

Regressive Sin Taxes, with an Application to the Optimal Soda Tax

Hunt Allcott, Benjamin B. Lockwood, and Dmitry Taubinsky*

First version: January 2017

This version: August 2018

Abstract

A common objection to “sin taxes”—corrective taxes on goods that are thought to be overconsumed, such as cigarettes, alcohol, and sugary drinks—is that they often fall disproportionately on low-income consumers. This paper studies the interaction between corrective and redistributive motives in a general optimal taxation framework, and delivers empirically implementable sufficient statistics formulas for determining the optimal commodity tax. The optimal sin tax is increasing in the price elasticity of demand, increasing in the degree to which lower-income consumers are more biased or more elastic to the tax, decreasing in the extent to which consumption is concentrated among the poor, and decreasing in income effects, because income effects imply that commodity taxes create labor supply distortions. Contrary to common intuitions, stronger preferences for redistribution can *increase* the optimal sin tax, if lower-income consumers are more responsive to taxes or are more biased. As an application, we estimate the optimal nationwide tax on sugar-sweetened beverages, using Nielsen Homescan data and a specially designed survey measuring nutrition knowledge and self-control. Our estimates imply that current city-level taxes in Berkeley and elsewhere are somewhat lower than the social optimum.

*Allcott: New York University and NBER. hunt.allcott@nyu.edu. Taubinsky: Berkeley and NBER. dmitry.taubinsky@berkeley.edu. Lockwood: Wharton and NBER. ben.lockwood@wharton.upenn.edu. We thank Kelly Brownell, Gabriel Chodorow-Reich, Stefano DellaVigna, Rebecca Diamond, Jean-Pierre Dube, Matt Gentzkow, Anna Grummon, David Laibson, Alex Rees-Jones, Christina Roberto, Na’ama Shenhav, Claire Wang, Danny Yagan, and seminar participants at Berkeley, Bonne, Brown, Carnegie Mellon and the University of Pittsburgh, Columbia, Cologne, Davis, Hong Kong University of Science and Technology, Michigan State University, the National Tax Association, National University of Singapore, the NBER Public Economics 2018 Spring Meetings, NYU, Princeton, Stanford, the Society for Benefit-Cost Analysis, University of British Columbia, University of Michigan, University of Virginia, and Wharton for helpful feedback. We are grateful to the Sloan Foundation for grant funding. The survey was determined to be exempt from review by the Institutional Review Boards at the University of Pennsylvania (protocol number 828341) and NYU (application FY2017-1123). Nielsen requires the following text: *This paper reflects the authors’ own analyses and calculations based in part on data reported by Nielsen through its Homescan, RMS, and PanelViews services for beverage categories over 2006-2016, for all retail channels in the U.S. market. Calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kiltz Center for Marketing Data Center at The University of Chicago Booth School of Business. Copyright © 2018 The Nielsen Company (US), LLC. All Rights Reserved. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.* Replication files are available from <https://sites.google.com/site/allcott/research>. This paper subsumes and replaces Lockwood and Taubinsky (2017).

“The only way to protect all of us, including the poor, from further harm is through a sugary drink tax...”

– *Forbes* magazine article (Huehnergath, 2016)

“A tax on soda and juice drinks would disproportionately increase taxes on low-income families in Philadelphia.”

– U.S. Senator Bernie Sanders (2016)

“They’ve [big soda] made their money of the backs of poor people, but this money [soda tax revenue] will stay in poor neighborhoods.”

– Philadelphia Mayor Jim Kenney (quoted in Blumgart (2016))

I Introduction

A large literature in behavioral economics suggests that biases such as self-control problems, inattention, and incorrect beliefs can lead to over-consumption of “sin goods” such as cigarettes, alcohol, unhealthy foods, and energy-inefficient durable goods. Consumption of these goods can also generate externalities, in the form of health care costs or pollution. Consequently, “sin taxes” that discourage consumption of such goods could increase social welfare.¹ This argument has led to widespread taxation of cigarettes and alcohol, as well as newer taxes on sugar-sweetened beverages in seven U.S. cities and 34 countries around the world (GFRP 2018).

What is the optimal level of a sin tax? The existing literature frequently invokes a corrective logic dating to Pigou (1920) and Diamond (1973): the optimal corrective tax equals the sum of the externality and the average mistake (or “internality”) of marginal consumers.² This principle, however, assumes that consumers do not vary in their marginal utility of money, and thus that policymakers care equally about the poor versus the rich. This assumption is starkly out of sync with public debates about sin taxes. As highlighted by the above quote from Senator Bernie Sanders, a common objection to sin taxes is that they are regressive.³

In response to such objections, however, others argue that the harms caused by overconsumption from behavioral biases are themselves regressive, so a corrective tax might confer greater benefits on the poor than the rich. For example, smoking and sugary drink consumption cause lung cancer, diabetes, and other health problems that disproportionately affect the poor.⁴ Furthermore, as

¹See, e.g., Chetty (2015), Mullainathan, Schwartzstein, and Congdon (2012), and Bernheim and Rangel (2009) for an overview of consumer biases and their implications for welfare calculations.

²For Pigouvian taxation of externalities, see, e.g., O’Donoghue and Rabin (2006), Gruber and Kőszegi (2001), Allcott, Mullainathan, and Taubinsky (2014), Allcott and Taubinsky (2015), Heutel (2015), Tsvetanov and Segerson (2013), and Mullainathan, Schwartzstein, and Congdon (2012).

³The poor disproportionately consume cigarettes and sugary drinks, while the rich disproportionately take up energy efficiency subsidies. See, e.g., Gruber and Kőszegi (2004) and Goldin and Homonoff (2013) on cigarettes, and Allcott, Knittel, and Taubinsky (2015), Davis and Borenstein (2016), and Davis and Knittel (2016) on energy efficiency subsidies and fuel economy standards.

⁴Gruber and Kőszegi (2004) make this argument formally in the context of cigarette taxes. Less formally, Huehnergath (2016) writes, “Here’s why Sanders’s position on sugary drink taxes plays right into the hands of Big Soda, and hurts America’s most vulnerable citizens: Big Soda’s targeted marketing to communities of color and low-income communities is regressive.... Type 2 diabetes, linked to excessive sugary drink consumption, is regressive....”

emphasized in the above quote from Philadelphia Mayor Jim Kenney, regressivity can be reduced by “recycling” sin tax revenues to programs that benefit the poor.⁵

In this paper, we present a general theoretical model that delivers the first explicit formulas for an optimal commodity tax that accounts for each of the three central considerations in this policy debate: correction of externalities and/or consumer bias, regressivity, and revenue recycling. We then use the theoretical formulas to estimate the optimal nationwide tax on sugar sweetened beverages, using new data and empirical techniques to estimate elasticities and biases.

Our theoretical model in Section II builds on Saez’s (2002a) extension of Atkinson and Stiglitz (1976) by considering an economy of consumers with heterogeneous earning abilities and tastes, who choose labor supply and a consumption bundle that exhausts their after-tax income. The policymaker chooses a set of linear commodity taxes and a nonlinear income tax, which can be used to provide transfers to poor consumers, to raise money for commodity subsidies, or to distribute commodity tax revenue (in a progressive way, if desired). Unlike those models, however, we allow for a corrective motive of taxation, driven by externalities or internalities from consumer mistakes.

In developing our theoretical results, we fill two gaps in the literature. First, even in the absence of internalities or externalities, there was no known formula for optimal commodity taxes in the presence of nonlinear income taxation and preference heterogeneity.⁶ As a result, tax economists’ policy recommendations have relied on the canonical result of Atkinson and Stiglitz (1976) that commodity tax rates should be uniform. This extreme result, however, was proven under a homogeneous preferences assumption that is likely unrealistic in many settings—and which we show is strongly rejected in the case of sugary drinks. Second, there has been no general account of how redistributive motives and behavioral biases jointly shape optimal commodity taxes. Although this is not the first paper to study internality-correcting commodity taxes, it is the first to embed internalities in the dominant optimal taxation framework of public economics, which allows for redistributive motives and both commodity taxes and a nonlinear income tax. Most previous papers studying internality taxes (e.g., the papers in footnote 2) abstract from redistributive motives. Those that do not (such as Gruber and Kőszegi, 2004; Bernheim and Rangel, 2004; Farhi and Gabaix, 2015) focus on specialized models of bias and/or consider simplified tax environments that do not permit redistribution through nonlinear income taxation, means-tested transfer programs, or progressive revenue recycling.⁷

⁵Philadelphia directs a share of soda tax revenues to pre-kindergarten education initiatives in low income neighborhoods. As a result, Blumgart (2016) writes that “By earmarking the revenues to programs that are at the top of the progressive agenda, [Mayor] Kenney has provided a strong counterargument to those concerned by the regressive nature of the tax: the money will go right back into those very neighborhoods that are hit hardest.”

⁶Saez (2002a) answered the qualitative question of when a “small” commodity tax can increase welfare in the presence of preference heterogeneity, but left to future work the task of deriving an expression for the optimal tax, writing “It would of course be extremely useful to obtain optimal commodity tax formulas” in such a general framework.

⁷Gruber and Kőszegi (2004) study the incidence of cigarette taxes on low- and high-income consumers with self-control problems, and show that low-income consumers with sufficiently elastic consumption can benefit from such taxes, but they do not characterize the optimal tax implications. Bernheim and Rangel (2004) study the optimal tax on addictive substances in a dynamic framework in which environmental cues trigger mindless consumption, and differences in the marginal utility from wealth arise from differences in spending on the addictive good or on

Our optimal commodity tax formula decomposes into two terms that address these two gaps. The first term represents the “redistributive motive”: the desire to use the commodity tax to transfer money from the rich to the poor by subsidizing goods consumed by low earners. Although Atkinson and Stiglitz (1976) show that it is optimal to carry out all redistribution solely through the income tax when preferences are homogeneous, this is no longer true when heterogeneous preferences are correlated with earning ability. Differential commodity taxes *are* optimal in the presence of preference heterogeneity, with the tax on a given commodity increasing with the inverse of the elasticity and the extent to which high-income consumers prefer it more. We derive a novel sufficient statistic that quantifies this preference heterogeneity: the difference between the cross-sectional variation in consumption of the sin good across incomes and the (causal) income effect.⁸

The second term in our optimal commodity tax formula represents the “corrective motive”: the desire to reduce over-consumption arising from internalities and externalities by imposing taxes on harmful goods. To incorporate internalities, we follow the sufficient statistics approach to behavioral public finance (Farhi and Gabaix, 2015; Chetty, 2015; Mullainathan, Schwartzstein, and Congdon, 2012) by adopting a money-metric definition of bias, which can transparently accommodate many specific behavioral biases and lends itself to direct empirical quantification.

The price elasticity of demand determines the relative importance of the corrective versus redistributive motives. High demand elasticity implies a large change in sin good consumption for a given degree of redistribution, and thus that the effects of the tax are primarily corrective rather than redistributive. Conversely, low demand elasticity implies that the effects of the tax are mostly redistributive rather than corrective.

In contrast to the conventional wisdom that inequality aversion unambiguously reduces optimal taxes on sin goods heavily consumed by the poor, we show that it can either decrease or increase the optimal sin tax. Although inequality aversion does magnify the redistributive motive, pushing toward a lower sin tax, it also amplifies the corrective motive when poor consumers are relatively more biased or more elastic, which pushes toward a higher sin tax.

rehabilitation. Farhi and Gabaix (2015) focus on an application of their general framework that assumes two types of consumers—behavioral and rational—with potential differences in the marginal utility from income. Both Farhi and Gabaix (2015) and Bernheim and Rangel (2004) assume that commodity taxes are the sole source of redistribution: the commodity tax revenue has to be distributed lump-sum and cannot, for example, be spent on transfers to the poor or programs that benefit the poor.

Also closely related is the work on health insurance by Baicker, Mullainathan, and Schwartzstein (2015), Spinnewijn (2017), and Handel et al. (2016). See also Mullainathan, Schwartzstein, and Congdon (2012) for a model that links the insights from the health insurance literature to commodity tax results; they derive formulas similar to those in Farhi and Gabaix (2015) for an economy in which consumers make a binary choice.

While some papers studying externalities allow for nonlinear income taxation, they still make the restrictive assumption of homogeneous preferences (Jacobs and de Mooij, 2015; Kaplow, 2012).

⁸Moreover, our results also generalize Jacobs and Boadway (2014), who study optimal commodity taxation in the absence of any market failure under the assumption that consumption preferences are homogeneous while allowing for non-separability between consumption and labor in the utility function, and also Jacobs and de Mooij (2015) and Kaplow (2012), who study the taxation of environmental externalities under the assumptions of homogeneous preferences. Our optimal tax formulas could also have important applications in environmental taxation. Our work is related to the “double dividend” literature that analyzes the interaction between carbon taxes and income tax distortions (e.g. Bovenberg and Goulder 1996; Goulder 1995; Goulder and Williams 2003; Goulder 2013), but this literature assumes linear income taxes and does not consider redistributive motives.

A key contribution of the theoretical work is that it delivers formulas for optimal commodity taxes as a function of sufficient statistics that can be estimated in a wide variety of empirical applications. In Section III, we apply the theory by estimating the necessary statistics for the optimal nationwide tax on sugar-sweetened beverages (hereafter, “SSBs”). We use Nielsen Homescan, a 60,000-household, nationally representative panel dataset of grocery purchases, and Nielsen Retail Measurement Services (RMS), a panel dataset from 37,000 stores covering about 40 percent of all U.S. grocery purchases. A basic plot of the Homescan data in Figure 1 illustrates how SSB taxes could be regressive: households with annual income below \$10,000 purchase about 99 liters of SSBs per adult each year, while households with income above \$100,000 purchase only half that amount.

To identify the price elasticity of demand, we develop an instrument that exploits retail chains’ idiosyncratic pricing decisions for the UPCs that a household usually buys at the retailers where the household usually shops. For example, if retailer X puts Gatorade on sale, people who often buy Gatorade at that retailer face a lower price for their SSBs compared to people who don’t buy Gatorade and/or don’t shop at that retailer. We ensure that the instrument is not contaminated by local or national demand shocks by using the retailer’s average price charged outside of a household’s county and by using only deviations from each UPC’s national average price. Because retailers regularly vary prices independently of each other, while keeping prices fairly rigid across their stores (DellaVigna and Gentzkow, 2017), the instrument delivers higher power while relaxing stronger assumptions required for traditional instruments such as those introduced by Hausman (1996) and Nevo (2001).

We estimate the income elasticity of demand using within-county income variation. We find a small positive income elasticity, which means that the downward-sloping consumption-income profile illustrated in Figure 1 is driven by strong between-income preference heterogeneity, not by causal income effects. This strong preference heterogeneity—but not the declining consumption-income profile *per se*—reduces the socially optimal SSB tax.

To quantify consumer bias, we designed and implemented a survey of 18,000 Homescan households to measure their nutrition knowledge and self-control. Panel (a) of Figure 2 shows that nutrition knowledge is strongly negatively associated with SSB purchases: households whose primary shoppers are in the lowest decile of nutrition knowledge purchase more than twice as many SSBs as households in the highest decile. Panel (b) shows that reported self-control is also strongly negatively associated with SSB consumption: households whose primary shoppers answer that they “definitely” drink SSBs “more often than I should” purchase more than twice as many SSBs as households whose respondents answer “not at all.” These striking relationships raise the possibility that lack of nutrition knowledge and self-control could be important contributors to SSB consumption. Furthermore, these factors might be regressive: people with household incomes below \$10,000 have 0.82 standard deviations lower nutrition knowledge and report that they have 0.40 standard deviations lower self-control than people with household incomes above \$100,000.

Extending the intuition from these figures, we formally quantify consumer bias using what we call the “counterfactual normative consumer” strategy, which builds on Bronnenberg et al. (2015)

and Handel and Kolstad (2015).⁹ We estimate the relationships in Figure 2 after conditioning on a rich set of preference measures and demographics and correcting for measurement error. We then predict “normative” consumption—that is, consumption if people had the nutrition knowledge of dietitians and nutritionists and perfect self-control. To interpret this prediction, we assume that any unobserved preferences are conditionally independent of bias; this unconfoundedness assumption is the key weakness of the “counterfactual normative consumer” approach.

For the average American household, predicted “normative” SSB consumption is only 52 to 62 percent of actual consumption. Put differently, we predict that the average American household would consume 38 to 48 percent fewer SSBs if they had the nutrition knowledge of dietitians and nutritionists and perfect self-control. This estimated overconsumption due to bias is higher among the poor, accounting for 47 to 58 percent of consumption for households with incomes below \$10,000, compared to 34 to 42 percent of consumption for households with incomes above \$100,000. This regressive bias implies a higher optimal soda tax.

In Section IV, we implement our optimal tax formulas using our empirical results. For a broad range of specifications, the optimal soda tax lies in the range of 1.5 to 2.8 cents per ounce—somewhat higher than the levels set in most U.S. cities to date.¹⁰ The welfare benefits from implementing the optimal tax are estimated to lie between \$3.6 billion and \$12.1 billion annually. And although SSB consumption is highly concentrated among low earners, the overall welfare effects are much less regressive, since our estimates imply that the internality corrections are also greatest at low incomes. The welfare gains are about \$400 million per year higher than what would be realized by imposing a one cent per ounce tax—currently the modal policy among U.S. cities that have implemented SSB taxes. Finally, we emphasize the importance of accounting for behavioral biases when designing policy: a tax designed without accounting for behavioral biases foregoes nearly \$1.6 billion per year in potential welfare gains.

In addition to its contributions to optimal tax theory and behavioral public economics, our work also connects to a large and growing empirical literature on SSB taxes. One set of papers estimates the price elasticity of SSB demand and/or the effect of SSB taxes on consumption.¹¹ Our work contributes transparent estimates in a large nationwide sample, whereas most previous papers require more restrictive identifying assumptions or deliver less precise estimates, for example

⁹Our approach is inspired by Bronnenberg et al. (2015), who show that sophisticated shoppers—in their application, doctors and pharmacists—are more likely to buy generic instead of branded drugs. They conduct welfare analysis assuming that only sophisticates’ choices are welfare-relevant. Analogously, we show that “normative” consumers—in our application, people who have as much nutrition knowledge as nutritionists and dietitians and report no self-control problems—purchase fewer SSBs, and we assume that only choices predicted in the absence of bias are welfare-relevant. Bartels (1996), Handel and Kolstad (2015), Johnson and Rehavi (2016), and Levitt and Syverson (2008) similarly compare informed to uninformed agents to identify the effects of imperfect information. All of these papers require the same identifying assumption: that preferences are conditionally uncorrelated with measures of informedness.

¹⁰This represents the range of optimal tax estimates under the prevailing U.S. income tax; under a reformed optimal U.S. income tax, the values span a wider range. See Section IV for details.

¹¹This includes Bollinger and Sexton (2017), Duffey et al. (2010), Finkelstein et al. (2013), Fletcher, Frisvold, and Teft (2010), Rojas and Wang (2017), Silver et al. (2017), Smith, Lin, and Lee (2010), Tiffin, Kehlbacher, and Salois (2015), Zhen et al. (2011), and others; see Andreyeva, Long, and Brownell (2010), Powell et al. (2013), and Thow, Downs, and Jan (2014) for reviews.

because they exploit only one specific SSB tax change. A second set of papers additionally estimates how SSB taxes would affect consumer surplus, including Dubois, Griffith, and O’Connell (2017), Harding and Lovenheim (2015), Wang (2015), and Zhen et al. (2014). These papers do not attempt to quantify consumer bias, making it difficult to use the estimates to evaluate a policy motivated by consumer bias.¹² The most important difference between our paper and the existing soda tax literature is that, to our knowledge, ours is the only one that attempts to answer the following basic question: what is the optimal soda tax? More broadly, this paper provides a theoretical and empirical framework for calculating optimal commodity taxes that can be applied in a wide variety of contexts.

II Deriving the Optimal Sin Tax

We begin with a conventional static optimal taxation setting: individuals have multidimensional heterogeneous types $\theta \in \Theta \subset \mathbb{R}_+^n$, distributed with measure $\mu(\theta)$. They supply labor to generate pre-tax income z , which is subject to a nonlinear tax $T(z)$. Net income is spent on a two goods: a numeraire consumption good c and a “sin good” s , with pre-tax price p , which is subject to a linear commodity tax t . Therefore the individual’s budget constraint is $c + (p + t)s \leq z - T(z)$.

Individuals make their choices to maximize “decision utility” $U(c, s, z; \theta)$ subject to this budget constraint. U is assumed to be increasing and weakly concave in its first two arguments, and decreasing and strictly concave in the third. The policymaker believes consumers should instead maximize “normative utility” $V(c, s, z; \theta)$, which may differ from U . The policymaker selects taxes $T(\cdot)$ and t to maximize aggregate normative utility, subject to a government budget constraint and to individual optimization. That is, the policymaker’s problem is to maximize aggregate normative utility, $W(t, T)$, weighted by type-specific Pareto weights $\alpha(\theta)$, less any externalities E :

$$\max_{T, t} \left[\int_{\Theta} \alpha(\theta) V(c(\theta), s(\theta), z(\theta); \theta) \mu(\theta) - E \cdot s(\theta) \right] \mu(\theta), \quad (1)$$

subject to the government budget constraint

$$\int_{\Theta} (ts(\theta) + T(z(\theta))) \mu(\theta) \geq R \quad (2)$$

and to individual optimization

$$\{c(\theta), s(\theta), z(\theta)\} = \arg \max_{\{c, s, z\}} U(c, s, z; \theta) \quad \text{s.t.} \quad c + (p + t)s \leq z - T(z) \text{ for all } \theta. \quad (3)$$

The difference between U and V can capture a variety of different psychological biases. For example, consumers may have incorrect beliefs about certain attributes of s , such as calorie content, future health costs, or energy efficiency, as documented for food choices and energy-efficiency choices by

¹²Dubois, Griffith, and O’Connell (2017) take a next step by asking whether consumer subgroups that they hypothesize to have larger bias (young people and higher sugar consumers) are more price elastic.

Allcott (2013), Attari et al. (2010), Bollinger, Leslie, and Sorensen (2011), and Schofield (2015). As another, example, consumers may have limited attention or salience bias with respect to certain attributes of s , as modeled in Gabaix and Laibson (2006), DellaVigna (2009), Gabaix (2014), and Bordalo, Gennaioli, and Shleifer (2013), and as documented by Allcott and Taubinsky (2015) for energy-efficient appliances. Finally, present focus may lead consumers to underweight the future health costs of some goods (e.g, potato chips or cigarettes) as in O’Donoghue and Rabin (2006). Our framework allows us to treat present focus as a bias. However, our framework also allows us to study other welfare criteria that may be applied to the model—for example, the policy might place some normative weight both on the “future-oriented self” and on the “in the moment self.”¹³

A key goal of our theoretical analysis is to derive optimal tax formulas that can accommodate a variety of possible consumer biases while still being empirically implementable. We do this by constructing a price metric for consumer bias.

Formally, let $s(p + t, y, z, \theta)$ be the sin good consumption chosen at total price $p + t$ by a type θ individual who earns z and has disposable income y .¹⁴ Analogously, define $s^V(p + t, y, z, \theta)$ to be the amount of s that would be chosen if the individual were maximizing V instead. We define the bias, denoted $\gamma(p + t, y, z, \theta)$, as the value for which $s(p + t, y, z, \theta) = s^V(p + t - \gamma, y - s\gamma, z, \theta)$. In words, γ is equal to the compensated price reduction that produces the same change in demand as the bias does. In terms of primitives, $\gamma = \frac{U'_s}{U'_c} - \frac{V'_s}{V'_c}$.¹⁵ (Throughout, we will use the notation f'_x to denote the derivative of $f(x, y)$ with respect to x , and f''_{xy} for the the cross-partial derivative with respect to x and y , etc. When no ambiguity arises, we sometimes suppress some arguments and write, for example, $s(\theta)$ for concision.) If $\gamma(\theta) > 0$, this means that type θ consumers over-consume relative to if they were maximizing normative utility, whereas $\gamma(\theta) < 0$ means that type θ consumers under-consume.¹⁶

The statistic can be quantified by comparing consumers’ choices in “biased” and “debiased” states, as we will do in our empirical application. Other examples that informally employ our definition of bias include Chetty, Looney, and Kroft (2009) and Taubinsky and Rees-Jones (2018), who quantify the (average) value of tax salience as the change in up-front prices that would alter demand as much as a debiasing intervention that displays tax-inclusive prices. Similarly, Allcott and Taubinsky (2015) run an experiment that directly estimates γ based on consumer’s valuation of energy efficient compact fluorescent lightbulbs (CFLs) with and without an informational intervention meant to eliminate informational and attentional biases. By being tied directly to choice, the bias measure γ is not tied to the policymaker’s aversion to wealth inequality and is invariant to monotonic transformations of the utility functions U and V .

¹³See Bernheim and Rangel (2009) and Bernheim (2016) for a discussion of why treating $\beta < 1$ fully as a bias may be problematic.

¹⁴Recall that because our model allows utility to not be weakly separable in leisure and consumption of s , s depends not only on disposable income but also on earned income z .

¹⁵To see this, note that by definition, $U'_s(c, s, z)/U'_c(c, s, z) = p + t$ with $s \cdot (p + t) + c = z - T(z)$, and $V'_s(c, s, z)/V'_c(c, s, z) = p + t - \gamma$ with $s \cdot (p + t - \gamma) + c = z - T(z) - s\gamma$. This also shows that the equivalence only holds under the assumption that the tax t is fully salient; otherwise, γ may deviate from $\frac{U'_s}{U'_c} - \frac{V'_s}{V'_c}$.

¹⁶Throughout the paper, we will assume that the sole source of disagreement between the consumer and policy maker is about the merits of s ; we do not focus on labor supply misoptimization.

To represent the policymaker’s inequality aversion concisely, we employ the notion, common in the optimal taxation literature, of “social marginal welfare weights”—the social value (from the policymaker’s perspective) of a marginal unit of consumption for a particular individual, measured in terms of public funds. We therefore define $g(\theta) = \alpha(\theta)V'_c/\lambda$, where V'_c represents the derivative of V with respect to its first argument, and λ is the marginal value of public funds (i.e., the multiplier on the government budget constraint at the optimum).¹⁷ These weights are endogenous to the tax system, but are useful for characterizing the necessary conditions that must hold at the optimum. We use $\bar{g} = \int_{\Theta} g(\theta)d\mu(\theta)$ to denote the average marginal social welfare weight. If there are no income effects on consumption and labor supply, then $\bar{g} = 1$ by construction.¹⁸

II.A Two Simple Examples with Discrete Types

We first use two examples to illustrate a key insight: when considering sin taxes which are “regressive”—i.e., taxes on goods consumed predominately by poorer consumers—the optimal tax depends crucially on *why* sin good consumption decreases with income. These insights, and the resulting tax formulas, extend transparently to settings where consumption rises with income, but for concreteness we discuss a declining consumption profile throughout this sub-section.

There are two broad reasons for why consumption of s can decrease with income. First, individuals with different ability levels (and thus different incomes) may have different preferences for the sin good. Second, the sin good may be inferior, so that individuals choose to consume less as their budget increases.¹⁹ In what follows, we describe a simple model with two types of consumers to illustrate the implications of preference heterogeneity vs. income effects for the optimal sin tax.

Our model economy consists of two consumer types, $\theta \in \{L, H\}$, in equal proportion, with wages $w(L) < w(H)$ and utility functions given by

$$U(c, s, z; \theta) = G \left(c + v(s; c, \theta) - \psi \left(\frac{z}{w(\theta)} \right) \right) \quad (4)$$

and

$$V(c, s, z; \theta) = G \left(c + v(s; c, \theta) - \psi \left(\frac{z}{w(\theta)} \right) - \gamma s \right). \quad (5)$$

$v(\cdot)$ is the sub-utility of sin good consumption, $\psi(\cdot)$ is the disutility of labor supply, and $G(\cdot)$

¹⁷Note that this definition implies that g represents the social value of a unit of marginal composite consumption c , rather than sin good consumption. When agents make rational decisions about consumption of s , this distinction is immaterial, since the policymaker places the same value on marginal spending on either good due to the envelope theorem.

¹⁸Note that because the Pareto weights $\alpha(\theta)$ are exogenous, and because U and V produce identical behavior (and identical choice-based measures of bias) up to monotonic transformations, the social marginal welfare weights reflect a policymaker’s or society’s normative preference for reducing wealth inequality—they cannot be inferred by observing behavior. As in the rest of the optimal taxation literature, our formulas for optimal taxes will thus depend both on observable behavior (and people’s quantifiable mistakes) and on the policy maker’s (or society’s) inequality aversion, as encoded by these weights.

¹⁹This distinction between preference heterogeneity and income effects is separate from the question of bias. If poor consumers’ higher consumption is driven by higher bias, the question remains whether that heterogeneous bias is invariant to ability (like preference heterogeneity) or causally affected by income (like an income effect).

is concave to reflect declining marginal utility of consumption. Note that because utility from consumption and labor are separable here, each type's choice of s will depend only on the tax t and after-tax income y . We consider two special cases:

- Pure income effects: $v(s; c, \theta) = v_{inc}(s; c)$ doesn't depend directly on type, only on numeraire consumption. Assume $\frac{\partial^2 v_{inc}}{\partial s \partial c} < 0$, so $s(L) > s(H)$, meaning that s is an inferior good.
- Pure preference heterogeneity: $v(s; c, \theta) = v_{pref}(s; \theta)$ doesn't depend on numeraire consumption, only on type. Assume $v'_s(s; L) > v'_s(s; H)$ for all s , so $s(L) > s(H)$, meaning that high earnings ability is correlated with a lower taste for s .

The optimal policy can be described by two allocations of gross and net-of-tax earnings, with one allocation intended for each type, denoted $(z^\theta, z^\theta - T^\theta)$, for $\theta \in \{L, H\}$. The income taxes T^L and T^H redistribute as much as possible from H to L subject to an incentive compatibility (IC) constraint, which requires that each type prefer their intended allocation to those intended for the other type. As is standard in two-type models, only the downward IC constraint for type H binds at the optimum.²⁰ The policymaker's optimization problem can be solved using a Lagrangian $\mathcal{L} = W + \lambda \cdot RC + \nu \cdot IC$, where RC denotes the nonnegative resource constraint shown in Equation (2), and IC denotes the incentive compatibility constraint, formally written in Footnote 20. Under the optimal policy, $\frac{d\mathcal{L}}{dt} = \frac{d\mathcal{L}}{dT^L} = \frac{d\mathcal{L}}{dT^H} = 0$, and the multiplier λ on RC is the marginal value of public funds.

To solve for the optimal commodity tax, we consider a perturbation of the commodity and income taxes that preserves the IC constraint for labor supply with equality. At the optimum, any such joint reform must have no first-order effect on social welfare, allowing us to write down a first-order condition for the optimal commodity tax. As we shall show, the income tax compensation necessary to preserve the IC constraint depends on whether sin good demand differences are driven by preference heterogeneity or income effects. We consider these cases in turn.

First, consider the case where sin good consumption varies solely due to income effects: $v(s; c, \theta) = v_{inc}(s; c)$. In this case, $s(\theta, y, t) = s(y, t)$: sin good demand depends only on y , and not on θ . Consider now how a sin tax perturbation dt alters the high type consumers' desire to select z^L instead of z^H . An increase dt in the commodity tax increases the relative tax burden of choosing z^L instead of z^H by $[s(y^L, t) - s(y^H, t)]dt$: because s is an inferior good, increasing the tax on it makes low levels of earnings less appealing. Without an income tax reform, this would leave the IC constraint slack,

²⁰Formally, this binding IC constraint is

$$\underbrace{y^H - (1+t)s(H, y^H, t)}_c + v\left(s(H, y^H, t); \underbrace{y^H - (1+t)s(H, y^H, t)}_c, H\right) - \psi\left(\frac{z^H}{w(H)}\right) - \left[\underbrace{y^L - (1+t)s(H, y^L, t)}_c + v\left(s(H, y^L, t); \underbrace{y^L - (1+t)s(H, y^L, t)}_c, H\right) - \psi\left(\frac{z^L}{w(H)}\right)\right] \geq 0$$

where $s(H, y^\theta, t)$ denotes the level of s which maximizes type H 's utility under tax t , conditional on having net-of-tax income y^θ .

implying that more redistribution could be carried out while still ensuring high types choose earnings z^H . A joint income tax reform that preserves equality of the IC constraint is $dT^L = -dt \cdot s(L)$ and $dT^H = -dt \cdot s(H)$. That is, the income tax is reformed to exactly rebate to each type the additional amount that they pay in sin taxes. Therefore, the total effect on the Lagrangian of this joint reform is equal to

$$\frac{d\mathcal{L}}{dt} = \frac{dW}{dt} + \lambda \cdot \frac{dRC}{dt} = \lambda \sum_{\theta} \left[-g(\theta)\gamma \left. \frac{\partial s(\theta)}{\partial t} \right|_u + t \cdot \left. \frac{\partial s(\theta)}{\partial t} \right|_u \right], \quad (6)$$

where the notation $\left. \frac{\partial s(\theta)}{\partial t} \right|_u$ denotes the compensated sin good demand response to a change in t .²¹ The optimality condition $\frac{d\mathcal{L}}{dt} = 0$ implies

$$t = \frac{\sum_{\theta} g(\theta)\gamma \left. \frac{ds(\theta)}{dt} \right|_u}{\sum_{\theta} \left. \frac{ds(\theta)}{dt} \right|_u}. \quad (7)$$

Now consider the second alternative that sin good consumption varies solely due to preference heterogeneity: $v(s; c, \theta) = v_{inc}(s; \theta)$. In this case, $s(\theta, y, t) = s(\theta, t)$: sin good demand depends only on θ , and not on y . Consequently, each type's choice of s will be independent of their choice of earnings, and thus an increase in the commodity tax will not affect the IC constraint in Footnote 20. Intuitively, when preference heterogeneity drives differences in $s(\theta)$, a sin tax perturbation dt requires no income tax reform to preserve labor supply choices. In this case, the perturbation has the following total effect on the Lagrangian:

$$\frac{d\mathcal{L}}{dt} = \frac{dW}{dt} + \lambda \cdot \frac{dRC}{dt} = \lambda \sum_{\theta} \left[-g(\theta) \left(s(\theta) + \gamma \cdot \frac{\partial s(\theta)}{\partial t} \right) + s(\theta) + t \cdot \frac{\partial s(\theta)}{\partial t} \right]. \quad (8)$$

The optimality condition $\frac{d\mathcal{L}}{dt} = 0$ implies

$$t = \underbrace{\frac{\sum_{\theta} g(\theta)\gamma \frac{ds(\theta)}{dt}}{\sum_{\theta} \frac{ds(\theta)}{dt}}}_{\text{corrective motive}} - \underbrace{\frac{\sum_{\theta} s(\theta)(1 - g(\theta))}{\sum_{\theta} \frac{ds(\theta)}{dt}}}_{\text{redistributive motive}}. \quad (9)$$

The first term represents the corrective motive to reduce s . This correction is more valuable to the extent that the consumers who respond most to the tax have higher welfare weight $g(\theta)$. The second term represents the redistributive motive to subsidize commodities consumed primarily by the poor. This term formalizes the common intuition that corrective taxes on goods consumed by the poor should be reduced for redistributive reasons.

