Journal of Environmental Economics and Management 103 (2020) 102356

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

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

Dynamic corrective taxes with time-varying salience[☆]

Ben Gilbert [a,*], Joshua S. Graff Zivin [b]

[a] Colorado School of Mines, 1500 Illinois St, Golden CO, 80401, United States
[b] University of California, San Diego and NBER, 9500 Gilman Dr, La Jolla CA, 92093-0519, USA

ARTICLE INFO

Article history:
Received 28 July 2019
Revised 6 June 2020
Accepted 18 June 2020
Available online 15 July 2020

JEL classification:
D03
D11
D62
D91
H21
H23
L97
Q40
Q41
Q50

Keywords:
Salience
Inattention
Optimal taxes
Energy demand
Consumption persistence

ABSTRACT

Economies across the globe are becoming increasingly cashless and many payment systems have become automated, driving a temporal wedge between consumption and payment and generally making the costs of consumption intermittently salient. Since this inconsistent price salience alters demand elasticities, it is a particular concern for goods that generate externalities and the price-based policies deployed to address them. This paper derives optimal dynamic corrective taxes for suboptimal and persistent consumption decisions. These taxes depend on the agent's ability to commit to a future consumption path. We also characterize a second-best constant tax and the excess burden from time-invariant tax rates. When calibrated to U.S. residential electricity consumption, the model shows that the second-best constant tax is more than twice the marginal external cost of carbon emissions.

1. Introduction

The standard economic remedy to externalities has been well known since the 1920s and has been a regular fixture in the policy maker's toolkit since at least the 1980s. The idea is simple, yet elegant. Impose a tax equal to marginal external damages in order to 'internalize' externalities and generate private decision making that is socially optimal (Pigou, 1920). But in an increasingly cashless society where the financial consequences of one's daily choices may only be experienced monthly or even quarterly, is it realistic to expect consumers to optimally perform this calculation? Indeed, an emerging literature that examines the impacts of price salience on purely private decisions suggests that this is unlikely to be the case across a range

☆ We thank Hunt Allcott, Terry Iverson, Jacob LaRiviere, Andrea La Nauze, Pete Maniloff, Chuck Mason, Kevin Novan, Linda Thunstrom, Klaas van 't Veld, and numerous seminar participants for helpful comments. Any errors are our own.

* Corresponding author.
E-mail addresses: bgilbert@mines.edu (B. Gilbert), jgraffzivin@ucsd.edu (J.S. Graff Zivin).

https://doi.org/10.1016/j.jeem.2020.102356
0095-0696/

of contexts.[1] In addition, many consumers use credit cards or other delayed payment mechanisms for everyday purchases, which means that they are in effect billed monthly. A large literature shows that using such payment modes increases spending (see, e.g., Hirschman, 1979; Feinberg, 1986; Prelec and Simester, 2001; and Thomas et al., 2011). These mechanisms can cause overconsumption of energy, water, food, and other resources with significant external social costs.

The so-called inattention problem raises the specter of additional policy interventions that increase salience in order to fix the 'internality' from privately suboptimal decision making.[2] The implications for a simple static setting are straightforward – impose a super tax that forces consumers to face the full internal (i.e. salient-equivalent) and external costs. This logic becomes more complex in a dynamic setting with a more realistic representation of price salience as intermittent. The optimal (time-varying) tax will now depend on both the persistence of consumption across periods and the sophistication consumers have about their own price inattention. A large empirical literature shows that consumption decisions are highly persistent for a broad class of goods such as residential electricity consumption (Costa and Gerard, 2018; Deryugina et al., 2020; Filippini et al., 2018; Heien and Durham, 1991; Leth-Petersen, 2007; and Sexauer, 1977, to name a few), food (Thunstrom, 2010), gasoline (Scott, 2012), and cigarette consumption (Becker et al., 1994). This type of persistence can arise if a good is habit-forming (Becker and Murphy, 1988; Landry, 2019), when consumption is based on defaults or status quo bias (Samuelson and Zeckhauser, 1988) or when consumption depends on some complementary durable good stock (Flavin and Nakagawa, 2008). Moreover, recent research has shown that consumption persistence is an important consideration in the welfare effects of a corrective policy (Costa and Gerard, 2018).

