varepsilon + \rho \left(\frac{d[wl_{i,t}^*]}{dY_{i,t}} \right)^2 \frac{A_{i,t}}{wl_{i,t}^*} \quad (22)$$

where ρ is the elasticity of intertemporal substitution (EIS), $\frac{d[wl_{i,t}^*]}{dY_{i,t}}$ measures the marginal propensity to earn out of unearned income (the income effect), and $\frac{A_{i,t}}{wl_{i,t}^*}$ is the ratio of assets to wage income. This equation implies $\varepsilon^F > \varepsilon$.

One can obtain more information about ε^F by calibrating the other parameters in (22). The ratio of assets to wage earnings was approximately $\frac{A_{i,t}}{wl_{i,t}^*} = 1.26$ for the median individual in the U.S. in 2008 (Dynan 2009, Table 1). Table III shows the values of the Frisch elasticity implied by a Hicksian elasticity of $\varepsilon = 0.33$ and $\frac{A_{i,t}}{wl_{i,t}^*} = 1.26$ for various combinations of ρ and $-\frac{d[wl_{i,t}^*]}{dY_{i,t}}$. To calibrate these two parameters, note that balanced growth requires that income and substitution effects cancel, implying $\frac{d[wl_{i,t}^*]}{dY_{i,t}} = -\varepsilon \Rightarrow \frac{d[wl_{i,t}^*]}{dY_{i,t}} = -0.33$. Both micro and macro studies find an EIS of $\rho \leq 1$ (Hall 1988, Vissing-Jorgensen 2002, Guvenen 2006). The largest Frisch elasticity consistent with these parameters is $\varepsilon^F = 0.47$. Intuitively, the Frisch elasticity cannot be much larger than the Hicksian elasticity for plausible values of the income effect because $\varepsilon^F - \varepsilon$ is proportional to the income effect squared and the ratio of assets to earnings is not very high for most households in the U.S.

An intensive margin Frisch elasticity of 0.47 is roughly consistent with the macro evidence on business cycle fluctuations in hours of work conditional on employment.³⁰ However, Chetty et al. (2011b) find that fluctuations in employment rates over the business cycle imply Frisch elasticities that are an order of magnitude larger than micro estimates. Unfortunately, this discrepancy between micro and macro estimates of the Frisch elasticity on the extensive margin cannot be explained by optimization frictions.

5 Conclusion

There are many frictions that induce agents to deviate from the optimal choices predicted by standard economic models. This paper has shown that the model mis-specification that

³⁰An interesting question left for future work is whether the structural or observed elasticity is more relevant for business cycle fluctuations. If there are small fluctuations in wage rates over the business cycle, the observed elasticity (attenuated by frictions) may be the better predictor of behavioral responses. But if a small group of individuals face very large wage shocks, then frictions may be overcome and the structural elasticity may be more relevant.

arises from the omission of these frictions can be handled using the tools of set identification. Abstractly, I exchange the standard orthogonality condition on the error term for a bounded support condition based on the utility costs of errors. I derive an analytical representation for bounds on structural price elasticities that is a function of the observed elasticity, size of the price change used for identification, and the degree of optimization frictions.

Applying the bounds to studies of taxation and labor supply offers a critique and synthesis of this literature. The critique is that many microeconomic studies of labor supply are uninformative about intensive margin elasticities because they cannot reject large values of ε with frictions of even 1% of earnings in choosing labor supply. The synthesis is that several patterns in this literature can be reconciled by allowing for such small frictions. Combining estimates from several studies, my preferred point estimates of structural Hicksian elasticities are 0.33 on the intensive margin, 0.25 on the extensive margin, and roughly 0.5 for aggregate hours. I also find that Frisch elasticities cannot be much larger than Hicksian elasticities given plausible income effects.

Both the methodology and application in this paper could be improved in several respects in future work. Methodologically, it is important to extend the bounds to settings beyond the binary treatment effect estimator considered here. Natural extensions include difference-in-difference estimates and regression models that allow for continuous price variation. It would also be interesting to explore whether the bounds can be sharpened by imposing additional restrictions from theory, such as a requirement that agents converge to unconstrained optima over time. In the labor supply application, it would be useful to calculate bounds in modern structural models of labor supply that incorporate factors such as human capital accumulation, credit constraints, and uncertainty. Because full identification of these models is challenging, bounding the structural elasticity may be a particularly fruitful approach in such cases.

Finally, the bounding methodology developed here can be applied to estimate a variety of other critical parameters such as the elasticity of intertemporal substitution, the marginal propensity to consume out of income, or the effects of the minimum wage on employment. Such analyses would shed light on which disagreements are economically significant and which can be reconciled simply by allowing for small frictions.

References

- Akerlof, G. A. and J. L. Yellen (1985): "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?" *American Economic Review*, 75, 708-20.
- Alesina, A. F., E. L. Glaeser, and B. Sacerdote (2005): "Work and Leisure in the U.S. and Europe: Why So Different?" in: Kenneth Rogoff and Mark Gertler (eds.), *NBER Macroeconomics Annual 2005*, 1-100. Cambridge: MIT Press
- Altonji, J. G. and C. H. Paxson (1992): "Labor Supply, Hours Constraints, and Job Mobility," *Journal of Human Resources*, 27, 256-278.
- Altonji, J. G. and J. Oldham (2003): "Vacation laws and annual work hours," *Economic Perspectives*, Federal Reserve Bank of Chicago Q III, 19-29.
- Attanasio, O. P. (2000): "Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S,s) Rules for Automobile Purchases," *Review of Economic Studies*, 67, 667-696.
- Auten, G. and R. Carroll (1999): "The Effect of Income Taxes on Household Behavior," *Review of Economics and Statistics*, 81, 681-693.
- Auten, G. and R. Carroll (1994): "Taxpayer Behavior and the 1986 Tax Reform Act," paper presented at NBER Summer Institute, Cambridge, Mass., July 1994.
- Bianchi, M., B. R. Gudmundsson, and G. Zoega (2001): "Iceland's Natural Experiment in Supply-Side Economics," *American Economic Review*, 91, 1564-1579.
- Blau, F. D. and L. M. Kahn (2000): "Gender Differences in Play," *The Journal of Economic Perspectives*, 14, 75-99.
- Blau, F. D. and L. M. Kahn (2007): "Changes in the Labor Supply Behavior of Married Women: 1980-2000," *Journal of Labor Economics*, 25, 393-438.
- Blomquist, S. and U. Hansson-Brusewitz (1990): "The Effect of Taxes on Male and Female Labor Supply in Sweden," *Journal of Human Resources*, 25, 317-357.
- Blundell, R., A. Bozio, and G. Laroque (2011): "Labour Supply Responses and the Extensive Margin: The US, UK and France," University College London working paper.
- Blundell, R., A. Duncan, and C. Meghir (1998): "Estimating Labor Supply Responses Using Tax Reforms," *Econometrica*, 66, 827-862.
- Browning, M. (2005): "A working paper from April 1985: Which demand elasticities do we know and which do we need to know for policy analysis?" *Research in Economics*, 59, 293-320.
- Card, D. and J. E. DiNardo (2002): "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, 20, 733-783.
- Chernozhukov, V., H. Hong, and E. Tamer (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75, 1243-1284.
- Chetty, R., A. Looney and K. Kroft (2009): "Salience and Taxation: Theory and Evidence," *American Economic Review*, 9, 1145-1177.
- Chetty, R. and E. Saez (2009): "Teaching the Tax Code: Earnings Responses to a Field Experiment with EITC Recipients," NBER Working Paper No. 14836.

- Chetty, R., J. Friedman, T. Olsen, and L. Pistaferri (2011a): “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *Quarterly Journal of Economics*, 126, 749–804.
- Chetty, R., A. Guren, D. Manoli, A. Weber (2011b): “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities,” NBER Working Paper No. 16729.
- Cochrane, J. (1989): “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *American Economic Review*, 79, 319-337.
- Davis, S. and M. Henrekson (2005): “Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich-Country Comparisons,” in R. G. Salvador, A. Lamo, B. Petrongolo, M. Ward, and E. Wasmer (eds), *Labour Supply and Incentives to Work in Europe*, Edward Elgar Press.
- DellaVigna, S. (2009): “Psychology and Economics: Evidence from The Field,” *Journal of Economic Literature*, 47, 315-372.
- Dynan, K. (2009): “Changing Household Financial Opportunities and Economic Security,” *Journal of Economic Perspectives*, 23, 49-68.
- Eberly, J. C. (1994): “Adjustment of Consumers’ Durables Stocks: Evidence from Automobile Purchases,” *Journal of Political Economy*, 102, 403-436.
- Eissa, N. and H. Hoynes (1998): “The Earned Income Tax Credit and the Labor Supply of Married Couples,” NBER Working Paper 6856.
- Eissa, N. and H. Hoynes (2004): “Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit,” *Journal of Public Economics*, 88, 1931-1958.
- Eissa, N. and H. Hoynes (2006): “Behavioral Responses to Taxes: Lessons from the EITC and Labor Supply,” *Tax Policy and the Economy*, 20, 74-110.
- Eissa, N. and J. B. Liebman (1996): “Labor Supply Response to the Earned Income Tax Credit,” *Quarterly Journal of Economics*, 111, 605-37.
- Feldstein, M. (1995): “The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act,” *Journal of Political Economy*, 103, 551-57.
- Fujii, E. T. and C. B. Hawley (1988): “On the Accuracy of Tax Perceptions,” *Review of Economics and Statistics*, 70, 344-347.
- Gelber, A. M. (2010): “Taxation and the Earnings of Husbands and Wives,” University of Pennsylvania working paper.
- Goolsbee, A. (1999): “Evidence on the High-Income Laffer Curve From Six Decades of Tax Reform,” *Brookings Papers on Economic Activity*, 2, 1-47.
- Graversen, E. K. (1998): “Labor Supply and Work Incentives,” PhD diss., University of Aarhus.
- Grossman, S. J. and G. Laroque (1990): “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 58, 25-51.
- Gruber, J. and D. A. Wise (1999): *Social Security and Retirement Around the World*, National Bureau of Economic Research, Inc.