²¹To derive (6), note that the mechanical welfare effect of the commodity tax reform is exactly offset by the mechanical welfare effect of the income tax reform and leaves the IC constraint unchanged (and binding with equality). Thus the remaining effect is the (compensated) substitution effect: $\sum_{\theta} (-(p+t)V'_c(c(\theta), s(\theta), z(\theta); \theta) + V'_s(c(\theta), s(\theta), z(\theta); \theta)) \left. \frac{\partial s(\theta)}{\partial t} \right|_u + \lambda \sum_{\theta} t \left. \frac{\partial s(\theta)}{\partial t} \right|_u$. Note that $-(p+t)V'_c(c(\theta), s(\theta), z(\theta); \theta) + V'_s(c(\theta), s(\theta), z(\theta); \theta) = -\lambda g(\theta)\gamma$, yielding the right hand side of 6.

Note that the welfare weights encoding inequality aversion appear in both terms. Although their appearance in the redistributive motive term is intuitive, their presence in the corrective motive is perhaps more surprising. This result is notable because it stands in contrast to the familiar “additivity result” for externalities (Sandmo, 1975), more generally called the “principle of targeting” by Kopczuk (2003), which demonstrates that under general assumptions, redistribution and externality correction are largely separable problems. The analogue of this logic fails for (heterogeneous) internalities because the optimal corrective commodity tax generally depends on the profile of social marginal welfare weights.

In the benchmark case of a constant elasticity of demand for s , inequality aversion pushes the corrective motive in the opposite direction from the redistributive motive: greater inequality aversion pushes toward a *larger* corrective tax. Intuitively, the corrective benefit is proportional to the change in s consumption, and constant elasticity of demand implies that those who consume the most sin good (poor consumers) reap the largest benefits from correction. A more inequality-averse policymaker values this concentration of corrective benefits on the poor more heavily, and thus perceives a larger corrective benefit from taxation.

II.B Optimal Sin Tax in a General Setting

The insights from the preceding examples extend to a more general setting with a continuum of types and multidimensional heterogeneity in ability, bias, demand elasticities, and preferences, under some weak assumptions. This section presents that more general setup, and the resulting optimal tax formulas.

II.B.1 Elasticity Concepts, Sufficient Statistics, and Simplifying Assumptions

The optimal sin tax depends on three types of sufficient statistics: elasticities, money-metric bias, and the “progressivity of bias correction.” We begin by defining the elasticities. These statistics are understood to be endogenous to the tax regime (t, T) , though we suppress those arguments for notational simplicity.

- $\zeta(\theta)$: the price elasticity of demand for s from type θ , formally equal to $-\left(\frac{ds(\theta)}{dt}\right)\frac{p+t}{s(\theta)}$.
- $\zeta^c(\theta)$: the compensated price elasticity of demand for s , equal to $-\left(\frac{ds(\theta)}{dt}\Big|_u\right)\frac{p+t}{s(\theta)}$.
- $\eta(\theta)$: the income effect on s expenditure, equal to $\zeta - \zeta^c$.
- $\xi(\theta)$: the causal income elasticity of demand for s , equal to $\frac{d}{dz}s(p+t, z - T(z), z; \theta) \cdot \frac{z}{s}$.²²

²²A change in net earnings may come from either a change in labor supply or a change in virtual income (e.g., due to a tax level reduction). If sin good consumption and labor are weakly separable in the utility function, sin good consumption will respond identically to either type of change in earnings. In that case, ξ is equal to $\frac{\eta(\theta)}{p+t} \cdot \frac{z(\theta)}{s(\theta)}(1 - T'(z(\theta)))$. More generally, ξ quantifies the sin good response to an increase in earnings *from labor*, for example due to a local reduction in marginal income tax rates. To test for weak separability, note that when s changes in response to a household income shock dz , weak separability implies that the response $\frac{ds}{dz}$ will depend on whether dz comes from a change in hours or nonlabor earnings, and so the estimated relationship $\frac{ds}{dz}$ will be

Additionally, we represent the labor supply response to tax reforms using the following parameters, which are defined formally in Appendix A. These responses are defined to include the full sequence of adjustments due to any nonlinearities in the income tax. (See Jacquet and Lehmann (2014) for an extensive discussion of such “circularities”).

- $\zeta_z^c(\theta)$: the compensated elasticity of taxable income with respect to the marginal income tax rate.
- $\eta_z(\theta)$: the income effect on labor supply.

It will prove convenient to average these statistics across all agents. We denote such averages using “bar” notation; for example, average consumption of s is denoted $\bar{s} = \int_{\Theta} s(\theta) d\mu(\theta)$, with aggregate elasticity of demand $\bar{\zeta} = -\left(\frac{d\bar{s}}{dt}\right) \frac{p+t}{\bar{s}}$. Similarly, we denote average consumption among consumers with a given income z as $\bar{s}(z)$, with income-conditional elasticities denoted by $\bar{\zeta}(z) = -\left(\frac{d\bar{s}(z)}{dt}\right) \frac{p+t}{\bar{s}(z)}$. The income distribution is denoted $H(z) = \int_{\Theta} 1\{z(\theta) \leq z\} d\mu(\theta)$, with income density denoted $h(z)$.

As discussed in Section II.A, it is helpful to distinguish between two sources of cross-sectional variation in $\bar{s}(z)$: income effects and preference heterogeneity. Let $\bar{s}'(z)$ denote the cross-sectional change in s with respect to income z at a particular point in the income distribution. This total derivative can be decomposed into two partial derivatives: the (causal) income effect, $s'_{inc}(z)$, and between-income preference heterogeneity $s'_{pref}(z)$. The causal income effect can be defined in terms of the empirically estimable income elasticity of s , as $s'_{inc}(z) = \mathbb{E}\left[\xi(\theta) \frac{s(\theta)}{z} \mid z(\theta) = z\right]$, and between-income preference heterogeneity as the residual: $s'_{pref}(z) = \bar{s}'(z) - s'_{inc}(z)$.²³ Using these definitions, we define our key sufficient statistic for preference heterogeneity, “cumulative between-income preference heterogeneity”:

$$s_{pref}(z) := \int_{x=z_{min}}^z s'_{pref}(x) dx. \quad (10)$$

This term quantifies the amount of sin good consumption at income z , relative to the lowest income level z_{min} , that can be attributed to preference heterogeneity rather than income effects.²⁴

To aggregate bias across individuals, we will follow Allcott, Mullainathan, and Taubinsky (2014) and Allcott and Taubinsky (2015) in defining the *average marginal bias*

insensitive to the inclusion of controls for hours worked. Using this test, we find support for weak separability in our empirical application, as discussed in footnote 35.

²³We describe this computation of $s'_{pref}(z)$ as a residual, because that is the estimation strategy we use in Section III. Alternatively, one could attempt to quantify $s'_{pref}(z)$ directly using data on preference heterogeneity, letting $s'_{inc}(z)$ be a residual.

²⁴One question that arises here is the relevant time horizon for income effects, since preferences may themselves be endogenous to income over long periods. For example, consumption patterns of children in poor households—including those causally driven by low income—may affect their preferences later in life, even if their income has increased. Conceptually, the relevant statistic for the optimal tax formula is the labor supply distortion generated by the commodity tax, which is proportional to the income effect on that commodity, measured over the *same time horizon as labor supply decisions*. Therefore, to the extent that preferences are endogenous to income patterns at longer horizons than labor supply decisions, these should be treated as preference heterogeneity rather than income effects. In practice, our empirical estimates of income effects will be estimated using annual income variation.

$$\bar{\gamma} = \frac{\int_{\Theta} \gamma(\theta) \left(\frac{ds(\theta)}{dt} \Big|_u \right) \mu(\theta)}{\int_{\Theta} \left(\frac{ds(\theta)}{dt} \Big|_u \right) \mu(\theta)}. \quad (11)$$

Intuitively, this aggregation represents the marginal bias weighted by individuals' marginal responses to a tax reform that raises t while reforming T to offset the average effect on wealth at each income z . In other words, if a tax perturbation causes a given change in total consumption of s , $\bar{\gamma}$ is the average amount by which consumers over- or under-estimate the change in utility from that change in consumption. We additionally define $\bar{\gamma}(z)$ as the response-weighted bias conditional on income.²⁵

Because our framework considers redistributive motives, unlike Allcott, Mullainathan, and Taubinsky (2014) and Allcott and Taubinsky (2015), we must also account for the progressivity of bias correction:

$$\sigma := Cov \left[g(z), \frac{\bar{\gamma}(z)}{\bar{\gamma}} \frac{\bar{\zeta}^c(z)}{\bar{\zeta}^c} \frac{\bar{s}(z)}{\bar{s}} \right]. \quad (12)$$

The term σ is the covariance of welfare weight with the product of consumption-weighted bias and (compensated) elasticity. If this term is positive, it indicates that bias reductions in response to a tax increase are concentrated among individuals with high welfare weights, i.e., those with lower incomes.

We impose the following assumptions, common in the optimal commodity taxation literature, in order to focus on the interesting features of sin taxes in a tractable context.

Assumption 1. *Constant social marginal welfare weights conditional on income: $g(\theta) = g(\theta')$ if $z(\theta) = z(\theta')$.*

This assumption is analogous to Assumption 1 in Saez (2002a). It holds immediately if types are homogeneous conditional on income. More generally, Saez (2002a) argues this is a reasonable normative requirement even under heterogeneity “if we want to model a government that does not want to discriminate between different consumption patterns...” Therefore we sometimes write $g(z)$ to denote the welfare weight directly as a function of earnings.

Assumption 2. *U and V are smooth functions that are strictly concave in c , s , and z , and μ is differentiable with full support.*

Assumption 3. *The optimal income tax function $T(\cdot)$ is twice differentiable, and each consumer's choice of income z admits a unique global optimum, with the second-order condition holding strictly at the optimum.*

²⁵Formally,

$$\bar{\gamma}(z) = \frac{\int_{\Theta} \gamma(\theta) \left(\frac{ds(\theta)}{dt} \Big|_u \right) 1\{z(\theta) = z\} d\mu(\theta)}{\int_{\Theta} \left(\frac{ds(\theta)}{dt} \Big|_u \right) 1\{z(\theta) = z\} d\mu(\theta)}.$$

Assumptions 2 and 3 ensure that the income distribution does not exhibit any atoms and individuals' labor supply and consumption decisions respond smoothly to perturbations of the tax system (Jacquet and Lehmann, 2014).

Assumption 4. *One of the two following conditions hold: (a) heterogeneity is unidimensional, so that individuals with a given ability all have the same preferences and behavioral responses, or (b) the sin good s accounts for a small share of all consumers' budgets (so that terms of order $\frac{(p+t)s}{z}$ are negligible) and s is orthogonal to η_z and ξ_z conditional on income.²⁶*

Assumption 4, which is realistic for sin goods such as sugary drinks that are a relatively small share of expenditures, simplifies the exposition of the formula for the optimal commodity tax. However, it is not necessary, and the full optimal commodity tax without this assumption is derived in Appendix C.A, Proposition 5. In the absence of this assumption, income-conditional covariances between the elasticities affect the income and substitution effects on the sin good, as reflected in the general formula.

II.B.2 General Expression for the Optimal Sin Tax

To characterize the optimal commodity tax, it is helpful to define the *social marginal utility of income*, denoted $\hat{g}(z)$, which is defined (as in Farhi and Gabaix, 2015) as the average welfare effect of marginally increasing the incomes of consumers currently earning income z . The weights $\hat{g}(z)$ incorporate any fiscal externalities resulting from income effects, and also the social welfare affect from mis-spending this marginal income due to bias.²⁷ In Appendix B.C, we provide formulas expressed entirely in terms of the standard social marginal welfare weights $g(z)$. All proofs are contained in Appendix C.A. Letting $e = E/\lambda$ denote externality costs from s in terms of public funds, we obtain the following proposition.

Proposition 1. *The commodity tax t and the income tax T satisfy the following at the optimum:*

$$t = \frac{\bar{\gamma}(\bar{g} + \sigma) + e - \frac{p}{\bar{s}\bar{\zeta}^c} \text{Cov}[\hat{g}(z), s_{pref}(z)]}{1 + \frac{1}{\bar{s}\bar{\zeta}^c} \text{Cov}[\hat{g}(z), s_{pref}(z)]} \quad (13)$$

$$T'(z^*) = \frac{\mathbb{E} \left[(g(z^*)\gamma(\theta) + e - t) \frac{\xi(\theta)s(\theta)\zeta_z^c(\theta)}{\zeta_z^c(z^*)z^*} \mid z(\theta) = z^* \right] + \frac{1}{\zeta_z^c zh(z^*)} \mathbb{E} [(1 - \hat{g}(z)) \mid z \geq z^*]}{1 + \frac{1}{\zeta_z^c(z^*)zh(z^*)} \mathbb{E} [(1 - \hat{g}(z)) \mid z \geq z^*]} \quad (14)$$

When income effects on labor supply are small and s is a small share of expenditures, the commodity tax is approximated by

$$t = \frac{\bar{\gamma}(1 + \sigma) + e - \frac{p}{\bar{s}\bar{\zeta}^c} \text{Cov}[g(z), s_{pref}(z)]}{1 + \frac{1}{\bar{s}\bar{\zeta}^c} \text{Cov}[g(z), s_{pref}(z)]} \quad (15)$$

²⁶Formally, this assumption implies that though the demand elasticity of s with respect to labor earnings, denoted ξ , say be non-negligible, the corresponding income effect $\xi \cdot \frac{s}{z} = \frac{ds(z, \theta)}{dz}$ is negligible, as is $\eta(\theta)$. The optimal tax condition in Proposition 1 below then represents the limiting optimal tax as these values go to zero.

²⁷Formally, $\hat{g}(z^*) = g(z^*) + \mathbb{E} \left[\eta_z(\theta) \frac{T'}{1-T'} \mid z(\theta) = z^* \right] + \mathbb{E} \left[(t - g(\theta)\gamma(\theta) - e) \left(\frac{\eta(\theta)}{p+t} + \frac{\eta_z(\theta)}{1-T'} \frac{\xi(\theta)s(\theta)}{z(\theta)} \right) \mid z(\theta) = z^* \right]$. See our derivation of the optimal income tax in Appendix C.A for intuition about this. When η is negligible (guaranteed by a small expenditure share $(p+t)s/z$) and when η_z is negligible, $\hat{g} = g$.

In our empirical implementation of the formula, and in our discussion of the intuition of the formula, we rely on the approximation that income effects on labor supply are small and that s is a small share of expenditures. The negligible income effects assumption is supported by Gruber and Saez (2002), who find small and insignificant income effects on labor supply and by Saez, Slemrod, and Giertz (2012), who review the empirical literature on labor supply elasticities and argue that “in the absence of compelling evidence about significant income effects in the case of overall reported income, it seems reasonable to consider the case with no income effects.”

Equation (15) shows that the optimal tax is the combination of two main terms. The numerator’s first term, $\bar{\gamma}(1 + \sigma) + e$, corresponds to the corrective motive of the tax. This term is increasing in the average marginal bias $\bar{\gamma}$, the extent to which bias correction is concentrated on the low-income consumers, and the size of the externality. This term illustrates a key difference between externalities and internalities. In the case of consumer bias, it is necessary to know whether the bias is bigger for the rich or for the poor, and whether the rich or the poor are more elastic. This is because the costs from a given individual’s overconsumption fall back on that individual, and thus are scaled by the social marginal welfare weight on the individual. The externality generated by any given individual, however, is borne by the whole population, and thus receives the same weight regardless of whose consumption generates it.

In the numerator’s second term, $Cov[g(z), s_{pref}(z)]$ scales the redistributive effects of the tax. This term depends on the extent to which the variation of sin good consumption by income is due to preference heterogeneity rather than income effects. The intuition for this result stems from an observation highlighted in the above two-type example: the labor supply distortion from a commodity tax is proportional to the income effect on the taxed good. This labor supply distortion is neutralized with an offsetting income tax reform (proportional to $s_{inc}(z)$) that adjusts marginal income tax rates in proportion to the income effect on s . Therefore, the net distributional effect of the commodity tax is proportional to the remaining, un-neutralized residual $s_{pref}(z)$. This is the component of the commodity tax that serves as a tag for ability, the distributional effects of which cannot already be achieved via an income tax reform.

II.B.3 General Expression for the Optimal Sin Tax at a Fixed Income Tax

Tax authorities may not be able to re-optimize the income tax system at the same time that a sin tax is imposed. For example, in the U.S., many soda taxes are set by cities that do not control the income tax structure. What is the optimal sin tax in that case? The causal income elasticity always plays a key role, even when the income tax is not necessarily optimal. As before, we define $g(z)$ to be the marginal utility from the numeraire, normalized by the marginal value of public funds.

Proposition 2. *Assuming s is a small share of expenditures and that η_z is negligible, the optimal commodity tax satisfies:*

$$t \approx \frac{\bar{s}\bar{\zeta}^c (\bar{\gamma}(\bar{g} + \sigma) + e) - p (Cov [g(z), s(z)] + \mathcal{A})}{\bar{s}\bar{\zeta}^c + Cov [g(z), s(z)] + \mathcal{A}}, \quad (16)$$

where $\mathcal{A} = \mathbb{E} \left[\frac{T'(z(\theta))}{1-T'(z(\theta))} \zeta_z(\theta) s(\theta) \xi(\theta) \right]$.

The expression for the constrained optimal SSB tax resembles that in Proposition 1, but replaces the term $Cov [g(z), s_{pref}(z)]$ with $Cov [g(z), s(z)] + \mathcal{A}$. This reflects the fact that under the optimal income tax, the fiscal externality from income adjustments in response to the tax (captured by the term \mathcal{A} in Proposition 2) is exactly equal to $Cov [g(z), s_{inc}(z)]$, leaving $Cov [g(z), s_{pref}(z)]$ as the residual term. Note that \mathcal{A} is proportional to the income elasticity ξ . Intuitively, when the sin good is a normal good, a higher sin tax is equivalent to an increase in the marginal income tax rate. The converse obtains when the sin good is an inferior good. When the income tax is set optimally, the second and third terms combine to produce the formula in Proposition 1.

II.C Interpretation and Implications

The optimal tax formulas above illustrate a number of important insights about the forces which affect the optimal sin tax. We highlight four key implications and special cases. To streamline the discussion, we focus on the slightly simpler expression from Equation (15).

1. Special case: no corrective concerns. An important special case of our framework is the case in which there are no corrective concerns. In this case,

$$\frac{t}{p+t} = -\frac{Cov [g(z), s_{pref}(z)]}{\bar{s}\bar{\zeta}^c}. \quad (17)$$

Equation (17) bears a striking resemblance to Diamond's (1975) "many-person Ramsey tax rule." Diamond (1975) studies a Ramsey framework in which the income tax is constrained to be a lump-sum transfer, and obtains almost the same expression as Equation (17), except with $s_{pref}(z)$ replaced by $\bar{s}(z)$. Equation (17) shows that the optimal commodity tax still adheres to the familiar inverse elasticity rule, and it clarifies that instead of taxing goods that high-income consumers *consume*, we should tax goods that high-income consumers *prefer*.

Equation (17) generalizes the seminal Atkinson-Stiglitz theorem to the case of arbitrary preference heterogeneity. The Atkinson-Stiglitz theorem itself obtains as a special case of (17) when all variation in s consumption is driven by income effects, which then implies that $s_{pref} \equiv 0$ and thus $t = 0$. Additionally, Equation (17) generalizes the qualitative analyses about the usefulness of commodity taxes by Saez (2002b) and Edwards, Keen, and Tuomala (1994), the results for two-type models by Edwards, Keen, and Tuomala (1994) and Nava, Shroyen, and Marchand (1996), and results for homogeneous preferences with consumption-labor complementarities by Jacobs and de Mooij (2015).

More generally, this formula holds when $\bar{\gamma} = 0$, which may arise even with non-zero internalities if the consumers who over-consume (or under-consume) the good are inelastic to the tax (or

subsidy). A stark implication of the result then is that when lower-income consumers prefer the good more, the optimal sin tax will be negative (a sin subsidy). This captures the spirit of a key result of Bernheim and Rangel (2004) about the optimality of subsidizing addictive goods when the marginal utility of income is increasing with the consumption of the addictive good.²⁸ Although the Bernheim and Rangel (2004) result that the sin good should be subsidized is seemingly in stark contrast to the sin tax results in O’Donoghue and Rabin (2006) and elsewhere, our general tax formula clarifies the economic forces that lead to each result.

2. Special case: no inequality aversion. If social marginal welfare weights are constant (implying the policymaker has no desire to redistribute marginal resources from high to low income consumers) then $t = \bar{\gamma} + e$. This matches the core principal of Pigouvian taxation and the typical sin tax results in the behavioral economics literature (e.g., O’Donoghue and Rabin 2006; Mullainathan, Schwartzstein, and Congdon 2012; Allcott, Mullainathan, and Taubinsky 2014; Allcott and Taubinsky 2015). This special case could arise either if welfare weights are constant across incomes (e.g., if V is linear in c) or if there is no income inequality (so that all individuals have the same marginal utility of consumption). In both cases, the optimal commodity tax must exactly offset the average marginal bias plus the externality.

3. Special case: no preference heterogeneity. When differences in consumption are due purely to differences in income, then regressive consequences of a sin tax can be perfectly offset by modifications to the income tax, so that there are no fiscal externalities resulting from the full policy reform. In this case the optimal sin tax is

$$t = \bar{\gamma} (\bar{g} + \sigma) + e. \tag{18}$$

This formula generalizes Equation (7) from the two-type case. Note that inequality aversion still plays a role in the size of the optimal commodity tax, as reflected by the σ term. The more price elastic low-income consumers are, and the more biased they are relative to the high income consumers, then the higher is σ and thus the higher is the benefit of bias correction. When bias and elasticity are constant across incomes, $\sigma > 0$, and the size of the optimal tax will exceed the optimal Pigouvian tax that prevails absent any inequality aversion.

4. A higher sin good demand elasticity reduces the redistributive motive relative to the corrective motive. When the demand elasticity grows large, the redistributive motive becomes small and the optimal tax t approaches $\bar{\gamma}(\bar{g} + \sigma) + e$, which is special case 3 above. At the opposite extreme, when the elasticity grows small, the corrective motive becomes negligible and the optimal tax approaches $-\frac{Cov[g(z), s_{pref}(z)]}{\bar{s}\zeta^c}$, which is special case 1. More generally, if preference heterogeneity

²⁸Because in Bernheim and Rangel (2004) over-consumption of the good is a consequence of cue-triggered neural processes that render the individual inelastic to prices, the average bias of consumers who are elastic to the tax is zero. Thus, their subsidy result is a direct consequence of the positive covariance between consumption of the addictive good and the marginal utility from income—analogous to our redistributive effect term.

accounts for *any* share of the decrease in s consumption across incomes, then for a sufficiently low elasticity the optimal tax becomes negative (a subsidy). Intuitively, if individuals do not respond to commodity taxes, then such taxes become a powerful instrument to enact redistribution through targeted subsidies.

These results clarify that the role of consumer bias in determining both the sign and magnitude of the optimal commodity tax is modulated by the price elasticity of demand. But perhaps most importantly, the demand elasticity also provides practical guidance on how *sensitive* the optimal tax is to different values of the bias $\bar{\gamma}$ and the externality e . A lower elasticity dampens the responsiveness of the optimal tax to the bias $\bar{\gamma}$, since the corrective benefits in Propositions 1 and 2 depend on the products $\bar{\zeta}^e \bar{\gamma}$ and $\bar{\zeta}^e e$. Simply put, learning that the average marginal bias or the externality are \$1 per unit higher than previously thought does not imply that the optimal tax should increase by \$1—the optimal adjustment could be higher or lower, depending on the demand elasticity.

II.D Extensions

Our theoretical results can be extended in a number of ways, which we do in Appendix B. In Appendix B.A, we derive a general condition for the optimal sin tax when there are more dimensions of consumption, allowing for substitution across goods with various degrees of marginal bias. To the extent to that consumers substitute from the sin good to another good that they also over-consume, the optimal sin tax on the good in question should be lowered. However, the above optimal sin tax formula remains unchanged if either there are no cross-price effects or if the goods consumers substitute to do not generate internalities or externalities.

In Appendix B.B, we also characterize results when the sin good s is a composite good consisting of several different items, such as soft drinks of various sizes. In our companion paper, Allcott, Lockwood, and Taubinsky (2018), we also study the case in which individuals are inattentive to the implicit labor supply incentives generated by the commodity tax on s .

III Estimating Key Parameters for the Optimal Soda Tax

In this section, we gather the empirical parameters needed to calibrate the optimal nationwide tax on sugar-sweetened beverages. First, we describe the data sources. Second, we estimate the price and income elasticities, and how elasticity varies by income. Third, we decompose the SSB consumption vs. income relationship into causal income effects $s'_{inc}(z)$ vs. between-income preference heterogeneity $s'_{pref}(z)$. Fourth, we estimate bias γ , and how this varies by income. Fifth, we discuss the externality e . In a final subsection, we assess the strengths and weaknesses of our empirical approaches.

III.A Data

III.A.1 Nielsen Retail Measurement Services and Homescan data

The Nielsen Retail Measurement Services (RMS) data include sales volumes and sales-weighted average prices at the UPC-by-store-by-week level at about 37,000 stores each year from 106 retail chains for 2006–2016.²⁹ RMS includes 53, 32, 55, 2, and 1 percent of national sales in the grocery, mass merchandiser, drug, convenience store, and liquor channels, respectively. For a rotating subset of stores that Nielsen audits each week, we also observe merchandising conditions: whether each UPC was “featured” by the retailer in the market where each store is located (through newspaper or online ads and coupons) and whether the UPC was on “display” inside each store.³⁰

To measure household grocery purchases, we use the Nielsen Homescan Panel for 2006-2016.³¹ Homescan includes about 38,000 households in 2006, and about 61,000 households each year for 2007-2016.

Each year, Homescan households report demographic variables such as household income (in 16 bins), educational attainment, household composition, race, binary employment status, and weekly hours worked (in three bins). Panel (a) of Table 1 presents descriptive statistics for Homescan households at the household-by-year level. All households report either one or two heads. If there are two heads, we use their average age, years of education, employment status, and weekly work hours. The U.S. government Dietary Guidelines provide calorie needs by age and gender; we combine that with Homescan household composition to get each household member’s daily calorie need. Household size in “adult equivalents” is the number of household heads plus the total calorie needs of all other household members divided by the nationwide average calorie consumption of household heads. In all tables and figures, we weight the sample for national representativeness.³²

Nielsen groups UPCs into product modules. We define sugar-sweetened beverages (SSBs) as the product modules that have typically been included in existing SSB taxes: fruit drinks (which includes sports drinks and energy drinks), pre-made coffee and tea (for example, bottled iced coffee and iced tea), carbonated soft drinks, and non-carbonated soft drinks (which includes cocktail mixes, breakfast drinks, ice pops, and powdered soft drinks). Fruit and vegetable juice and artificially sweetened drinks such as diet soda are not included. The bottom two rows of Panel (a) of Table 1 show that the average Homescan household purchases 153 liters of SSBs per year,

²⁹We define a retail chain by the RMS “parent code,” with the exception that we define a separate retail chain for any RMS “retailer code” within a parent code that has more than 100 store-by-year observations. We rely primarily on parent codes because there are occasional errors in assigning retailer codes to RMS stores.

³⁰The raw data include feature and display information for one original set of audited stores, but a retailer’s advertising captured by the feature variable is generally market-wide. Following standard practice with the RMS data (Kilts Center 2018), we thus impute the feature variable from audited stores to all other stores in a UPC-retailer-market-week cell, where “markets” are defined using Nielsen’s Designated Market Areas. After this imputation, we observe the feature and display variables in approximately 8.5 and 81 percent of store-by-week observations, respectively.

³¹See Einav, Leibtag, and Nevo (2010) for a validation study of Homescan.

³²For analyses using the full sample, we use the “projection factors” provided by Nielsen. For analyses using the subsample of PanelViews survey respondents, we construct our own nationally representative sample weights using the first seven variables presented in Panel (a) of Table 1.

at an average price of \$1.12 per liter.³³ (Average price paid is undefined for the 3.3 percent of household-by-year observations with no SSB purchases.) We deflate all prices and incomes to real 2016 dollars.

There are two important differences between Homescan purchase data and actual SSB consumption. First, Homescan does not include data on beverages purchased and consumed away from home, such as at restaurants and vending machines. Second, people might give soda to others or throw it out instead of drinking it themselves. We discuss the implications of these differences in Section III.F.

III.A.2 Homescan PanelViews Survey

We designed a special survey to measure biases and preferences affecting SSB consumption. Using its PanelViews survey platform, Nielsen fielded the survey in October 2017 to all adult heads of the approximately 60,000 eligible households that were in the 2015 or 2016 Homescan data. We have complete responses from 20,640 people at 18,159 households; there are 2,481 households where both heads responded. Panel (b) of Table 1 summarizes the respondent-level data. Appendix D gives the exact text of the survey questions described below.

We attempt to quantify two classes of consumer bias that might drive a wedge between consumers' decisions and normative utility: imperfect nutrition knowledge and imperfect self-control. To measure nutrition knowledge, we delivered 28 questions from the General Nutrition Knowledge Questionnaire (GNKQ).³⁴ The GNKQ is widely used in the public health literature; see Kliemann, Wardle, Johnson, and Croker (2016) for a validation study. The nutrition knowledge variable is the share correct of the 28 questions; the average score was approximately 0.70 out of 1.

To measure self-control, we asked respondents to state their level of agreement with the following statements: "I drink soda pop or other sugar-sweetened beverages more often than I should," and, if the household has a second head, "The other head of household in my house drinks soda pop or other sugar-sweetened beverages more often than they should." There were four responses: "Definitely," "Mostly," "Somewhat," and "Not at all." To construct the self-control variable, we code those responses as 0, 1/3, 2/3, and 1, respectively.

To measure taste and preference heterogeneity, we asked, "Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking the following?" We asked this question for five types of SSBs (sweetened juice drinks, regular soft drinks, pre-made coffee and tea, sports drinks, and caffeinated energy drinks) and two non-SSBs

³³We measure quantity interchangeably in liters and ounces, under the standard approximation that all drinks have the density of water. For liquid drinks, the weight is reported by Nielsen directly from the package label. For powdered drinks, we transform to the weight when consumed as liquid, i.e. with water added.

³⁴One example question is, "If a person wanted to buy a yogurt at the supermarket, which would have the least sugar/sweetener?" The four possible responses were "0% fat cherry yogurt," "Plain yogurt," "Creamy fruit yogurt," and "Not sure." A second example is, "Which is the main type of fat present in each of these foods?" The five possible responses were "Polyunsaturated fat," "Monounsaturated fat," "Saturated fat," "Cholesterol," and "Not sure." This question was asked about olive oil (correct answer: monounsaturated), butter (saturated), sunflower oil (polyunsaturated), and eggs (cholesterol).

(100% fruit juice and diet soft drinks). To measure health preferences, we asked, “In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?” Responses to each question were originally on a scale from 0 to 10, which we rescale to between 0 and 1. Finally, we asked gender, age, occupation, and whether the respondent makes the majority of the grocery purchase decisions.

III.B Price and Income Elasticities

III.B.1 Empirical Model

In this section, we estimate the price and income elasticities of demand, ζ and ξ . Let s_{it} denote Homescan SSB purchases (in liters per adult equivalent) by household i in quarter t . Let p_{it} denote the average price paid per liter of SSB purchased, and let \mathbf{f}_{it} denote a vector of feature and display variables, whose construction we detail below. z_{ct} is the mean per capita income reported by the Bureau of Economic Analysis (2017) for county c in the calendar year that contains quarter t , ω_t is a vector of quarter of sample indicators, and μ_{ic} is a household-by-county fixed effect. We estimate the following regression:

$$\ln s_{it} = -\zeta \ln p_{it} + \xi \ln z_{ct} + \nu \mathbf{f}_{it} + \omega_t + \mu_{ic} + \varepsilon_{it}, \quad (19)$$

with standard errors clustered by county and with $\ln p_{it}$ instrumented in a manner we describe below.³⁵

Because SSBs are storable, previous purchases could affect current stockpiles and thus current purchases, and Hendel and Nevo (2006b) and others document stockpiling in weekly data. In our quarterly data, however, there is no statistically detectable effect of lagged prices and merchandising conditions on current purchases, and it is statistically optimal to not include lags in Equation (19) according to both the Akaike and Bayesian information criteria. See Appendix E for details.