- Gruber, J. and E. Saez (2002): “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, 84, 1–32.
- Guvenen, F. (2006): “Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective,” *Journal of Monetary Economics*, 53, 1451-1472.
- Hall, R. E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96, 339-357.
- Hansen, L. P. and T. J. Sargent (2007): *Robustness*, Princeton: Princeton University Press.
- Hausman, J. A. (1981): “Labor Supply,” in: Henry J. Aaron and Joseph A. Pechman (eds.), *How Taxes Affect Economic Behavior*, 27-72. Washington, DC: The Brookings Institution.
- Heckman, J. (1993): “What Has Been Learned About Labor Supply in the Past Twenty Years?” *American Economic Review*, 83, 116-21.
- Helms, L. J. (1985): “Expected Consumer’s Surplus and the Welfare Effects of Price Stabilization,” *International Economic Review*, 26, 603-17.
- Imbens, G. W., and C. F. Manski (2004): “Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 72, 1845-1857.
- Imbens, G. W., D. B. Rubin, and B. Sacerdote (1999): “Estimating the Effect of Unearned Income on Labor Supply, Earnings, Savings, and Consumption: Evidence From a Survey of Lottery Players,” NBER Working Paper No. 14836.
- Imbens, G. W., D. B. Rubin, and B. Sacerdote (2001): “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players,” *American Economic Review*, 91, 778-794.
- Jacob, B. A. and J. Ludwig (2008): “The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery,” NBER Working Paper No. 14570.
- Jones, A. D. (2008): “Information, Inertia, and Public Benefit Participation: Experimental Evidence from the Advance EITC,” UC-Berkeley working paper.
- Kopczuk, W. (2010): “Polish Business Flat Tax and its Effect on Reported Incomes: a Pareto Improving Tax Reform?” Columbia University working paper.
- Liebman, J., and E. Saez (2006): “Earnings Responses to Increases in Payroll Taxes,” UC-Berkeley working paper.
- Ljungqvist, L. and T. J. Sargent (2006): “Do Taxes Explain European Employment? Indivisible Labor, Human Capital, Lotteries, and Savings,” in: D. Acemoglu, K. Rogoff and M. Woodford (eds.), *NBER Macroeconomics Annual 2006*, 181-246. Cambridge: MIT Press.
- Low, H., C. Meghir, and L. Pistaferri (2010): “Wage Risk and Employment Risk Over the Life Cycle,” *American Economic Review*, 100, 1432-1467.
- MaCurdy, T. (1981): “An Empirical Model of Labor Supply in a Life-Cycle Setting,” *Journal of Political Economy*, 89, 1059-1085.
- MaCurdy, T., D. Green and H. Paarsche (1990): “Assessing Empirical Approaches for Analyzing Taxes and Labor Supply,” *Journal of Human Resources*, 25, 415-490.
- Mankiw, N. G. (1985): “Small Menu Costs and Large Business Cycles: A Macroeconomic

- Model,” *Quarterly Journal of Economics*, 100, 529–38.
- Manski, C. F. (2007): *Identification for Prediction and Decision*, Cambridge: Harvard University Press.
- Meyer, B. and D. Rosenbaum (2000): “Making Single Mothers Work: Recent Tax and Welfare Policy and its Effects,” *National Tax Journal*, 53, 1027-59.
- Meyer, B. and D. Rosenbaum (2001): “Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers,” *Quarterly Journal of Economics*, 116, 1063-1114.
- Nevo, A. and A. Rosen (2008): “Identification with Imperfect Instruments,” NBER Working Paper No. 14434.
- Nickell, S. (2003): “Employment and Taxes,” CESifo Working Paper No. 1109.
- Prescott, E. (2004): “Why Do Americans Work So Much More Than Europeans?” *Federal Reserve Bank of Minneapolis Quarterly Review*, 28, 2–13.
- Rogerson, R. (1988): “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 21, 3–16.
- Rogerson, R. and J. Wallenius (2009): “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory*, 144, 2277-2292.
- Rothstein, J. (2008): “The Unintended Consequences of Encouraging Work: Tax Incidence and the EITC,” CEPS Working Paper No. 165.
- Saez, E. (2002): “Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses,” *Quarterly Journal of Economics*, 117, 1039-1073.
- Saez, E. (2004): “Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications,” *Tax Policy and the Economy*, 18, 117-174. Cambridge: MIT Press.
- Saez, E. (2010): “Do Taxpayers Bunch at Kink Points?” *American Economic Journal: Economic Policy*, 2, 180-212.
- Saez, E., J. Slemrod, and S. Giertz (2011): “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” forthcoming, *Journal of Economic Literature*.
- Slemrod, J. B. (1995): “Income Creation or Income Shifting? Behavioral Responses to the Tax Reform Act of 1986,” *American Economic Review*, 82, 175-180.
- Vissing-Jorgensen, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110, 825-853.
- Yelowitz, A. S. (1995): “The Medicaid Notch, Labor Supply, and Welfare Participation: Evidence from Eligibility Expansions,” *Quarterly Journal of Economics*, 110, 909-39.
- Ziliak, J. P. and T. J. Kniesner (1999): “Estimating Life Cycle Labor Supply Tax Effects,” *Journal of Political Economy*, 107, 326-359.

Appendix A: Theoretical Derivations

(i) Bounds on Intensive Margin Elasticities with Income Effects and Stochastic Prices

This section establishes two results. First, the bounds in Proposition 1 apply to the Hicksian elasticity when the quasilinearity assumption in (2) is relaxed. Second, allowing for stochastic prices p_t does not affect the bounds. To simplify notation, I ignore heterogeneity across agents and assume all agents have a flow utility function $v(x_t, y_t)$. Heterogeneity does not affect the result under the assumption that the structural elasticity does not vary locally across agents, as discussed below.

Let \mathbb{E}_t denote the conditional expectation operator over prices given information available in period t and $p = (p_1, \dots, p_T)$ the realized price vector. To account for stochastic prices, I redefine the nominal model so that the agent maximizes expected lifetime utility

$$\mathbb{E}_t \sum_{s=t}^T v(x_s, y_s) \quad (23)$$

subject to the dynamic budget constraint $Z_{t+1} = Z_t - p_t x_t - y_t$ and the terminal condition $Z_{T+1} = 0$.

Let $V_t(p, Z_t) = \sum_{s=t}^T v(x_s(p), y_s(p))$ denote the utility the agent attains from periods t to T with a realized price vector of p and wealth Z_t . Following Helms (1985), I define the agent's expenditure function with stochastic prices as the minimum wealth required to attain expected utility above a given threshold U . The agent's partial expenditure function (on all other goods) conditional on consuming \tilde{x}_t units of good x_t in period t is

$$\tilde{e}(\tilde{x}_t, U) = \min_Z Z - p_t \tilde{x}_t \text{ s.t. } \mathbb{E}_t V_t(p, Z) \geq U \text{ and } x_t = \tilde{x}_t$$

and hence the total expenditure function can be written as

$$E(p_t, U) = \min_{x_t} p_t x_t + \tilde{e}(x_t). \quad (24)$$

Let the expenditure-minimizing choice of x_t be denoted by $x_t^{c,*}(p_t, U_t)$, the structural Hicksian demand function under the nominal model in (23). Let $x_t^c(p_t, U_t)$ denote the observed Hicksian demand function with frictions. Let $\varepsilon(p_t) = -\frac{\partial x_t^{c,*}}{\partial p_t} \frac{p_t}{x_t^c}$ denote the structural Hicksian price elasticity of demand at price p_t . When utility is not quasilinear, identifying $\varepsilon(p_t)$ requires variation in prices *within* period t because price changes across periods conflate the Frisch and Hicksian elasticities (MaCurdy 1981). Consider an experiment in which some agents face a price of p_A and others face a price of p_B in period t and let

$$\hat{\varepsilon}(p_A, p_B) = -\frac{\log x_B^c(p_B) - \log x_A^c(p_A)}{\log(p_B) - \log(p_A)}$$

denote the observed elasticity from this experiment. Our objective is to identify $\varepsilon(p_t)$ from estimates of $\hat{\varepsilon}$ in an environment with frictions.

In this setting, the δ class of models is defined by the condition:

$$[p_t x_t^c + \tilde{e}(x_t^c)] - [p_t x_t^{c,*} + \tilde{e}(x_t^{c,*})] \leq \delta p_t x_t^{c,*} \quad (25)$$

I first establish an analog of Lemma 1 to characterize the choice set with frictions.

Lemma A1. For small δ , the set of observed Hicksian demands is approximately

$$X_t^c(p_t, \delta) = \{x_t^c : |\log x_t^c - \log x_t^{c,*}| \leq [2\varepsilon(p_t)\delta]^{1/2}\} \quad (26)$$

Proof. The first order condition for (24) is

$$\tilde{e}_x(x_t^{c,*}) = -p_t \quad (27)$$

Using a quadratic approximation to the partial expenditure function, we can exploit this first order condition to obtain

$$[p_t x_t^c + \tilde{e}(x_t^c)] - [p_t x_t^{c,*} + \tilde{e}(x_t^{c,*})] \simeq \frac{1}{2}(x_t^{c,*})^2 (\log x_t^c - \log x_t^{c,*})^2 \tilde{e}_{xx}(x_t^{c,*})$$

and hence we can rewrite (25) as

$$|\log x_t^c - \log x_t^{c,*}| \leq [2\delta \frac{p_t}{x_t^{c,*}} \frac{1}{\tilde{e}_{xx}(x_t^{c,*})}]^{1/2} \quad (28)$$

Differentiating (27) with respect to p_t implies $1/\tilde{e}_{xx}(x_t^{c,*}) = -\frac{\partial x_t^{c,*}}{\partial p_t}$ and substituting this equation into (28) completes the proof.

Next, I establish the analog of Proposition 1. When utility is not quasilinear, the structural elasticity $\varepsilon(p_t)$ varies with the price p_t . Let $\varepsilon(p_A)$ and $\varepsilon(p_B)$ denote the structural point elasticities at the initial and final prices and $\varepsilon(p_A, p_B) = -\frac{\log x_B^{c,*}(p_B) - \log x_A^{c,*}(p_A)}{\log(p_B) - \log(p_A)}$ denote the structural arc elasticity between the two prices. Then the upper bound on $\varepsilon(p_A, p_B)$ is characterized by an equation analogous to (12):

$$\widehat{\varepsilon}(p_A, p_B) = -\frac{\log x_B^c(p_B) - \log x_A^c(p_A)}{\log(p_B) - \log(p_A)} = \varepsilon(p_A, p_B) - \frac{2(2\varepsilon(p_B)\delta)^{1/2}}{\Delta \log p}$$

Solving this equation requires a parametric assumption about utility to relate the two point elasticities at p_A and p_B to the arc elasticity. I make the following local iso-elasticity assumption, which is analogous to assumption A2 in the extensive margin case.

A2' The structural Hicksian elasticity is constant between p_A and p_B : $\varepsilon(p_t) = -\frac{\partial x_t^{c,*}}{\partial p_t} \frac{p_t}{x_t^{c,*}} = \varepsilon(p_A, p_B)$ for $p_t \in [p_A, p_B]$

Under A2', the upper and lower bounds on the structural arc elasticity $\varepsilon(p_A, p_B)$ are characterized by the same equations as (12) and (13):

$$\widehat{\varepsilon} = \varepsilon \pm 2 \frac{(2\varepsilon\delta)^{1/2}}{\Delta \log p}$$

Proposition A1. Under assumption A2', for small δ , the range of structural Hicksian elasticities $\varepsilon(p_A, p_B)$ consistent with an observed Hicksian elasticity $\widehat{\varepsilon}(p_A, p_B)$ is approximately

$(\varepsilon_L, \varepsilon_U)$ where

$$\begin{aligned} \varepsilon_L &= \widehat{\varepsilon} + \frac{4\delta}{(\Delta \log p)^2}(1 - \rho) \text{ and } \varepsilon_U = \widehat{\varepsilon} + \frac{4\delta}{(\Delta \log p)^2}(1 + \rho) \\ \text{with } \rho &= \left(1 + \frac{1}{2} \frac{\widehat{\varepsilon}(p_t)}{\delta} (\Delta \log p)^2\right)^{1/2} \end{aligned}$$

Proof. Identical to the proof of Proposition 1.

In a model with heterogeneous utilities $v_i(x_s, y_s)$, Proposition A1 requires a stronger iso-elasticity assumption, namely that the structural elasticity $\varepsilon(p_t)$ does not vary across agents between p_A and p_B . It also requires an assumption analogous to A1, i.e. that tastes are orthogonal to the price change used for identification.

(ii) Bounds on Extensive Margin Elasticities

With quasilinear utility, the agent's flow utility in period t is $v_{i,t}(x, y) = y + b_{i,t}x$. Recognizing that the consumption path of y does not affect lifetime utility, the flow utility cost of choosing x suboptimally in period t is

$$u_{i,t}(x^*(p_t)) - u_{i,t}(x) = (x_{i,t}^* - x)(b_{i,t} - p_t)$$

I define a δ class of models around the nominal model by a condition analogous to (7):

$$(x_{i,t}^* - x)(b_{i,t} - p_t) \leq \delta_i p_t \text{ and } \frac{1}{N} \sum_i \delta_{i,t} \leq \delta \text{ and } F(b_{i,t} | \delta_{i,t}) = F(b_{i,t}) \quad (29)$$

The last condition in (29) – that the taste distribution cannot vary across agents with different frictions – is needed to ensure that the choice set has the same width for the marginal agents at each level of p .³¹ This condition was not necessary in the intensive margin case because the marginal agent did not vary with p there.

Proof of Lemma 2. Equation (29) implies that agent i 's observed demand for x is

$$x_{i,t} = \begin{cases} 1 & \text{if } b_{i,t} - p_t > \delta_{i,t} p_t \\ \{0, 1\} & \text{if } |b_{i,t} - p_t| \leq \delta_{i,t} p_t \\ 0 & \text{if } b_{i,t} - p_t < -\delta_{i,t} p_t \end{cases}$$

Let $\theta_{\delta_{i,t}}(p_t)$ denote the observed participation rate for agents who have frictions $\delta_{i,t}$ and $\theta_t = \mathbb{E}\theta_{\delta_{i,t}}(p_t)$ denote the observed participation rate in the aggregate economy. Under the condition that $F(b_{i,t} | \delta_{i,t}) = F(b_{i,t})$, it follows that $\theta_{\delta_{i,t}}(p_t)$ lies in the set:

$$\begin{aligned} & [1 - F((1 + \delta_{i,t})p_t), 1 - F((1 - \delta_{i,t})p_t)] \\ &= [\theta_t^* + F(p_t) - F((1 + \delta_{i,t})p_t), \theta_t^* + F(p_t) - F((1 - \delta_{i,t})p_t)] \\ &\simeq [\theta_t^* - f(p_t)p_t\delta_{i,t}, \theta_t^* + f(p_t)p_t\delta_{i,t}] \end{aligned}$$

³¹To see why this condition is needed, suppose agents with $b_{i,t}$ close to p_t have very large $\delta_{i,t}$ while those away from the margin have $\delta_{i,t} = 0$. This would result in a wide choice set for the participation rate at p_t even if $\mathbb{E}\delta_{i,t} < \delta$.

where the last line uses a first-order Taylor expansion of $F(p_t)$ around p_t . Under A1' and A2', $\eta = -\frac{d \log[1-F(p_t)]}{d \log p_t} = \frac{f(p_t)}{\theta^*(p_t)} p_t$. Hence

$$\begin{aligned} \theta_{\delta_{i,t}}(p_t) &\in [\theta_t^* \cdot (1 - \eta \delta_{i,t}), \theta_t^* \cdot (1 + \eta \delta_{i,t})] \\ &\Rightarrow \mathbb{E} \theta_{\delta_{i,t}}(p_t) \in [\theta_t^* \cdot (1 - \eta \mathbb{E} \delta_{i,t}), \theta_t^* \cdot (1 + \eta \mathbb{E} \delta_{i,t})] \\ &\Rightarrow \theta_t(p_t) / \theta_t^*(p_t) \in [1 - \eta \delta, 1 + \eta \delta] \end{aligned}$$

The approximation $\log(1 + \eta \delta) \simeq \eta \delta$ for small δ yields $|\log \theta_t - \log \theta_t^*| \leq \eta \delta$.

Proof of Proposition 2. Given a structural elasticity η , the maximal observed response to a price change of $\Delta \log p$ is $\Delta \log \theta = \eta \Delta \log p + 2\delta \eta$ and the minimal observed response is $\Delta \log \theta = \eta \Delta \log p - 2\delta \eta$. Therefore the observed elasticity $\hat{\eta} = \frac{\Delta \log \theta}{\Delta \log p}$ must satisfy

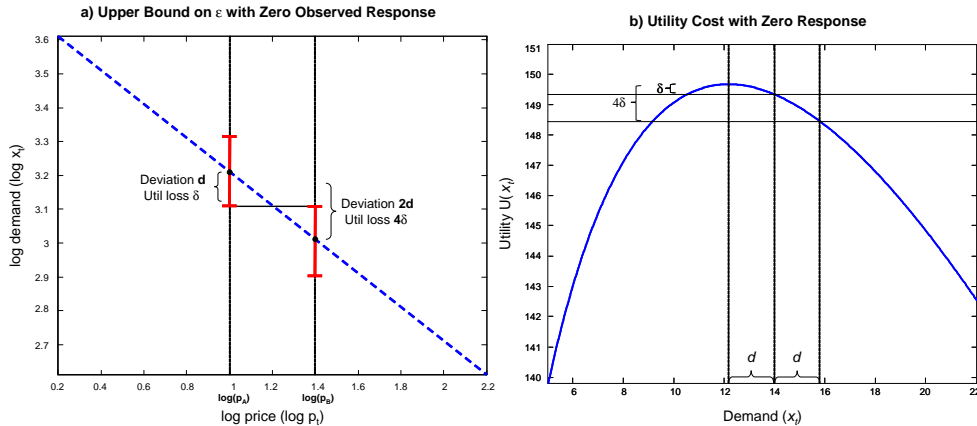
$$(1 - \rho_\eta) \eta \leq \hat{\eta} \leq (1 + \rho_\eta) \eta \quad (30)$$

where $\rho_\eta = \frac{2\delta}{\Delta \log p}$. If $\rho_\eta \geq 1$, η is unbounded above for a given value of $\hat{\eta}$ because both inequalities in (30) are satisfied for arbitrarily large η . If $\frac{2\delta}{\Delta \log p} < 1$, then the upper and lower bounds on η are obtained when (30) holds with equality. Solving these equations yields (16).

Proof of Corollary 2. Suppose $\hat{\eta} = 0$. Then $\rho_\eta < 1 \Rightarrow \eta_U = 0$. Hence a positive structural elasticity ($\eta > 0$) can only generate a 0 observed elasticity if $\rho_\eta = \frac{2\delta}{\Delta \log p} \geq 1 \Leftrightarrow \Delta u_{\text{ext},\%} = \Delta \log p \leq 2\delta$.

(iii) Intuition for 4δ threshold in Corollary 1

This section explains why $\Delta u_{\%}(\varepsilon)$ must be below 4δ in order to observe $\hat{\varepsilon} = 0$. Let $d = x_A^*(p_A) - \min(X_A(p_A, \delta))$ denote the difference between the mean optimal demand and the lowest mean demand in the initial choice set. Figure (a) below shows that at the upper bound ε_U , the difference between the optimal demands at the two prices is $x^*(p_A) - x^*(p_B) = 2d$. By definition, the percentage utility cost of choosing $\min(X_A(p_A, \delta))$ instead of $x^*(p_A)$ is δ . Given that the utility cost of deviating by d units is δ , the utility cost of deviating by $2d$ units is 4δ , as illustrated in Figure (b).



Appendix B: Sources and Calculations for Studies in Table I

This appendix describes how the values in columns 3-5 Table I are calculated. The papers used for the analysis along with comprehensive documentation of the calculations are available at http://obs.rc.fas.harvard.edu/chetty/bounds_opt_meta_analysis.zip

I use compensated intensive margin estimates reported in each paper when available and use the Slutsky equation to calculate compensated elasticities in cases where uncompensated elasticities are reported.

The studies do not always directly report the relevant inputs, especially the net-of-tax change $\Delta \log(1 - \tau)$. For studies whose estimates are identified from a single quasi-experiment (e.g. Feldstein 1995), I define $\Delta \log(1 - \tau)$ as the change in the marginal NTR for the group that the authors' define as the "treated" group. For studies that pool multiple tax or wage changes of different sizes and do not explicitly isolate a treatment group (e.g. Gruber and Saez 2002), I define $\Delta \log(1 - \tau)$ as twice the standard deviation of $\Delta \log(1 - MTR)$ in the sample. The logic for this approach is as follows. In a linear regression $Y_i = \alpha + \beta_1 X_i + u_i$, the standard error of $\hat{\beta}_1$ is the square root of $(var(u) / var(X)) / N$ where N is the sample size. Consider a second regression $Y_i = \alpha + \beta_2 Z_i + u_i$, where $Z_i = 0$ for half the observations (the "control group") and $Z_i = 2 \cdot SD(X)$ for the remaining observations (the "treatment group"). Setting the size of the single treatment to $2 \cdot SD(X)$ yields $var(Z) = var(X)$. Hence, the standard error of $\hat{\beta}_2$ equals the standard error of $\hat{\beta}_1$. A single tax change of $2 \cdot SD(\Delta \log(1 - MTR))$ therefore produces an estimate of $\hat{\varepsilon}$ with the same precision as the original variation in marginal tax rates used for identification.

I calculate the bounds by assuming that agents face a linear budget set whose slope is given by their marginal tax rate (MTR) and apply Proposition A1 using $\Delta \log(1 - MTR)$ in place of $\Delta \log p$. This yields valid bounds on ε for agents who remain in the interior of budget segments in a progressive tax system. However, the bounds cannot be applied to agents who locate at kinks. Given that most of the studies in Table I estimate elasticities from changes in the behavior of agents away from kinks, this is not a serious limitation.³²

The remainder of the appendix describes how I calculate $\hat{\varepsilon}$, the standard error of $\hat{\varepsilon}$, and $\Delta \log(1 - MTR)$ for each study in Table I.

A. Hours Elasticities

1. MaCurdy (1981): $\hat{\varepsilon}$: reported in text on page 1083. $s.e.(\hat{\varepsilon})$: imputed from the t-statistic for δ reported in row 5 of Table 1 as 0.15/0.98 because the estimate of compensated elasticity is approximately equal to δ . $\Delta \log(1 - \tau)$: the relevant within-person annual wage variation is not reported in the paper, so I use $2 \times SD = 2 \times (0.152^2 + 2 \cdot 0.086^2)^{1/2}$ from Table 1, column 4 of Low, Meghir and Pistaferri (2010) who estimate the standard deviation of changes in log wages. Note that this is likely an overestimate of the size of $\Delta \log(1 - \tau)$, resulting in bounds that are too tight, because MaCurdy uses family background characteristics, age, and year dummies as instruments for wage growth and does not use all elements of wage growth for identification.

³²Recent studies that identify observed elasticities from bunching at kinks (e.g. Saez 2009, Chetty et al. 2011) are an exception. I incorporate these studies into the linear-demand framework by exploiting the fact that they also study *movements* in the kinks over time, which create reductions in marginal rates for the subgroup of individuals located between the old and new bracket cutoffs. These studies imply that these individuals do not increase labor supply significantly when their marginal tax rates are lowered. This constitutes an observed elasticity estimate based on choices at interior optima, permitting application of Proposition 1.

2,3. Eissa and Hoynes (1998): $\hat{\varepsilon}$: authors report for men an intensive margin “wage elasticity” of 0.07 and an income elasticity of -0.03 in Table 8, col. 3. This “wage elasticity” uses the total hours change, which includes the hours change induced by the increased EITC rebate, which raised the *average* net of tax rate by 0.042 for a couple earning \$15,000 with two children (for whom the average net-of-tax rate changed from 107.5% in 1993 to 112.1% in 1994 computed using TAXSIM). This rebate should have changed hours (in log terms) by -0.03×0.042 , giving an uncompensated elasticity of 0.069. The compensated elasticity is $\hat{\varepsilon}_{l,w}^{men} = \hat{\varepsilon}_{l,w} - \frac{wl}{y}\hat{\varepsilon}_{l,y} = 0.200$, with w , l , and y from Table 3, column 4. A parallel calculation using Table 9 gives $\hat{\varepsilon}_{l,w}^{women} = 0.088$. $s.e.(\hat{\varepsilon})$: assume that w , l , y and the change in income from the EITC expansion are measured without error. Then, using the t-statistics from the coefficients on $\ln(\text{wage})$ and virtual inc to impute the standard errors for the elasticities, yields $SE(\hat{\varepsilon}_{l,w}^{men}) = \{SE(\hat{\varepsilon}_{l,w})^2 + [\frac{wl}{y}SE(\hat{\varepsilon}_{l,y})]^2\}^{1/2} = 0.074$ and $SE(\hat{\varepsilon}_{l,w}^{women}) = 0.067$. Note that this calculation is limited because the full variance-covariance matrix for the regression coefficients is not reported. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of log net-of-tax-rate in the phase out EITC rates listed in Table 1 for 1984–96 because most married couples that receive the EITC are in the phase-out region (Table 2).