We use county mean income z_{ct} instead of Homescan panelists’ self-reported income because we are concerned about measurement error in the within-household self-reported income variation, because there is uncertainty as to the time period for which the self-reported incomes apply, and because variation in income that is not due to variation in labor market conditions is less likely to be exogenous to preferences for SSB; see Appendix F for details. We think of Equation (19) as the reduced form of an instrumental variable (IV) regression with time-varying household income as an endogenous variable and county mean income as an instrument, where the first stage should have a coefficient of one.³⁶

³⁵As we show in Appendix Table A7, employment status and weekly hours worked are not statistically significantly associated with SSB consumption when included in Equation (2), so we cannot reject weak separability of SSB consumption and labor. Thus, it does not matter whether variation in z_{ct} results from non-labor windfalls such as government benefits or from wage changes, nor does it matter whether such a wage change results from a shift in local labor supply or demand.

³⁶This approach would be biased if migration causes changes in county mean income that are correlated with within-household SSB demand changes and with income changes for continuing residents. However, migration is likely responsible for only a small share of year-to-year county income variation. Bias could also arise if county income changes are correlated with unobserved within-household preference changes, for example if high-socioeconomic status

III.B.2 Price Instrument Using Chain-Level Time-Varying Prices

A key challenge in demand estimation is addressing simultaneity bias: omitted variables bias generated by a potential correlation between price and unobserved demand shifters. We address simultaneity bias using a price instrument leveraging two facts documented by DellaVigna and Gentzkow (2017) and Hitsch, Hortacsu, and Lin (2017). First, retail chains vary prices over time in a highly coordinated way across their stores: if retailer X is offering Gatorade on sale right now in Toledo, it’s probably also offering Gatorade on sale in Topeka. Second, different chains vary their prices independently of each other over time: retailer X’s current sale has little relationship to what other retailers are doing. Appendix G illustrates these patterns in more detail.

To construct the instrument, define $\ln p_{jkw}$ as the natural log of the price charged at store j for UPC k in week w . Further define $\ln p_{kw}$ as the national average of natural log price of UPC k in week w , unweighted across stores. Then, let $\ln p_{krt,-c}$ denote the unweighted average of $\ln p_{jkw} - \ln p_{kw}$ at all of retail chain r ’s stores *outside* of county c during quarter t .³⁷ The leave-out construction guarantees that our instrument is not contaminated by store-specific responses to local demand shocks, although in practice the leave-out construction makes little difference because price variation is so coordinated within chains. Differencing out the national average price helps remove responses to national-level demand shocks that might influence the price of the specific UPC k , which could still be a concern even after we condition on time fixed effects ω_t that soak up shocks to overall SSB demand.

To construct an instrument for the average SSB price faced by each household, we fit the leave-out price deviations $\ln p_{krt,-c}$ to the household’s average purchasing patterns. Define s_{ijk} as household i ’s total purchases of UPC k at store j while living in county c , and define s_{ic} as household i ’s total SSB purchases while living in county c , both measured in liters. Then $\pi_{ijk} = s_{ijk}/s_{ic}$ is the share of household i ’s SSB liters purchased while in county c that are of UPC k at store j . Household i ’s predicted local price deviation in quarter t is

$$Z_{it} = \sum_{k,j \in RMS} \pi_{ijk} \ln p_{krt,-c}. \quad (20)$$

Price deviations $\ln p_{krt,-c}$ are only observed at RMS stores, so Z_{it} sums only over purchases at RMS stores; approximately 34 percent of SSB purchases are at RMS chains. Because π_{ijk} is the purchase share across all SSB purchases (at both RMS and non-RMS stores), each household’s quantity-weighted prices paid $\ln p_{it}$ should vary approximately one-for-one with Z_{it} .

counties experience both faster income growth and faster declines in preferences for SSBs. However, Appendix Table A7 shows that the estimated $\hat{\xi}$ is unaffected by controlling for a linear time trend in county sample average income.

³⁷In constructing $\ln p_{krt,-c}$, we include only the 81 percent of store-week observations where the feature variable is observed. This is important because the majority of features are associated with a price discount, and Appendix Table A7 shows that omitting the feature variable does change the estimated price elasticity. Therefore, we could introduce bias by including weekly price observations where feature is unobserved. By contrast, in-store displays are less frequently associated with a price decrease, and Appendix Table A7 shows that omitting display does not generate significant omitted variables bias. Thus, to increase power in the instrument, we construct $\ln p_{krt,-c}$ with weekly price observations regardless of whether display is observed.

The feature and display variables are constructed analogously. Let the 2-vector \mathbf{f}_{jkt} denote the number of weeks in which UPC k is featured at store j in quarter t , divided by the number of weeks in which feature is observed for that store in that quarter, as well as the analogous share of weeks in which UPC k is observed to be on display at store j . The feature and display variables we use in the household-by-quarter regressions are

$$\mathbf{f}_{it} = \sum_{k,j \in RMS} \pi_{ijkc} \mathbf{f}_{jkt}. \quad (21)$$

The exclusion restriction is that the local price instrument Z_{it} is uncorrelated with demand shifters ε_{it} , conditional on the set of controls in Equation (19). The key economic content of this assumption is that when retail chains vary prices across weeks and quarters, they do not observe and respond to *chain-specific* demand shocks. One threat to this assumption would be price cuts coordinated with retailer-specific advertising, but retailers do little advertising beyond the newspaper and online ads and coupons that are already captured by RMS feature variable (DellaVigna and Gentzkow, 2017). Furthermore, we show below that the estimates are largely unaffected by alternative instrument constructions and fixed effect controls that address other types of regional and city-specific demand shocks.

III.B.3 Estimation Results

Panels (a) and (b) of Figure 3 present binned scatterplots of the first stage and reduced form of our instrumental variables estimates, conditioning on the other controls in Equation (19). Dividing the reduced form slope by the first stage slope implies a price elasticity of approximately $1.59/1.14 \approx 1.40$.

Table 2 presents estimates of Equation (19). Column 1 presents our primary IV estimates, which give estimated price elasticity $\hat{\zeta} \approx 1.40$ and income elasticity $\hat{\xi} \approx 0.18$.

The exclusion restriction would be violated if chains vary prices in response to chain-specific demand shocks. For example, retailers might respond to local economic downturns in cities where they operate, or to seasonal variation in soft drink demand that could vary across warm and cold cities. Column 2 addresses these concerns by adding city-by-quarter fixed effects. Demand shocks could also vary across chains serving different demographic groups, for example if an economic downturn primarily affects low-income households that shop at some retailers more than others. Column 3 addresses this allowing the city-by-quarter fixed effects to differ for above- versus below-median household income. In both columns, the point estimates move slightly but are statistically indistinguishable.

While these control strategies can address demand shocks that are common across SSB UPCs, they cannot address UPC-specific demand shocks. For example, a warm spring on the east coast might increase demand for soft drinks more than it increases demand for bottled coffee. If retailers were to recognize and respond to this, then the subgroup of east coast households that often buy soft drinks would have a positive demand shock and an instrument Z_{it} that is correlated with that

shock, even conditional on city-by-time fixed effects. Column 4 addresses this concern by using an instrument constructed with deviations from the census region average log price instead of the national average log price. The estimates are again very similar. Thus, for the exclusion restriction to be violated, there must be some specific form of endogeneity not addressed by these multiple alternative specifications.³⁸

To measure whether elasticities vary by income, Appendix Table A8 presents separate estimates of Equation (19) for households with above vs. below \$50,000 income. The price elasticity point estimates are statistically indistinguishable but differ somewhat: 1.48 for below \$50,000 income and 1.35 for above \$50,000 income. The income elasticities are very similar and also statistically indistinguishable.

III.C Causal Income Effects vs. Between-Income Preference Heterogeneity

The second key empirical statistic needed to determine the optimal sin tax is between-income preference heterogeneity $s'_{pref}(z)$. The dark circles on Figure 4 repeat the consumption-income relationship from Figure 1; this is now compressed due to an expanded y-axis range. The curve at the top of the figure uses the income elasticity estimate $\hat{\xi} \approx 0.18$ from Table 2 to predict the causal effects of income increases on the SSB consumption of households earning less than \$10,000 per year: $s_{inc}(z) = \bar{s}(z_{min}) \left(\frac{z}{z_{min}}\right)^{\hat{\xi}}$. Between-income preference heterogeneity is $s_{pref}(z) = \bar{s}(z) - s_{inc}(z)$, the difference between actual consumption and consumption predicted by income effects. On the graph, $s_{pref}(z)$ is thus the vertical difference between the dark circles and the curve, which is drawn in x's.

The estimate of $s_{pref}(z)$ indicates large between-income preference heterogeneity. If all households were exogenously re-assigned to earn the same income, households currently making over \$100,000 per year would purchase 128 liters fewer SSBs than households currently making under \$10,000 per year. This difference is about 1.9 times average consumption. This result that lower-income households have stronger *preferences* for SSBs—regardless of whether they have higher *consumption*—means that in the absence of internalities and externalities, a policymaker would want to subsidize SSBs.

III.D Measuring Bias

III.D.1 The Counterfactual Normative Consumer Estimation Strategy

In section II, we defined bias γ as the compensated price cut that would induce the counterfactual normative self to consume as much of the sin good as the actual biased self. Our *counterfactual*

³⁸The estimates include only observations with positive SSB consumption, as price paid p_{it} is undefined for the 15 percent of quarterly observations with no SSB purchases. In theory, this can bias our estimates, as high prices are more likely to cause zero-purchase observations. Appendix Table A6 addresses this by presenting Tobit estimates (thereby formally accounting for latent demand that is censored at zero) of the reduced form (thereby giving an instrumented price for every observation), with SSB purchases in levels instead of logs (thereby giving a dependent variable for every observation). Price elasticity estimates are economically similar and statistically indistinguishable.

normative consumer empirical strategy directly implements this definition, using an approach that builds on Bronnenberg et al. (2015) and Handel and Kolstad (2015). The process is to use surveys to elicit proxies of bias, estimate the relationship between bias proxies and quantity consumed, use that relationship to predict the counterfactual quantity consumed if consumers instead maximized normative utility, and transform the quantity difference into dollar units using the price elasticity.

To formalize the approach, recall that money-metric bias γ is defined to satisfy $s(p, y, \theta) = s^V(p - \gamma, y - s\gamma, \theta)$. We now log-linearize this demand equation and use an i subscript for each household in the data, recognizing that each i maps to a (p, y, θ) triple. This gives

$$\ln s_i = \ln s_i^V + \zeta_i^c \gamma_i / p_i, \quad (22)$$

where $\ln s_i^V$ denotes the log of the quantity that household i would consume in the absence of bias, s_i and p_i are observed in the Homescan data, and ζ_i^c is the compensated price elasticity of demand, which we obtain from the Slutsky equation using our estimates of the uncompensated price elasticity and income effects.³⁹ As an example, imagine that bias increases quantity demanded by 15% and that the compensated demand elasticity is 1.5. Then, the impact of bias on consumption is the same as a 10% price reduction: $\gamma_i = p_i \cdot 15\% / 1.5 = 10\% \cdot p_i$.

Let $\mathbf{b}_i = [b_{ki}, b_{si}]$ denote a vector of indices measuring household i 's bias: nutrition knowledge b_{ki} and self-control b_{si} . Let $\mathbf{b}^V = [b_k^V, b_s^V]$ denote the value of \mathbf{b} for a “normative” consumer that maximizes V . \mathbf{a}_i is the vector of preferences (beverage tastes and health preferences) measured in the PanelViews survey, \mathbf{x}_i is the vector of household characteristics introduced in Table 1, and μ_c is a county fixed effect.

We assume that a household's SSB purchases depend on the average biases, preferences, and demographics of all (one or two) household heads. Appendix Table A13 presents estimates under the alternative assumption that a household's SSB purchases depend on the biases and preferences of the primary shopper only. The pattern of results is very similar, but the coefficient estimates and resulting bias magnitudes are attenuated by about 15 percent.

Our empirical strategy for estimating γ_i requires the following three assumptions:

Assumption 5. *Normative consumers:* $b_k^V = \mathbb{E}[b_{ki} | \text{dietitian}, \text{nutritionist}]$, $b_s^V = 1$.

For nutrition knowledge, we set b_k^V equal to the average nutrition knowledge score of the 24 dietitians and nutritionists in the PanelViews survey, which is 0.92. For self-control, we set $b_s^V = 1$:

³⁹To derive Equation (22), we take a first-order Taylor expansion of $s^V(p - \gamma, y - s\gamma, \theta)$:

$$\ln s^V(p - \gamma, y - s\gamma, \theta) \approx \ln s(p, y, \theta) - \frac{d \ln s^V(p, y, \theta)}{d \ln p} \Big|_u \cdot (\ln(p + \gamma(\theta)) - \ln p) \approx \ln s(p, y, \theta) - \zeta^{c,V}(\theta) \gamma(\theta) / p,$$

where $\zeta_i^{c,V}$ is the compensated price elasticity of normative demand. We assume that $\zeta_i^{c,V} \approx \zeta_i^c$. This holds in many standard cases—for example, if marginal health costs are locally constant and are underweighted by some constant amount due to present bias or incorrect beliefs. In general, a sufficient condition is that γ_i is locally constant in p . Substituting the i subscript for the (p, y, θ) triple and re-arranging gives $\ln s_i = \ln s_i^V + \zeta_i^c \gamma_i / p_i$.

Compensated elasticity is $\zeta_i^c = \zeta_i - \frac{\xi_i}{1 - T'(z_i)} \cdot \frac{p_i s_i}{z_i}$, because $\zeta^c = \zeta - \frac{ds}{dy} p$ per the Slutsky equation, and $\frac{ds}{dy} p = \frac{ds}{dz} \frac{dz}{dy} p = \frac{ds}{dz} \frac{1}{1 - T'(z)} p = \frac{\xi}{1 - T'(z)} \frac{ps}{z}$. We use the ζ_i and ξ_i estimated at household i 's income from Appendix Table A8 and the $T'(z_i)$ implied by the U.S. income tax schedule, as described in Appendix K.B.

normative consumers are those for whom “not at all” is the correct response to the statement, “I drink soda pop or other sugar-sweetened beverages more often than I should.”

Assumption 6. *Linearity:* $\zeta_i^c \gamma_i / p_i = \boldsymbol{\tau} \cdot (\mathbf{b}^V - \mathbf{b}_i)$, where $\boldsymbol{\tau}$ comprises two parameters scaling the effects of nutrition knowledge and self-control.

In our data, linearity is a realistic assumption, as demonstrated in Appendix Figure A5.⁴⁰

Assumption 7. *Unconfoundedness:* $\mathbf{b}_i \perp (\ln s_i^V | \mathbf{a}_i, \mathbf{x}_i, \mu_c)$.

In words, bias is conditionally independent of normative consumption. While such unconfoundedness assumptions are often unrealistic, this is more plausible in our setting because of our tailor-made survey measures of beverage tastes and health preferences.

Equation 22 and Assumptions 6 and 7 imply the following linear model:

$$\ln(s_i + 1) = \boldsymbol{\tau} \mathbf{b}_i + \beta_a \mathbf{a}_i + \beta_x \mathbf{x}_i + \mu_c + \varepsilon_i. \quad (23)$$

We add 1 to SSB purchases before taking the natural log so as to include households with zero purchases.

The regression we estimate differs from this model because \mathbf{b}_i and \mathbf{a}_i are measured imperfectly, for two reasons. First, household head average biases and preferences are unobserved for two-head households in which only one head responded to the survey. Second, the PanelViews survey delivers bias *proxies* that may contain measurement error. The observed bias proxies are as follows. For two-head households, \tilde{b}_{1si} is the average of the primary shopper’s self-control assessments for herself and the other head, and in the 2,481 households where both heads responded, \tilde{b}_{2si} is the analogous average constructed from the other head’s responses.⁴¹ In one-head households, \tilde{b}_{1si} is simply the household head’s self-control assessment for herself. Similarly, \tilde{b}_{1ki} and \tilde{b}_{2ki} are the nutrition knowledge scores for the primary shopper and other head, if observed, and \tilde{b}_{ki} is the household head average nutrition knowledge score.

We make two additional assumptions:

Assumption 8. *Measurement error in self-control:* For self-control, $\tilde{b}_{1si} = b_{si} + \nu_{1si}$ and $\tilde{b}_{2si} = b_{si} + \nu_{2si}$, with $(\nu_{1si}, \nu_{2si}) \perp b_{si}$ and $\nu_{1si} \perp \nu_{2si}$. Moreover, the distribution of ν_{1si} does not depend on whether the household has one or two heads and on whether one or both heads responded to the survey. We assume no measurement error in nutrition knowledge.

⁴⁰Linearity is also theoretically plausible, as it results from any “structural” behavioral model in which \mathbf{b}_i scales the share of costs that are misperceived. For example, consider a model in which utility of soda consumption is $v(s) - ps - \beta_i hs$, where h is the future health cost per unit of soda consumption and β is the self-control parameter in the standard β, δ model (Laibson, 1997; O’Donoghue and Rabin, 1999). If the policymaker defines normative utility using the long-run criterion, then $\gamma_i = (1 - \beta_i)h$. Thus, if $\beta_i = b_i/b^V$, i.e. b_i is proportional to self-control and b^V represents the value for time-consistent consumers, then linearity holds, with $\tau = \frac{h\zeta_i^c}{b^V p}$.

⁴¹We follow Bronnenberg et al. (2015) in defining each household’s “primary shopper” by selecting the survey respondent who makes the majority of the grocery purchase decisions, then the female, then the older person, until we arrive at a single primary shopper in each household.

It is reasonable to make the approximation that nutrition knowledge is measured without error, because the General Nutrition Knowledge Questionnaire is a many-question scale with high sensitivity to change and construct validity and a high test-retest reliability of 0.89 (Kliemann, Wardle, Johnson, and Croker, 2016).

Assumption 9. *Missing at random: in two-head households, biases and preferences are not correlated with whether both heads responded to the survey.*

To address non-response in two-head households where only one head responded, we impute household average nutrition knowledge nutrition knowledge \tilde{b}_{ki} and preferences \mathbf{a}_i based on the observed head’s bias proxies and preferences.⁴² We denote household i ’s observed or imputed nutrition knowledge and preferences as $\hat{\tilde{b}}_{ki}$ and $\hat{\mathbf{a}}_i$. To address measurement error in self-control, we use two-sample two-stage least squares (2SLS). In the first stage, we regress \tilde{b}_{2si} on \tilde{b}_{1si} and the other covariates from Equation (23) in the subset of households with two survey respondents, and we construct fitted values $\hat{\tilde{b}}_{2si}$ for all households.⁴³ For the second stage, we estimate Equation (23) using $\hat{\tilde{\mathbf{b}}}_i = [\hat{\tilde{b}}_{ki}, \hat{\tilde{b}}_{2si}]$ and $\hat{\mathbf{a}}_i$ in place of \mathbf{b}_i and \mathbf{a}_i . We calculate standard errors using the procedure of Chodorow-Reich and Wieland (2016), which accounts for heteroskedasticity and the interdependence between the first- and second-stage samples. Assumptions (8) and (9) guarantee that the imputation and measurement error correction procedures yield consistent estimates of τ in Equation (23).

Inserting our parameter estimates into Equation 22 and Assumption 6, we obtain estimates of counterfactual normative consumption and money-metric bias:

$$\log \hat{s}_i^V = \log s_i - \hat{\tau}(\mathbf{b}^V - \hat{\tilde{\mathbf{b}}}_i) \tag{24}$$

$$\hat{\gamma}_i = \hat{\tau}(\mathbf{b}^V - \hat{\tilde{\mathbf{b}}}_i)p_i/\hat{\zeta}_i^c. \tag{25}$$

For all empirical analyses, we use each household’s most recent year in the Homescan data, which for about 98 percent of households is 2016.⁴⁴

III.D.2 Descriptive Facts

Figure 2, discussed in the introduction, shows that there is a strong unconditional relationship between our bias proxies and SSB purchases. After conditioning on other controls, this is the variation that identifies τ in Equation (23).

⁴²Specifically, for each variable v in \tilde{b}_{ki} and \mathbf{a}_i , we impute with the predictions from a regression of household average v on \mathbf{x}_i , μ_c , and the primary shopper’s \tilde{b}_{ki} and \mathbf{a}_i , in the sample of two-head households where both heads responded. Appendix Table A10 presents these regressions. There is little uncertainty in these imputations: the adjusted R^2 ’s are around 0.7 to 0.8.

⁴³Appendix Table A11 presents this first stage regression.

⁴⁴As we use more historical consumption data, the correlations with bias proxies measured in the 2017 survey become weaker, suggesting that biases may change over time.

Figure 5 shows that both nutrition knowledge and self-control are strongly correlated with income. Panel (a) shows that people with household income above \$100,000 score 0.12 higher (0.82 standard deviations) than people with income below \$10,000 on the nutrition knowledge questionnaire. Panel (b) shows that people with income above \$100,000 also report about 0.14 higher (0.40 standard deviations) self-control. These relationships suggest that bias is regressive, which augments the corrective benefits of SSB taxes.

Figure 6 shows that preferences entering normative utility also differ systematically by income. Panel (a) shows that relative to people with household income above \$100,000, people with income below \$10,000 average about 0.09 higher (0.24 standard deviations) in terms of how much they “like the taste and generally enjoy drinking” regular soft drinks. Panel (b) shows that relative to that highest-income group, the lowest-income group averages about 0.06 points lower (0.36 standard deviations) in their reported importance of staying healthy. Both results imply that lower-income consumers have stronger normative preferences for SSBs. This corroborates the result illustrated in Figure 4 that the declining consumption-income relationship is driven by preference heterogeneity, not income effects.

III.D.3 Regression Results

Table 3 presents results. Column 1 is our primary specification. Columns 2-4 show how the $\hat{\tau}$ estimates change with the exclusion of particular controls. Preferences, income, and education are correlated with both SSB purchases and bias proxies, so it is unsurprising that their exclusion increases the $\hat{\tau}$ in column 2 relative to the primary estimates in column 1. Other demographics (age, race, the presence of children, household size, employment status, and weekly work hours) and county indicators, however, have relatively little effect on $\hat{\tau}$ in columns 3 and 4. While certainly not dispositive, this limited coefficient movement in columns 3 and 4 is consistent with the idea that unobservables also have limited impacts on $\hat{\tau}$.⁴⁵

Column 5 uses self-control $\tilde{b}_{1,si}$ directly instead of the two-sample 2SLS measurement error correction. Comparing to column 1, we see that the measurement error correction addresses what would otherwise be substantial attenuation bias in the self-control coefficient. Because knowledge and self-control are positively correlated, this measurement error correction reduces the knowledge coefficient, because some of the self-control effect otherwise loads onto knowledge.⁴⁶

⁴⁵To quantify this argument, imagine that the additional unobserved confounders have the same correlation with bias proxies and consumption as the demographic controls excluded in column 3, and thus would also decrease the nutrition knowledge coefficient by about 0.11 and increase the self-control coefficient by 0.10. This would decrease the average bias $\tau \cdot (\mathbf{b}^V - \mathbf{b}_i)$ by about one percent relative to the estimates from column 1. However, the low R^2 of these regressions means that there are important unobservables, and our estimates could have large bias if these unobservables are correlated with bias (Oster, 2017).

⁴⁶Appendix Table A12 presents additional estimates including an interaction term between knowledge and self-control. This interaction term is highly significant, perhaps because it takes knowledge of health damages to believe that one “should” consume less. Including this interaction in the model, however, does not materially change the estimates of average marginal bias.

III.D.4 Estimates of Bias

For the average American household, predicted “normative” SSB consumption is only $\frac{\hat{s}^V}{s_i} \approx 62$ percent of actual consumption. Put differently, we predict that the average American household would consume $\frac{s_i - \hat{s}_i^V}{s_i} \approx 38$ percent fewer SSBs if they had the nutrition knowledge of dietitians and nutritionists and no self-control problems. Figure 7 plots the share of consumption attributable to bias, i.e. the unweighted average of $\frac{s_i - \hat{s}_i^V}{s_i}$ across households, by income. Predicted overconsumption is much larger for low-income households: it is 47 and 34 percent, respectively, for households with income below \$10,000 and above \$100,000.

Figure 8 plots the demand-slope weighted average marginal bias $\hat{\gamma} = \frac{\sum_i \hat{\zeta}_i^c \frac{s_i}{p_i} \hat{\gamma}_i}{\sum_i \hat{\zeta}_i^c \frac{s_i}{p_i}}$ by income. The average marginal bias across all American households is 1.14 cents per ounce. Since nutrition knowledge and self-control increase with income and elasticities and prices do not differ much by income, we know from Equation (25) that money-metric bias $\hat{\gamma}$ will decline in income. Indeed, average marginal biases are 1.38 and 1.04 cents per ounce, respectively, for households with income below \$10,000 and above \$100,000.

III.E Externalities

We import an externality estimate from outside sources. Using epidemiological simulation models, Wang et al. (2012) estimate that one ounce of soda consumption increases health care costs by an average of approximately one cent per ounce. Yong, Bertko, and Kronick (2011) estimate that for people with employer-provided insurance, about 15 percent of health costs are borne by the individual, while 85 percent are covered by insurance. Similarly, Cawley and Meyerhoefer (2012) estimate that 88 percent of the total medical costs of obesity are borne by third parties, and obesity is one of the primary diseases thought to be caused by SSB consumption. Accordingly, we approximate the health system externality at $e \approx 0.85$ cents per ounce.

There are two caveats to this calculation. First, Bhattacharya and Bundorf (2009) find that obese people in jobs with employer-provided health insurance bear the full health costs of obesity through lower wages. However, this result may or may not apply to the other diseases caused by SSB consumption, including diabetes and cardiovascular disease, and it does not apply to people with government-provided health insurance through Medicaid or Medicare. Second, the diseases caused by SSBs might decrease life expectancy, reducing the amount of social security benefits that people claim and thereby imposing a positive fiscal externality (Bhattacharya and Sood, 2011; Fontaine et al., 2003). Accounting for these two factors would reduce the externality estimate. Section IV presents optimal tax estimates under alternative assumptions that illustrate the impact of externalities.

III.F Assessment

We now assess the strengths and weaknesses of the empirical strategies employed above.

Identifying the price elasticity of demand. Our instrumental variables approach has advantages relative to alternative approaches to identifying price elasticity. One alternative is to use recent SSB tax changes in U.S. cities as an instrument. This would require the assumption that no other factors affected SSB demand at the time the SSB tax was implemented. Rees-Jones and Rozema (2018) document one potential violation of this assumption: local public debates about sin taxes also affect demand, even in cities where prices never changed because the tax never went into effect.⁴⁷ Furthermore, because there have been only a small handful of city-level SSB taxes implemented in the U.S., this approach will tend to deliver less precise estimates for the foreseeable future.

A second alternative is to use pricing instruments such as those introduced by Hausman (1996) and Nevo (2001), who instrument for the price of a product at time t using the average price of that product in all other cities. This requires the assumption that time-varying demand shocks for a given product are uncorrelated across cities, which assumes away possibilities such as national advertising campaigns. Our instrument also delivers more power, because it exploits the extensive idiosyncratic chain-specific price variation instead of the more limited variation in national average prices.

A third alternative is to regress quantities on prices without an instrument, as in Dubois, Griffith, and O’Connell (2017), Tiffin, Kehlbacher, and Salois (2015), Zhen et al. (2011), and others. With our particular vector of controls, Appendix Table A7 shows that the OLS estimates understate price elasticity by about a factor of two, which is suggestive of either a standard simultaneity bias or measurement error in prices.⁴⁸

Substitution to un-taxed goods. As discussed in Section II.D, optimal commodity taxes depend on substitution patterns across all goods that cause internalities or externalities. There is a robust debate in the public health literature about the health risks of diet drinks and of sugar consumed through solid foods compared to drinks, and we do not attempt to quantify internalities or externalities for drinks other than SSBs. Appendix Table A9 presents estimates of substitution between SSBs and nine groups of non-SSB beverages and high-sugar foods. Unsurprisingly, we find that SSBs and diet drinks are substitutes. There is no statistically significant cross-price elasticity with other beverages. One of the four high-sugar foods actually is actually estimated to be a complement to SSBs; if correct, this would slightly increase the optimal SSB tax.

Long-run versus short-run elasticities. Because the commodity taxes and income taxes we model are long-lasting, we ideally want long-run elasticities to calibrate optimal taxes. In the absence of plausibly exogenous long-run price and income variation, we estimate the price and income elasticities of demand for SSBs using quarterly prices and annual income. To the extent that there is habit formation in SSB consumption, this variation would cause us to understate elasticities. In Appendix Table A2, we see no impact of lagged quarterly prices on current consumption, although

⁴⁷In our theory model, a tax affects demand only through prices, not through signaling or public awareness. The tax elasticity of demand estimated in these papers would be one crucial input into a richer model.

⁴⁸Dubois, Griffith, and O’Connell (2017) and the other papers adopting this approach have different data and control strategies, so their estimates certainly need not be biased in this same way.

this does not rule out habit formation over a longer time horizon.

Measuring SSB consumption. As discussed in Section III.A, scanner purchase datasets such as Homescan provide imperfect measures of actual SSB intake. Furthermore, Homescan measures consumption at the household level, whereas our survey measures of bias and preferences are at the individual level and include only household heads.

To address these issues, we carried out a beverage intake questionnaire as part of the PanelViews survey, and we analyze these data in Appendix H. We estimate larger $\hat{\tau}$ —that is, stronger associations between SSB consumption and bias proxies—using the PanelViews self-reports compared to the Homescan purchase data. As a result, the PanelViews bias estimates are larger: 48 of the average household’s consumption is attributable to bias (58 and 42 percent, respectively, for household incomes below \$10,000 and above \$100,000), and average marginal bias is 2.78 cents per ounce. We present optimal tax estimates based on these consumption self-reports in Section IV.

While purchase data are incomplete for the two reasons discussed above, self-reported intake data have other limitations that could bias $\hat{\tau}$ up or down. For example, self-reports of consumption and self-control could both suffer from recency or salience bias: the judgments may be constructed based on only the most recent or salient consumption episodes. The resulting correlated measurement error would generate an upward-biased relationship between self-reported consumption and self-control. Additionally, social desirability bias could cause people with more health knowledge or less perceived self-control to under-report consumption, which could bias $\hat{\tau}$ either up or down. Using scanner-based purchase data avoids these potential problems.

Homescan’s incomplete consumption measures could also affect our price elasticity estimate. If households substitute between scanned and non-scanned SSBs, for example by buying less soda at a restaurant when grocery store prices are lower, the estimated price elasticity of scanned purchases will likely overstate the true SSB price elasticity. We explore the policy implications of a more inelastic alternative elasticity assumption in Section IV.

Other types of consumer bias. We base optimal taxes only on the two biases that we measure in the PanelViews survey: imperfect nutrition knowledge and self-control. If SSB consumption is affected by other biases—for example, “non-rational” (i.e. unanticipated) habit formation or inattention to health harms—or if our surveys provide incomplete measures of the types of imperfect knowledge and self-control that affect SSB consumption, then we will understate the optimal SSB tax.⁴⁹

Measurement error correction. Our paper is one of only a small number to correct for measurement error when estimating the relation between experimental or survey measures and outcomes; see Gillen, Snowberg, and Yariv (forthcoming) for evidence on the importance of this type of correction.

Comparison to other bias measurement approaches. Bernheim and Taubinsky (2018),

⁴⁹We do not consider rational habit formation in the sense of Becker and Murphy (1988) to be a behavioral bias that would drive a wedge between decision utility and normative utility. We note that our self-control survey question was designed to reflect both naive and sophisticated self-control problems by asking whether people drink too many SSBs in general, not whether they will drink too many SSBs in the future.

Mullainathan, Schwartzstein, and Congdon (2012), and others review the possible approaches to estimating biases for public policy applications. We briefly discuss the merits of these other approaches in the context of SSB consumption.

One approach is to calibrate or estimate models of “bias” parameters. For example, in the older working paper version of this paper, Lockwood and Taubinsky (2017) assume that all bias is a consequence of present bias, and use evidence from behavioral economics and public health to calibrate a present bias parameter β and a marginal health cost parameter h to arrive to an estimate of bias $\gamma = (1 - \beta)h$. This strategy rests on many assumptions: that measures of present bias from laboratory experiments about effort or money can be extrapolated to a representative population and to the alternate domain of SSB consumption; that the measures of marginal health costs (obtained from mostly correlational public health studies) are accurate and independent of the types of SSBs purchased; and that bias is homogeneous across the population; and that there are no biases besides present bias. Preference reversal experiments with groceries such as Sadoff, Samek, and Sprenger (2015) could be used to provide qualitative within-domain evidence of present bias, but such designs do not deliver money-metric bias estimates and would be prohibitively difficult to carry out on a nationally representative sample.

Mistakes arising from misinformation could also be measured using survey-elicited beliefs. This approach of eliciting beliefs about health costs is similar in spirit to our approach, but unlike our approach requires strong assumptions about the “true” marginal health costs as well as reasonably strong structural assumptions about how health costs affect decision and normative utility. Furthermore, consumers may or may not act on the beliefs that they state in a survey.