4. Blundell, Duncan, and Meghir (1998): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 4, row 1. I interpret this estimate as an intensive margin elasticity because the variation in wages from the grouping estimator does not appear to affect participation, based on the discussion on page 845. $\Delta \log(1 - \tau)$: defined as $2 \times SD(\log \hat{w}_{gt} - \log \hat{w}_g - \log \hat{w}_t) = 0.23$ reported in Table 9 because the variation arises from group-time interactions in wages.

5. Ziliak and Kniesner (1999): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 1, column 3. $\Delta \log(1 - \tau)$: the study effectively uses within-person annual wage variation because lagged wage growth is included as an instrument. Since within-person annual wage variation is not reported in the paper, I again use $2 \times SD = 2 \times (0.152^2 + 2 \cdot 0.086^2)^{1/2}$ from Table 1, column 4 of Low, Meghir and Pistaferri (2010).

B. Taxable Income Elasticities

6. Bianchi, Gudmundsson, and Zoega (2001): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: average percent change in earnings for men and women weighted by observations (columns 1–4 of Table 6) divided by the percent change in the net-of-tax rate. Standard error computed from the standard errors reported for the changes in earnings. I interpret this estimate as an intensive margin elasticity because Table 6 conditions on work in 1986 and tax rates were generally lower in 1987 and 1988 than in 1986. I take this to be a compensated elasticity because Bianchi, Gudmundsson, and Zoega argue that income effects are small on page 1565–6, although this is somewhat tenuous. Note that the elasticity estimates provided by the authors are computed using *average* rather than marginal tax rates, necessitating use of the computation described above. $\Delta \log(1 - \tau)$: log change from tax rate of 0 in 1987 to 0.3875, which is an average of the flat tax in 1988 and the mean of the top marginal tax rate and bottom marginal tax rate in 1986 reported in Table 1 because the change in earnings estimate compares 1987 to the average earnings in 1986 and 1988.

7. Gruber and Saez (2002): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: average of the estimates in column 2 of Table 9 for individuals with taxable income between \$10,000 and \$50,000 and those with taxable income between \$50,000 and \$100,000. These estimates are compensated elasticities, as Gruber and Saez note on page 20 that income effects are essentially zero in their sample. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the change in log net-of-tax-rate and computed separately for columns 3

and 4 of Table 3 using the means and standard deviations for each year. The two estimates of $\Delta \log(1 - \tau)$ are then averaged in the same way as in the elasticity calculation described above.

8. Saez (2004): $\hat{\varepsilon}$, s.e.($\hat{\varepsilon}$): Table 7B, column 6 for the top 5 to 1 percent of tax units. Note that Saez uses gross income, not taxable income. I interpret his estimate as an intensive margin elasticity because his sample consists of repeated cross sections of workers and because the extensive margin is unlikely to be important for the top 5 to 1 percent of tax payers. I interpret this estimate as a compensated elasticity following the aforementioned evidence from Gruber and Saez (2002) that income effects are small. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the log net-of-tax-rate for the top 5 to 1 percent of tax units listed in column 8 of Table 5.

9. Jacob and Ludwig (2008): $\hat{\varepsilon}$: authors report that head of households' quarterly earnings conditional on working changed by \$228 from a control mean of \$5,558 in Table 3. As with Eissa and Hoynes, I calculate how much income would have changed absent the grant worth \$6,860 (page 9) in order to compute a compensated wage elasticity. Jacob and Ludwig do not report the effect of unearned income on earnings, so I use an estimate from Imbens, Rubin, and Sacerdote (2001), who report in Table 4, specification V, col. 1 a marginal propensity to earn out of unearned income (MPE) of -0.114 with a standard error of 0.015. In an earlier version, Imbens, Rubin, and Sacerdote (1999) reported earnings and participation elasticities of "around" -0.20 and -0.14 respectively, so I assume an intensive MPE of $\frac{d[wl]}{dY} = -0.114\{1 - (0.14/0.20)\} = -0.034$. On a quarterly basis, the grant should have lowered earnings by $-0.034 \cdot (6860/4) = 58.65$. Dividing the change in earnings absent the grant by the tax change gives an uncompensated elasticity of $\{\log(5558 - 228 + 58.65) - \log(5558)\} / \{\log(1) - \log(1 - 0.3)\} = 0.086$. Finally, the elasticity is $\hat{\varepsilon} = \hat{\varepsilon}^u - \frac{d[wl]}{dY} = 0.086 + 0.034 = 0.121$. s.e.($\hat{\varepsilon}$): Assuming that the standard error on the intensive MPE is proportional to the error on the total MPE, and that the change in income due to the grant is measured without error, then the standard error is 0.031. $\Delta \log(1 - \tau)$: MTR changed from 0 to 0.30 for those receiving the housing voucher as described in footnote 29 so that $\log(1) - \log(1 - 0.3) = 0.36$.

10, 11. Gelber (2010): $\hat{\varepsilon}$, s.e.($\hat{\varepsilon}$): Table 3, column 1 for men and column 2 for women. These estimates use earned income since it is less susceptible to manipulation than taxable labor income. These estimates presumably reflect primarily intensive margin responses since the extensive margin is unlikely to be important for the high income group affected by the change in top bracket tax rates. $\Delta \log(1 - \tau)$: percent change in net-of-tax rate from 1989 to 1991 for the highest tax brackets reported in Table 1.

12. Saez (2010): $\hat{\varepsilon}$, s.e.($\hat{\varepsilon}$): Table 2, row 1 of column 6 for wage earners with two or more children. $\Delta \log(1 - \tau)$: change in NTR at first kink in the EITC benefit schedule from 1995 to 2004.

13, 14. Chetty et al. (2011a): $\hat{\varepsilon}$, s.e.($\hat{\varepsilon}$): observed elasticities at middle and top kinks, calculated using equation 6 in Chetty et al. (2011a) as $b/K\Delta \log(1 - \tau)$. In this equation, K is the location of the tax bracket cutoff (DKr 164,300 for the middle tax and DKr 267,600 for the top tax). The estimated excess mass at the kink (b) is 1.79 (s.e. 0.05) for married women at the top kink (Figure IIIb) and 0.06 (s.e. 0.03) at the middle kink (Figure VIa). $\Delta \log(1 - \tau)$: size of tax changes at the middle and top tax kinks as reported in Figure II.

15. Chetty et al. (2011a): $\hat{\varepsilon}$, s.e.($\hat{\varepsilon}$): Table 2, col 1. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the changes in the log net-of-tax rate reported in the last row of Table 1, col 1.

C. Top Income Elasticities

16. Feldstein (1995): $\hat{\varepsilon}$: high minus medium tax rate specification in Table 2. For this and other studies based on TRA86, I follow the literature in interpreting elasticities as compensated elasticities because the reform was revenue neutral (Feldstein 1999). $s.e.(\hat{\varepsilon})$: not reported by Feldstein (1995). For a rough estimate, rescaling the standard error cited by Feldstein on page 566 for Auten and Carroll (1994) by the ratio of sample sizes in the two studies yields: $s.e.(\hat{\varepsilon}) = 0.15\sqrt{14425/3735} = 0.295$. $\Delta \log(1 - \tau)$: reported in Table 2 for the high tax rate group.

17. Auten and Carroll (1999): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 2, Col 6. $\Delta \log(1 - \tau)$: reported by Goolsbee (1999) for the highest income group in Table 3, row C for 1985 to 1989 because TRA86 “provided tax variation mostly at the top of the income scale, so that their overall estimates are identified primarily by reactions of high income taxpayers” (Gruber and Saez 2002, pg 24-25).

18. Goolsbee (1999): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 4, column 1. $\Delta \log(1 - \tau)$: Table 3, row C for 1985 to 1989 based on the quote above.

19. Saez (2004): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 3C, column 3 for the top 1 percent of tax units. Note that Saez uses gross income, not taxable income. I interpret his estimate as an intensive margin elasticity because his sample consists of repeated cross sections of workers and because the extensive margin is unlikely to be important for the top 1 percent of tax payers. I interpret this estimate as a compensated elasticity following the aforementioned evidence from Gruber and Saez (2002) that income effects are small. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the log net-of-tax-rate for the top 1 percent of tax units listed in column 3 of Table 5.

20. Kopczuk (2010): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: Table 9, second panel, col (1), 2002-2005, with standard error imputed from the reported t-statistic. This is a compensated elasticity following Gruber and Saez (2002) equation 2. $\Delta \log(1 - \tau)$: reported on page 17.

D. Macro/Cross-Sectional

21. Prescott (2004): $\hat{\varepsilon}$, $s.e.(\hat{\varepsilon})$: calculated by regressing log hours per worker on log net-of-tax rates using OECD data reported by Prescott in Table 2 on hours per adult, which are converted to hours per worker using labor force participation rates from OECD Stat Extracts.³³ The data on labor force participation rates are missing for Canada and the U.K. in the 1970’s and these observations are therefore excluded. The elasticity estimate can be interpreted as a compensated labor supply elasticity if government expenditure is viewed as unearned income in the aggregate. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the change in log net-of-tax rate for the 12 observations with non-missing data on hours per employed person.

22. Davis and Henrekson (2005): $\hat{\varepsilon}$: computed using log differences in annual hours per employed adult based on the slope coefficient in Table 2.3 (middle panel, Sample C) and the sample means of annual hours per employed person and tax rates in Table 2.1 for the corresponding sample. The elasticity estimate can be interpreted as a compensated labor supply elasticity if government expenditure is viewed as unearned income in the aggregate. $s.e.(\hat{\varepsilon})$: calculated from the standard error reported for the slope coefficient in Table 2.3 (middle panel, Sample C). $\Delta \log(1 - \tau)$: computed as $2 \times SD$ of log one minus sum of tax rates for the 19 countries in Sample C.³⁴

³³Data are for men and women aged 15-64 for 1970-1974 and 1993-1996 in order to match Prescott’s data. Data are available from OECD Stat Extracts at the following URL: http://stats.oecd.org/Index.aspx?DataSetCode=LFS_SEXAGE_I_R

³⁴Data are for 1995 for all countries except New Zealand and Australia for which I use 1986 and 1985 values

23. Blau and Kahn (2007): $\hat{\varepsilon}$: computed from intensive margin (with selection correction) elasticities reported in Table 6, defining the income elasticity as the elasticity of women’s hours with respect to husband’s wages and using the Slutsky equation to compute compensated elasticities in corresponding fashion. Mean values of wl and y are from Table A2 and A3. I report an unweighted average of the elasticities from Model 1 for each of the three time periods. $s.e.(\hat{\varepsilon})$: calculated from the standard error reported for the regression coefficients in Table 7 of NBER Working Paper 11230. I assume that the covariance between the coefficient estimates is zero because the full variance-covariance matrix for the regression coefficients is not reported. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of log wage rates because the study effectively exploits cross-sectional variation in wage rates for identification; the instruments used in Table 6 correct only for measurement error. The standard deviation of log wages for married women is not reported, and is therefore taken from Rothstein (2008), who reports a value of 0.50 in col. 4 of Table 1 for married women in 1992/3. This estimate is consistent with other published estimates of the standard deviations of women’s log wages in the CPS (e.g., Blau and Kahn 2000, Card and DiNardo 2002).

Appendix C: Sources and Calculations for Extensive Margin Studies in Table II

This appendix describes the sources of the values in Columns 3-5 of Table II for each study. For studies 1-7, the elasticity estimates ($\hat{\eta}$) and standard errors in columns 3-4 are taken from Table 1 in Chetty et al. (2011b), and details on the sources of these estimates are given in Appendix B of that paper. Studies 8-10 are also from Chetty et al. (2011b) and details on these estimates can be found in Appendix C of that paper. I follow the same methods as in Appendix B to calculate $\Delta \log(1 - \tau)$, defined here as the change in the net-of-average tax rate. The papers used for the analysis along with comprehensive documentation of the calculations are available at:

http://obs.rc.fas.harvard.edu/chetty/bounds_opt_meta_analysis.zip

A. Quasi-Experimental Elasticities

1. Eissa and Liebman (1996): $\Delta \log(1 - \tau)$: Meyer and Rosenbaum (2000) use the same data source and in Table 2 calculate the financial gain from working for single mothers in 1990 as \$8,458, compared with \$7,469 in 1984. I therefore define $\Delta \log(1 - \tau) = \log(8458) - \log(7469)$.

2. Graversen (1998): $\Delta \log(1 - \tau)$: Table 3 reports level changes in employment rates and participation elasticities, from which I back out $\Delta \log(1 - \tau) = (\Delta\theta/\bar{\theta})/\hat{\eta}$, where $\Delta\theta = -0.031$ is the estimated change in employment rates for single women, $\bar{\theta} = 0.7$ is the mean employment rate for single women using an average of the six participation rates in Table 2 weighted by sample sizes, and $\hat{\eta} = -0.174$ is the elasticity estimate reported in Table 3.

3. Devereux (2004): $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the deviations from the mean log wage change for each region/age/education group in Table A1 for women because the variation used for identification is across region and time by education/age group. Note that this table conditions on some work whereas in the sample used to estimate $\hat{\eta}$, nonparticipants’ wages are imputed as the average for their group.

following Davis and Henrekson’s data appendix. Austria is excluded because Davis and Henrekson exclude it from Sample C. The variable of interest in the data set is tw which stands for “tax wedge.” See Davis and Henrekson for more details. The mean (0.496 vs. 0.500) and standard deviation (0.14 vs. 0.133) reported for Sample C in Table 2.1 differs slightly from those used in this calculation. The data were accessed from the .zip appendix at the following URL: <http://cep.lse.ac.uk/pubs/number.asp?number=502>

4. Meyer and Rosenbaum (2001): $\Delta \log(1 - \tau)$: see discussion of study 4 in Chetty et al. (2011b), who define $\Delta \log(1 - \tau) = 45\%$ after accounting for taxes and transfers as in Meyer and Rosenbaum (2000, pg. 1043).

5. Eissa and Hoynes (2004): $\Delta \log(1 - \tau)$: Meyer and Rosenbaum (2000, pg. 1043) report a tax change of 45% from 1984 to 1996 for the group studied by Eissa and Hoynes.

6. Liebman and Saez (2006) : $\Delta \log(1 - \tau)$: defined as $\log(1 - 0.419) - \log(1 - 0.31)$ based on the change in tax rates reported on pages 10–11 for OBRA93.

7. Blundell, Bozio, and Laroque (2011): $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of log net-of-tax-rates for participation. Standard deviation of 0.37 obtained from personal correspondence with authors.

B. Macro/Cross-Sectional Elasticities

8. Nickell (2003): $\hat{\eta}$: computed using the average point estimate of 2 percent (reported on page 8) and the sample means of employment rates and tax rates from Tables 1 and 2, respectively. $s.e.(\hat{\eta})$: not reported because Nickell does not report standard errors for the studies in Table 4 on which his point estimate is based. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of log net-of-tax-rates using values listed in Table 2 because most of the studies in Table 4 used in Nickell’s estimate of the effect of taxation on employment use panel or cross sectional data for OECD countries.


9. Prescott (2004): $\hat{\eta}$, $s.e.(\hat{\eta})$: calculated by regressing log labor force participation rates from OECD Stat Extracts on log net-of-tax rates using the same sample of countries and years as Prescott.³⁵ The data on tax rates is taken from Table 2 of Prescott (2004). The data on labor force participation rates are missing for Canada and the U.K. in the 1970’s and these observations are therefore excluded. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of the change in log net-of-tax rate for the 12 observations with non-missing data on labor force participation rates.

10. Davis and Henrekson (2005): $\hat{\eta}$: computed using the log difference in employment based on the slope coefficient in Table 2.3 (bottom panel, Sample C) and the sample means of labor force participation and tax rates in Table 1 for the corresponding sample. $s.e.(\hat{\eta})$: calculated from the standard error reported for the slope coefficient in Table 2.3 (bottom panel, Sample C). $\Delta \log(1 - \tau)$: computed as $2 \times SD$ of log one minus sum of tax rates for the 19 countries in Sample C.

11. Blau and Kahn (2007): $\hat{\eta}$: I report an unweighted average of the own wage participation elasticities for each of the three time periods in Table 6, Model 1. $s.e.(\hat{\eta})$: The standard error is calculated from the standard error reported for own log wage in Table 7 of NBER Working Paper 11230. I assume that the covariance between the coefficient estimates is zero because the full variance-covariance matrix for the parameters in the probit model is not reported. $\Delta \log(1 - \tau)$: defined as $2 \times SD$ of log wages, calculated as described in study 23 in Appendix B above.


³⁵Data are for men and women aged 15-64 for 1970-1974 and 1993-1996 in order to match Prescott’s data. Data are available from OECD Stat Extracts at the following URL: http://stats.oecd.org/Index.aspx?DataSetCode=LFS_SEXAGE_I_R

FIGURE 1
Choice Set in a δ Class of Models




NOTE—This figure illustrates the choice set $X(p_t, \delta)$ in a δ class of models when there is no heterogeneity across agents and $\delta = 1\%$ and $a_{i,t} = a = \exp(3.5)$. The blue curve plots flow utility $u(x_t) = 100 - p_t x_t + a \log x_t$ with $\log p_t = 1$. The set of demand levels that yield utility within $0.01 p_t x^*(p_t)$ dollars of the maximum is shown by the red interval on the x axis.

FIGURE 2
Identification with Optimization Frictions



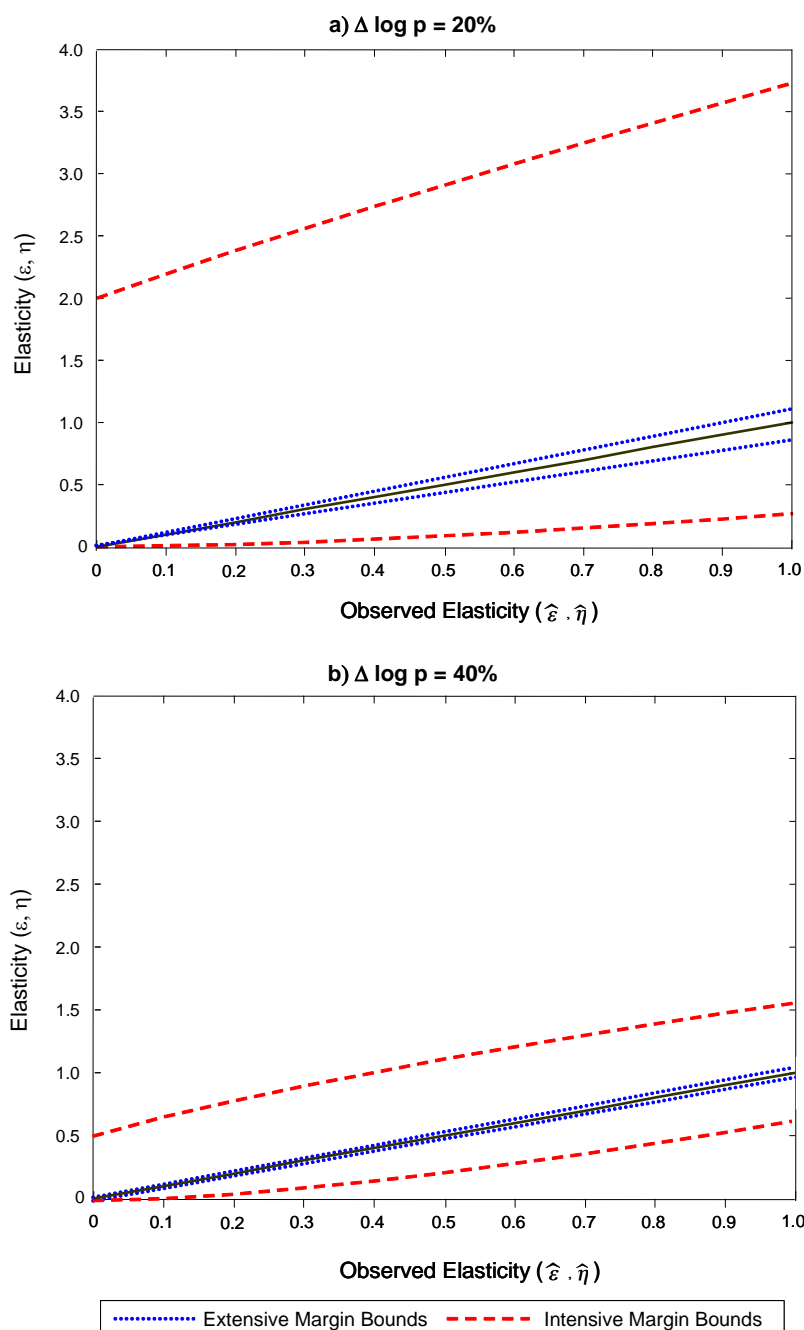
NOTE—This figure plots the choice sets at two price levels, $X(p_A, \delta)$ and $X(p_B, \delta)$, with $\log p_A = 1$ and $\log p_B = 1.2$. In Panel A, $\varepsilon = 1$; in Panel B, $\varepsilon = 0$. All other parameters are specified as in Figure 1. The dashed blue line shows the optimal demand $x^*(p_t)$. The black lines in Panel A illustrate some of the responses $(\log x_B(p_B) - \log x_A(p_A))$ that may be observed for a price increase from p_A to p_B with a structural elasticity of $\varepsilon = 1$ and frictions of $\delta = 1\%$.