Finally, one might estimate informational and attentional biases through an information provision field experiment, as in Allcott and Taubinsky (2015). This strategy requires the assumption that the information provided fully debiases the information treatment group, which could be unrealistic in the case of complex information around nutrition and health. Our nutrition knowledge scores capture information gathered over a long period, which would probably be prohibitively difficult to induce experimentally. Moreover, if the dependent variable is naturally-occurring decisions such as those in Homescan, this strategy would also require that the information provision is processed by all individuals, is remembered, and is salient at the time of choice—conditions that seem implausible outside a controlled, online shopping experiment as in Allcott and Taubinsky (2015). At the same time, a controlled online shopping experiment would not lead to plausible estimates of demand curves (and therefore bias) in our context because SSBs are easy to purchase, and these outside options would be perfect substitutes for the products in the experiment.

In sum, while our approach requires a strong unconfoundedness assumption and has other weaknesses, other possible approaches to estimating bias involve arguably stronger assumptions and are, in our view, even more likely to suffer from confounds. Moreover, and in contrast to the requirements of some other approaches, the dataset we use to estimate demand elasticities is essentially the same one we use for estimating bias—a key advantage for internal consistency.

IV Computing the Optimal Sugar-Sweetened Beverage Tax

We now combine the theoretical results from Section II with the empirical estimates from Section III to compute the optimal nationwide tax on SSBs. We compute the optimal tax across a range of specifications, under two different assumptions about the income tax. First, we compute the optimal SSB tax assuming the income tax is held fixed at the current status quo in the U.S., using Proposition 2. Second, we compute the optimal SSB tax assuming the income tax is also reformed to be optimal, using Equation (15) in Proposition 1.

These computations require an assumption about inequality aversion. We employ social marginal welfare weights proportional to $y_{US}^{-\nu}$, where y_{US} is post-tax income in the U.S., and ν is a parameter which governs the strength of inequality aversion. We use $\nu = 1$ as our baseline, and $\nu = 0.25$ and $\nu = 4$ as our “weak” and “strong” redistributive preferences, respectively. Calibrations of the status quo U.S. income distribution and income tax are drawn from Piketty, Saez, and Zucman (2018); see Appendix K.B for details.

The sufficient statistics formulas for the optimal tax depend on a number of population statistics, as well as key covariances involving welfare weights, for which we need estimates of the statistics by income. The necessary statistics, based on our empirical estimates described in Section III, are reported in Table 4. Panel (a) presents estimates of population-level statistics, while Panel (b) presents estimates of statistics within each Homescan income bin, from which we can compute the welfare weight covariances.⁵⁰

Using the distribution of statistics by income bin, we calculate $\sigma = Cov \left[g(z), \frac{\bar{\gamma}(z)}{\bar{\gamma}} \frac{\bar{\zeta}^c(z)}{\bar{\zeta}^c} \frac{\bar{s}(z)}{\bar{s}} \right] \approx 0.22$. The corrective motive is thus equal to $\bar{\gamma}(1 + \sigma) + e \approx 1.17(1 + 0.22) + 0.85 \approx 2.27$ cents per ounce. This highlights the importance of accounting for internalities when setting optimal corrective taxes. Less than half of the corrective motive is driven by externality correction. Indeed, the average marginal bias $\bar{\gamma}$ is larger than the externality, and the internality correction is further inflated by about 27 percent due to bias correction progressivity, reflecting the fact that the benefits of bias correction accrue disproportionately to poorer consumers.

Counteracting this corrective motive, the regressivity of the tax pushes toward a smaller optimal SSB tax, the poor have much stronger preferences for SSBs than the wealthy. In the presence of the optimal income tax, the impact of the redistributive motive is proportional to $-Cov [g(z), s_{pref}(z)]$. This statistic can be computed directly from Table 4, Panel B, where $s_{pref}(z)$ is constructed as the difference between observed SSB consumption $\bar{s}(z)$ and consumption predicted from estimated income elasticities, similarly to Figure 4. In this case, the estimated covariance is $Cov [g(z), s_{pref}(z)] = 13.59$. Proposition 1 therefore implies an optimal sin tax of 1.26 cents per ounce.

When the income tax is constrained to the status quo, the redistributive motive term is in-

⁵⁰In principle all statistics can be estimated separately within each income group; however in the case of the elasticity and income effects, we do not have sufficient statistical power to precisely estimate full interactions with income. Therefore we estimate these statistics in only two income bins—above and below median.

stead proportional to $-Cov[g(z), s(z)] - \mathbb{E} \left[\frac{T'(z)}{1-T'(z)} \bar{\zeta}_z \bar{s}(z) \bar{\xi}(z) \right]$.⁵¹ Using Table 4, we can calculate $Cov[g(z), s(z)] = 6.07$ and $\mathbb{E} \left[\frac{T'(z)}{1-T'(z)} \bar{\zeta}_z \bar{s}(z) \bar{\xi}(z) \right] = 0.23$. Proposition 2 then implies an optimal SSB tax of 1.75 cents per ounce. The optimal SSB tax is higher under the status quo income tax than under the optimal income tax because if the status quo marginal income tax rates are “too low” relative to the optimum, as is the case here, then the fiscal externalities from labor supply distortions due to the sin tax will be lower than under the optimal income tax.

These estimates of the optimal SSB tax are reported in Table 5, along with analogous calculations under several alternative assumptions that we now summarize. The “Pigouvian” specification reports the optimal tax in the absence of any inequality aversion, in which case the tax is simply equal to $\bar{\gamma} + e$. Weaker preferences for redistribution lead to a higher tax than in the status quo, since the reduced corrective motive (due to smaller bias correction progressivity) is more than offset by a weaker redistributive motive. Similarly, stronger inequality aversion leads to a lower tax. In both cases, the deviation from baseline is greater under the optimal income tax than under the status quo, for the same reason that the redistributive motive is stronger under the optimal income tax than in the status quo.

A higher SSB demand elasticity scales down the redistributive motive term, raising the optimal tax. Conversely, a lower elasticity reduces the optimal tax. The fourth and fifth rows of Table 5 report the optimal tax assuming a (constant) demand elasticity of either 2 or 1, rather than our heterogeneous empirical estimate.

The next two rows report the optimal tax under two alternative assumptions about the roles of preference heterogeneity versus income effects in accounting for cross-sectional variation in SSB consumption. The first case assumes all SSB consumption differences are driven by between-income preference heterogeneity. In this case, the optimal SSB tax is independent of the income tax, and so it is the same in both columns. The “Pure income effects” case instead assumes preferences are homogeneous, implying that SSBs are highly inferior goods. In this case, redistribution is more efficiently carried out through the optimal income tax. This does not substantially alter the optimal SSB tax under the (suboptimal) status quo income tax. If the income tax is optimal, however, then the optimal SSB tax is *higher* than the Pigouvian case of no inequality aversion, as in special case 3 in Section II.C.

The row labeled “No internality” reports the optimal tax if consumer bias is assumed to be zero, which substantially reduces the optimal tax. The specification “No corrective motive” also assumes that *externalities* (in addition to internalities) are zero. In this case, only the redistributive motive is active, resulting in an optimal subsidy for SSBs.

The final row of Table 5 reports the optimal tax estimated using self-reported SSB consumption from our PanelViews survey, rather than data captured by Homescan. This specification results in a higher optimal SSB tax of 2.82 cents per ounce.

We compute the welfare gains from SSB taxes in Appendix J. In our baseline specification,

⁵¹Recall that the first term represents the mechanical effect of the tax based on actual SSB consumption, while the second represents the change in income tax revenues due to the effect of the SSB tax on labor supply.

the optimal tax generates an estimated increase in social welfare of \$11.48 per adult equivalent consumer per year, or about \$3.6 billion in aggregate across the U.S. This generates \$400 million more annual welfare gains than a nationwide one-cent-per-ounce tax—the modal policy among U.S. cities with SSB taxes—and \$1.6 billion more than the tax that would incorrectly be deemed “optimal” if policymakers do not account for consumer bias. These gains from the optimal tax highlight the importance of using empirical estimates of internalities and externalities, rather than assuming them away, or using round number “rules of thumb” to design policy.

The welfare gains can be decomposed into four distinct components, plotted in Figure 9 and described in Appendix J. The net gains vary across income groups for two competing reasons. First, groups that consume more SSBs have a larger decision utility Equivalent Variation loss due to the financial burden of the tax. Second, groups that are more biased and more elastic experience a larger benefit from bias correction due to the tax. As a result, the profile of net benefits from a seemingly “regressive” sin tax can be increasing or decreasing with income. The former force tends to dominate in our baseline estimate, generating an upward-sloping profile of net benefits across the income distribution. That slope is modest, however, and a similar exercise using the PanelViews self-reports—which give a steeper negative slope of bias in income—suggests that the poor benefit somewhat more than the rich; see Appendix Figure A6.

A potential concern about the estimates in Table 5, and the corresponding welfare calculations, is that sufficient statistics formulas yield only an approximation to the optimal tax. This is for two reasons. First, the formulas we use assume that income effects on labor supply and the budget share of SSBs are assumed to be small enough to be negligible. Second, the formulas rely on statistics that we estimate from the status quo equilibrium, rather than under the optimal tax—that is, we do not account for how a new tax regime would change the consumption of s by income, or how it would affect the elasticities. To explore the importance of these sources of error, Appendix K presents computations of the optimal tax using two different structural models, with taxes computed using sufficient statistics at the optimum and fully accounting for all behavioral responses. Those estimates exhibit the same qualitative patterns and are quantitatively close to the values reported here—particularly in the case where the income tax is held fixed—providing additional evidence that our sufficient statistics formulas provide a close approximation of the optimal tax.

V Conclusion

This paper provides a tractable theoretical and empirical framework for setting and evaluating optimal commodity taxes in the presence of internalities and externalities. We provide the first optimal commodity tax formula that takes account of the three key elements of public policy debates around sin taxes: correcting consumer bias and externalities, distributional concerns, and revenue recycling through income taxes or income-targeted transfers. Prior work in behavioral economics and public economics had considered only subsets of these three issues or imposed unrealistic assumptions around preference heterogeneity and other parameters.

We demonstrate the usefulness of the theoretical results by focusing on a particularly timely and controversial public policy question: what is the optimal soda tax? Our PanelViews survey data provide novel insights about the relationship between nutrition knowledge, self-control, income and SSB consumption, and we provide a credible estimate of the price elasticity of demand for SSBs using a new and broadly usable instrument. Our results suggest that the socially optimal nationwide SSB tax is between 1.26 and 2.82 cents per ounce. Thus, most existing city-level taxes may be too low. Our model suggests that the optimal federal tax would increase welfare by \$3.6 billion per year.

Our theoretical results and empirical methodology could immediately be applied to study optimal taxes (or subsidies) on cigarettes, alcohol, unhealthy foods such as sugar or saturated fat, and consumer products such as energy-efficient appliances. Leaving aside internalities, our results can be used to clarify active debates about the regressivity of externality correction policies such as carbon taxes and fuel economy standards. Our theory is also applicable to questions about capital income taxation or subsidies on saving, and with some appropriate modification our empirical methods could be extended to quantify taxes in those domains as well.

As we discuss throughout the paper, our approach has its weaknesses. One should be cautious about advocating any particular optimal tax estimate too strongly, and we encourage further work extending, generalizing, and critiquing our approach. But by leveraging robust economic principles tied closely to data, our methods almost surely provide valuable input into thorny public policy debates that often revolve around loose intuitions, unsubstantiated assumptions, personal philosophies, or political agendas.

References

- ALLCOTT, H. (2013): “The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market,” *American Economic Journal: Economic Policy*, 5, 30–66.
- ALLCOTT, H., C. KNITTEL, AND D. TAUBINSKY (2015): “Tagging and Targeting of Energy Efficiency Subsidies,” *American Economic Review Papers and Proceedings*, 112, 72–88.
- ALLCOTT, H., B. B. LOCKWOOD, AND D. TAUBINSKY (2018): “Ramsey Strikes Back: Optimal Commodity Taxes and Redistribution in the Presence of Salience Effects,” *American Economic Association Papers and Proceedings*, 108.
- ALLCOTT, H., S. MULLAINATHAN, AND D. TAUBINSKY (2014): “Energy Policy with Externalities and Internalities,” *Journal of Public Economics*, 112, 72–88.
- ALLCOTT, H. AND D. TAUBINSKY (2015): “Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market,” *American Economic Review*, 105, 2501–2538.
- ANDREYEVA, T., M. W. LONG, AND K. D. BROWNELL (2010): “The Impact of Food Prices on Consumption: A Systematic Review of Research on the Price Elasticity of Demand for Food,” *American Journal of Public Health*, 100, 216–222.
- ATKINSON, A. B. AND J. STIGLITZ (1976): “Design of Tax Structure - Direct Versus Indirect Taxation,” *Journal of Public Economics*, 6, 55–75.
- ATTARI, S., M. DEKAY, C. DAVIDSON, AND W. B. DE BRUIN (2010): “Public Perceptions of Energy Consumption and Savings,” *Proceedings of the National Academy of Sciences*, 37, 16054–16059.
- BAICKER, K., S. MULLAINATHAN, AND J. SCHWARTZSTEIN (2015): “Behavioral Hazard in Health Insurance,” *Quarterly Journal of Economics*, 130, 1623–1667.
- BARTELS, L. M. (1996): “Uninformed Votes: Information Effects in Presidential Elections,” *American Journal of Political Science*, 40, 194–230.
- BECKER, G. S. AND K. M. MURPHY (1988): “A theory of rational addiction,” *The Journal of Political Economy*, 675–700.
- BERNHEIM, B. D. (2016): “The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics,” *Journal of Cost Benefit Analysis*, 1–57.
- BERNHEIM, B. D. AND A. RANGEL (2004): “Addiction and Cue-triggered Decision Processes,” *American Economic Review*, 1558–1590.
- (2009): “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics,” *The Quarterly Journal of Economics*, 124, 51–104.

- BERNHEIM, B. D. AND D. TAUBINSKY (2018): “Behavioral Public Economics,” in *The Handbook of Behavioral Economics*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson.
- BHATTACHARYA, J. AND M. K. BUNDORF (2009): “The Incidence of the Healthcare Costs of Obesity,” *Journal of Health Economics*, 28, 649–658.
- BHATTACHARYA, J. AND N. SOOD (2011): “Who Pays for Obesity?” *Journal of Economic Perspectives*, 25, 139–158.
- BLUMGART, J. (2016): “A Soda Tax is Fundamentally Regressive. What if It’s Philadelphia’s Best Hope for Progressive Change?” .
- BOLLINGER, B., P. LESLIE, AND A. SORENSEN (2011): “Calorie Posting in Chain Restaurants,” *American Economic Journal: Economic Policy*, 3, 91–128.
- BOLLINGER, B. AND S. SEXTON (2017): “Local Excise Taxes, Sticky Prices, and Spillovers: Evidence from Berkeley’s Soda Tax,” *Working Paper*.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2013): “Salience and Consumer Choice,” *Journal of Political Economy*, 121, 803–843.
- BOVENBERG, A. L. AND L. H. GOULDER (1996): “Optimal Environmental Taxation in the Presence of Other Taxes: General-Equilibrium Analyses,” *American Economic Review*, 86, 985–1000.
- BRONNENBERG, B., J.-P. DUBÉ, M. GENTZKOW, AND J. SHAPIRO (2015): “Do Pharmacists Buy Bayer? Sophisticated Shoppers and the Brand Premium,” *Quarterly Journal of Economics*, 130.
- CAWLEY, J. AND C. MEYERHOEFER (2012): “The Medical Care Costs of Obesity: An Instrumental Variables Approach,” *Journal of Health Economics*, 31, 219–230.
- CHETTY, R. (2012): “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply,” *Econometrica*, 80, 969–1018.
- (2015): “Behavioral Economics and Public Policy: A Pragmatic Perspective,” *American Economic Review Papers and Proceedings*, 105, 1–33.
- CHETTY, R., A. LOONEY, AND K. KROFT (2009): “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 99, 1145–1177.
- CHODOROW-REICH, G. AND J. WIELAND (2016): “Secular Labor Reallocation and Business Cycles,” *NBER Working Paper No. 21864*.
- DAVIS, L. AND S. BORENSTEIN (2016): “The Distributional Effects of U.S. Clean Energy Tax Credits,” *NBER Tax Policy and the Economy*, 30, 191–234.
- DAVIS, L. AND C. KNITTEL (2016): “Are Fuel Economy Standards Regressive?” *Working Paper*.

- DELLAVIGNA, S. (2009): “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 47, 315–372.
- DELLAVIGNA, S. AND M. GENTZKOW (2017): “Uniform Pricing in US Retail Chains,” *NBER Working Paper 23996*.
- DIAMOND, P. A. (1973): “Consumption Externalities and Imperfect Corrective Pricing,” *The Bell Journal of Economics and Management Science*.
- (1975): “A Many-Person Ramsey Tax Rule,” *Journal of Public Economics*, 4, 335–342.
- DUBOIS, P., R. GRIFFITH, AND M. O’CONNELL (2017): “How Well Targeted Are Soda Taxes?” *Working Paper*.
- DUFFEY, K. J., P. GORDON-LARSEN, J. M. SHIKANY, D. GUILKEY, D. R. JACOBS, AND B. M. POPKIN (2010): “Food Price and Diet and Health Outcomes,” *Archives of Internal Medicine*, 170, 420–426.
- EDWARDS, J., M. KEEN, AND M. TUOMALA (1994): “Income Tax, Commodity Taxes and Public Good Provision: A Brief Guide,” *FinanzArchiv*, 51, 83–95.
- EINAV, L., E. LEIBTAG, AND A. NEVO (2010): “Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?” *Quantitative Marketing and Economics*, 8, 207–239.
- FARHI, E. AND X. GABAIX (2015): “Optimal Taxation with Behavioral Agents,” Working Paper 21524, National Bureau of Economic Research.
- FINKELSTEIN, E. A., C. ZHEN, M. BILGER, J. NONNEMAKER, A. M. FAROOQUI, AND J. E. TODD (2013): “Implications of a sugar-sweetened beverage (SSB) tax when substitutions to non-beverage items are considered,” *Journal of Health Economics*, 32, 219–239.
- FLETCHER, J. M., D. E. FRISVOLD, AND N. TEFT (2010): “The effects of soft drink taxes on child and adolescent consumption and weight outcomes,” *Journal of Public Economics*, 94, 967–974.
- FONTAINE, K. R., D. T. REDDEN, C. WANG, A. O. WESTFALL, AND D. B. ALLISON (2003): “Years of Life Lost Due to Obesity,” *Journal of the American Medical Association*, 289, 187–193.
- GABAIX, X. (2014): “A Sparsity Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 129, 1661–1710.
- GABAIX, X. AND D. LAIBSON (2006): “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121.
- GILLEN, B., E. SNOWBERG, AND L. YARIV (forthcoming): “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economy*.

- GOLDIN, J. AND T. HOMONOFF (2013): “Smoke Gets In Your Eyes: Cigarette Tax Salience and Regressivity,” *American Economic Journal: Economic Policy*, 5, 302–336.
- GOULDER, L. H. (1995): “Environmental taxation and the double dividend: A reader’s guide,” *International Tax and Public Finance*, 2, 157–183.
- (2013): “Climate change policy’s interactions with the tax system,” *Energy Economics*, 40, 3–13.
- GOULDER, L. H. AND R. C. WILLIAMS (2003): “The Substantial Bias from Ignoring General Equilibrium Effects in Estimating Excess Burden, and a Practical Solution,” *Journal of Political Economy*, 111, 898–927.
- GRUBER, J. AND B. KŐSZEGI (2001): “Is Addiction “Rational”? Theory and Evidence,” *The Quarterly Journal of Economics*, 116, 1261–1303.
- (2004): “Tax Incidence When Individuals Are Time-Inconsistent: The Case of Cigarette Excise Taxes,” *Journal of Public Economics*, 88, 1959–1987.
- GRUBER, J. AND E. SAEZ (2002): “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, 84, 1–32.
- HANDEL, B. R. AND J. T. KOLSTAD (2015): “Health Insurance for “Humans”: Information Frictions, Plan Choice, and Consumer Welfare,” *American Economic Review*, 105, 2449–2500.
- HANDEL, B. R., J. T. KOLSTAND, AND J. SPINNEWIJN (2016): “Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets,” *Working Paper*.
- HARDING, M. AND M. LOVENHEIM (2015): “The effect of prices on nutrition: Comparing the impact of product- and nutrient-specific taxes,” *Journal of Health Economics*.
- HAUSMAN, J. (1996): *Valuation of New Goods Under Perfect and Imperfect Competition*, National Bureau of Economic Research, vol. 58.
- HEDRICK, V., J. SAVLA, D. COMBER, K. FLACK, P. ESTABROOKS, P. NSIAH-KUMI, S. ORTMEIER, AND B. DAVY (2012): “Development of a Brief Questionnaire to Assess Habitual Beverage Intake (BEVQ-15): Sugar-Sweetened Beverages and Total Beverage Energy Intake,” *Journal of the Academy of Nutrition and Dietetics*, 112, 840–849.
- HENDEL, I. AND A. NEVO (2006a): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74, 1637–1673.
- (2006b): “Sales and Consumer Inventory,” *RAND Journal of Economics*, 37, 543–561.
- (2010): “A Simple Model of Demand Anticipation,” *Working Paper*.

- HEUTEL, G. (2015): “Optimal Policy Instruments for Externality-Producing Durable Goods under Present Bias,” *Journal of Environmental Economics and Management*, 72, 54–70.
- HIRTH, R. A., M. E. CHERNEW, E. MILLER, M. FENDRICK, AND W. G. WEISSERT (2000): “Willingness to Pay for a Quality-adjusted Life Year: In Search of a Standard,” *Medical Decision Making*, 20, 332–342.
- HITSCH, G. J., A. HORTACSU, AND X. LIN (2017): “Prices and Promotions in U.S. Retail Markets: Evidence from Big Data,” *Working Paper*.
- HUEHNERGARTH, N. F. (2016): “Senator Sanders Is Wrong: A Soda Tax Is Progressive, Not Regressive,” .
- JACOBS, B. AND R. BOADWAY (2014): “Optimal Linear Commodity Taxation under Optimal Non-Linear Income Taxation,” *Journal of Public Economics*, 117, 201–210.
- JACOBS, B. AND R. A. DE MOOIJ (2015): “Pigou Meets Mirrlees: On the Irrelevance of Tax Distortions for the Second-Best Pigouvian Tax,” *Journal of Environmental Economics and Management*.
- JACQUET, L. AND E. LEHMANN (2014): “Optimal Nonlinear Income Taxation with Multidimensional Types: The Case with Heterogeneous Behavioral Responses,” .
- JOHNSON, E. M. AND M. M. REHAVI (2016): “Physicians Treating Physicians: Information and Incentives in Childbirth,” *American Economic Journal: Economic Policy*, 8, 115–141.
- KAPLOW, L. (2012): “Optimal Control of Externalities in the Presence of Income Taxation,” *International Economic Review*, 53, 487–509.
- KIT, B., T. FAKHOURI, S. PARK, S. J. NIELSEN, AND C. OGDEN (2013): “Trends in Sugar-Sweetened Beverage Consumption Among Youth and Adults in the United States: 1999-2010,” *American Journal of Clinical Nutrition*, 98, 180–188.
- KLIEMANN, N., J. WARDLE, F. JOHNSON, AND H. CROKER (2016): “Reliability and validity of a revised version of the General Nutrition Knowledge Questionnaire,” *European Journal of Clinical Nutrition*, 70, 1174–1180.
- KOPCZUK, W. (2003): “A note on optimal taxation in the presence of externalities,” *Economics Letters*, 80, 81–86.
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, 112, 443–478.
- LEVITT, S. D. AND C. SYVERSON (2008): “Market Distortions when Agents are Better Informed: The Value of Information in Real Estate Transactions,” *Review of Economics and Statistics*, 90, 599–611.

- LOCKWOOD, B. B. AND D. TAUBINSKY (2017): “Regressive Sin Taxes,” *NBER Working Paper No. 23085*.
- LONG, M. W., S. L. GORTMAKER, Z. J. WARD, S. C. RESCH, M. L. MOODIE, G. SACKS, B. A. SWINBURN, R. C. CARTER, AND Y. C. WANG (2015): “Cost Effectiveness of a Sugar-Sweetened Beverage Excise Tax in the U.S.” *American Journal of Preventive Medicine*, 49, 112–123.
- GFRP (GLOBAL FOOD RESEARCH PROGRAM) (2018): “Sugary Drink Taxes Around the World,” .
- KILTS CENTER FOR MARKETING (2018): “Retail Scanner Dataset Manual,” Tech. rep.
- MULLAINATHAN, S., J. SCHWARTZSTEIN, AND W. J. CONGDON (2012): “A Reduced-Form Approach to Behavioral Public Finance,” *Annual Review of Economics*, 4, 1–30.
- NAVA, M., F. SHROYEN, AND M. MARCHAND (1996): “Optimal Fiscal and Public Expenditure Policy in a Two-Class Economy,” *Journal of Public Economics*, 61, 119–137.
- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- O’DONOGHUE, T. AND M. RABIN (1999): “Doing It Now or Later,” *American Economic Review*, 89, 103–124.
- (2006): “Optimal Sin Taxes,” *Journal of Public Economics*.
- OF ECONOMIC ANALYSIS”, B. (2017): “Personal Income Summary: Personal Income, Population, Per Capita Personal Income,” .
- OSTER, E. (2017): “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics*, 0, 1–18.
- PIGOU, A. C. (1920): *The Economics of Welfare*.
- PIKETTY, T., E. SAEZ, AND G. ZUCMAN (2018): “Distributional National Accounts: Methods and Estimates for the United States,” *Quarterly Journal of Economics*.
- POWELL, L. M., J. F. CHRIQUI, T. KHAN, R. WADA, AND F. J. CHALOUPKA (2013): “Assessing the Potential Effectiveness of Food and Beverage Taxes and Subsidies for Improving Public Health: A Systematic Review of Prices, Demand and Body Weight Outcomes,” *Obesity Reviews*.
- REES-JONES, A. AND K. ROZEMA (2018): “On the Non-Price Effects of Sin Taxes,” Tech. rep., Wharton.
- ROJAS, C. AND E. WANG (2017): “Do Taxes for Soda and Sugary Drinks Work? Scanner Data Evidence from Berkeley and Washington,” *Working Paper*.

- SADOFF, S., A. SAMEK, AND C. SPRENGER (2015): “Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert,” .
- SAEZ, E. (2001): “Using Elasticities to Derive Optimal Income Tax Rates,” *The Review of Economic Studies*, 68, 205–229.
- (2002a): “The Desirability of Commodity Taxation Under Non-Linear Income Taxation and Heterogeneous Tastes,” *Journal of Public Economics*, 83, 217–230.
- (2002b): “Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses,” *The Quarterly Journal of Economics*, 117, 1039–1073.
- SAEZ, E., J. SLEMROD, AND S. H. GIERTZ (2012): “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Literature*, 50, 3–50.
- SANDERS, B. (2016): “A Soda Tax Would Hurt Philly’s Low-Income Families,” .
- SANDMO, A. (1975): “Optimal Taxation in the Presence of Externalities,” *The Swedish Journal of Economics*, 77, 86–98.
- SCHOFIELD, H. (2015): “The Economic Costs of Low Caloric Intake: Evidence from India,” *Working Paper*.
- SILVER, L. D., S. W. NG, S. RYAN-IBARRA, L. S. TAILLIE, M. INDUNI, D. R. MILES, J. M. POTI, AND B. M. POPKIN (2017): “Changes in Prices, Sales, Consumer Spending, and Beverage Consumption One Year After a Tax on Sugar-Sweetened Beverages in Berkeley, California, US: A Before-And-After Study,” *PLoS: Medicine*.
- SMITH, T. A., B.-H. LIN, AND J.-Y. LEE (2010): “Taxing Caloric Sweetened Beverages: Potential Effects on Beverage Consumption, Calorie Intake, and Obesity,” *USDA Economic Research Report Number 100*.
- SPINNEWIJN, J. (2017): “Heterogeneity, Demand for Insurance and Adverse Selection,” *American Economic Journal: Economic Policy*.
- TAUBINSKY, D. AND A. REES-JONES (2018): “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” *Review of Economic Studies*.
- THOW, A. M., S. DOWNS, AND S. JAN (2014): “A systematic review of the effectiveness of food taxes and subsidies to improve diets: Understanding the recent evidence,” *Nutrition Reviews*.
- TIFFIN, R., A. KEHLBACHER, AND M. SALOIS (2015): “The Effects of A Soft Drink Tax in the UK,” *Health Economics*.
- TSVETANOV, T. AND K. SEGERSON (2013): “Re-Evaluating the Role of Energy Efficiency Standards: A Behavioral Economics Approach,” *Journal of Environmental Economics and Management*.

- WANG, E. Y. (2015): “The Impact of Soda Taxes on Consumer Welfare: Implications of Storability and Taste Heterogeneity,” *The RAND Journal of Economics*.
- WANG, Y. C., P. COXSON, Y.-M. SHEN, L. GOLDMAN, AND K. BIBBINS-DOMINGO (2012): “A Penny-Per-Ounce Tax on Sugar-Sweetened Beverages Would Cut Health and Cost Burdens of Diabetes,” *Health Affairs*, 31, 199–207.
- YONG, P. L., J. BERTKO, AND R. KRONICK (2011): “Actuarial Value and Employer-Sponsored Insurance,” .
- ZHEN, C., E. A. FINKELSTEIN, J. NONNEMAKER, S. KARNS, AND J. E. TODD (2014): “Predicting the Effects of Sugar-Sweetened Beverage Taxes on Food and Beverage Demand in a Large Demand System,” *American Journal of Agricultural Economics*.
- ZHEN, C., M. K. WOHLGENANT, S. KARNS, AND P. KAUFMAN (2011): “Habit Formation and Demand for Sugar-Sweetened Beverages,” *American Journal of Agricultural Economics*.

Table 1: **Descriptive Statistics**

	Obs.	Mean	Std. dev.	Min	Max
Household income (\$000s)	653,554	68.00	47.59	3	167
Years education	653,554	13.92	2.06	6	18
Age	653,554	52.24	14.41	18	95
1(White)	653,554	0.77	0.42	0	1
1(Black)	653,554	0.12	0.33	0	1
1(Have children)	653,554	0.33	0.47	0	1
Household size (adult equivalents)	653,554	2.48	1.36	1	11
1(Employed)	653,554	0.61	0.44	0	1
Weekly work hours	653,554	22.84	16.73	0	40
SSBs purchased (liters)	653,554	153.14	190.70	0	12,582
Average price (\$/liter)	632,073	1.12	1.38	0	228

(a) **Homescan Household-by-Year Data**

	Obs.	Mean	Std. dev.	Min	Max
Nutrition knowledge	20,640	0.70	0.15	0	1
Self-control	20,640	0.77	0.34	0	1
Other head self-control	13,066	0.67	0.38	0	1
Taste for juice drinks	20,640	0.49	0.32	0	1
Taste for soda	20,640	0.52	0.36	0	1
Taste for tea/coffee	20,640	0.45	0.36	0	1
Taste for sports drinks	20,640	0.29	0.32	0	1
Taste for energy drinks	20,640	0.17	0.28	0	1
Taste for fruit juice	20,640	0.72	0.29	0	1
Taste for diet drinks	20,640	0.32	0.37	0	1
Health importance	20,640	0.84	0.18	0	1
SSB consumption (liters)	20,640	87.65	146.05	0	1735
1(Male)	20,640	0.28	0.45	0	1
1(Primary shopper)	20,640	0.88	0.33	0	1

(b) **PanelViews Respondent-Level Data**

Notes: Panel (a) presents descriptive statistics on the Nielsen Homescan data, which are at the household-by-year level for 2006-2016. If there are two household heads, we use the two heads' mean age, education, employment status, and weekly work hours. We code weekly work hours as zero for people who are not employed. For people who are employed, weekly work hours is reported in three bins: <30 , $30-34$, and ≥ 35 , which we code as 24, 32, and 40, respectively. The U.S. government Dietary Guidelines include calorie needs by age and gender; we combine that with Homescan household composition to get each household member's daily calorie need. Household size in "adult equivalents" is the number of household heads plus the total calorie needs of all other household members divided by the nationwide average calorie need of household heads. Prices and incomes are in real 2016 dollars. Panel (b) presents descriptive statistics on the Homescan PanelViews data, with one observation for each respondent. See Appendix D for the text of the PanelViews survey questions. Observations are weighted for national representativeness.