FIGURE 3
Bounding the Structural Elasticity with Optimization Frictions



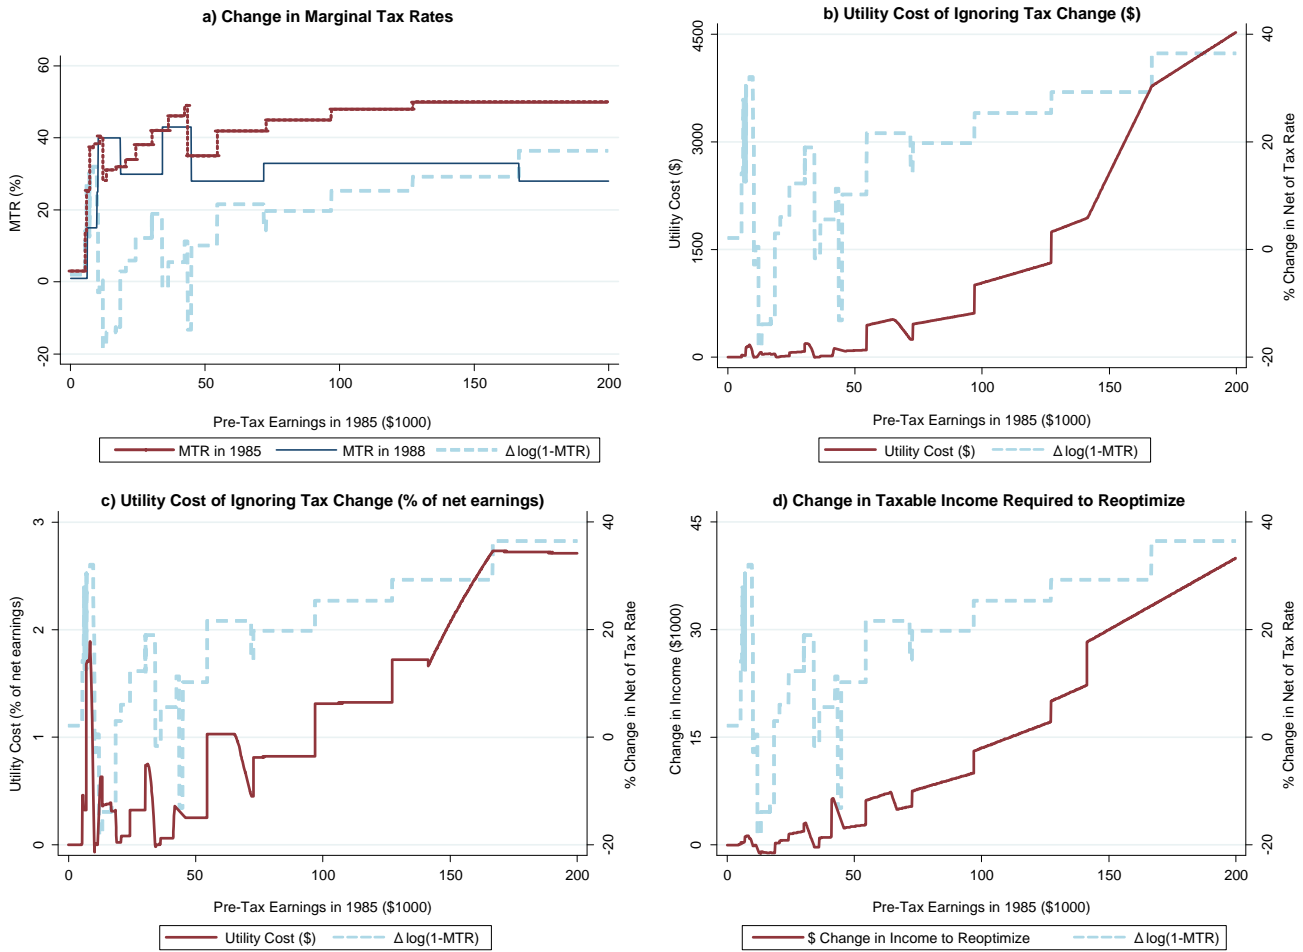
NOTE—The solid black line in each panel depicts the observed demand response for a price increase from p_A to p_B with an observed elasticity $\hat{\varepsilon} = 0.3$, $\log p_A = 1$, and $\log p_B = 1.4$. Panel A depicts the highest structural elasticity, $\varepsilon_U = 1$, that could have generated this observed response with $\delta = 1\%$. The dashed blue line depicts the optimal demand $x^*(p_t)$ with $\varepsilon = 1$. Panel B analogously depicts the lowest structural elasticity, $\varepsilon_L = 0.1$, that could have generated the same observed response.

FIGURE 4
 Bounds on Structural Elasticities as a Function of Observed Elasticities



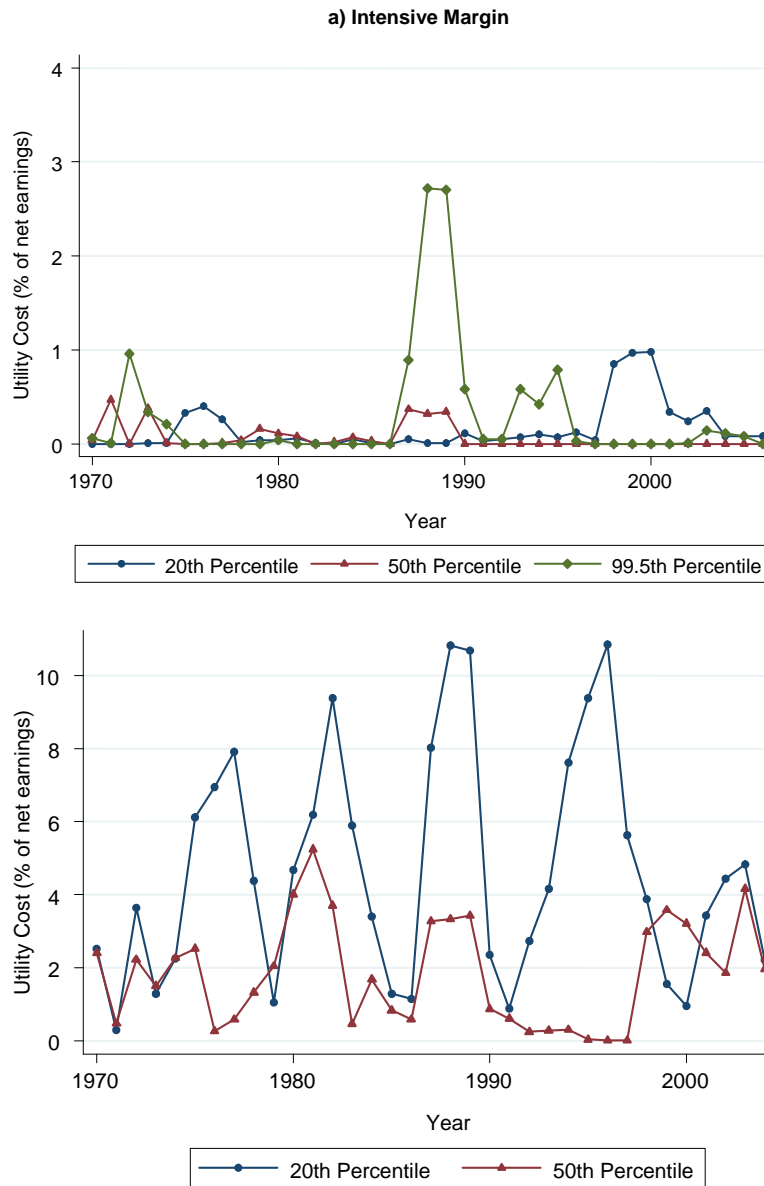
NOTE—The dashed (red) lines show the bounds on the intensive margin structural elasticity $(\varepsilon_L, \varepsilon_U)$ vs. the observed intensive margin elasticity $\hat{\varepsilon}$, computed using Proposition 1. The dotted (blue) lines show the bounds on the extensive margin structural elasticity (η_L, η_U) vs. the observed extensive margin elasticity $\hat{\eta}$, computed using Proposition 2. The solid black line is the 45 degree line. The bounds are computed with $\delta = 1\%$ frictions and $\Delta \log p = 20\%$ (Panel A) and $\Delta \log p = 40\%$ (Panel B).

FIGURE 5
Tax Reform Act of 1986



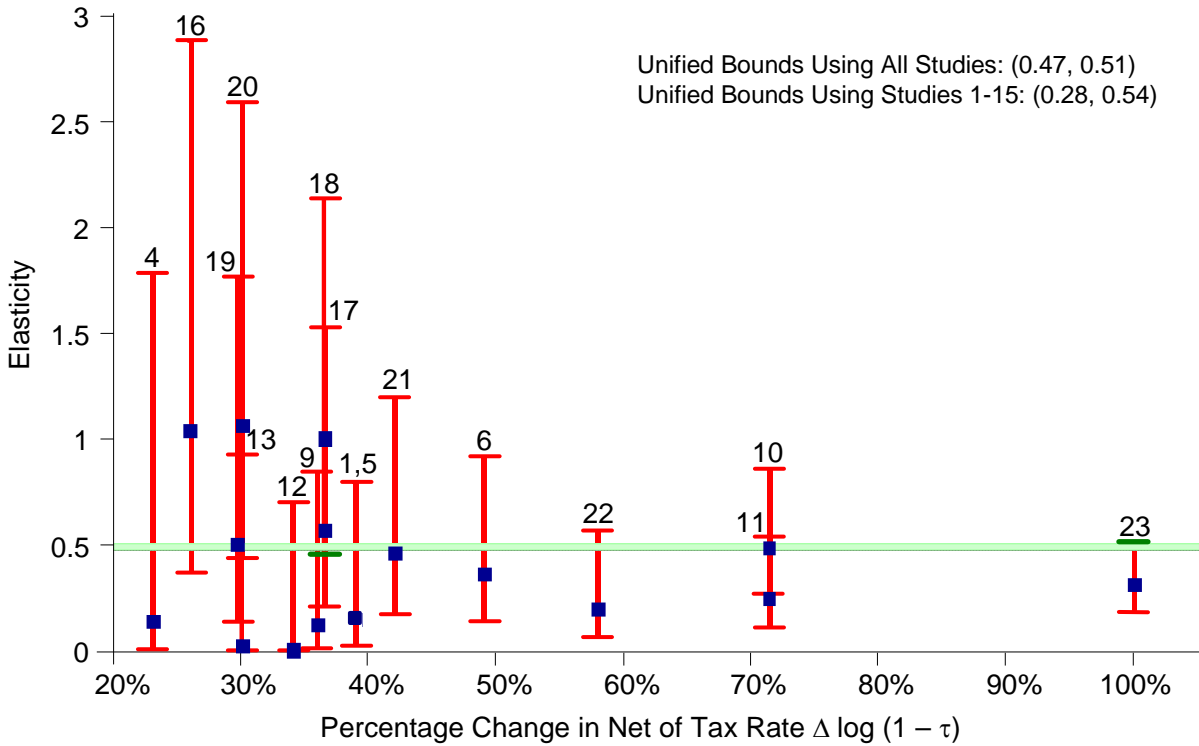
NOTE—These figures are based on the Tax Reform Act of 1986. The x axis in all the figures is gross earnings in the year prior to the reform. Panel A shows how marginal tax rates changed between 1985 and 1988 for single filers with two children. Panel B plots the utility cost Δu_i , measured in dollars, from failing to reoptimize labor supply on the intensive margin in response to the tax change with $\varepsilon = 0.5$. Panel C plots the same utility cost as a percentage of net-of-tax earnings ($\Delta u_{i,\%}$), defined as the dollar cost in Panel B divided by the agent’s optimal net-of-tax earnings in 1988. Panel D shows the change in gross earnings ($wl_{i,1988}^* - wl_{i,1985}^*$) required to reoptimize relative to the tax change. In Panels B-D, the dashed blue line (right y axis) replicates the log change in the net-of-tax rate (1-MTR) shown in Panel A.

FIGURE 6
Utility Cost of Ignoring Tax Changes by Year




NOTE—These figures plot the utility cost of ignoring changes in taxes over three-year periods from the 1970 to 2006 for selected percentiles of the household income distribution. Panel A shows the utility cost of failing to reoptimize labor supply on the intensive margin ($\Delta u_{i,t,\%}$) with a structural intensive-margin elasticity of $\varepsilon = 0.5$, calculated as in Figure 5c. In each year y , the point that is plotted shows the utility loss (as a percentage of optimal net-of-tax earnings in year y) from choosing labor supply optimally according to the tax system in year $y - 3$ instead of year y . Panel B depicts the percentage utility cost ($\Delta u_{i,t,ext,\%}$) of failing to reoptimize labor supply on the extensive margin in year y for the marginal agent in year $y - 3$. This is the agent whose disutility of working b_i made him indifferent between working and not working in $y - 3$. The utility cost $\Delta u_{i,t,ext,\%}$ is measured as a percentage of net-of-tax earnings when working in year y , as in Corollary 2.

FIGURE 7
 Bounds on Intensive Margin Hicksian Labor Supply Elasticities with $\delta = 1\%$



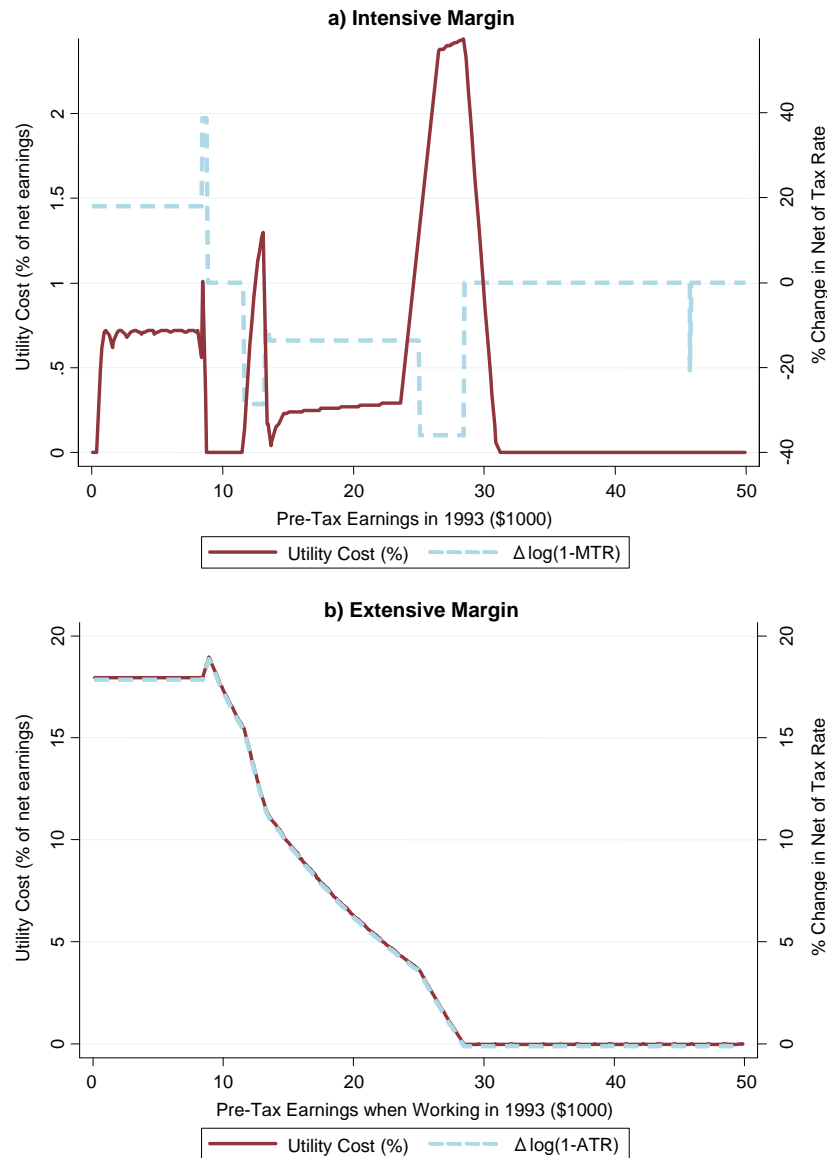
NOTE—The red intervals show the bounds on the structural intensive margin Hicksian elasticity ε implied by each of the studies with corresponding numbers listed in Table I. The blue squares show the point estimate of each study. The x axis is the log change in the net of tax rate ($\Delta \log(1 - \tau)$) used for identification in each study. Studies with $\Delta \log(1 - \tau) < 20\%$ are excluded from this figure for scaling purposes. The shaded region shows the range of structural elasticities consistent with all the observed elasticities in Panels A-D of Table I, the unified bounds of (0.47, 0.51). When only studies 1-15 (Panels A and B) are used, the unified bounds are (0.28, 0.54).

FIGURE 8
 Unified Bounds on Intensive Margin Hicksian Elasticity vs. Degree of Frictions



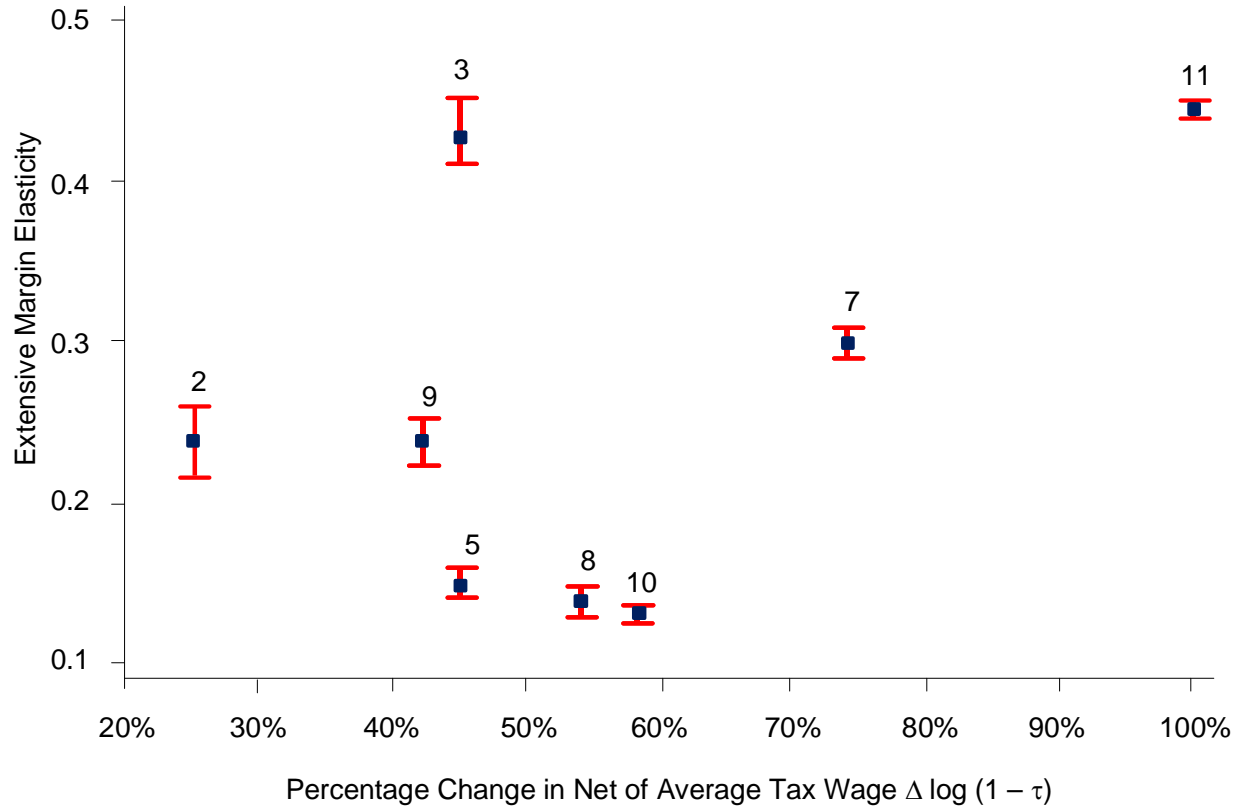
NOTE—This figure shows how the unified bounds on the structural intensive-margin elasticity ε vary with the level of frictions δ . The solid red lines plot the unified bounds implied by the studies in Panels A and B of Table I. These unified bounds are defined only for $\delta > \delta_{\min} = 0.5\%$ because δ 's below this threshold cannot reconcile the observed elasticities. The dashed red lines show a 95% confidence interval for the unified bounds.

FIGURE 9
Utility Costs of Ignoring the Clinton EITC Expansion



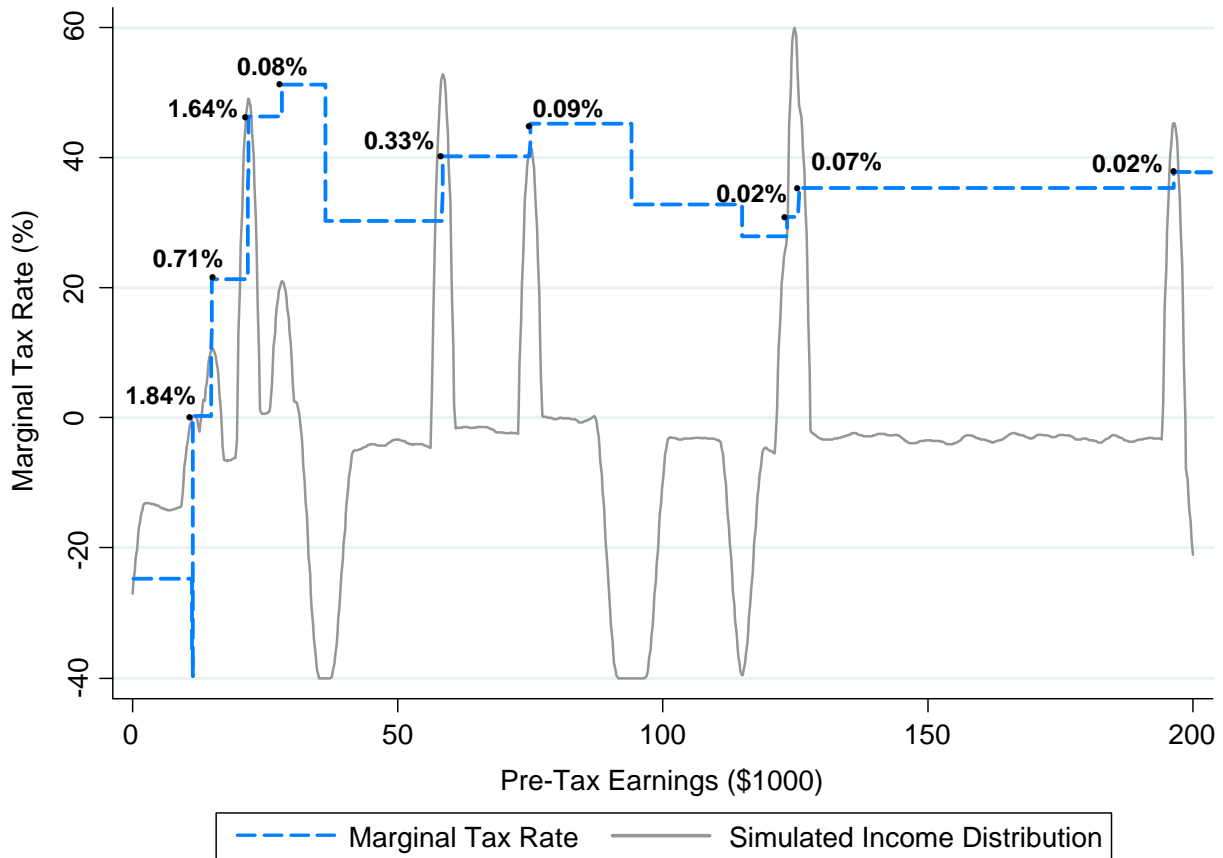
NOTE—These figures are based on the Clinton EITC Expansion enacted between 1993 and 1996. Panel A considers the intensive margin. The x axis is gross earnings in the year prior to the reform. The dashed blue line (right y axis) shows the log change in the net-of-marginal-tax rate (1-MTR) from 1993-1996 for single filers with two children. The solid red line plots the utility cost as a percentage of optimal net-of-tax earnings in 1996 ($\Delta u_{i,1996,\%}$) from failing to reoptimize hours of work in response to the tax change when $\varepsilon = 0.5$. Panel B considers the extensive margin. The dashed blue line (right y axis) shows the log change in the net-of-average-tax rate (1-ATR) from 1993-1996 for single filers with two children. The solid red line plots the utility cost ($\Delta u_{i,1996,ext,\%}$) of failing to enter the labor force in 1996 for the marginal agent who chose not to work at each earnings level in 1993. This is this agent whose disutility of working b_i made him indifferent between working and not working in 1993 at the gross earnings level shown on the x axis. The utility cost $\Delta u_{i,1996,ext,\%}$ is measured as a percentage of net-of-tax earnings when working in 1996.