Table 2: **Instrumental Variables Estimates of Price and Income Elasticities**

	(1)	(2)	(3)	(4)
ln(Average price/liter)	-1.401*** (0.096)	-1.513*** (0.104)	-1.529*** (0.104)	-1.371*** (0.096)
ln(County income)	0.183** (0.071)	0.195** (0.078)	0.158** (0.078)	0.179** (0.071)
Feature	1.145*** (0.069)	1.137*** (0.074)	1.125*** (0.073)	1.156*** (0.071)
Display	0.492*** (0.087)	0.514*** (0.087)	0.503*** (0.087)	0.494*** (0.086)
Market-quarter fixed effects	No	Yes	No	No
Market-quarter-income fixed effects	No	No	Yes	No
Kleibergen-Paap first stage F stat	191.1	187.1	189.1	214.0
N	2,205,632	2,205,632	2,205,584	2,205,632

Notes: This table presents estimates of Equation (19). All regressions include quarter of sample indicators and household-by-county fixed effects. Columns 1-3 present instrumental variables estimates using the primary IV. Column 4 constructs the instrument using deviations from regional average prices instead of national average prices. In columns 2 and 3, “market” is Nielsen’s Designated Market Area (DMA). Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table 3: **Regressions of Sugar-Sweetened Beverage Consumption on Bias Proxies**

	(1)	(2)	(3)	(4)	(5)
Nutrition knowledge	-0.730*** (0.086)	-0.987*** (0.084)	-0.836*** (0.085)	-0.781*** (0.085)	-0.927*** (0.089)
Self-control	-1.449*** (0.070)	-1.847*** (0.058)	-1.353*** (0.070)	-1.460*** (0.068)	-0.850*** (0.043)
Taste for soda	0.444*** (0.047)		0.445*** (0.047)	0.443*** (0.048)	0.618*** (0.046)
Health importance	-0.194*** (0.074)		-0.033 (0.074)	-0.286*** (0.076)	-0.273*** (0.078)
ln(Household income)	-0.024 (0.018)		-0.057*** (0.016)	-0.054*** (0.018)	-0.045** (0.019)
ln(Years education)	-0.701*** (0.096)		-0.772*** (0.096)	-0.846*** (0.100)	-0.727*** (0.101)
Other beverage tastes	Yes	No	Yes	Yes	Yes
Other demographics	Yes	Yes	No	Yes	Yes
County indicators	Yes	Yes	Yes	No	Yes
Self-control 2SLS	Yes	Yes	Yes	Yes	No
R^2	0.287	0.252	0.276	0.169	0.287
N	18,568	18,568	18,568	18,568	18,568

Notes: This table presents estimates of Equation (23). Data are at the household level, and the dependent variable is the natural log of SSB purchases per adult equivalent in the most recent year that the household was in Homescan. Columns 1-4 correct for measurement error in self-control using two-sample 2SLS, with standard errors calculated per Chodorow-Reich and Wieland (2016). Taste for soda is the response to the question, “Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking [Regular soft drinks (soda pop)]?” “Other beverage tastes” are the responses to parallel questions for other beverages. Health importance is the response to the question, “In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?” Responses to each question were originally on a scale from 0 to 10, which we rescale to between 0 and 1. “Other demographics” are natural log of age, race, an indicator for the presence of children, household size in adult equivalents, employment status, and weekly work hours. Observations are weighted for national representativeness. Robust standard errors are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table 4: **Baseline Optimal Tax Calculation**

	Value
SSB consumption (ounces per week): \bar{s}	41.21
SSB price (cents per ounce): p	3.17
SSB demand elasticity: $\bar{\zeta}^c$	1.44
Elasticity of taxable income: $\bar{\zeta}_z$	0.33
Average marginal bias (cents per ounce): $\bar{\gamma}$	1.17
Externality (cents per ounce): e	0.85

(a) **Population Sufficient Statistics**

z	f	$\bar{s}(z)$	$\bar{\zeta}^c(z)$	$\xi(z)$	$\bar{\gamma}(z)$	$g(z)$	$T'(z)$	$s_{pref}(z)$
5000	0.12	55.6	1.48	0.15	1.37	2.70	-0.19	0.0
15000	0.15	50.4	1.48	0.15	1.11	1.40	-0.05	-15.2
25000	0.13	47.5	1.48	0.15	1.13	1.01	0.08	-23.3
35000	0.10	41.7	1.48	0.15	1.10	0.81	0.15	-32.8
45000	0.08	39.7	1.48	0.15	1.13	0.68	0.19	-37.6
55000	0.07	38.1	1.35	0.17	1.17	0.59	0.21	-41.8
65000	0.09	34.4	1.35	0.17	1.18	0.52	0.21	-47.8
85000	0.09	30.7	1.35	0.17	1.17	0.42	0.22	-55.4
125000	0.14	26.4	1.35	0.17	1.06	0.31	0.23	-65.7

$\sigma = 0.22$, $Cov[g(z), s_{pref}(z)] = 13.59$, $\mathbb{E}\left[\frac{T'(z)}{1-T'(z)}\bar{\zeta}_z\bar{s}(z)\bar{\xi}_{inc}(z)\right] = 0.23$

(b) **Calculating Covariances**

Under existing income tax	$t \approx \frac{\bar{s}\bar{\zeta}^c(\bar{\gamma}(\bar{g}+\sigma)+e)-p(Cov[g(z),s(z)]+\mathcal{A})}{\bar{s}\bar{\zeta}^c+Cov[g(z),s(z)]+\mathcal{A}} \approx 1.75$
Under optimal income tax	$t \approx \frac{\bar{s}\bar{\zeta}^c(\bar{\gamma}(1+\sigma)+e)-pCov[g(z),s_{pref}(z)]}{\bar{s}\bar{\zeta}^c+Cov[g(z),s_{pref}(z)]} \approx 1.26$

(c) **Optimal Sugar-Sweetened Beverage Tax**

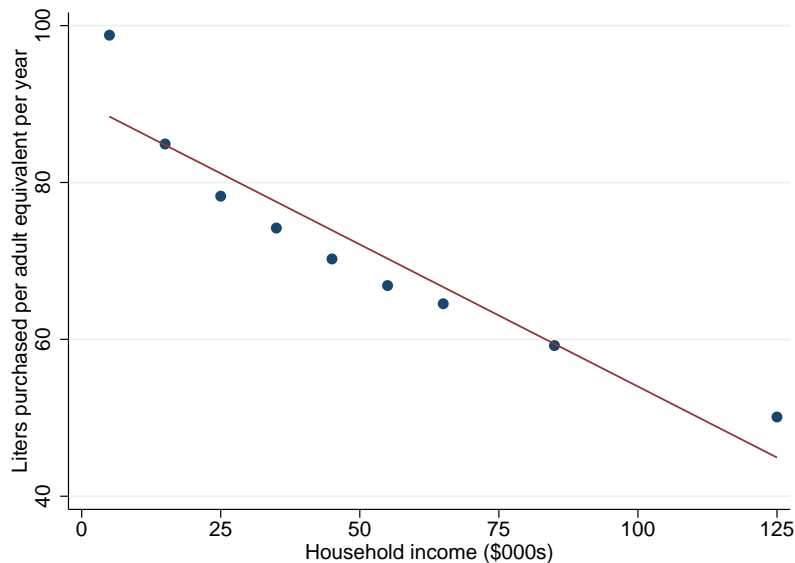
Notes: Panel (a) reports estimates of population-level sufficient statistics required to compute the optimal SSB tax. All statistics are computed using the data described in Section III, except for the externality e , the calculation of which is described in Section III.E, and the elasticity of taxable income $\bar{\zeta}_z^c$, which is drawn from Chetty (2012). Panel (b) reports sufficient statistics by income bin, which are used to compute covariances using the formula $Cov[a, b] = \sum_n f_n(a_n b_n) - \sum_n f_n a_n \sum_n f_n b_n$, where n indexes rows. Income bins (z) are those recorded in the Homescan data discussed in III.A; f represents the U.S. population share with pre-tax incomes in ranges bracketed by midpoints between each income bin, according to Piketty, Saez, and Zucman (2018). The statistics $\bar{s}(z)$, $\bar{\zeta}^c(z)$, $\xi(z)$, and $\bar{\gamma}(z)$ represent SSB consumption (in ounces per week), the compensated SSB demand elasticity, the SSB income elasticity, and average marginal money-metric bias estimated within each income bin, as described in Sections III.B to III.D. The column $g(z)$ reports our assumed marginal social welfare weights, while $T'(z)$ represents estimated net marginal tax rates (see text for discussion of each). The column $s_{pref}(z)$ is computed as $\bar{s}(z) - s_{inc}(z)$, where $s_{inc}(z_n) = \bar{s}(z_1) \prod_{j=2}^n \left(\frac{z_j}{z_{j-1}}\right)^{\frac{\xi_j + \xi_{j-1}}{2}}$. Panel (c) uses each of these values to compute the optimal tax estimates using Equation (15) and Proposition 2, where $\mathcal{A} = \mathbb{E}\left[\frac{T'(z)}{1-T'(z)}\bar{\zeta}_z\bar{s}(z)\bar{\xi}(z)\right]$.

Table 5: **Optimal Sugar-Sweetened Beverage Tax Under Alternative Assumptions**

	Existing income tax	Optimal income tax
Baseline	1.75	1.26
Pigouvian (no redistributive motive)	2.02	-
Weaker redistributive preferences	1.92	1.79
Stronger redistributive preferences	1.58	0.64
Higher demand elasticity ($\zeta^c(\theta) = 2$)	1.86	1.48
Lower demand elasticity ($\zeta^c(\theta) = 1$)	1.53	0.90
Pure preference heterogeneity	1.77	1.77
Pure income effects	1.82	2.27
No internality	0.46	0.10
No corrective motive	-0.30	-0.59
Self-reported SSB consumption	2.82	2.21

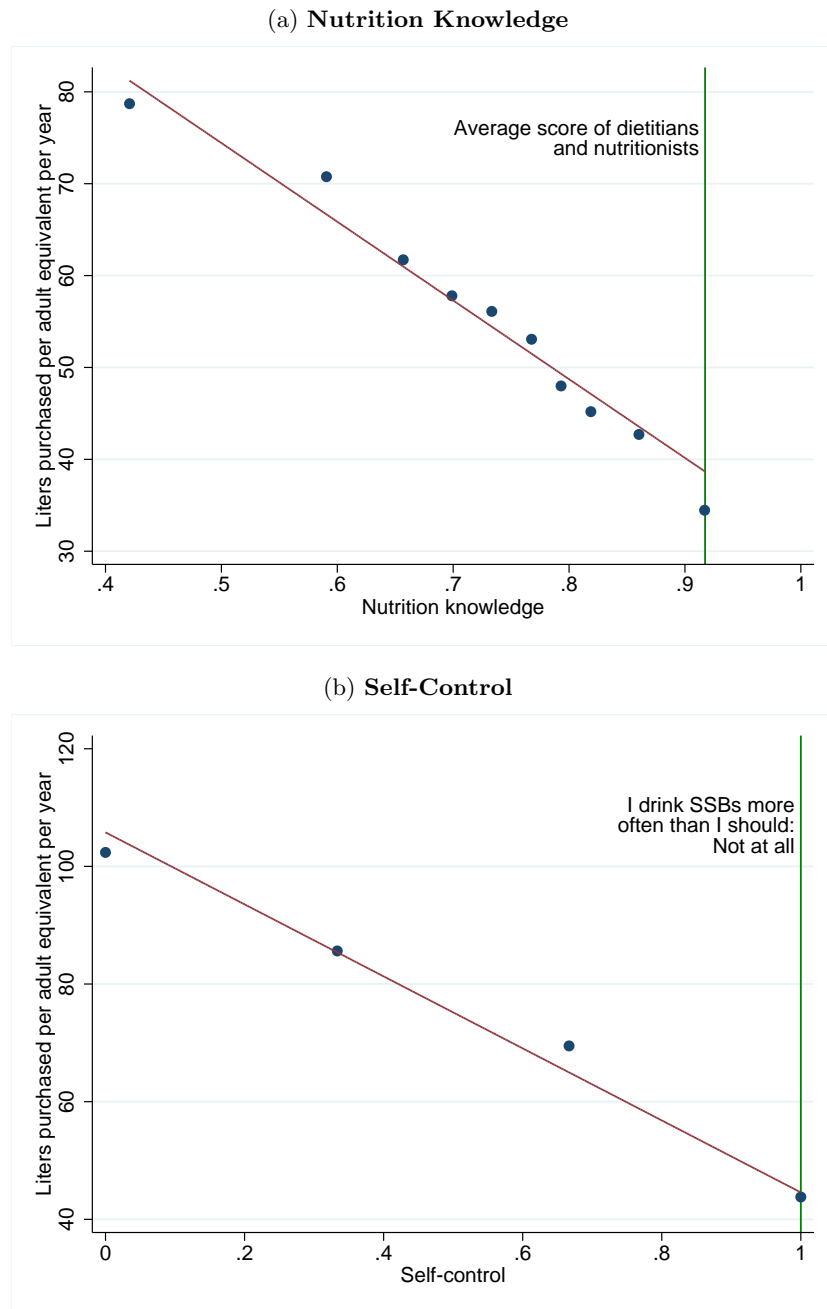
Notes: This table reports the optimal sweetened beverage tax, as computed using the sufficient statistics formulas for t^* under the status quo U.S. income tax (using Equation (15)) and under the optimal income tax (using Proposition 2) across a range of assumptions. The first row reports our baseline calculations, which employ the sufficient statistics by income bin displayed in Table 4. The second row reports the Pigouvian optimal tax, equal to $\bar{\gamma} + e$. The following two rows report the optimal tax under weaker and and stronger redistributive social preferences than the baseline. (Social marginal welfare weights are computed to be proportional to $y_{US}^{-\nu}$, where y_{US} is post-tax income in each bin—see Appendix K.B for details—with $\nu = 1$, $\nu = 0.25$, and $\nu = 4$ in the baseline, and under “weaker” and “stronger” redistributive preferences, respectively.) “Self-reported SSB consumption” reports results using SSB consumption data from our PanelViews survey, rather than from Homescan. “No internality” is computed assumes zero bias for all individuals, while “No corrective motive” assumes zero bias and zero externality.

Figure 1: **Homescan Sugar-Sweetened Beverage Purchases by Income**




Notes: This figure presents the average purchases of sugar-sweetened beverages by household income, using Nielsen Homescan data for 2006-2016. Purchases are measured in liters per “adult equivalent,” where household members other than the household heads are rescaled into adult equivalents using the recommended average daily consumption for their age and gender group. Observations are weighted for national representativeness.

Figure 2: Nutrition Knowledge and Self-Control vs. Sugar-Sweetened Beverage Consumption




Notes: These figures present average purchases of sugar-sweetened beverages for each household’s most recent year in the Nielsen Homescan data against the primary shopper’s nutrition knowledge (in Panel (a)) and self-control (in Panel (b)). Nutrition knowledge is the share correct out of 28 questions from the General Nutrition Knowledge Questionnaire (Kliemann, Wardle, Johnson, and Croker, 2016). Self-control is level of agreement with the statement, “I drink soda pop or other sugar-sweetened beverages more often than I should,” with answers coded as “Definitely”=0, “Mostly”=1/3, “Somewhat”= 2/3, and “Not at all”=1. Purchases are measured in liters per “adult equivalent,” where household members other than the household heads are rescaled into adult equivalents using the recommended average daily consumption for their age and gender group. Observations are weighted for national representativeness.

Figure 3: **Contemporaneous First Stage and Reduced Form of the Local Price Instrument**




Notes: These figures present binned scatterplots of the first stage (in Panel (a)) and reduced form (in Panel (b)) of the instrumental variables estimates of Equation (19). Both relationships are residual of the other variables in Equation (19): feature and display, natural log of county mean income, quarter of sample indicators, and household-by-county fixed effects. Purchases are measured in liters per “adult equivalent,” where household members other than the household heads are rescaled into adult equivalents using the recommended average daily consumption for their age and gender group. Observations are weighted for national representativeness.

Figure 4: Causal Income Effects and Between-Income Preference Heterogeneity



Notes: The dark circles plot average purchases of sugar-sweetened beverages by household income, using Nielsen Homescan data for 2006-2016. Purchases are measured in liters per “adult equivalent,” where household members other than the household heads are rescaled into adult equivalents using the recommended average daily consumption for their age and gender group. The curve represents predicted average purchases at each income level that would result from applying the estimated income elasticity $\hat{\xi} \approx 0.18$ to the average consumption of households with annual income under \$10,000: $s(z_{min}) \left(\frac{z}{z_{min}}\right)^{\hat{\xi}}$. The x’s are $\bar{s}_{pref}(z)$, the difference between actual consumption and consumption predicted only using income elasticity. Observations are weighted for national representativeness.


Figure 5: Nutrition Knowledge and Self-Control by Income




Notes: These figures present average nutrition knowledge (in Panel (a)) and self-control (in Panel (b)) by household income. Nutrition knowledge is the share correct out of 28 questions from the General Nutrition Knowledge Questionnaire (Kliemann et al., 2016). Self-control is level of agreement with the statement, “I drink soda pop or other sugar-sweetened beverages more often than I should,” with answers coded as “Definitely”=0, “Mostly”=1/3, “Somewhat”= 2/3, and “Not at all”=1. Observations are weighted for national representativeness.

Figure 6: Preferences by Income

(a) Taste for Soda




(b) Health Importance




Notes: These figures present average taste for soda (in Panel (a)) and health importance (in Panel (b)) by household income. Taste for soda is the response to the question, “Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking [Regular soft drinks (soda pop)]?” Health importance is the response to the question, “In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?” Responses to each question were originally on a scale from 0 to 10, which we rescale to between 0 and 1. Observations are weighted for national representativeness.

Figure 7: **Share of Consumption Attributable to Bias by Income**




Notes: This figure presents the share of sugar-sweetened beverage consumption attributable to bias, i.e. the unweighted average of $\frac{s_i - \hat{s}_i^V}{s_i}$, by income, using Homescan purchase data. Observations are weighted for national representativeness.

Figure 8: **Average Marginal Bias by Income**



Notes: This figure presents the demand slope-weighted average marginal bias by income, using Homescan purchase data. Observations are weighted for national representativeness.

Figure 9: Welfare Consequences of Optimal Sugar-Sweetened Beverage Tax



Notes: This figure plots the decomposition of welfare changes resulting from the baseline optimal sugar-sweetened beverage tax, across income bins. “Redistributed revenues” and “Externality correction” are money-metric values, assumed to be distributed equally across the income distribution. “Internality correction” is the increase in (money-metric) welfare due to the change in consumption resulting from the tax, at each income level. “Decision utility EV” is the value dy such that consumers are indifferent between a change in net income dy and the introduction of the optimal SSB tax.

Online Appendix: Not for Publication

Regressive Sin Taxes, with an Application to the Optimal Soda Tax

Hunt Allcott, Benjamin B. Lockwood, Dmitry Taubinsky

A Income Elasticity Definitions

We define labor supply responses to include any “circularities” due to the curvature of the income tax function, which is assumed to be differentiable. Thus, following Jacquet and Lehmann (2014), we define a tax function \hat{T} which has been locally perturbed around the income level z_0 by raising the marginal tax rate by τ and reducing the tax level by ν :

$$\hat{T}(z; z_0, \tau, \nu) = T(z) + \tau(z - z_0) - \nu. \quad (26)$$

Let $z^*(\theta)$ denote a type θ 's choice of earnings under the status quo income tax T , and let $\hat{z}(\theta; \tau, \nu)$ denote θ 's choice of earnings under the perturbed income tax $\hat{T}(z; z^*(\theta), \tau, \nu)$. Then the compensated elasticity of taxable income is defined in terms of the response of \hat{z} to τ , evaluated at $\tau = \nu = 0$:

$$\zeta_z^c(\theta) := \left(- \frac{\partial \hat{z}(\theta; \tau, 0)}{\partial \tau} \Big|_{\tau=0} \right) \frac{1 - T'(z^*(\theta))}{z^*(\theta)}. \quad (27)$$

The income tax is similarly defined in terms of the response of \hat{z} to a tax credit ν (this statistic will be nonpositive if leisure is a non-inferior good):

$$\eta_z(\theta) := \left(\frac{\partial \hat{z}(\theta; 0, \nu)}{\partial \nu} \Big|_{\nu=0} \right) (1 - T'(z^*(\theta))). \quad (28)$$

These definitions are comparable to those in Saez (2001), except that they include circularities and thus permit a representation of the optimal income tax in terms of the actual earnings density, rather than the “virtual density” employed in that paper.

B Extensions and Additional Results

In this appendix we present additional theoretical results. To facilitate clarity, we restrict to the case of unidimensional types θ . Unless otherwise noted, the proofs of these results are contained in Appendix C.

B.A Many Dimensions of Consumption

Our results readily generalize to $N > 2$ dimensions of consumption. Let utility be given by $U(c_1, c_2, \dots, c_n, z; \theta)$, with c_1 representing the numeraire, and assume that $U_{c_1}/U_z = V_{c_1}/V_z$. Let \mathbf{S} denote the Slutsky matrix of compensated demand responses in which the j, i entry is the compensated demand response of c_j with respect to t_i .

Let $\gamma_i(\theta)$ denote the bias in consumption dimension i , with $\gamma_1(\theta) = 0$ for all θ . Let $\bar{\gamma}_{ij}$ denote the average marginal bias from consumption of j with respect to tax i :

$$\bar{\gamma}_{ij} = \frac{\int_{\Theta} \gamma_j(\theta, t, T) \left(\frac{dc_j(\theta, t, T)}{dt_i} \Big|_u \right) \mu(\theta)}{\int_{\Theta} \left(\frac{dc_j(\theta, t, T)}{dt_i} \Big|_u \right) \mu(\theta)}, \quad (29)$$

where $\frac{dc_j(z)}{dt_i}\Big|_u$ denotes the (compensated) derivative of θ 's consumption of c_j with respect to tax t_i .

We also define $\sigma_{ij} := Cov_H \left[g(z), \frac{\bar{\gamma}_j(z)}{\bar{\gamma}_j} \frac{\frac{dc_j(z)}{dt_i}\Big|_u}{\frac{d\bar{c}_j}{dt_i}\Big|_u} \right]$, where \bar{c}_j denotes population average consumption of c_j . Finally, let e_j denote the marginal externality from consumption of c_j . We define \mathbf{R} to be a column vector in which the i th entry is $R_i = -\sum_j \frac{dc_j(\theta, t, T)}{dt_i}\Big|_u (\bar{\gamma}_{ij}(\bar{g} + \sigma_{ij}) + e_j)$. That is, R_i is the total corrective benefit from increasing tax t_i , keeping income constant.

Employing notation similar to Section II, we define $c_{inc,j}(z)$ to be the portion of c_j consumption explained by cumulative income effects, and we define $c_{pref,j}(z) := (c_j(z) - c_{inc,j}(z))$ to be cumulative between-income preference heterogeneity. We define $\boldsymbol{\rho}$ to be the column vector in which the j th entry is $Cov[\hat{g}(z), c_{pref,j}(z)]/\bar{c}_j$.

Proposition 3. *The optimal commodity taxes $\mathbf{t} = (t_1, \dots, t_N)$ satisfy*

$$\mathbf{tS} = \boldsymbol{\rho} - \mathbf{R} \quad (30)$$

and the optimal income tax satisfies

$$\frac{T'(z)}{1 - T'(z)} = \sum_i \frac{g(z)\bar{\gamma}_i(z) + e_i - t_i}{1 - T'(z)} \frac{dc_i(z, \theta)}{dz} + \frac{1}{\zeta_z^c h(z)} \int_z^\infty (1 - \hat{g}(x)) dH(x). \quad (31)$$

B.B Composite Goods

We can also allow for the possibility that the sin good s is a composite of several sin goods, (s_1, s_2, \dots, s_n) with pre-tax prices (p_1, p_2, \dots, p_n) , and with marginal (money metric) bias of $\gamma_i(\theta) = \frac{U'_c}{U'_c} - \frac{V'_c}{V'_c}$ at the optimum. Assume all sin goods are subject to a common *ad valorem* tax τ , so that the post-tax price of each s_i is $p_i(1 + \tau)$. Under this composite setup, the optimal sin tax τ still satisfies the general expression from Proposition 1, provided the parameters therein are reinterpreted as follows:

- $s(\theta) := \sum_i p_i s_i(\theta)$, representing total revenues (net of taxes) from sin good spending by θ
- $\gamma(\theta) := \sum_i \gamma_i(\theta) \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i / \sum_i \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i$, representing the response-weighted bias, where $q_i := (1 + \tau)p_i$
- $\zeta(\theta) := \sum_i \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i \cdot (1 + \tau) / s(\theta)$, representing the aggregate elasticity of sin good spending
- $p := \sum p_i \frac{s_i}{s}$, representing the average price of the composite

B.C Computing the Optimal Tax in Terms of $g(z)$ rather than $\hat{g}(z)$

Here we solve for the optimal tax system in terms of the social marginal welfare weights g rather than \hat{g} .

The key variable needed for this solution is the following: “When we raise the commodity tax by dt , how much money can we give back to each person to offset the tax so as to keep the average

labor supply choices of each z -earner constant?” Roughly, we seek to characterize the extent to which we can offset the regressivity costs of the commodity tax by making the income tax more progressive. We begin with a lemma providing this characterization:

Lemma 1. *Let $\chi(z) := s_{inc}(z) - \int_0^z w(x, z) \frac{\eta_z}{\zeta_x^c} (s(x) - s_{inc}(x)) dx$, where $w(x, z) = e^{\int_{z'=x}^{z'=z} -\frac{\eta_z}{\zeta_x^c} dx'}$. Then increasing the commodity tax by dt and decreasing the income tax by $\chi(z)dt$ leaves the average labor supply of z -earners unchanged.*

The intuition for the “income effect adjustment” term is that the higher the income effect, the more the commodity tax increases labor supply through the income effect, and thus the higher χ can be. Note that $\chi(z) = s_{inc}(z)$ when $\eta_z = 0$. That is, in the absence of labor supply income effects, the term $s_{inc}(z)$ —which we defined as the portion of s consumption explained by income effects—is the extent to which the income tax should be optimally lowered for z -earners. When $\eta_z \neq 0$ and $s_{inc}(z) < \bar{s}(z)$, $\chi(z)$ will be higher than $s_{inc}(z)$. Intuitively, this is because giving back $s_{inc}(z)dt$ to consumers offsets the effective impact on the marginal keep-rate from before-tax earnings, but does not fully offset the income effect as it leaves consumers poorer by an amount $(s - s_{inc})dt$.

Proposition 4. *The optimal commodity tax t satisfies*

$$t = \underbrace{\bar{s}\bar{\gamma}(\bar{g} + \sigma)}_{\text{corrective effect}} + \underbrace{\frac{p+t}{\bar{\zeta}^c} \mathbb{E}[(1-g(z))(s(z) - \chi(z))]}_{\text{redistributive effect}} - \underbrace{\frac{1}{\bar{\zeta}^c} \mathbb{E}[s(z) - \chi(z)]\eta(z)(t - g(z)\bar{\gamma}(z))}_{\text{additional impact from income effect}}. \quad (32)$$

Proposition 4 provides a commodity tax formula similar to Equation (15), with two differences. First, the covariance term $Cov[g(z), s_{pref}(z)]$, where $s_{pref}(z) = s(z) - s_{inc}(z)$, corresponding to regressivity costs is written more generally as $E[(g(z)-1), s(z)-\chi(z)]$. In the absence of labor supply income effects, $\chi(z) = s_{inc}(z)$ and thus the two terms are equivalent. Second, the presence of income effects η means that a tax reform that decreases consumers’ after-tax income by $s(z) - \chi(z)$ also impacts consumption of s , beyond the compensated elasticity with respect to prices. This generates fiscal externalities proportional to $t\eta(s(z) - \chi(z))$ and additional corrective benefits proportional to $g\bar{\gamma}(s(z) - \chi(z))$. Combined, this generates the third term in Equation (32) above. In the absence of both kinds of income effects it is easy to see that the formula reduces to

$$t = \bar{\gamma}(\bar{g} + \sigma) - \frac{p+t}{\bar{s}\bar{\zeta}^c} Cov[g(z), s(z) - s_{inc}(z)], \quad (33)$$

which can be solved for t to get the formula in Equation (15).

C Proofs of propositions

C.A Proof of Proposition 1

We prove Proposition 1 in three steps. First, we characterize the optimal income tax. Second, we characterize the optimal commodity tax under fully general heterogeneity, which constitutes Proposition 5 below. Third, we invoke Assumption 4 from Section II.B.1, which simplifies the commodity tax expression in Proposition 5 to yield the optimal commodity tax formula in Proposition 1.

For the general result, we define two new elasticity concepts. First, we define the elasticity of s with respect to a *windfall* of pre-tax income, $\xi_\nu := \frac{\partial s}{\partial \nu} \cdot \frac{z}{s} \cdot (1 - T')$, where ν is defined as in Appendix A. Note that ξ_ν relates to the standard income effect η as follows: $\eta = \xi_\nu \frac{(p+t)s}{z} \cdot \frac{1}{1-T'}$. We analogously define $\eta_\xi := \xi \cdot \frac{(p+t)s}{z} \cdot \frac{1}{1-T'}$. In words, η_ξ is the change in expenditure on s from an extra \$1 of post-tax income earned.

C.A.1 Proof of Proposition 1, Equation 14 (optimal income tax)

Following Saez (2001), we have the following effect from increasing the marginal tax rate between z^* and $z^* + dz$:

1. *Direct effect* (fiscal and welfare): $\int_{x=z^*}^{\infty} (1 - g(x)) dH(x)$

2. *Compensated elasticity effect* (including both fiscal and welfare components):

$$\begin{aligned} & - \bar{\zeta}_z^c(z^*) \cdot z^* \cdot \frac{T'}{1 - T'} h(z^*) - \mathbb{E} \left[\frac{\xi(\theta) \zeta_z^c(\theta) s(\theta)}{1 - T'} (t - g(z) \gamma(\theta) - e) \mid z(\theta) = z^* \right] h(z^*) = \\ & - \bar{\zeta}_z^c(z^*) \cdot z^* \cdot \frac{T'}{1 - T'} h(z^*) - \mathbb{E} \left[\frac{\xi(\theta) \zeta_z^c(\theta) s(\theta)}{1 - T'} (t - g(z) \gamma(\theta) - e) \mid z(\theta) = z^* \right] h(z^*) \quad (34) \end{aligned}$$

3. *Income effect* (fiscal and welfare):

$$\begin{aligned} & - \mathbb{E} \left[\eta_z(\theta) \frac{T'}{1 - T'} \mu(\theta) \mid z(\theta) > z^* \right] \\ & - \mathbb{E} \left[(t - g(\theta) \gamma(\theta) - e) \left(\frac{\eta(\theta)}{p + t} + \frac{\eta_z(\theta)}{1 - T'} \frac{\xi(\theta) s(\theta)}{z(\theta)} \right) \mu(\theta) \mid z(\theta) > z^* \right], \end{aligned}$$

where the second term combines the income effect on s from the change in after-tax income with the income effect on s from the adjustment in earnings z .

These three terms must sum to zero, as the net welfare effect of this reform should be zero at the optimum. Using the definition of $\hat{g}(z)$ from footnote 27, the direct effect from 1 combines with the income effects from 3 to become $\int_{x=z^*}^{\infty} (1 - \hat{g}(x)) dH(x)$. Therefore, the optimal income tax satisfies

$$\frac{T'(z^*)}{1 - T'(z^*)} = \frac{\int_{x=z^*}^{\infty} (1 - \hat{g}(x)) dH(x)}{\bar{\zeta}_z^c(z^*) z^* h(z^*)} - \mathbb{E} \left[\frac{(t - g(\theta) \gamma(\theta) - e)}{1 - T'(z(\theta))} \cdot \frac{\xi(\theta) \zeta_z^c(\theta) s(\theta)}{\bar{\zeta}_z^c(z^*) z(\theta)} \mid z(\theta) = z^* \right] \quad (35)$$

Re-solving for $T'(z^*)$ yields the desired result.

C.A.2 Expression for the optimal commodity tax under full heterogeneity

Let $\Sigma_{a,b}^{(z^*)} := \text{Cov}_{\mu|z^*} [a(\theta), b(\theta) | z(\theta) = z^*]$ denote the income-conditional covariance between variables $a(\theta)$ and $b(\theta)$ at the optimum. Then we have the following proposition.

Proposition 5. *The optimal commodity tax satisfies*

$$t = \frac{\bar{s}\bar{\zeta}^c (\bar{\gamma}(\bar{g} + \sigma) + e) + A_2 + eA_1 - p\text{Cov}[\hat{g}(z), s_{pref}(z)] - pA_3}{\bar{s}\bar{\zeta}^c + \text{Cov}[\hat{g}(z), s_{pref}(z)] + A_1 + A_3}, \quad (36)$$

where

$$\begin{aligned} A_1 &:= \int_{z^*=0}^{\infty} \left(\Sigma_{s,\eta}^{(z^*)} + \Sigma_{\xi s, \zeta_z^c \eta \xi}^{(z^*)} + \Sigma_{s\eta z, \eta \xi}^{(z^*)} \right) dH(z^*) \\ A_2 &:= \int_{z^*=0}^{\infty} g(z^*) \left(\Sigma_{s,\eta\gamma}^{(z^*)} + \Sigma_{\xi s, \zeta_z^c \eta \xi \gamma}^{(z^*)} + \Sigma_{s\eta z, \eta \xi \gamma}^{(z^*)} \right) dH(z^*) \\ A_3 &:= \int_{z^*=0}^{\infty} \frac{\int_{x=z^*}^{\infty} (1 - \hat{g}(x)) dH(x)}{\bar{\zeta}_z^c(z^*) \cdot z^*} \Sigma_{\eta z, s}^{(z^*)} dz^*. \end{aligned}$$

Term A_1 concerns income-conditional covariances that affect substitution, and thus affect the fiscal externality and the impact on externality reduction. Term A_2 is similar: it involves income-conditional covariances between the substitution and the bias, which may also be heterogeneous conditional on income. To begin the proof, consider the total effect on welfare of a marginal increase in the commodity tax dt . The total welfare effect, written in terms of the marginal value of public funds, can be decomposed into the following components:

- *Mechanical revenue effect:* the reform mechanically raises revenue from each consumer by $dt \cdot s(\theta)$, for a total of $dt \cdot S$.
- *Mechanical welfare effect:* the reform mechanically reduces each consumer's net income by $dt \cdot s(\theta)$. To isolate the mechanical effect, we compute the loss in welfare as if this reduction all comes from composite consumption c for a welfare loss of $dt \cdot s(\theta)g(\theta)$, and we account for adjustments in s and z in the behavioral effects below. Thus the total mechanical welfare effect is $-dt \int_{\Theta} s(\theta)g(\theta)d\mu(\theta)$.
- *Direct effect on sin good consumption:* the reform causes each consumer to decrease their s consumption by $dt \cdot \zeta(\theta) \frac{s(\theta)}{p+t}$. This generates a fiscal externality equal to $-dt \int_{\Theta} t\zeta(\theta) \frac{s(\theta)}{p+t} \mu(\theta)$. It also generates a behavioral welfare effect equal to $dt \int_{\Theta} g(\theta)\gamma(\theta)\zeta(\theta) \frac{s(\theta)}{p+t} \mu(\theta)$ and an externality effect equal to $dt \int_{\Theta} e\zeta(\theta) \frac{s(\theta)}{p+t} \mu(\theta)$.
- *Effect on earnings:* The reform causes a change in income tax revenue collected from type θ equal to $\frac{dz(\theta)}{dt} T'(z(\theta))$. To compute $\frac{dz(\theta)}{dt}$, we use Lemma 1 from Saez (2002a), which carries through in this context: the change in earnings for type θ due to the increase in the

commodity tax dt is equal to the change in earnings that would be induced by imposing a type-specific income tax reform, raising the income tax by $dT^\theta(z) = dt \cdot s(p+t, z - T(z), z; \theta)$. Intuitively, an increase in the commodity tax alters the utility that type θ would gain from selecting each possible level of earnings, in proportion to the amount of s that θ would consume at each income. This alternative income tax reform reduces the utility θ would realize at each possible income by instead altering the income tax. The resulting adjustment in earnings can be decomposed into a compensated elasticity effect and an income effect. The compensated elasticity effect depends on the change in the marginal income tax rate, which equals $dt \cdot \frac{ds(z; \theta)}{dz}$. (We adopt the shorthand $\frac{ds(z; \theta)}{dz} := \frac{ds(p+t, z - T(z), z; \theta)}{dz}$. Because this derivative holds type constant, it is different from $\bar{s}'(z)$ even with unidimensional types, which also incorporates how types vary with z .) The income effect depends on the change in the tax level, which equals $dt \cdot s(\theta)$. Combining these effects, we have $\frac{dz(\theta)}{dt} = -\zeta_z^c(\theta) \left(\frac{z(\theta)}{1 - T'(z(\theta))} \right) \frac{ds(z; \theta)}{dz} - \frac{\eta_z(\theta)}{1 - T'(z(\theta))} s(\theta)$. This generates a total fiscal externality through the income tax equal to $-dt \cdot \int_{\Theta} \frac{T'}{1 - T'} \left(\zeta_z^c(\theta) z(\theta) \frac{ds(z; \theta)}{dz} + \eta_z(\theta) s(\theta) \right) \mu(\theta)$.

- *Indirect effects on sin good consumption:* The change in earnings affects consumption indirectly. This generates additional fiscal externalities and welfare effects. These total to

$$\begin{aligned} dt \cdot \int_{\Theta} \frac{dz(\theta)}{dt} \frac{ds(z; \theta)}{dz} (t - g(\theta)\gamma(\theta) - e) d\mu(\theta) \\ = -dt \cdot \int_{\Theta} \frac{1}{1 - T'(z(\theta))} \left(\frac{ds(z; \theta)}{dz} \zeta_z^c(\theta) z(\theta) + \eta_z(\theta) s(\theta) \right) \frac{ds(z; \theta)}{dz} (t - g\gamma - e). \end{aligned}$$

Combining these components, the total welfare effect of the tax reform dt is equal to

$$\frac{dW}{dt} = \overbrace{\int_{\Theta} \left(s(\theta)(1 - g(\theta)) - s(\theta) \left(\frac{\zeta^c(\theta)}{p+t} + \frac{\eta(\theta)}{p+t} \right) (t - g(\theta)\gamma(\theta) - e) \right) \mu(\theta)}^{\text{Direct effect on } s \text{ consumption}} \quad (37)$$

$$- \overbrace{\int_{\Theta} \frac{T'(z(\theta))}{1 - T'} \left(\zeta_z^c(\theta) z(\theta) \frac{ds(z; \theta)}{dz} + \eta_z(\theta) s(\theta) \right) \mu(\theta)}^{\text{Effect on labor supply}} \quad (38)$$

$$+ \overbrace{\int_{\Theta} \frac{dz(\theta)}{dt} \frac{ds(z; \theta)}{dz} (t - g(\theta)\gamma(\theta) - e) \mu(\theta)}^{\text{Indirect effects on } s \text{ consumption}} \quad (39)$$

$$= \int_{\Theta} \left(s(\theta)(1 - g(\theta)) - s(\theta) \left(\frac{\zeta^c(\theta)}{p+t} + \frac{\eta(\theta)}{p+t} \right) (t - g(\theta)\gamma(\theta) - e) \right) d\mu(\theta) \quad (40)$$

$$- \int_{\Theta} \frac{T'(z(\theta))}{1 - T'} \left(\zeta_z^c(\theta) z(\theta) \frac{ds(z; \theta)}{dz} + \eta_z(\theta) s(\theta) \right) d\mu(\theta) \quad (41)$$

$$- \int_{\Theta} \frac{1}{1 - T'(z(\theta))} z \zeta_z^c(\theta) \left(\frac{ds(z; \theta)}{dz} \right)^2 (t - g\gamma - e) \mu(\theta) \quad (42)$$

$$- \int_{\Theta} \frac{1}{1 - T'(z(\theta))} \eta_z(\theta) s(\theta) \frac{ds(z; \theta)}{dz} (t - g\gamma - e) \mu(\theta) \quad (43)$$

$$= \int_{\Theta} \left(s(\theta)(1 - g(\theta)) - s(\theta) \left(\frac{\zeta^c(\theta)}{p+t} + \frac{\eta(\theta)}{p+t} \right) (t - g(\theta)\gamma(\theta) - e) \right) \mu(\theta) \quad (44)$$

$$- \int_{\Theta} \zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1 - T'} + \frac{t - g(z)\gamma(\theta) - e}{1 - T'(z(\theta))} \frac{ds(z, \theta)}{dz} \right) \left(\frac{ds(z; \theta)}{dz} + \frac{\eta_z(\theta) s(\theta)}{\zeta_z^c(\theta) z(\theta)} \right) \mu(\theta). \quad (45)$$

If the income tax is optimal, any small perturbation to it generates no first-order change in welfare. In particular, this statement must hold for a reform which raises the income tax at each z in proportion to average s consumption at that z , i.e., for the perturbation $dT(z) = dt \cdot \bar{s}(z)$. Letting $\frac{dW}{dT}$ denote the total welfare effect of this marginal reform, we have

$$\frac{dW}{dT} = \int_{\Theta} \left(\bar{s}(z(\theta))(1 - g(\theta)) - \bar{s}(z(\theta)) \frac{\eta(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \right) \mu(\theta) \quad (46)$$

$$- \int_{\Theta} \frac{T'(z(\theta))}{1 - T'} \left(\zeta_z^c(\theta) z(\theta) \bar{s}'(z(\theta)) + \eta_z \bar{s}(z(\theta)) \right) \mu(\theta) \quad (47)$$

$$- \int_{\Theta} z \zeta_z^c \bar{s}'(z(\theta)) \frac{t - g(z)\gamma(\theta) - e}{1 - T'(z(\theta))} \frac{ds(z, \theta)}{dz} \mu(\theta) \quad (48)$$

$$- \int_{\Theta} \eta_z(\theta) \bar{s} \frac{t - g(z)\gamma(\theta) - e}{1 - T'(z(\theta))} \frac{ds(z, \theta)}{dz} \mu(\theta) \quad (49)$$

$$= \int_{\Theta} \left(\bar{s}(z(\theta))(1 - g(\theta)) - \bar{s}(z(\theta)) \frac{\eta(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \right) \mu(\theta) \quad (50)$$

$$- \int_{\Theta} \zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1 - T'} + \frac{t - g(z)\gamma(\theta) - e}{1 - T'(z(\theta))} \frac{ds(z, \theta)}{dz} \right) \left(\bar{s}'(z(\theta)) + \frac{\eta_z(\theta) \bar{s}(z(\theta))}{\zeta_z^c(\theta) z(\theta)} \right) \mu(\theta). \quad (51)$$

Since the income tax is optimal by assumption, we can simplify $\frac{dW}{dt}$ by subtracting $\frac{dW}{dT} = 0$. Using the fact that $g(\theta)$ is constant conditional on income, we have

$$\frac{dW}{dt} = - \int_{\Theta} s(\theta) \frac{\zeta^c(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \mu(\theta) \quad (52)$$

$$- \int_{\Theta} (s(\theta) - \bar{s}(z(\theta))) \frac{\eta(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \mu(\theta) \quad (53)$$

$$- \int_{\Theta} \zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1-T'} + \frac{t - g(z)\gamma(\theta) - e}{1-T'(z(\theta))} \frac{ds(z, \theta)}{dz} \right) \left(\frac{ds(z; \theta)}{dz} - \bar{s}'(z(\theta)) \right) \mu(\theta) \quad (54)$$

$$- \int_{\Theta} \zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1-T'} + \frac{t - g(z)\gamma(\theta) - e}{1-T'(z(\theta))} \frac{ds(z, \theta)}{dz} \right) \left(\frac{\eta_z(\theta) s(\theta)}{\zeta_z^c(\theta) z(\theta)} - \frac{\eta_z(\theta) \bar{s}(z(\theta))}{\zeta_z^c(\theta) z(\theta)} \right) \mu(\theta). \quad (55)$$

The term in line (52) can be written as

$$\int_{\Theta} s(\theta) \frac{\zeta^c(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \mu(\theta) = \frac{\bar{s} \bar{\zeta}^c}{p+t} (t - \bar{\gamma}(\bar{g} + \sigma) - e). \quad (56)$$

The term in line (53) can be rewritten as

$$\begin{aligned} & \int_{z^*=0}^{\infty} \mathbb{E} \left[(s(\theta) - \bar{s}(z(\theta))) \frac{\eta(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \mid z(\theta) = z^* \right] dH(z^*) \\ &= \int_{z^*=0}^{\infty} Cov \left[s(\theta), \frac{\eta(\theta)}{p+t} (t - g(\theta)\gamma(\theta) - e) \mid z(\theta) = z^* \right] dH(z^*) \\ &= \int_{z^*=0}^{\infty} \left[\frac{t}{p+t} \Sigma_{s, \eta}^{(z^*)} - \frac{1}{p+t} g(z^*) \Sigma_{s, \eta \gamma}^{(z^*)} - \frac{e}{p+t} \Sigma_{s, \eta}^{(z^*)} \right] dH(z^*). \end{aligned} \quad (57)$$

The term in line (54) can be written as

$$\begin{aligned} & \int_{z^*=0}^{\infty} \mathbb{E} \left[\zeta_z^c(\theta) z^* \left(\frac{T'(z^*)}{1-T'} + \frac{t - g(z)\gamma(\theta) - e}{1-T'(z^*)} \frac{ds(z, \theta)}{dz} \right) \left(\frac{ds(z; \theta)}{dz} - \bar{s}'(z^*) \right) \mid z(\theta) = z^* \right] dH(z^*) \\ &= \int_{z^*=0}^{\infty} \bar{\zeta}_z^c(z^*) z^* \mathbb{E} \left[\frac{T'(z^*)}{1-T'(z^*)} + \frac{t - g(z^*)\gamma(\theta) - e}{1-T'(z^*)} \frac{ds(z, \theta)}{dz} \frac{\zeta_z^c(\theta)}{\bar{\zeta}_z^c(z^*)} \mid z(\theta) = z^* \right] (s'_{inc}(z^*) - \bar{s}'(z^*)) h(z^*) dz^* \\ &+ \int_{z^*=0}^{\infty} z^* Cov \left[\zeta_z^c(\theta) \left(\frac{T'(z^*)}{1-T'(z^*)} + (t - g(z^*)\gamma(\theta) - e) \frac{\eta_{\xi}(\theta)}{p+t} \right), \xi(\theta) \frac{s(\theta)}{z(\theta)} \mid z(\theta) = z^* \right] \\ &= \int_{z^*=0}^{\infty} \left\{ \int_{x=z^*}^{\infty} (1 - \hat{g}(x)) dH(x) (s'_{inc}(z^*) - \bar{s}'(z^*)) \right\} dz^* \\ &+ \int_{z^*=0}^{\infty} \left[\frac{t}{p+t} \Sigma_{\zeta_z^c \eta_{\xi}, \xi s}^{(z^*)} - \frac{e}{p+t} \Sigma_{\zeta_z^c \eta_{\xi}, \xi s}^{(z^*)} - \frac{g(z^*)}{p+t} \Sigma_{\zeta_z^c \gamma \eta_{\xi}, \xi s}^{(z^*)} \right] dH(z^*). \end{aligned} \quad (58)$$

The term in line (55) can be written as

$$\begin{aligned}
& \int_{z^*=0}^{\infty} \mathbb{E} \left[\zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1-T'} + \frac{t-g(z)\gamma(\theta)-e}{1-T'(z(\theta))} \frac{ds(z,\theta)}{dz} \right) \left(\frac{\eta_z(\theta)s(\theta)}{\zeta_z^c(\theta)z(\theta)} - \frac{\eta_z(\theta)\bar{s}(z(\theta))}{\zeta_z^c(\theta)z(\theta)} \right) \mid z(\theta) = z^* \right] dH(z^*) \\
&= \int_{z^*=0}^{\infty} \mathbb{E} \left[\left(\frac{T'(z^*)}{1-T'(z^*)} + \frac{t-g(z^*)\gamma(\theta)-e}{1-T'(z^*)} \frac{ds(z,\theta)}{dz} \right) \eta_z(\theta) (s(\theta) - \bar{s}(z^*)) \mid z(\theta) = z^* \right] dH(z^*) \\
&= \int_{z^*=0}^{\infty} h(z^*) \mathbb{E} \left[\frac{T'(z^*)}{1-T'(z^*)} + \frac{t-g(z^*)\gamma(\theta)-e}{1-T'(z^*)} \frac{ds(z,\theta)}{dz} \mid z(\theta) = z^* \right] \mathbb{E} [\eta_z(\theta) (s(\theta) - \bar{s}(z^*)) \mid z(\theta) = z^*] dz^* \\
&\quad + \int_{z^*=0}^{\infty} Cov \left[\frac{t-g(z^*)\gamma(\theta)-e}{1-T'(z^*)} \frac{ds(z,\theta)}{dz}, \eta_z(\theta) (s(\theta) - \bar{s}(z^*)) \mid z(\theta) = z^* \right] dH(z^*) \\
&= \int_{z^*=0}^{\infty} \frac{\int_{x=z^*}^{\infty} (1-\hat{g}(x))dH(x)}{\bar{\zeta}_z^c(z^*) \cdot z^*} \Sigma_{\eta_z, s}^{(z^*)} dz^* \\
&\quad + \int_{z^*=0}^{\infty} Cov \left[(t-g(z^*)\gamma(\theta)-e) \frac{\eta_\xi(\theta)}{p+t}, \eta_z(\theta) (s(\theta) - \bar{s}(z^*)) \mid z(\theta) = z^* \right] dH(z^*) \\
&= \int_{z^*=0}^{\infty} \frac{\int_{x=z^*}^{\infty} (1-\hat{g}(x))dH(x)}{\bar{\zeta}_z^c(z^*) \cdot z^*} \Sigma_{\eta_z, s}^{(z^*)} dz^* \\
&\quad + \int_{z^*=0}^{\infty} \left[\frac{t}{p+t} \Sigma_{s\eta_z, \eta_\xi}^{(z^*)} - \frac{1}{p+t} g(z^*) \Sigma_{s\eta_z, \eta_\xi \gamma}^{(z^*)} - \frac{e}{p+t} \Sigma_{s\eta_z, \eta_\xi}^{(z^*)} \right] dH(z^*). \tag{59}
\end{aligned}$$

Finally, integration by parts, using $\Delta\bar{s}(z)$ to denote $\bar{s}(z) - \bar{s}(z_{min})$, yields

$$\begin{aligned}
& \int_{z^*=0}^{\infty} \left\{ \left(\int_{x=z^*}^{\infty} (1-\hat{g}(x))dH(x) \right) (s'_{inc}(z^*) - \bar{s}'(z^*)) \right\} dz^* \\
&= \int_0^{\infty} (1-\hat{g}(z^*))h(z^*) \left(\int_0^z s'_{inc}(x)dx - \Delta\bar{s}(z^*) \right) dz^* = Cov[\hat{g}(z), s_{pref}(z)]. \tag{60}
\end{aligned}$$

We now employ the terms A_1 , A_2 , and A_3 , defined in Proposition 5. Using the simplifications in Equations (56), (57), (58), and (59), and the requirement that $\frac{dW}{dt} = 0$ at the optimum, we have

$$0 = \frac{t}{p+t} (\bar{s}\bar{\zeta}^c + A_1) - \frac{1}{p+t} (\bar{s}\bar{\zeta}^c(\bar{\gamma}(\bar{g} + \sigma) + e) + A_2 + eA_1) + Cov[\hat{g}(z), s_{pref}(z)] + A_3.$$

Rearranging this expression yields Equation (36).

C.A.3 Proof of Proposition 1, Equation 13 (optimal commodity tax)

Equation (36) characterizes the optimal tax under arbitrary multidimensional heterogeneity. We now show how Assumption 4 simplifies this expression to the one presented in Proposition 1. Under condition (a) types are unidimensional, in which case there is no heterogeneity conditional on income, and all of the Σ terms in A_1 , A_2 , and A_3 are equal to zero. Then Equation (36) immediately reduces to the expression in Proposition 1.

Under condition (b), terms of order $\frac{(p+t)s}{z}$ are negligible, and so we find the optimal tax by taking the limit of Equation (36) as $\frac{(p+t)s(\theta)}{z(\theta)} \rightarrow 0$ for all θ . Note that by the definitions $\eta_\xi := \xi \cdot \frac{(p+t)s}{z} \cdot \frac{1}{1-T'}$

and $\eta = \xi \nu \frac{(p+t)s}{z} \cdot \frac{1}{1-T'}$, the income effect terms η_ξ and η go to zero as $\frac{(p+t)s(\theta)}{z(\theta)} \rightarrow 0$, and therefore A_1 and A_2 go to zero. Condition (b) further assumes that η_z is orthogonal to s conditional on earnings, implying A_3 is equal to zero. Therefore under this assumption, the limit of Equation (36) is the simplified expression in Proposition 1.

C.B Proof of Proposition 2

We now characterize the optimal commodity tax when the income tax is suboptimal at some fixed $T(z)$, assuming η and η_z are negligible, and that ζ_z , ξ , and s are uncorrelated conditional on income. We can begin with the derivation for $\frac{dW}{dt}$ on lines (44) and (45) above, before we made use of the optimality of T' :

$$\begin{aligned} \frac{dW}{dt} &= \int_{\Theta} \left(s(\theta)(1-g(\theta)) - s(\theta) \left(\frac{\zeta^c(\theta)}{p+t} + \frac{\eta(\theta)}{p+t} \right) (t-g(\theta)\gamma(\theta) - e) \right) \mu(\theta) \\ &\quad - \int_{\Theta} \zeta_z^c(\theta) z(\theta) \left(\frac{T'(z(\theta))}{1-T'} + \frac{t-g(\theta)\gamma(\theta) - e}{1-T'(z(\theta))} \frac{ds(z, \theta)}{dz} \right) \left(\frac{ds(z; \theta)}{dz} + \frac{\eta_z(\theta)s(\theta)}{\zeta_z^c(\theta)z(\theta)} \right) \mu(\theta) \\ &= \int_{\Theta} \left(s(\theta)(1-g(\theta)) - s(\theta) \left(\frac{\zeta^c(\theta)}{p+t} \right) (t-g(\theta)\gamma(\theta) - e) \right) \mu(\theta) \\ &\quad - \int_{\Theta} \zeta_z^c(\theta) \left(\frac{T'(z(\theta))}{1-T'} + \frac{t-g(\theta)\gamma(\theta) - e}{p+t} \eta_\xi(\theta) \right) \xi(\theta)s(\theta)\mu(\theta) \\ &= \mathbb{E}[s(z)(1-g(z))] - \frac{\bar{s}\bar{\zeta}^c}{p+t} (t - \bar{\gamma}(\bar{g} + \sigma) - e) - \mathbb{E} \left[\bar{\zeta}_z^c(z) \left(\frac{T'(z)}{1-T'} \right) \bar{\xi}(z)\bar{s}(z) \right]. \end{aligned}$$

Where we have invoked the assumption of negligible income effects the assumption in Proposition 2, and the orthogonality of ζ_z , ξ , and s conditional on income. Noting that absent income effects, $\mathbb{E}[g(z)] = 1$, setting this expression equal to zero and yields Proposition 2 in the paper.

C.C Extension to Many Dimensions of Consumption: Proof of Proposition 3

Extending Equation (35) from the proof of Proposition 1, the income tax must satisfy

$$\frac{T'(z)}{1-T'(z)} = \frac{1}{\bar{\zeta}_z^c(z) \cdot z \cdot h(z)} \int_{x=z}^{\infty} (1 - \hat{g}(x)) dH(x) - \sum_i \mathbb{E} \left[\frac{t_i - g(z)\gamma_i(\theta) - e_i}{1-T'(z(\theta))} \cdot \frac{dc_i(z; \theta)}{dz} \mid z(\theta) = z \right]$$

Next, we have that $\frac{dz(\theta)}{dt_i} = -\zeta_z^c(\theta) \left(\frac{z(\theta)}{1-T'(z(\theta))} \right) \frac{dc_i(z; \theta)}{dz}$

The welfare impact of increasing commodity tax t_i is given by

$$\frac{dW}{dt_i} = \overbrace{\int_{\Theta} \left(c_i(\theta)(1 - g(\theta)) + \sum_j \frac{dc_j(\theta)}{dt_i} (t_j - g(\theta)\gamma_j(\theta) - e_j) \right) \mu(\theta)}^{\text{Direct effect on consumption}} \quad (61)$$

$$- \overbrace{\int_{\Theta} \frac{T'(z)}{1 - T'} \left(\zeta_z^c(\theta)z(\theta) \frac{dc_i(z; \theta)}{dz} + \eta_z(\theta)c_i(\theta) \right) \mu(\theta)}^{\text{Direct effect on labor supply}} \quad (62)$$

$$- \overbrace{\sum_j \int_{\Theta} \frac{dz(\theta)}{dt_i} \frac{dc_j(z; \theta)}{dz} (t_j - g(\theta)\gamma_j(\theta) - e_j) \mu(\theta)}^{\text{Indirect effect on consumption through changes in } z} \quad (63)$$

$$= \sum_j \int_{\Theta} \left(c_i(\theta)(1 - g(\theta)) + \frac{dc_j(\theta)}{dt_i} (t_j - g(\theta)\gamma_j(\theta) - e_j) \right) \mu(\theta) \quad (64)$$

$$- \int_{\Theta} \frac{T'(z)}{1 - T'} \left(\zeta_z^c(\theta)z(\theta) \frac{dc_i(z; \theta)}{dz} + \eta_z(\theta)c_i(\theta) \right) \mu(\theta) \quad (65)$$

$$+ \sum_j \int_{\Theta} \frac{1}{1 - T'(z(\theta))} z \zeta_z^c(\theta) \frac{dc_i(z; \theta)}{dz} \frac{dc_j(z; \theta)}{dz} (t_j - g\gamma_j - e_j) \mu(\theta) \quad (66)$$

The welfare impact of an income tax perturbation $dT(z) = dt\bar{c}_i(z)$ is

$$\frac{dW}{dT} = \sum_j \int_{\Theta} \left(\bar{c}_i(z(\theta))(1 - g(\theta)) - \bar{c}_i(z(\theta)) \frac{\eta_j(\theta)}{p_j + t_j} (t_j - g(\theta)\gamma_j(\theta) - e_j) \right) \mu(\theta) \quad (67)$$

$$- \int_{\Theta} \frac{T'(z)}{1 - T'} \left(\zeta_z^c(\theta)z(\theta) \bar{c}_i'(z(\theta)) + \eta_z(\theta)\bar{c}_i(z(\theta)) \right) \mu(\theta) \quad (68)$$

$$+ \sum_j \int_{\Theta} z \zeta_z^c \bar{c}_i'(z(\theta)) \frac{t - g(z)\gamma(\theta) - e}{1 - T'(z(\theta))} \frac{dc_j(z, \theta)}{dz} \mu(\theta). \quad (69)$$

Proceeding exactly as in the proof of Proposition 1 (keeping in mind the unidimensional types assumption), we have

$$\begin{aligned} \frac{dW}{dt_i} &= \sum_j \int_0^\infty \frac{c_j(z)}{dt_i} \Big|_u (t_j - g(z)\bar{\gamma}_j(z) - e) dH(z) \\ &\quad - \int_0^\infty \int_{x \geq z} (1 - \hat{g}(x)) dx (s'_{inc}(z) - \bar{s}'(z)) dH(z). \end{aligned} \quad (70)$$

We can transform Equation (70) using integration by parts as in the proof of Proposition 1, exactly

as we did in Equation (60). This yields, for each i ,

$$\begin{aligned} \sum_j \int_0^\infty t_j \frac{\bar{c}_j}{dt_i} \Big|_u dH(z) &= \sum_j \int_0^\infty (g(z)\gamma_j(z) - e) \frac{c_j(z)}{dt_i} \Big|_u dH(z) + Cov[\hat{g}(z), c_{pref,j}(z)] \\ &\quad - \sum_j \frac{c_j(\theta, t, T)}{dt_i} \Big|_u (\bar{\gamma}_{ij}(\bar{g} + \sigma_{ij}) + e) + Cov[\hat{g}(z), c_{pref,j}(z)] \\ &= -R_i + \rho_i. \end{aligned}$$

C.D Extension to a Composite Sin Good

A composite sin good can be represented in the utility function by assuming that individuals maximize decision utility $U(c, s_1, s_2, \dots, s_n, z; \theta)$ subject to their budget constraint $c + (1 + \tau) \sum_{i=1}^n p_i s_i \leq z - T(z)$, while the policymaker seeks to maximize aggregate normative utility $V(c, s_1, s_2, \dots, s_n, z; \theta)$.

As in the proof of Proposition 1 above, consider the implications of a small joint perturbation to the commodity tax $d\tau$, combined with an offsetting compensation through the income tax which preserves labor supply choices for all individuals. This reform has the following effects:

- *Mechanical revenue effect:* $d\tau \cdot \sum_i p_i s_i(\theta) = d\tau \cdot p \cdot s(\theta)$.
- *Mechanical welfare effect:* the reform mechanically reduces each consumer's net income by $d\tau \cdot \sum_i p_i s_i(\theta)$, for a mechanical welfare loss of $d\tau \cdot \sum_i p_i s_i(\theta) g(\theta)$. Thus the total mechanical welfare effect is $-d\tau \int_\Theta \sum_i p_i s_i(\theta) g(\theta) \mu(\theta) = -d\tau \int_\Theta p s(\theta) g(\theta) \mu(\theta)$.
- *Direct effect on sin good consumption:* each consumer alters their consumption of each sin good i by $d\tau \cdot \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i$, where $q_i := (1 + \tau)p_i$, generating a fiscal externality of $d\tau \int_\Theta \tau p_i \frac{\partial s_i(q_i, \theta)}{\partial q_i} \mu(\theta)$ and a behavioral effect (in terms of public funds) of $\frac{\alpha(\theta)}{\lambda} \sum_i (V'_{s_i}(\theta) - V'_c(\theta) p_i (1 + \tau)) \left(d\tau \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i \right)$, or $d\tau \int_\Theta g(\theta) \sum_i \gamma_i(\theta) \frac{\partial s_i(q_i, \theta)}{\partial q_i} p_i d\mu(\theta) = d\tau \int_\Theta g(\theta) \sum_i \gamma(\theta) \frac{\partial s(q, \theta)}{\partial q} p d\mu(\theta)$. Then the direct effect on sin good consumption can be written $d\tau \int_\Theta g(\theta) \gamma(\theta) \zeta(\theta) \frac{s(\theta)}{p(1 + \tau)} \mu(\theta)$, equivalent to the expression from C.A.

C.E Computing the Optimal Tax in Terms of $g(z)$ rather than $\hat{g}(z)$

Proof of Lemma 1

Proof. We can instead ask the following intuitive question: “When we raise the commodity tax by dt , how much money can we give back to each person to offset the tax so as to keep the average labor supply choices of each z -earner constant?” Call this quantity $\chi(z)$. This term satisfies the following differential equation:

$$\bar{\zeta}_z^c z \chi'(z) + \eta_z \chi(z) = \bar{\zeta}_z^c z s'_{inc}(z) + \eta_z s(z) \quad (71)$$

The right-hand side is the impact of the commodity-tax, as shown in Appendix C.A. The left-hand side follows straightforwardly: the first term is just compensated effect of increasing the marginal

tax rate, while the second term is the income effect. Rearranging yields

$$\chi'(z) + \frac{\eta_z}{\zeta_z^c z} \chi(z) = s'_{inc}(z) + \frac{\eta_z}{\zeta_z^c z} s(z) \quad (72)$$

The solution to this first-order differential equation is

$$\chi(z) = \frac{\int_{x=0}^z e^{\int_{\zeta_z^c x'}^{\eta_z} dx'} \left(s'_{inc}(x) + \frac{\eta_z}{\zeta_z^c x} s(x) \right) dx + K}{e^{\int_{\zeta_z^c z}^{\eta_z} dx'}} \quad (73)$$

where K is some integration constant.

Now

$$\frac{d}{dx} s_{inc}(x) e^{\int_{\zeta_z^c x'}^{\eta_z} dx'} = e^{\int_{\zeta_z^c x'}^{\eta_z} dx'} \left(s'_{inc}(x) + \frac{\eta_z}{\zeta_z^c x} s_{inc}(x) \right) \quad (74)$$

and thus

$$\chi(z) = s_{inc}(z) + \frac{K - s_{inc}(0)}{e^{\int_{\zeta_z^c z}^{\eta_z} dx'}} + \int_0^z w(x, z) \frac{\eta_z}{\zeta_z^c x} (s(x)) dx \quad (75)$$

where $w(x, z) = e^{\int_{x'=x}^{x'=z} -\frac{\eta_z}{\zeta_z^c z} dx'}$. And to get the initial conditions $\chi(0) = s_{inc}(0)$, we must have $K = s_{inc}(0)$, so that

$$\chi(z) = s_{inc}(z) + \underbrace{\int_0^z w(x, z) \frac{\eta_z}{\zeta_z^c x} (\bar{s}(x) - s_{inc}(x)) dx}_{\text{labor supply income effect adjustment}} \quad (76)$$

□

Proof of Proposition 4

Proof. Consider now the welfare impacts of a reform that increase the commodity tax by dt and the income tax $\chi(z)$ at each point z . Because by Lemma 1 this has no effect on labor supply, it has the following impact on welfare⁴:

$$dW = \int_z \left(\underbrace{\left((1 - g(z))(s(z) - \chi(z)) \right)}_{\text{Mechanical revenue and welfare}} - \underbrace{s(z)(t - \gamma) \frac{\zeta^c(z)}{p + t}}_{\text{Corrective benefit I}} - \underbrace{\frac{\eta(z)}{p + t} (t - g(z)\gamma(z))(s(z) - \chi(z))}_{\text{Corrective benefit II}} \right) dH(\bar{z}) \quad (77)$$

Equation (77) follows from the following effects. The first effect is the direct revenue and welfare effect of decreasing each individual's income by $s(z) - \chi(z)$ and transferring that to public funds. The second effect corresponds to the compensated demand response, which generates both a welfare

effect and a fiscal externality from substitution. The third effect comes from the effect that a wealth decrease of $s(z) - \chi(z)$ has on s consumption; again, this generates both a fiscal externality and a welfare effect. Setting $dW = 0$, we have

$$t = \bar{\gamma}(\bar{g} + \sigma) - \frac{1}{\bar{s}\bar{\zeta}^c} \int (s(z) - \chi(z))\eta(z)(t - g(z)\bar{\gamma}(z)) + \frac{p+t}{\bar{s}\bar{\zeta}^c} E[(1 - g(z))(s(z) - \chi(z))]. \quad (78)$$

$$= \bar{\gamma}(\bar{g} + \sigma) + \frac{p+t}{\bar{s}\bar{\zeta}^c} E[(1 - g(z))(s(z) - \chi(z))] \quad (79)$$

$$+ \frac{1}{\bar{s}\bar{\zeta}^c} E[(s(z) - \chi(z))\eta(z)g(z)\bar{\gamma}(z)] - \frac{t}{\bar{s}\bar{\zeta}^c} E[(s(z) - \chi(z))\eta(z)] \quad (80)$$

Resolving for t yields

$$t = \frac{\bar{s}\bar{\zeta}^c(\bar{g} + \sigma) + pE[(1 - g(z))(s(z) - \chi(z))] + E[(s(z) - \chi(z))\eta(z)g(z)\bar{\gamma}(z)]}{\bar{s}\bar{\zeta}^c - E[(1 - g(z))(s(z) - \chi(z))] + E[(s(z) - \chi(z))\eta(z)]}.$$

□

D PanelViews Survey Questions

This appendix presents the text of the PanelViews survey questions used for this project.

D.A Self-Reported Beverage Intake

For each of the following types of drinks, please tell us how many 12-ounce servings you drink in an average week. (A normal can of soda is 12 ounces.)

- 100% fruit juice
- Sweetened juice drinks (for example, fruit ades, lemonade, punch, and orange drinks)
- Regular soft drinks (soda pop)
- Diet soft drinks and all other artificially sweetened drinks
- Pre-packaged (i.e. canned or bottled) tea or coffee (for example, iced tea, iced coffee, and flavored tea)
- Sports drinks
- Caffeinated energy drinks

D.B Self-Control

Please indicate how much each of the following statements reflects how you typically are: I drink soda pop or other sugar-sweetened beverages more often than I should.