FIGURE 10
 Bounds on Extensive Margin Labor Supply Elasticities with $\delta = 1\%$



NOTE—The red intervals show the bounds on the structural extensive margin elasticity η implied by each of the studies with corresponding numbers listed in Table II. The blue squares show the point estimate of each study. The x axis is the log change in the net of tax rate ($\Delta \log(1 - \tau)$) used for identification in each study. Papers with $\Delta \log(1 - \tau) < 20\%$ are excluded from this figure for scaling purposes.

FIGURE 11
Gains from Bunching at Kinks in 2006 Tax Schedule



NOTE—The dashed blue curve shows the 2006 marginal tax rate schedule in the U.S. The solid grey curve shows the distribution of taxable income predicted by the frictionless labor supply model with $\varepsilon = 0.5$. This curve assumes a uniform distribution of a_i and plots an Epanechnikov kernel density of the simulated earnings distribution with a bandwidth of \$1000. The numbers near each convex kink are the utility gain as a percentage of optimal net-of-tax earnings ($\Delta u\%$) from locating at that kink when $\varepsilon = 0.5$. To compute $\Delta u\%$ at a given kink, I first define $\Delta u_{i,\%}$ as the utility gain for an individual with taste parameter a_i from locating at that kink relative to optimizing under the (incorrect) assumption that the tax rate in the previous bracket continues into the next bracket. I then define $\Delta u\%$ as the unweighted mean of $\Delta u_{i,\%}$ over all individuals whose a_i would make it optimal for them to locate at that kink. The first two kinks (1.84% and 0.71%) correspond to the end of the phase-in and start of the phase-out regions of the EITC.

TABLE 1
Bounds on Intensive Margin Hicksian Labor Supply Elasticities with $\delta = 1\%$ Frictions

Study (1)	Identification (2)	$\hat{\varepsilon}$ (3)	s.e. ($\hat{\varepsilon}$) (4)	$\Delta \log(1-\tau)$ (5)	ε_L (6)	ε_U (7)	95% CI		
							ε_L (8)	ε_U (9)	
<i>A. Hours Elasticities</i>									
1. MaCurdy (1981)	Lifecycle wage variation, 1967-1976	0.15	0.15	0.39	0.03	0.80	0.04	1.20	
2. Eissa and Hoynes (1998)	U.S. EITC Expansions, 1984-1996, Men	0.20	0.07	0.07	0.00	15.29	0.00	15.51	
3. Eissa and Hoynes (1998)	U.S. EITC Expansions, 1984-1996, Women	0.09	0.07	0.07	0.00	15.07	0.00	15.30	
4. Blundell, Duncan, and Meghir (1998)	U.K. Tax Reforms, 1978-1992	0.14	0.09	0.23	0.01	1.78	0.00	2.04	
5. Ziliak and Kniesner (1999)	Lifecycle wage, tax variation 1978-1987	0.15	0.07	0.39	0.03	0.80	0.00	0.99	
	Mean observed elasticity	0.15							
<i>B. Taxable Income Elasticities</i>									
6. Bianchi, Gudmundsson, and Zoega (2001)	Iceland 1987 Zero Tax Year	0.37	0.05	0.49	0.15	0.92	0.10	1.04	
7. Gruber and Saez (2002)	U.S. Tax Reforms 1979-1991	0.14	0.14	0.14	0.00	4.42	0.00	4.84	
8. Saez (2004)	U.S. Tax Reforms 1960-2000	0.09	0.04	0.15	0.00	3.51	0.00	3.64	
9. Jacob and Ludwig (2008)	Chicago Housing Voucher Lottery	0.12	0.03	0.36	0.02	0.84	0.01	0.92	
10. Gelber (2010)	Sweden, 1991 Tax Reform, Women	0.49	0.02	0.71	0.28	0.86	0.25	0.91	
11. Gelber (2010)	Sweden, 1991 Tax Reform, Men	0.25	0.02	0.71	0.12	0.54	0.10	0.59	
12. Saez (2010)	U.S., 1st EITC Kink, 1995-2004	0.00	0.02	0.34	0.00	0.70	0.00	0.77	
13. Chetty et al. (2011a)	Denmark, Married Women, Top Kinks, 1994-2001	0.02	0.00	0.30	0.00	0.93	0.00	0.94	
14. Chetty et al. (2011a)	Denmark, Middle Kinks, 1994-2001	0.00	0.00	0.11	0.00	6.62	0.00	6.62	
15. Chetty et al. (2011a)	Denmark Tax Reforms, 1994-2001	0.00	0.00	0.09	0.00	9.88	0.00	9.89	
	Mean observed elasticity	0.15							
<i>C. Top Income Elasticities</i>									
16. Feldstein (1995)	U.S. Tax Reform Act of 1986	1.04		0.26	0.37	2.89			
17. Auten and Carroll (1999)	U.S. Tax Reform Act of 1986	0.57	0.12	0.37	0.21	1.53	0.11	1.81	
18. Goolsbee (1999)	U.S. Tax Reform Act of 1986	1.00	0.15	0.37	0.47	2.14	0.32	2.47	
19. Saez (2004)	U.S. Tax Reforms 1960-2000	0.50	0.18	0.30	0.14	1.77	0.03	2.21	
20. Kopczuk (2010)	Poland, 2002 Tax Reform	1.07	0.22	0.30	0.44	2.58	0.24	3.09	
	Mean observed elasticity	0.84							
<i>D. Macro/Cross-Sectional</i>									
21. Prescott (2004)	Cross-country Tax Variation, 1970-1996	0.46	0.09	0.42	0.18	1.20	0.10	1.41	
22. Davis and Henrekson (2005)	Cross-country Tax Variation, 1995	0.20	0.08	0.58	0.07	0.57	0.01	0.76	
23. Blau and Kahn (2007)	U.S. wage variation, 1980-2000	0.31	0.004	1.00	0.19	0.51	0.18	0.52	
	Mean observed elasticity	0.32							
					Unified Bounds Using Panels A and B:	0.28	0.54	0.23	0.61
					Minimum- δ Estimate ($\varepsilon_{\delta-\min}$):	0.33			
					Unified Bounds Using All Panels:	0.47	0.51	0.23	0.53
					Minimum- δ Estimate ($\varepsilon_{\delta-\min}$):	0.50			

Note: This table shows bounds on structural intensive margin Hicksian elasticities using estimates from existing studies. Column 3 shows the point estimate of the observed elasticity, column 4 shows the associated standard error, and column 5 shows the size of the net-of-marginal-tax wage change used for identification. Columns 6 and 7 show the lower and upper bounds on the structural elasticity, calculated using Proposition A1. Columns 8 and 9 give a 95% confidence interval for ε , constructed as in Imbens and Manski (2004). See Appendix B for sources and details underlying calculations in columns 3-5.

TABLE 2
 Bounds on Extensive Margin Labor Supply Elasticities with $\delta = 1\%$ Frictions

Study (1)	Identification (2)	$\hat{\eta}$ (3)	s.e.($\hat{\eta}$) (4)	$\Delta\log(1-\tau)$ (5)	η_L (6)	η_U (7)	95% CI	
							η_L (8)	η_U (9)
<i>A. Quasi-Experimental Estimates</i>								
1. Eissa and Liebman (1996)	U.S. EITC Expansions 1984-1990, Single Mothers	0.30	0.10	0.12	0.26	0.36	0.12	0.56
2. Graversen (1998)	Denmark 1987 Tax Reform, Women	0.24	0.04	0.25	0.22	0.26	0.16	0.33
3. Meyer and Rosenbaum (2001)	U.S. Welfare Reforms 1985-1997, Single Women	0.43	0.05	0.45	0.41	0.45	0.33	0.53
4. Devereux (2004)	U.S. Wage Trends 1980-1990, Married Women	0.17	0.17	0.12	0.14	0.20	0.00	0.53
5. Eissa and Hoynes (2004)	U.S. EITC expansions 1984-1996, Low-Income Married Men and Women	0.15	0.07	0.45	0.14	0.16	0.03	0.28
6. Liebman and Saez (2006)	U.S. Tax Reforms 1991-1997, Women Married to High Income Men	0.15	0.30	0.17	0.13	0.17	0.00	0.72
7. Blundell, Bozio, and Laroque (2011)	U.K. Tax Reforms 1978-2007, Prime-age Men and Women	0.30	n/a	0.74	0.29	0.31		
	Mean observed elasticity	0.25						
<i>B. Macro/Cross-Sectional</i>								
8. Nickell (2003)	Cross-country Tax Variation, 1961-1992	0.14	n/a	0.54	0.13	0.15		
9. Prescott (2004)	Cross-country Tax Variation, 1970-1996	0.24	0.14	0.42	0.22	0.25	0.00	0.50
10. Davis and Henrekson (2005)	Cross-country Tax Variation, 1995	0.13	0.11	0.58	0.13	0.13	0.00	0.33
11. Blau and Kahn (2007)	U.S. Wage Variation 1989-2001, Married Women	0.45	0.004	1.00	0.44	0.45	0.43	0.46
	Mean observed elasticity	0.24						

Note: This table shows bounds on structural extensive margin elasticities using estimates from existing studies. Column 3 shows the point estimate of the observed elasticity, column 4 shows the associated standard error, and column 5 shows the size of the net-of-average-tax wage change used for identification. Columns 6 and 7 show the lower and upper bounds on the structural elasticity, calculated using Proposition 2. Columns 8 and 9 give a 95% confidence interval for η , constructed as in Imbens and Manski (2004). See Appendix C for sources and details underlying calculations in columns 3-5.

TABLE 3
Frisch Elasticities Implied by Hicksian Elasticity of 0.33

		Income Effect: $-d[w/l^*]/dY$						
		0.00	0.11	0.22	0.33	0.44	0.55	0.66
EIS (ρ)	0.00	0.33	0.33	0.33	0.33	0.33	0.33	0.33
	0.20	0.33	0.34	0.35	0.36	0.38	0.41	0.44
	0.40	0.33	0.34	0.36	0.39	0.43	0.49	0.55
	0.60	0.33	0.34	0.37	0.42	0.48	0.56	0.66
	0.80	0.33	0.35	0.38	0.44	0.53	0.64	0.77
	1.00	0.33	0.35	0.39	0.47	0.58	0.71	0.88
	1.20	0.33	0.35	0.41	0.50	0.63	0.79	0.99
	1.40	0.33	0.35	0.42	0.53	0.67	0.87	1.10

Note: This table shows the intensive margin Frisch elasticity implied by various combinations of the EIS and income effect. The calculations assume that the ratio of wealth to earned income is $A/w^* = 1.26$ (Dynan 2009) and the intensive margin Hicksian (compensated) elasticity is $\epsilon = 1/3$ (Table 1). Values within the dashed lines are consistent with evidence that the $EIS \leq 1$ and the uncompensated labor supply elasticity is positive. The values are computed using the equation $\epsilon^F = \epsilon + \rho(d[w/l^*]/dY)^2(A/w^*)$.