- Not at all
- Somewhat
- Mostly
- Definitely

Please indicate how much each of the following statements describes the other head of household: The other head of household in my house drinks soda pop or other sugar-sweetened beverages more often than they should.

- Not at all
- Somewhat
- Mostly
- Definitely
- I am the only head of household

D.C Preferences: Taste for Beverages and Health Importance

Imagine for a moment that you could drink whatever beverages you want without any health or nutritional considerations. Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking the following? Please indicate your liking by selecting a number on a scale of 0-10, with 0 being “not at all” and 10 being “very much.”

- 100% fruit juice
- Sweetened juice drinks (for example, fruit ades, lemonade, punch, and orange drinks)
- Regular soft drinks (soda pop)
- Diet soft drinks and all other artificially sweetened drinks
- Pre-made tea or coffee
- Sports drinks
- Caffeinated energy drinks

In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.? Please indicate the importance by selecting a number on a scale of 0-10, with 0 being “not at all important” and 10 being “extremely important.”

D.D Nutrition Knowledge

Now we'd like to ask about nutrition knowledge. This is a survey, not a test. If you don't know the answer, mark "not sure" rather than guess. Your answers will help identify which dietary advice people find confusing.

Do health experts recommend that people should be eating more, the same amount, or less of the following foods? Please select one response for each. *Possible responses: More, Same, Less, Not sure.*

- Fruit
- Food and drinks with added sugar
- Vegetables
- Fatty foods
- Processed red meat
- Whole grains
- Salty foods
- Water

Which of these types of fats do experts recommend that people should eat less of? Please select one response for each. *Possible responses: Eat less, Not eat less, Not sure.*

- Unsaturated fats
- Trans fats
- Saturated fats

Do you think these foods and drinks are typically high or low in added sugar? Please select one response for each. *Possible responses: High in added sugar, Low in added sugar, Not sure.*

- Diet cola drinks
- Plain yogurt
- Ice cream
- Tomato ketchup
- Melon

Do you think these foods are typically high or low in salt? Please select one response for each. *Possible responses: High in salt, Low in salt, Not sure.*

- Breakfast cereals
- Frozen vegetables
- Bread
- Baked beans
- Red meat
- Canned soup

Do you think these foods are typically high or low in fiber? Please select one response for each. *Possible responses: High in fiber, Low in fiber, Not sure.*

- Oats
- Bananas
- White rice
- Eggs
- Potatoes with skin
- Pasta

Do you think these foods are a good source of protein? Please select one response for each. *Possible responses: Good source of protein, Not a good source of protein, Not sure.*

- Poultry
- Cheese
- Fruit
- Baked beans
- Butter
- Nuts

Which of the following foods do experts count as starchy foods? Please select one response for each. *Possible responses: Starchy food, Not a starchy food, Not sure.*

- Cheese
- Pasta
- Potatoes

- Nuts
- Plantains

Which is the main type of fat present in each of these foods? Please select one response for each.
Possible responses: Polyunsaturated fat, Monounsaturated fat, Saturated fat, Cholesterol, Not sure.

- Olive oil
- Butter
- Sunflower oil
- Eggs

Which of these foods has the most trans-fat? Please select one.

- Biscuits, cakes and pastries
- Fish
- Rapeseed oil
- Eggs
- Not sure

The amount of calcium in a glass of whole milk compared to a glass of skimmed milk is...? Please select one.

- Much higher
- About the same
- Much lower
- Not sure

Which one of the following nutrients has the most calories for the same weight of food? Please select one.

- Sugar
- Starchy
- Fiber/roughage
- Fat
- Not sure

If a person wanted to buy a yogurt at the supermarket, which would have the least sugar/sweetener? Please select one.

- 0% fat cherry yogurt
- Plain yogurt
- Creamy fruit yogurt
- Not sure

If a person wanted soup in a restaurant or cafe, which one would be the lowest fat option? Please select one.

- Mushroom risotto soup (field mushrooms, porcini mushrooms, arborio rice, butter, cream, parsley and cracked black pepper)
- Carrot butternut and spice soup (carrot , butternut squash, sweet potato, cumin, red chilies, coriander seeds and lemon)
- Cream of chicken soup (British chicken, onions, carrots, celery, potatoes, garlic, sage, wheat flour, double cream)
- Not sure

Which of these combinations of vegetables in a salad would give the greatest variety of vitamins and antioxidants? Please select one.

- Lettuce, green peppers and cabbage
- Broccoli, carrot and tomatoes
- Red peppers, tomatoes and lettuce
- Not sure

One healthy way to add flavor to food without adding extra fat or salt is to add...? Please select one.

- Coconut milk
- Herbs
- Soy sauce
- Not sure

Which of these diseases is related to how much sugar people eat? Please select one.

- High blood pressure

- Tooth decay
- Anemia
- Not sure

Which of these diseases is related to how much salt (or sodium) people eat? Please select one.

- Hypothyroidism
- Diabetes
- High blood pressure
- Not sure

Which of these options do experts recommend to prevent heart disease? Please select one.

- Taking nutritional supplements
- Eating less oily fish
- Eating less trans-fats
- Not sure

Which of these options do experts recommend to prevent diabetes? Please select one.

- Eating less refined foods
- Drinking more fruit juice
- Eating more processed meat
- Not sure

Which one of these foods is more likely to raise people's blood cholesterol? Please select one.

- Eggs
- Vegetable oils
- Animal fat
- Not sure

Which one of these foods is classified as having a high Glycemic Index (Glycemic Index is a measure of the impact of a food on blood sugar levels, thus a high Glycemic Index means a greater rise in blood sugar after eating)? Please select one.

- Wholegrain cereals

- White bread
- Fruit and vegetables
- Not sure

Fiber can decrease the chances of gaining weight. Please select one.

- Agree
- Disagree
- Not sure

If someone has a Body Mass Index (BMI) of 23, what would their weight status be? (BMI is conventionally measured in kg/m²) Please select one.

- Underweight
- Normal weight
- Overweight
- Obese
- Not sure

If someone has a Body Mass Index (BMI) of 31, what would their weight status be? (BMI is conventionally measured in kg/m²) Please select one.

- Underweight
- Normal weight
- Overweight
- Obese
- Not sure

D.E Other Questions

Are you the primary shopper? By this we mean, the one household member who makes the majority of your household's grocery purchase decisions.

Are you. . .

- Male
- Female

What is your primary occupation?

E Stockpiling and Optimal Lag Length for Price Elasticity Estimates

Because SSBs are storable, past prices and merchandising conditions that affect past purchases could in theory affect current stockpiles and thus current purchases. To address this, we estimate a version of Equation (19) that includes lags of the price and merchandising condition variables. Letting l index quarterly lags, we estimate the following regression:

$$\ln s_{it} = \sum_{l=0}^L \zeta_l \ln p_{i,t-l} + \sum_{l=0}^L \nu_l \mathbf{f}_{i,t-l} + \xi \ln z_{ct} + \omega_t + \mu_{ic} + \varepsilon_{it}, \quad (81)$$

with standard errors clustered by county.

In Table A1, we find that the local price IV Z_{it} is a powerful predictor of price paid in quarter t but is not strongly conditionally correlated with prices paid in the quarters before and after. This implies that we have strong first stages and do not have large serial autocorrelation problems.

A standard approach to determining the optimal number of lags in a distributed lag model is to choose the specification that minimizes the Aikake Information Criterion (AIC) or Bayesian Information Criteria (BIC). In essence, these criteria find the specification that best predicts purchases in period t . As shown in Table A2, the specification that minimizes AIC and BIC is to set $L = 0$, i.e. to include only contemporaneous quarter prices. Furthermore, the coefficients on lagged prices in the table are all statistically zero, implying no statistically detectable stockpiling of SSBs from quarter to quarter.⁵² For these reasons, we set $L = 0$ when estimating demand in the body of the paper, thereby including only contemporaneous prices and merchandising conditions.

⁵²This result is consistent with DellaVigna and Gentzkow (2017), who find that stores offering lower prices in a given week see little decrease in sales in future weeks and months, even for relatively storable goods. Prior work highlighting household stockpiling (e.g. Hendel and Nevo (2006a; 2006b; 2010) and Wang (2015)) typically uses weekly purchase data, and stockpiling is naturally more relevant across weeks than across quarters.

Table A1: Regressions of Price Paid on the Local Price IV

	(1)	(2)	(3)	(4)	(5)
	Price (t)	Price (t-1)	Price (t-2)	Price (t-3)	Price (t-4)
Local price IV (t)	1.138*** (0.075)	-0.087 (0.066)	-0.261*** (0.062)	0.068 (0.070)	0.258*** (0.089)
Local price IV (t-1)	-0.180*** (0.057)	1.282*** (0.080)	0.047 (0.068)	-0.246*** (0.068)	-0.051 (0.065)
Local price IV (t-2)	-0.179*** (0.069)	-0.155*** (0.057)	1.267*** (0.091)	0.035 (0.077)	-0.270*** (0.071)
Local price IV (t-3)	-0.001 (0.057)	-0.183*** (0.063)	-0.154** (0.060)	1.201*** (0.092)	0.047 (0.077)
Local price IV (t-4)	0.166** (0.072)	0.074 (0.055)	-0.088 (0.067)	-0.132** (0.065)	1.065*** (0.095)
N	2,205,632	2,065,554	1,911,355	1,721,779	1,620,371

Notes: This table presents regressions of price paid $\ln p_{i,t-l}$ on the local price instrumental variable Z_{it} and four additional quarterly lags. All regressions include the additional control variables in Equation (81): feature and display (and four lags thereof), natural log of county average per capita income, quarter of sample indicators, and a household-by-county fixed effect. Columns 1-5, respectively, use the contemporaneous price paid and then the first through fourth lags of prices paid as the dependent variable. Sample sizes vary across columns because using a longer lag of price paid $p_{i,t-l}$ reduces the number of observations. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A2: Determining Optimal Number of Lags

	(1)	(2)	(3)	(4)	(5)
$\ln(\text{Average price/liter})$	-1.453*** (0.146)	-1.413*** (0.121)	-1.406*** (0.122)	-1.373*** (0.115)	-1.334*** (0.129)
$\ln(\text{Average price/liter})$ (t-1)	0.047 (0.118)	0.080 (0.113)	0.086 (0.106)	0.067 (0.115)	
$\ln(\text{Average price/liter})$ (t-2)	-0.097 (0.108)	-0.119 (0.104)	-0.100 (0.112)		
$\ln(\text{Average price/liter})$ (t-3)	-0.069 (0.117)	0.006 (0.122)			
$\ln(\text{Average price/liter})$ (t-4)	0.170 (0.161)				
$\ln(\text{County income})$	0.145* (0.087)	0.147* (0.086)	0.144* (0.085)	0.133 (0.085)	0.133 (0.085)
N	1,226,185	1,226,185	1,226,185	1,226,185	1,226,185
Akaike Information Criterion	2,877,873	2,852,217	2,846,694	2,824,888	2,811,829
Bayesian Information Criterion	2,878,534	2,852,842	2,847,283	2,825,441	2,812,346

Notes: This table presents estimates of Equation (81), with $L = 4, 3, 2, 1, 0$ in columns 1-5, respectively. All regressions include feature and display (and L lags thereof), quarter of sample indicators, and household-by-county fixed effects. All regressions use a common sample: the set of observations in column 1. Observations are weighted for national representativeness. Robust standard errors in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

F Income Elasticity Estimated with Household-Level Variation

When estimating the income elasticity in Equation (19), our specifications in the body of the paper use county mean income. In this appendix, we present alternative specifications using household

income reported by Homescan households on annual surveys. One benefit of the household-level data is that there is substantially more variation to exploit, so the estimates can be more precise.

There are three problems, however. First, one might naturally expect more measurement error in the household-level survey data, especially because we are exploiting within-household variation. Second, using household-level income can generate additional endogeneity problems. In particular, changes in household composition (for example, a divorce or a adult moving in) can affect household income and the quantity of SSBs (and other groceries) purchased. Third, there is some uncertainty over the year for which Homescan panelists are reporting their income. The surveys reported for year t are taken in the fall of year $t - 1$. In the early years of the sample, panelists were asked to report total annual income as of year-end of the *previous* calendar year, i.e. year $t - 2$. Nielsen believes that panelists are actually reporting their annualized income at the time of the survey, and as of 2011, the instructions for income no longer mention the previous calendar year, but rather annualized income.

Table A3 explores this third issue by regressing the natural log of household income in year t on the natural log of current and lagged county mean income. The second lag of county income is most predictive, suggesting that the modal Homescan household is reporting their income for that year. However, other lags also have predictive power, suggesting that there will be some measurement error. The low t-statistics on county income also highlight that county income variation is not very predictive of reported household-level income variation.

Table A4 presents estimates of Equation (19). Column 1 reproduces the primary estimate from column 2 of Table 2. Columns 2–5 substitute household income and the first and second lead of household income in place of county income. Across the various columns, the estimated income elasticities $\hat{\xi}$ range from 0 to about 0.04. This is substantially smaller than the estimate of $\hat{\xi} \approx 0.18$ in the primary estimates reproduced in column 1. One natural explanation is that the difference is attenuation bias driven by measurement error in household income. Even though the magnitudes differ, we emphasize that the income elasticity estimates $\hat{\xi}$ are consistent in their key qualitative implication: they all imply that the steeply declining consumption-income profile is not driven by year-to-year causal income effects.

To address the endogeneity concern introduced above, columns 2-5 all include the full set of demographic controls used elsewhere in the paper: natural logs of education and age, race, an indicator for the presence of children, household size in adult equivalents, employment status, and weekly work hours. Column 6 demonstrates the endogeneity concern by omitting these variables. The estimated income elasticity $\hat{\xi}$ becomes slightly negative.

Table A3: Regressions of Household Income on County Income

	(1)	(2)	(3)	(4)	(5)
ln(County income)	0.056 (0.043)	0.124*** (0.036)			
ln(County income) (t-1)	0.037 (0.038)		0.179*** (0.037)		
ln(County income) (t-2)	0.104*** (0.036)			0.215*** (0.041)	
ln(County income) (t-3)	0.108*** (0.041)				0.190*** (0.044)
N	653,554	653,554	653,554	653,554	653,554

Notes: This table presents regressions of the natural log of household income on county income and its lags, using household-by-year Homescan data for 2006-2016. All regressions include year indicators and household-by-county fixed effects. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A4: Estimating Income Elasticity Using Household Income Instead of County Income

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Average price/liter)	-1.401*** (0.096)	-1.253*** (0.157)	-1.367*** (0.094)	-1.305*** (0.113)	-1.250*** (0.156)	-1.389*** (0.096)
ln(County income)	0.183** (0.071)					
ln(Household income)		0.004 (0.008)	0.013* (0.007)			-0.059*** (0.007)
ln(Household income) (year+1)		0.040*** (0.008)		0.040*** (0.007)		
ln(Household income) (year+2)		0.003 (0.007)			0.013 (0.008)	
Feature	1.145*** (0.069)	1.126*** (0.097)	1.185*** (0.069)	1.202*** (0.083)	1.128*** (0.097)	1.148*** (0.069)
Display	0.492*** (0.087)	0.429*** (0.129)	0.513*** (0.083)	0.472*** (0.108)	0.430*** (0.129)	0.493*** (0.086)
Other demographics	No	Yes	Yes	Yes	Yes	No
Kleibergen-Paap first stage F stat	191	85	193	124	85	193
N	2,205,632	1,232,316	2,205,632	1,620,361	1,232,316	2,205,632

Notes: This table presents estimates of Equation (19). Column 1 presents base IV estimates, while columns 2-6 present estimates using natural log of household income instead of county mean income. All regressions include quarter of sample indicators and household-by-county fixed effects. “Other demographics” are natural logs of education and age, race, an indicator for the presence of children, household size in adult equivalents, employment status, and weekly work hours. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.


G Example of Uniform Pricing

The power of our instrument derives from two facts documented by DellaVigna and Gentzkow (2017) and Hitsch, Hortacsu, and Lin (2017). First, chains vary their prices independently of each other over time. Second, chains set prices in a nearly uniform fashion across all their stores.


Figure A1 illustrates these patterns for a high-volume SSB UPC. Panel (a) presents the quarterly average price (unweighted across stores and weeks) for each retail chain that sold this UPC over the full sample, with darker shading indicating higher prices. Other than a nationwide price increase in early 2008, there is no clear pattern, illustrating that chains vary their prices independently of each other over time. Panel (b), by contrast, shows a clear pattern. This figure looks within one example retail chain, presenting the quarterly average price (again unweighted across weeks) for each store at that chain. The vertical patterns illustrate that this chain varies prices in a coordinated way across all of its stores. In mid-2008, for example, prices at all stores were relatively high, whereas in early 2007, prices at all stores were relatively low. The figure echoes a similar figure in DellaVigna and Gentzkow (2017), who document that this within-chain, across-store price coordination is not limited to the example chain we chose for this figure.

Figure A1: Intuition for the Price Instruments

(a) Relative Prices Vary Across Retail Chains



(b) Uniform Pricing Within an Example Retail Chain



Notes: This figure presents prices of an example high-volume sugar-sweetened beverage UPC. Panel (a) presents the quarterly average price (unweighted across stores and weeks) for each retail chain that sold this UPC over the full sample. Panel (b) presents the quarterly average price (again unweighted across weeks) for each store at an example retail chain.

H Bias Estimation Using PanelViews SSB Consumption Data

As introduced in Section III, the Homescan grocery purchase data are imperfect measures of SSB consumption, both because they do not measure away-from-home SSB consumption that panelists do not scan into the data and because the Homescan data are at the household level, while the bias proxies from PanelViews are at the individual level. For a comprehensive measure of individual-level SSB consumption, we therefore delivered a beverage intake frequency questionnaire as part of the PanelViews survey. In this appendix, we describe these SSB consumption data and use them to estimate average marginal bias.

We used a modified version of the BEVQ-15, a validated questionnaire that is standard in the public health literature (Hedrick et al., 2012). We asked, “For each of the following types of drinks, please tell us how many 12-ounce servings you drink in an average week,” for five types of SSBs (sweetened juice drinks, regular soft drinks, pre-packaged tea or coffee, sports drinks, and caffeinated energy drinks) and two non-SSBs (100% fruit juice and diet soft drinks).

How does Homescan consumption compare to the self-reports? Because average SSB consumption is declining over time in the U.S., we make this comparison using only the most recent year of Homescan data (2016) for closest comparability with the 2017 PanelViews survey. In 2016, Homescan purchases average 58 liters per adult equivalent, while annualized PanelViews self-reported intake averages 88 liters, or 58 percent larger. As a benchmark, Kit et al. (2013, Table 2) report that away from home SSB intake is 86 percent larger than intake at home for adults over our sample period, although away from home vs. at home intake is not the only difference between our two consumption measures. Up to some sampling error in the smaller PanelViews sample, Figure A2 shows that the two data sources are both consistent in showing a similar slope of consumption with respect to income.


At the individual household level, however, the two sources are not always consistent. The relationship between natural log of SSB purchases per adult equivalent in the most recent year of Homescan purchases and natural log of average self-reported SSB consumption per adult in the PanelViews survey has only $R^2 \approx 0.16$.⁵³ This relatively low value underscores the importance of having both consumption measures. Because each measure has different strengths (the PanelViews surveys are at the individual level, measure intake, and include consumption away from home, but they suffer from the noise and possible biases of self-reported data), readers may disagree over which measure should be prioritized, so we present optimal tax calculations in the introduction and in Section IV based on both.

Table A5 presents estimates of Equation (23) using PanelViews respondent-level data. This parallels Table 3, except that i indexes PanelViews respondents instead of Homescan Households, s_i , \tilde{b}_{ki} , \tilde{b}_{si} , and \mathbf{a}_i are the SSB consumption, nutrition knowledge, self-control, and preference measures

⁵³Only some of this low R^2 is from not having self-reports from the full household: when limiting to one-person households, the relationship has $R^2 \approx 0.21$. Some of the low R^2 is from variation within households over time. The PanelViews self-reports are from October 2017, whereas the most recent year of Homescan data is 2016, and a regression of household-level Homescan SSB purchases on its one year lag has only $R^2 \approx 0.62$.

for respondent i , \tilde{b}_{2si} is the self-control of respondent i as rated by the other household head, and \mathbf{x}_i are the household-level characteristics for respondent i 's household. We cluster standard errors at the household level. Figures A3 and A4 present the predicted quantity effect of bias and average marginal bias by income, re-creating Figures 7 and 8 in the body of the paper but adding the PanelViews estimates. The patterns of results are very similar to those from the Homescan data, except that the $\hat{\tau}$ estimates, and the resulting bias estimates, are materially larger.

Figure A2: **Homescan Purchases vs. Self-Reported Intake**



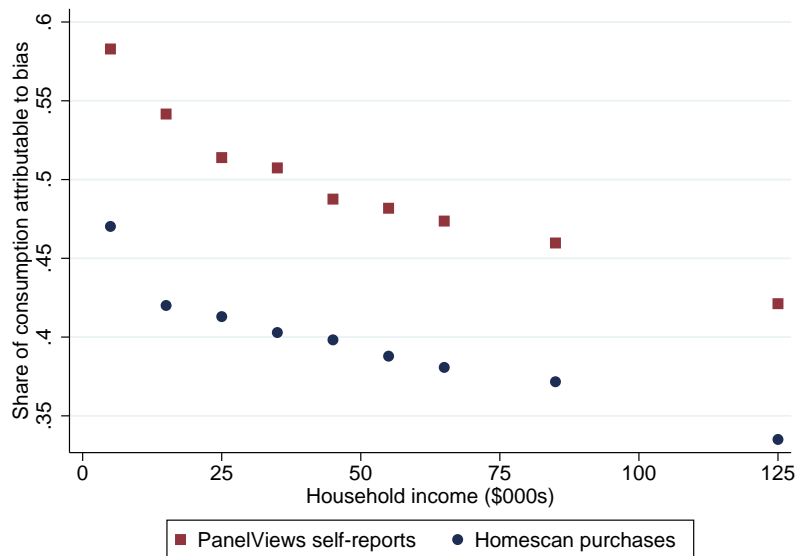
Notes: This figure presents the average purchases of sugar-sweetened beverages for each household's most recent year in the Nielsen Homescan data and the average annualized self-reported SSB consumption from the PanelViews survey, by income group. Homescan purchases are measured in liters per "adult equivalent," where household members other than the household heads are rescaled into adult equivalents using the recommended average daily consumption for their age and gender group. Observations are weighted for national representativeness.

Table A5: **Regressions of Sugar-Sweetened Beverage Consumption on Bias Proxies Using PanelViews Self-Reported Consumption**

	(1)	(2)	(3)	(4)	(5)
Nutrition knowledge	-0.428*** (0.123)	-0.637*** (0.133)	-0.487*** (0.123)	-0.551*** (0.130)	-0.948*** (0.122)
Self-control	-3.210*** (0.097)	-3.919*** (0.086)	-3.203*** (0.096)	-3.286*** (0.100)	-1.973*** (0.058)
Taste for soda	0.736*** (0.065)		0.735*** (0.065)	0.707*** (0.067)	1.121*** (0.059)
Health importance	-0.040 (0.105)		0.043 (0.104)	-0.081 (0.110)	-0.405*** (0.104)
ln(Household income)	-0.109*** (0.027)		-0.068*** (0.023)	-0.104*** (0.027)	-0.096*** (0.028)
ln(Years education)	-0.268* (0.147)		-0.416*** (0.140)	-0.409*** (0.146)	-0.391*** (0.150)
Other beverage tastes	Yes	No	Yes	Yes	Yes
Other demographics	Yes	Yes	No	Yes	Yes
County indicators	Yes	Yes	Yes	No	Yes
Self-control 2SLS	Yes	Yes	Yes	Yes	No
R^2	0.414	0.307	0.411	0.324	0.414
N	20,640	20,640	20,640	20,640	20,640

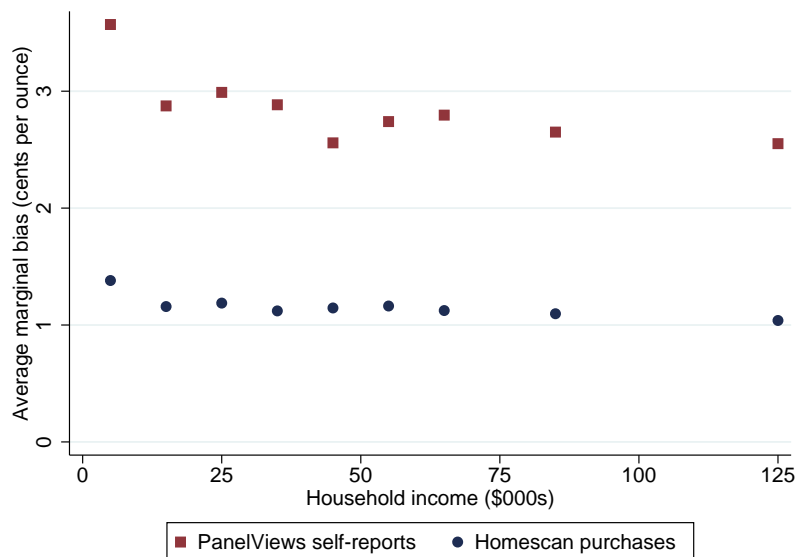
Notes: This table presents estimates of Equation (23). Data are at the PanelViews respondent level, and the dependent variable is the natural log of PanelViews self-reported SSB consumption. \tilde{b}_{ki} , \tilde{b}_{si} , and \mathbf{x}_i are the knowledge, self-control, and preference measures for respondent i , \tilde{b}_{2si} is the self-control of respondent i as rated by the other household head, and \mathbf{x}_i are the household-level characteristics for respondent i 's household. Columns 1-4 correct for measurement error in self-control using two-sample 2SLS, with standard errors calculated per Chodorow-Reich and Wieland (2016). Taste for soda is the response to the question, "Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking [Regular soft drinks (soda pop)]?" "Other beverage tastes" are the responses to parallel questions for other beverages. Health importance is the response to the question, "In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?" Responses to each question were originally on a scale from 0 to 10, which we rescale to between 0 and 1. "Other demographics" are natural log of age, race, an indicator for the presence of children, household size in adult equivalents, employment status, and weekly work hours. Observations are weighted for national representativeness. Robust standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Figure A3: Share of Consumption Attributable to Bias by Income Using PanelViews and Homescan



Notes: This figure presents the share of sugar-sweetened beverage consumption attributable to bias, i.e. the unweighted average of $\frac{s_i - s_i^V}{s_i}$, by income, using both PanelViews and Homescan data. Observations are weighted for national representativeness.

Figure A4: Average Marginal Bias by Income Using PanelViews and Homescan



Notes: This figure presents the demand slope-weighted average marginal bias by income, using both PanelViews and Homescan data. Observations are weighted for national representativeness.

I Additional Empirical Results

I.A Price and Income Elasticities

Table A6: **Estimates of Price and Income Elasticities to Address Censoring at Zero Consumption**

	(1) Primary IV	(2) Reduced form in IV sample	(3) Full sample	(4) Tobit
ln(Average price/liter)	-22.96*** (2.17)			
Local price IV		-26.11*** (2.94)	-21.16*** (2.33)	-26.08*** (2.65)
ln(County income)	2.52 (1.56)	0.10 (1.39)	-0.00 (1.24)	-0.47 (1.39)
Feature	14.46*** (2.19)	20.29*** (1.96)	18.46*** (1.54)	24.93*** (1.75)
Display	7.24*** (2.23)	7.92*** (2.32)	7.25*** (1.94)	9.40*** (2.20)
Price elasticity	1.47	1.68	1.36	1.67
SE(Price elasticity)	0.139	0.189	0.150	0.170
N	2,205,632	2,205,632	2,614,216	2,614,216

Notes: This table presents estimates of Equation (19). All regressions include quarter of sample indicators and a household-by-county fixed effect. Column 1 presents the primary instrumental variables estimates, except with quantity purchased in levels instead of natural logs. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A7: **Alternative Estimates of Price and Income Elasticities**

	(1)	(2)	(3)	(4)	(5)	(6)
	Primary	Employment controls	Drop feature	Drop display	Income control	OLS
ln(Average price/liter)	-1.400*** (0.095)	-1.399*** (0.095)	-1.686*** (0.098)	-1.418*** (0.096)	-1.399*** (0.095)	-0.747*** (0.003)
ln(County income)	0.182** (0.071)	0.184*** (0.071)	0.199*** (0.076)	0.184*** (0.071)	0.183** (0.072)	0.108 (0.066)
Feature	1.146*** (0.069)	1.147*** (0.069)		1.173*** (0.069)	1.151*** (0.069)	1.379*** (0.060)
Display	0.493*** (0.087)	0.493*** (0.087)	0.656*** (0.092)		0.488*** (0.087)	0.536*** (0.082)
1(Employed)		-0.005 (0.032)				
Weekly work hours		-0.001 (0.001)				
Year \times ln(County mean income)					-0.014*** (0.005)	
Kleibergen-Paap first stage F stat	192.0	191.7	220.2	192.1	191.7	
N	2,205,632	2,205,632	2,205,632	2,205,632	2,205,632	2,205,632

Notes: Column 1 reproduces the primary instrumental variables estimates of Equation (19) from column 2 of Table 2, while the other columns present alternative estimates. All regressions include quarter of sample indicators and household-by-county fixed effects. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A8: **Heterogeneous Price Elasticity by Income**

	(1)	(2)
	Income < 50,000	Income \geq 50,000
ln(Average price/liter)	-1.480*** (0.146)	-1.349*** (0.128)
ln(County income)	0.150 (0.102)	0.174* (0.090)
Feature	1.074*** (0.116)	1.202*** (0.083)
Display	0.520*** (0.129)	0.469*** (0.116)
Kleibergen-Paap first stage F stat	93.4	170.9
N	898,449	1,307,183

Notes: This table presents instrumental variables estimates of Equation (19). Column 1 includes only households with sample average income below \$50,000, while column 2 includes only households with sample average income above \$50,000. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A9: Estimates of Substitution to Other Product Groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Coffee and tea	Fruit juice	Diet drinks	Milk	Water	Baked goods	Candy	Desserts	Sweeteners
ln(Average price/liter) (SSBs)	-0.148 (0.091)	0.131 (0.090)	0.245* (0.139)	-0.037 (0.071)	0.037 (0.148)	-0.158 (0.114)	-0.227** (0.108)	-0.056 (0.098)	0.040 (0.130)
ln(Average price/kg)	-0.856*** (0.105)	-1.245*** (0.164)	-1.012*** (0.127)	-0.772*** (0.254)	-1.615*** (0.246)	-2.341*** (0.251)	-2.190*** (0.183)	-1.262*** (0.191)	-1.078*** (0.130)
ln(County income)	-0.028 (0.059)	0.125* (0.071)	0.111 (0.083)	0.036 (0.069)	0.078 (0.088)	0.201*** (0.071)	0.144** (0.069)	0.025 (0.064)	-0.130** (0.056)
Feature (SSBs)	-0.070 (0.070)	-0.001 (0.078)	-0.316*** (0.097)	0.180*** (0.049)	-0.018 (0.107)	0.037 (0.074)	-0.029 (0.084)	0.123* (0.073)	0.063 (0.071)
Display (SSBs)	-0.039 (0.084)	0.011 (0.111)	-0.147 (0.144)	0.152** (0.071)	-0.344** (0.152)	0.050 (0.122)	0.001 (0.117)	0.190** (0.094)	-0.001 (0.080)
Feature	1.188*** (0.090)	1.775*** (0.097)	0.809*** (0.065)	0.971*** (0.108)	1.086*** (0.126)	2.059*** (0.125)	3.094*** (0.162)	1.534*** (0.130)	0.761*** (0.199)
Display	0.673*** (0.146)	0.707*** (0.209)	0.323*** (0.080)	1.515 (0.959)	0.479*** (0.121)	1.173*** (0.183)	1.618*** (0.210)	0.929*** (0.235)	0.155 (0.142)
Kleibergen-Paap first stage F stat	43.5	145.5	69.0	50.3	41.1	68.7	114.1	49.9	37.1
N	1,489,464	1,680,153	1,106,323	2,053,424	1,131,988	1,919,468	1,915,485	1,911,514	1,566,451


Notes: This table presents instrumental variables estimates of the following equation:

$$\ln s_{itg} = \zeta_{SSB} \ln p_{itg} + \zeta_g \ln p_{itg} + \xi \ln z_{ct} + \nu_{SSB} \mathbf{f}_{it} + \nu_g \mathbf{f}_{itg} + \omega_t + \mu_{ic} + \varepsilon_{it}, \quad (82)$$

where s_{itg} , p_{itg} , and \mathbf{f}_{itg} are the consumption (in kilograms), price (in price per kilogram), and merchandising conditions of product group g . The product groups are as described in the column headers; columns 1-5 are non-SSB beverages, and columns 6-9 are foods with high sugar content. Observations are weighted for national representativeness. Robust standard errors, clustered by county, are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

I.B Measuring Bias

Figure A5: **Linearity of Relationship Between Consumption and Bias Proxies**



Notes: This figure presents binned scatterplots of the relationship between natural log of SSB consumption and bias proxies, residual of the other variables in Equation (23). In the top two panels, the dependent variable is the natural log of purchases per adult equivalent in the most recent year that the household was in Homescan. In the bottom two panels, the dependent variable is natural log of PanelViews self-reported consumption.

Table A10: **Example Regressions Used to Impute Nutrition Knowledge and Preferences**

	(1) Nutrition knowledge	(2) Taste for soda	(3) Health importance
Nutrition knowledge	0.826*** (0.019)	0.052 (0.041)	-0.023 (0.022)
Self-control	0.029*** (0.007)	-0.186*** (0.023)	0.042*** (0.011)
Taste for soda	0.006 (0.006)	0.630*** (0.017)	0.024*** (0.008)
Health importance	0.022 (0.015)	0.001 (0.036)	0.728*** (0.024)
ln(Household income)	0.000 (0.004)	-0.015 (0.010)	0.000 (0.005)
ln(Years education)	0.057*** (0.016)	-0.002 (0.046)	0.006 (0.021)
Other beverage tastes	Yes	Yes	Yes
Other demographics	Yes	Yes	Yes
County indicators	Yes	Yes	Yes
R^2	0.883	0.838	0.814
Adjusted R^2	0.811	0.732	0.701
N	2,809	2,576	2,809

Notes: In our primary estimates of Equation 23, nutrition knowledge and preference variables are missing for two-head households where only one head responded to the PanelViews survey. This table presents regressions used to impute nutrition knowledge, taste for soda, and health importance, using the sample of two-head households where both heads responded. The dependent variable in each column is the household average of the variable reported in the column header. The bias proxies and preference controls used as independent variables use only the primary shopper's survey responses. Observations are weighted for national representativeness. Robust standard errors are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A11: **First Stage of Two-Sample 2SLS Self-Control Measurement Error Correction**

	(1)
Nutrition knowledge	0.136*** (0.045)
Self-control	0.586*** (0.026)
Taste for soda	-0.120*** (0.023)
Health importance	0.054 (0.040)
ln(Household income)	0.015 (0.010)
ln(Years education)	0.018 (0.047)
Other beverage tastes	Yes
Other demographics	Yes
County indicators	Yes
R^2	0.757
N	2,812

Notes: This table presents the first stage of the two-sample 2SLS procedure used to correct for measurement error in self-control, using the sample of households where both heads responded to the PanelViews survey. Observations are weighted for national representativeness. Robust standard errors are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A12: **Regressions of Sugar-Sweetened Beverage Consumption on Bias Proxies and Their Interaction**

	(1)	(2)
	Homescan	PanelViews
Nutrition knowledge	0.180 (0.273)	1.613*** (0.318)
Self-control	-0.595** (0.257)	-1.242*** (0.293)
Nutrition knowledge \times self-control	-1.265*** (0.364)	-2.912*** (0.424)
Taste for soda	0.430*** (0.047)	0.722*** (0.065)
Health importance	-0.184** (0.074)	-0.031 (0.105)
ln(Household income)	-0.024 (0.018)	-0.106*** (0.027)
ln(Years education)	-0.692*** (0.096)	-0.238 (0.147)
Other beverage tastes	Yes	Yes
Other demographics	Yes	Yes
County indicators	Yes	Yes
Self-control 2SLS	Yes	Yes
R^2	0.288	0.415
N	18,568	20,640

Notes: This table presents estimates of Equation (23). The estimates replicate column 1 of Tables 3 and A5, adding the interaction between nutrition knowledge and self-control. Observations are weighted for national representativeness. Both columns correct for measurement error in self-control using two-sample 2SLS, with standard errors calculated per Chodorow-Reich and Wieland (2016). Robust standard errors are in parentheses, clustered by household in column 2. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

Table A13: **Regressions of Sugar-Sweetened Beverage Consumption on Primary Shopper’s Bias Proxies**

	(1)	(2)	(3)	(4)	(5)
Nutrition knowledge	-0.634*** (0.078)	-0.860*** (0.077)	-0.689*** (0.078)	-0.652*** (0.078)	-0.822*** (0.081)
Self-control	-1.215*** (0.062)	-1.533*** (0.054)	-1.172*** (0.062)	-1.203*** (0.061)	-0.712*** (0.038)
Taste for soda	0.350*** (0.038)		0.351*** (0.038)	0.345*** (0.039)	0.465*** (0.037)
Health importance	-0.143** (0.065)		-0.055 (0.065)	-0.181*** (0.067)	-0.251*** (0.068)
ln(Household income)	-0.048*** (0.018)		-0.061*** (0.016)	-0.077*** (0.018)	-0.040** (0.019)
ln(Years education)	-0.834*** (0.097)		-0.885*** (0.097)	-1.077*** (0.100)	-0.889*** (0.102)
Other beverage tastes	Yes	No	Yes	Yes	Yes
Other demographics	Yes	Yes	No	Yes	Yes
County indicators	Yes	Yes	Yes	No	Yes
Self-control 2SLS	Yes	Yes	Yes	Yes	No
R^2	0.272	0.238	0.264	0.148	0.272
N	18,568	18,568	18,568	18,568	18,568

Notes: This table presents estimates of Equation (23). Data are at the household level, and the dependent variable is the natural log of SSB purchases per adult equivalent in the most recent year that the household was in Homescan. This table parallels Table 3, except using the bias proxies and preferences for the primary shopper, instead of the average across all household heads. Columns 1-4 correct for measurement error in self-control using two-sample 2SLS, with standard errors calculated per Chodorow-Reich and Wieland (2016). Taste for soda is the response to the question, “Leaving aside any health or nutrition considerations, how much would you say you like the taste and generally enjoy drinking [Regular soft drinks (soda pop)]?” “Other beverage tastes” are the responses to parallel questions for other beverages. Health importance is the response to the question, “In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?” Responses to each question were originally on a scale from 0 to 10, which we rescale to between 0 and 1. “Other demographics” are natural log of age, race, an indicator for the presence of children, household size in adult equivalents, employment status, and weekly work hours. Observations are weighted for national representativeness. Robust standard errors are in parentheses. *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively.

J Equivalent Variation and Welfare Calculations

We can use the sufficient statistics discussed in Section II to approximate the change in welfare from the introduction of a sin tax. The net change in welfare is equal to the equivalent variation (measured under normative utility) of the sin tax, weighted by each individual’s welfare weight, plus the increase in tax revenues and externality reductions, measured in terms of public funds. We consider each in turn.

Equivalent variation (EV) for a consumer of type θ is defined as the change in wealth which would alter normative utility by the same amount as the introduction of a given tax t . For a small tax, this is equal to decision utility EV (the wealth change such that the consumer would express indifference between it and the introduction of the tax t) plus the change in welfare due


to internalities. We log-linearize demand around the (observed) no-tax equilibrium $t = 0$, so the demand for s from z -earners under a tax t is approximated by $\hat{s}(z, t) = \bar{s}(z) \left(\frac{p+t}{p}\right)^{-\bar{\zeta}^c(z)}$, and decision utility EV is approximated by $\frac{1}{1-\zeta^c(\theta)} ((p+t)\hat{s}(z, t) - p\bar{s}(z))$. The welfare change due to internalities is equal to $\gamma(\theta) \cdot (\hat{s}(z, t) - \bar{s}(z))$. The sum of these two figures gives the money-metric change in welfare for consumers earning z . Weighting this money-metric change by social marginal welfare weights (net of income effects, i.e., $\hat{g}(z)$), these welfare changes can be aggregated with the value of resulting tax revenues and reduced externalities (both measured in units of public funds) to approximate the total change in welfare from the tax, accounting for distributional concerns.

We can compute these gains for the case of the SSB tax using the estimates presented in Section IV. Our baseline optimal tax generates an estimated increase in social welfare of \$11.48 per adult equivalent consumer per year, or about \$3.6 billion in aggregate across the U.S.⁵⁴ Figure 9 plots the decomposition of these gains into the four component parts, for each of the income bins reported in Nielsen, assuming that tax revenues and externality reductions are distributed evenly across the population. Figure A6 plots analogous results using the specification based on self-reported SSB consumption from PanelViews, for which we estimate an annual increase in social welfare of \$39.06 per person, or about \$12.1 billion across population. Finally, if one did not believe consumers were biased (i.e., if internalities were zero, as in the “No externality” specification in Table 5), then the lower optimal tax would generate an increase in social welfare of \$1.04 per person year, or \$325 million nationally. (This decomposition is displayed in Figure A7.)


One may also be interested in the welfare gains which would be realized under various suboptimal SSB tax policies. We consider two in particular. First, consider the modal policy among U.S. cities which have passed SSB taxes: a tax of one cent per ounce. Under our baseline parameter estimates, this lower-than-optimal tax would generate a welfare gain of \$10.1 per person, or \$3.1 billion across the population—implying that the loss from this sub-optimal policy rather than the optimum is approximately \$400 million annually. Second, consider a policy maker who mistakenly believes consumers are unbiased and who therefore implements the “No externality” optimal tax estimate from Table 5, when in fact our estimates of bias are accurate. The resulting welfare gains from this suboptimal tax are only \$6.47 per person per year, or \$2 billion across the population, implying a loss of nearly \$1.6 billion per year relative to the optimal tax. These results emphasize the importance of accounting for behavioral biases when designing the optimal policy.

⁵⁴This calculation assumes 311 million adult equivalents in the U.S., based on a U.S. population of 326 million, 22.8% of whom are under 18, with an average recommended caloric intake equal to 80% of that for adults.

Figure A6: Welfare Consequences of Optimal SSB Tax Using PanelViews Data



Notes: This figure is identical to Figure 9, except the optimal tax is computed using self-reported SSB consumption from the PanelViews survey, rather than Homescan data. The figure plots the decomposition of welfare changes resulting from the baseline optimal sugar-sweetened beverage tax, across income bins. “Redistributed revenues” and “Externality correction” are money-metric values, assumed to be distributed equally across the income distribution. “Internality correction” is the increase in (money-metric) welfare due to the change in consumption resulting from the tax, at each income level. “Decision utility EV” is the value dy such that consumers are indifferent between a change in net income dy and the introduction of the optimal SSB tax.

Figure A7: **Welfare Consequences of Optimal SSB Tax if Consumers Are Unbiased**

Notes: This figure is identical to Figure 9, except the optimal tax and welfare changes are computed assuming consumers are unbiased (i.e., internalities are zero). The figure plots the decomposition of welfare changes resulting from the baseline optimal sugar-sweetened beverage tax, across income bins. “Redistributed revenues” and “Externality correction” are money-metric values, assumed to be distributed equally across the income distribution. “Internality correction” is the increase in (money-metric) welfare due to the change in consumption resulting from the tax, at each income level. “Decision utility EV” is the value dy such that consumers are indifferent between a change in net income dy and the introduction of the optimal SSB tax.

K Structural Simulations of Optimal SSB Tax

This appendix presents the optimal estimated SSB tax using structural models calibrated to the parameters estimated in Section III. This also provides a robustness test of the stability of the sufficient statistic policy estimates reported in Table 5.

K.A Functional Form Specifications

For the simulations to follow, we employ the following functional form for utility:

$$U(c, s, z; \theta) = c + \overbrace{v(s, c, \theta)}^{\text{SSB subutility}} - \overbrace{\Psi(z/w(\theta))}^{\text{Effort cost}} \quad (83)$$

$$V(c, s, z; \theta) = c + v(s, c, \theta) - \Psi(z/w(\theta)) - \tilde{\gamma}(\theta)s \quad (84)$$

We consider two different specifications for $v(s, c, \theta)$, to explore the robustness of these results to particular assumptions about functional forms. The first is

$$v(s, c, \theta) = a(c)b(\theta) \left(\frac{s^{1-k(\theta)}}{1-k(\theta)} \right). \quad (85)$$

An attractive feature of this functional form is that there is type-specific constant compensated price elasticity of demand, equal to $\zeta^c(\theta) = 1/k(\theta)$. (This may vary across types.) However a disadvantage of it is that it doesn't technically accommodate a strict present bias interpretation, since the bias $\tilde{\gamma}$ cannot be mapped into a specific value of β . To allow for this possibility, we employ a different functional form for our second specification:

$$v(s, c, \theta) = a(c)b(\theta) \left(\frac{s^{1-k(\theta)}}{1-k(\theta)} \right) - \iota(\theta)s. \quad (86)$$

Then $\iota(\theta) + \tilde{\gamma}(\theta)$ can be interpreted as the long run money-metric health costs of SSB consumption, of which the present biased agent internalizes $\iota(\theta)$, accounting for only a share $\beta(\theta) = \frac{\iota(\theta)}{\iota(\theta) + \tilde{\gamma}(\theta)}$, when making consumption decisions.

In both specifications, the product $a(c)b(\theta)$ controls the level of soda consumption, and can be calibrated nonparametrically using the observed cross-income profile of SSB consumption. The components $a(c)$ and $b(\theta)$ control whether that cross-sectional variation is driven by income effects or preference heterogeneity, respectively. We can nonparametrically calibrate $a(c)$ to generate the income elasticity estimated from our data, and we attribute the residual cross-sectional SSB variation to preference heterogeneity.

The term $\tilde{\gamma}(\theta)$ controls the marginal internality from SSB consumption. If some SSB variation is due to income effects (non-constant $a(c)$), then normative utility in Equation (84) is not quasilinear in numeraire consumption, and so $\tilde{\gamma}(\theta)$ is not exactly equal to $\frac{U'_s}{U'_c} - \frac{V'_s}{V'_c}$. Therefore we use the ‘‘tilde’’ notation to distinguish this structural internality parameter from the equilibrium money metric internality, which we will continue to denote γ .

We assume isoelastic disutility of labor effort, $\Psi(\ell) = \frac{1}{1+1/\zeta_z^c} \ell^{(1+1/\zeta_z^c)}$, where ζ_z^c is the Frisch elasticity of labor supply, which is assumed to be homogeneous conditional on income.

To compute the optimal SSB tax, we first calibrate the parameters in Equations (83) and (84) to match the patterns of SSB consumption documented in Section III, and other data from the U.S. economy described below. We then use a numerical solver to compute the optimal tax policies under those parameter values.

K.B Calibration Procedure and Data Sources

We draw from Piketty, Saez, and Zucman (2018) to calibrate the pre- and post-tax income distributions for the United States. We use their reported distribution for 2014, which includes percentiles 1 to 99, and much finer partitions within the 99th percentile. Each of these points in the income distribution is treated as an individual point in a discretized ability grid for the purposes of simulation.⁵⁵ Thus in the simulations there is a one-to-one mapping between type θ and ability $w(\theta)$. We also use these distributions to calibrate effective marginal tax rates in the U.S. status quo, computed as $1 - \frac{dy_{US}}{dz_{US}}$, where z_{US} and y_{US} represent the reported pre- and post-tax income distributions.⁵⁶

To encode a preference for redistribution, we assume type-specific Pareto weights given by $\alpha(\theta) = y_{US}(\theta)^{-\nu}$, where ν controls the degree of inequality aversion. Following Saez (2002b), we use a baseline value of $\nu = 1$, approximately corresponding to the weights which would arise under logarithmic utility from consumption, and we report alternative specifications with $\nu = 0.25$ and $\nu = 4$. We compute these weights using the status quo U.S. net income distribution, and they are held fixed during simulations in order to isolate the effects of other model parameters on optimal taxes without endogenously changing distributional motives.

We calibrate the status quo level of SSB consumption using the consumption estimates across incomes, as shown in Figure A2.⁵⁷ To interpolate consumption across the full ability distribution in our simulations, we find the percentile in our income grid which corresponds to each income point in Figure A2, and we then interpolate (and extrapolate) SSB consumption linearly across income percentiles in our ability grid.

The price elasticity of SSB demand, $\zeta^c(\theta)$, is calibrated to match the estimates described in Section III.B. We predict the uncompensated SSB elasticity as a linear function of (log) net income, using the interaction term on elasticity and log income reported in Table A8. We then convert these estimates into the compensated elasticity $\zeta^c(\theta)$ using the Slutsky equation and the estimated SSB income elasticity reported in Table 2.

In the second structural specification, we must also calibrate the health costs $\iota(\theta)$ to which the present-biased agent does attend. We calibrate those costs using estimates of the expected health costs from SSB consumption. Long et al. (2015) estimate the impact of SSB consumption on quality-adjusted life years (QALY) by tracing its impact on the likelihood of suffering a stroke, ischemic heart disease, hypertensive heart disease, diabetes mellitus, osteoarthritis, postmenopausal breast cancer, colon cancer, endometrial cancer, and kidney cancer. That paper estimates that a permanent 20% reduction in SSB consumption generates an average increase of 0.00278 QALYs per person in the first ten years. A “back of the envelope” calculation suggests this equates to a total

⁵⁵We drop the bottom three percentiles of the distribution, which have slightly negative reported incomes due to reported business or farm losses.

⁵⁶An alternative approach would be to use estimated effective marginal tax reported, for example from the NBER TAXSIM model or as computed by CBO. However since these estimates often omit some types of taxes, and fail to include many types of benefits, we have chosen to use the more comprehensive notion of taxes and benefits from Piketty, Saez, and Zucman (2018).

⁵⁷We use Homescan consumption as our baseline, though we also report results using self-reported consumption from our PanelViews survey below.

increase of 0.0021 QALYs per person from a one-year SSB reduction of 20%.⁵⁸ This translates to \$105 per person using the commonly used conservative estimate of \$50,000 for the value of a QALY (Hirth et al., 2000), corresponding to a marginal health cost of \$0.10 per ounce of SSB consumed.⁵⁹ Although these marginal health costs are likely to be heterogeneous in practice, for purposes of this calibration we assume they are constant across the population, so that $\iota(\theta) + \tilde{\gamma}(\theta) = \0.10 per ounce for all individuals. Using our estimates of bias from Section III.D, this implies that 5% to 15% of health costs are uninternalized due to present bias and/or incorrect beliefs.

We jointly calibrate the functions $a(c)$ and $b(\theta)$ nonparametrically to match both the level of SSB consumption (across the income distribution) and the estimated SSB income elasticity, using the following procedure. We first assume $a'(c) = 0$, and we calibrate the product $a(c)b(\theta)$ at each point in the income distribution—without yet worrying about the decomposition into separate terms—to match the observed schedule of SSB consumption via the first-order condition for SSB consumption. We then compute a path of $a'(c)$ which is consistent with our estimated SSB income elasticities, and we numerically integrate to find $a(c)$, which in turn identifies $b(\theta)$ at each ability gridpoint. Finally, we compute the implied ability distribution $w(\theta)$ which generates the observed income distribution. We then repeat this procedure with the new $a'(c)$ (which affects the FOC for SSB demand) and we iterate to convergence. In this manner, we find a non-parametric schedule of $a(c)$ and $b(\theta)$ which is consistent with both the level of estimated soda consumption at each income, and with the estimated SSB income elasticities in our data. During simulations, we compute the nonparametric function $a(c)$ using linear interpolation (extrapolation) over $\log(c)$.

For estimates which assume an externality cost from SSB consumption, we use a value of 0.85 cents per ounce, as calculated in Section III.E.

K.C Simulation Results

Our structural model estimates for the optimal SSB tax are presented in Table A14. As in Table 5, for each specification we report the optimal SSB tax under two income tax regimes. First, we compute the optimal SSB tax holding fixed the income tax at the current status quo, under the assumption that SSB tax revenues are distributed equally across all consumers as a lump-sum payment. Second, we jointly solve for the optimal income tax, which may be quite different from the U.S. status quo. The former exercise is most relevant if the SSB tax is viewed as an isolated policy over which policymakers much make choices, independent of overall income tax reforms. The latter is more relevant as an illustration of the theoretical results in Section II, which follow the

⁵⁸The simulated experiment in Long et al. (2015) is a 20% permanent reduction of SSB consumption, whereas we are interested in the QALY cost of each SSB consumed. To reach this, we first compute the approximate total QALYs gained from reducing SSB consumption for a single year. If the 10-year total of 0.00278 QALYs were composed of 10 years of exposure to the effects of a 20% reduction during year 1, 9 years of exposure from the reduction during year 2, etc., then the first year's reduction would account for $10/55 = 0.18$ of the 10 year total, or 0.00051 QALYs. This accounts only for the QALY difference in the first 10 years. If the per-year effect were the same in the ensuing years (a conservative assumption, given that health costs typically compound later in life), then this ten-year effect should be multiplied by four, given a median US age of 38 and life expectancy of 78. This yields a total approximate effect of 0.0021 QALYs for a one-year 20% SSB reduction.

⁵⁹Long et al.'s (2015) estimates imply that on average, people consume 5475 oz of SSBs per year which implies that a 20% reduction would reduce consumption by 1095 oz.

tradition in the optimal taxation literature of jointly characterizing optimal income and commodity taxes. To illustrate the difference between these two regimes, Figure A8 plots the mapping between pre- vs. post-tax income in the U.S. status quo and under the simulated optimal income tax.

K.C.1 The Optimal Sweetened Beverage Tax

Table A14 computes estimates for each of the specifications from Table 5 using this structural model. The values in the first column (under the prevailing income tax) are quantitatively very close to the sufficient statistics computations from Table 5. The values in the second column (computed under the optimal income tax) generate the same qualitative patterns, but they are generally lower than the values in Table 5 under the optimal income tax. This is driven primarily by the endogenous change in SSB consumption generated by the income tax reform in the structural model. Recall that the sufficient statistics formula (under the optimal income tax) is computed using the *observed* profile of soda consumption, whereas in the structural model, the tax is computed accounting to the change in consumption due to the income tax reform. The optimal income tax substantially raises the net income of the lowest earning households (see Figure A8) and, since SSB income effects are positive, this raises the amount of SSB consumption among low earners. Therefore, the demand response to a tax, $\frac{ds}{dt}$, is substantially larger in magnitude at low incomes in the structural model than is accounted for in the sufficient statistics calculation, and as a result, bias correction progressivity σ is accordingly higher in the structural model, implying a larger corrective motive and hence a higher optimal tax.

K.C.2 How the Optimal Internality-Correcting Tax Varies with Elasticity and Bias

The rows in Table A14 are intended to represent realistic alternative specifications for the purpose of understanding the robustness of our estimates. However, one might also wish to understand how more extreme variations in parameters affect the optimal sin tax. In particular, the analytic formulas in Section II emphasize the role of bias and elasticity of demand in determining the optimal tax. In order to illustrate those insights quantitatively, Figure A9 plots the optimal sin tax across a range of values of bias and elasticities of demand. To isolate the role of the bias and elasticity of demand, the simulations in this figure assume constant values of each parameter across the income distribution (akin to row 10 in Table A14), and zero externality cost. As a result, the optimal tax in the absence of distributional concerns (i.e., the optimal Pigouvian tax) lies on the 45-degree line, and departures from that line illustrate the effect of redistributive concerns on the optimal sin tax. The baseline (average) value of bias is represented in the figure by the vertical dashed line, and the baseline (average) demand elasticity is plotted by the red line. All values are computed holding fixed the status quo U.S. income tax.

Figure A9 illustrates an important theoretical insight from Section II. In particular, the regressivity costs term (which reduces the optimal SSB tax) rises with the demand elasticity. Quantitatively, these costs are modest for elasticities in the range estimated in our data, and the optimal SSB tax lies close to (although distinctly below) the 45 degree line representing the optimal Pigouvian

tax. However at lower elasticity values, regressivity costs become large. Indeed at a low elasticity value of 0.25, the optimal SSB tax is in fact a subsidy for the average bias computed in our data. The Optimal SSB Tax Under Constrained Targeted Revenue Recycling

Table A14: **Simulation Results: Optimal SSB Tax (Cents Per Ounce)**

	Existing income tax	Optimal income tax
1. Baseline	1.72	1.47
2. Pigouvian (no redistributive motive)	1.98	-
3. Weaker redistributive preferences	1.90	1.84
4. Stronger redistributive preferences	1.47	0.76
5. Higher demand elasticity ($\bar{\zeta}^c(\theta) = 2$)	1.82	1.66
6. Lower demand elasticity ($\bar{\zeta}^c(\theta) = 1$)	1.39	0.94
7. Constant bias and elasticity	1.61	1.35
8. Pure preference heterogeneity	1.63	1.63
9. Pure income effects	1.67	2.05
10. No internality	0.44	0.16
11. No corrective motive	-0.34	-0.56
12. Self-reported SSB consumption	2.77	2.54

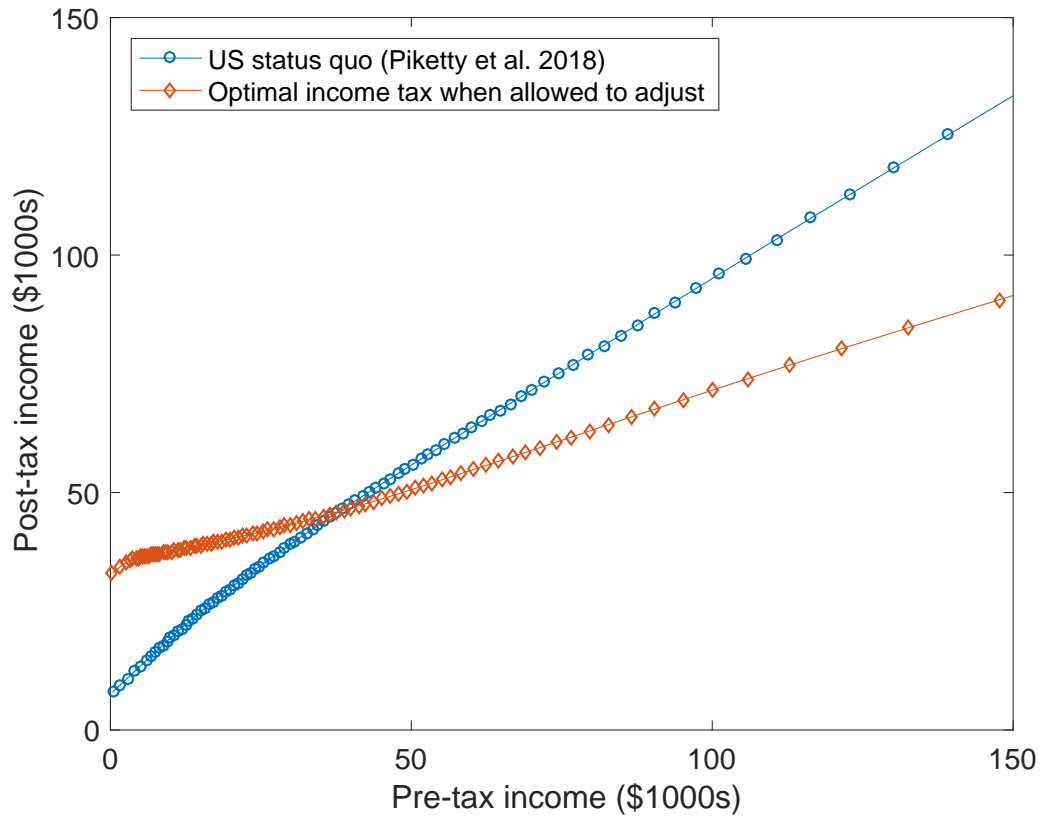
(a) **Specification 1 (Constant Elasticity)**

	Existing income tax	Optimal income tax
1. Baseline	1.86	1.67
2. Pigouvian (no redistributive motive)	2.00	-
3. Weaker redistributive preferences	1.96	1.91
4. Stronger redistributive preferences	1.71	1.04
5. Higher demand elasticity ($\bar{\zeta}^c(\theta) = 2$)	1.94	1.84
6. Lower demand elasticity ($\bar{\zeta}^c(\theta) = 1$)	1.57	1.21
7. Constant bias and elasticity	1.75	1.55
8. Pure preference heterogeneity	1.77	1.77
9. Pure income effects	1.80	2.06
10. No internality	0.47	0.17
11. No corrective motive	-0.37	-0.65
12. Self-reported SSB consumption	2.96	2.87

(b) **Specification 2 (Partially Internalized Linear Health Costs)**

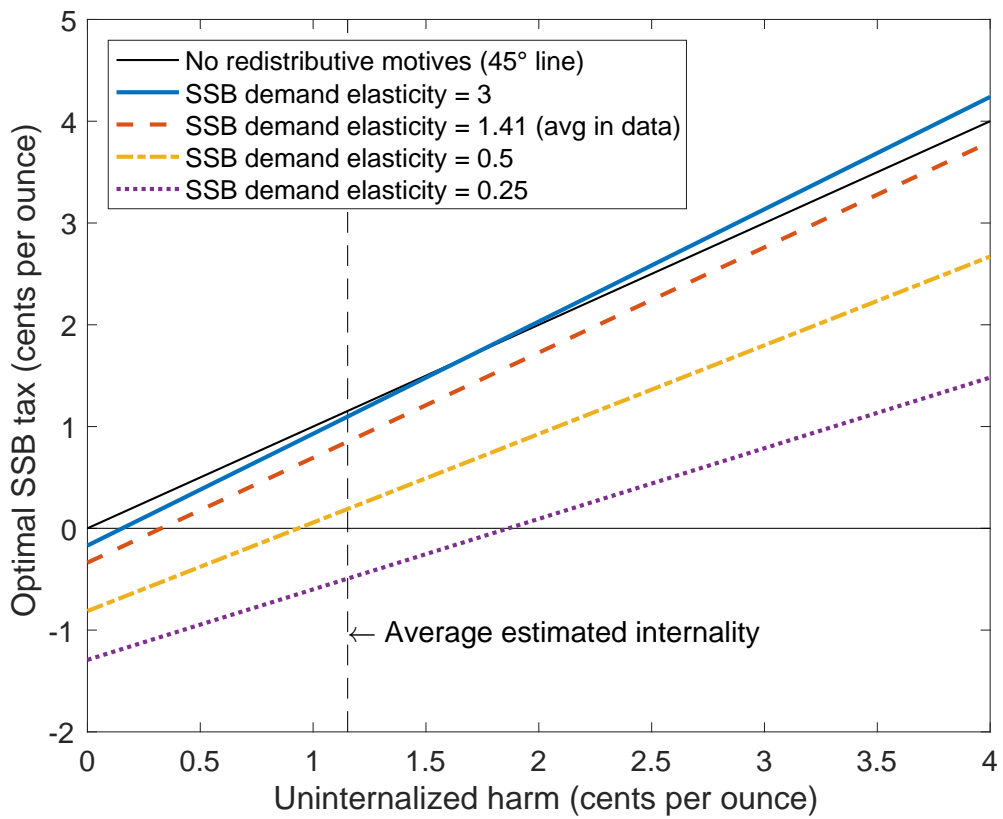
Notes: This table reports the optimal tax estimates using the calibrated structural models. Panel (a) uses a structural model with constant compensated elasticity of demand for SSBs. Panel (b) uses a specification in which a portion of linear health costs are internalized. In both panels, the first column displays the optimal SSB tax when the income tax is held fixed at the U.S. status quo; the second column displays the optimal SSB tax when the simulations also solve for the optimal income tax.

Figure A8: **Simulations: Status Quo Income Tax and Optimal Income Tax in Baseline Specification**



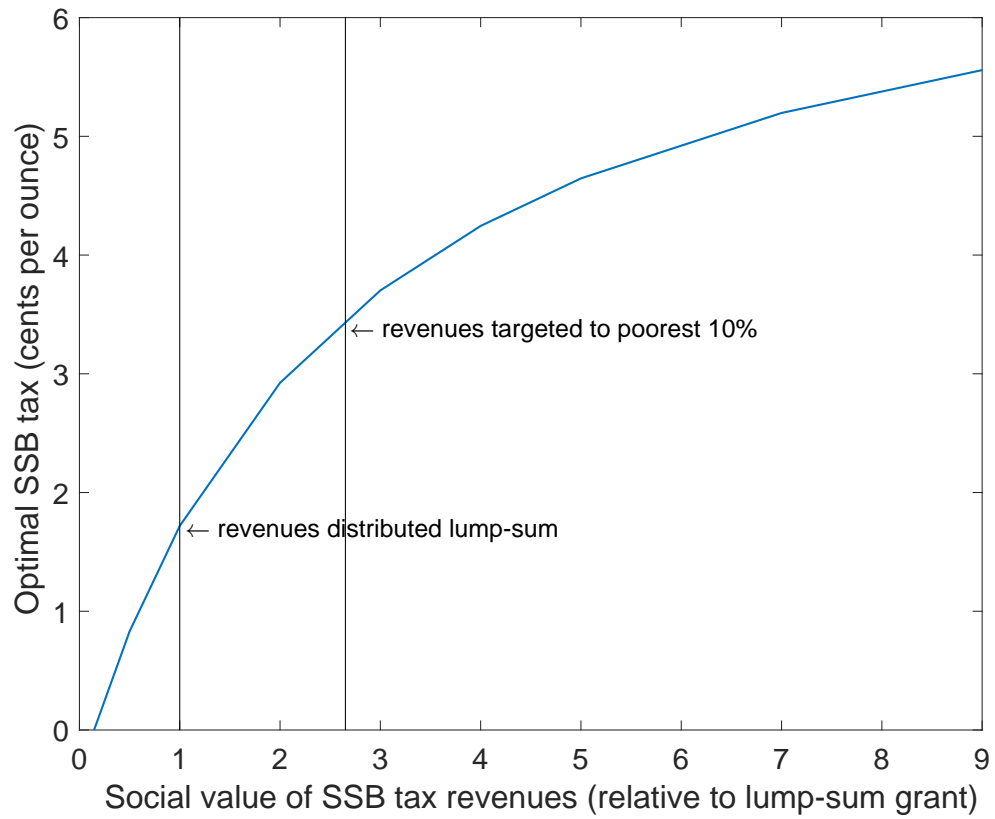
Notes: This figure compares the simulated optimal income tax to the status quo income tax in the U.S.

Figure A9: **Optimal Internality-Correcting Tax Across Values of Bias and Elasticities**



Notes: This figure plots the simulated optimal SSB tax in cents per ounce across values of internality bias (in cents per ounce) for different values of the SSB price elasticity of demand. To illustrate the effects parameters on the optimal internality-correcting tax, these simulations assume zero externality.

Figure A10: **Optimal SSB Tax for Different Marginal Social Values of SSB Tax Revenues**



Notes: This figure shows how the optimal SSB tax varies depending on the social usefulness of SSB tax revenues. A value of one on the horizontal axis implies that SSB tax revenues are valued the same as marginal funds raised via the income tax.