Our work extends the key insights from this literature by showing that welfare effects are compounded by time-varying inattention and by demonstrating the importance of both persistence and inattention for the design of optimal and second-best policy. We develop a model of consumer behavior when prices are intermittently salient, demand persists across periods, and consumption creates external damages. We then derive optimal dynamic tax rules for two distinct types of consumers: a *naive* agent who dynamically optimizes her consumption persistence under current perceived prices while unaware of future price inattention; and a *partially sophisticated* agent who anticipates her own future price inattention but does so imperfectly.[3] Importantly for policy design, persistence provides a vehicle through which partially sophisticated agents can commit to reduce future consumption when prices are salient.[4] Our work builds on the static models of corrective policy with externalities and internalities developed by Allcott et al. (2014) and Allcott and Taubinsky (2015) by evaluating a fully dynamic setting with time-varying inattention and consumption persistence. We represent persistence with a Becker and Murphy (1988) model of habit formation. This stylized framework produces consumption and taxation dynamics that are qualitatively in line with other common models of consumption persistence. Examples include durable goods investment by a representative household, nondurable consumption or durable goods utilization of an individual agent with status quo bias, or an agent that alternates between active and passive decision making because of external cues such as monthly bills.[5]

We find that the optimal time-varying tax depends upon the price-elasticity of demand and the size of external damages, as well as the price salience decay function, consumption persistence, and time preference parameters. During relatively price-salient periods, greater consumption persistence and lower future price salience lead to lower taxes for the partially sophisticated agent and have no effect on taxes for the naive agent. This occurs because the partially sophisticated agent's implied demand for consumption is diminished by their commitment to lower habits in the future. Given that time-varying taxes are difficult to implement in practice, we evaluate the welfare consequences of some plausible constant tax alternatives.

Our theoretical work is followed by a numerical calibration for U.S. residential electricity consumption. The optimal tax rises as salience decays within a billing cycle, until the tax is more than three times larger than marginal external damage. The second-best constant tax, which leads to lower-than-optimal tax rates during non-salient periods and higher-than-optimal tax rates during salient periods, is more than twice the marginal external damage. The excess burden from the second-best tax ranges from about 19 dollars per household for naive agents to about 50 dollars per household for partially sophisticated agents. With approximately 126 million households in the U.S., this amounts to between 2.4 billion and 6.3 billion dollars in deadweight

[1] For example, Finkelstein (2009) finds that toll road drivers become less price elastic following electronic toll collection, Sexton (2015) finds that residential electricity demand is less price elastic with automatic bill pay, and Grubb and Osborne (2015) find that cell phone consumers are inattentive to both the marginal price structure of their bills and to their own past usage. DellaVigna (2009) includes a thorough review of the empirical evidence for inattention to cost information of various types.

[2] For example, Chetty et al. (2009), Chetty (2009), Allcott et al. (2014), Allcott and Taubinsky (2015), and Farhi and Gabaix (2020) all derive optimal policies for internalities.

[3] In order to further evaluate the role of sophistication and commitment, an earlier version of this paper considers additional cases, including an agent who is not forward-looking about inattention or persistence (Gilbert and Graff Zivin, 2018). These cases may apply to related behavioral biases such as status quo bias, but not to alternative justifications for persistence such as durable goods complementarity.

[4] The problem is conceptually similar to the one that arises in models of taxation with present bias, in which agents' sophistication about their own time-inconsistency affects their demand for commitment. See Gruber and Koszegi (2004) or O'Donoghue and Rabin (1999, 2006) for a generic treatment of present bias, or Heutel (2015) for externality correction with present bias.

[5] Allcott et al. (2014) and Allcott and Taubinsky (2015) model adoption of a durable good whose utilization may generate externalities. Although durable goods have dynamic implications in the sense that the impacts of decisions will be felt across time, adoption and utilization are one-shot decisions in their models. In contrast, we model a recurrent purchase over an infinite time horizon with periodically salient price signals. A less stylized but fully dynamic representation of the durable goods problem would require an explicit representation of how investment timing is influenced by bill shocks versus routine replacement, and how this reinvestment motivation affects the salience of energy costs. Such a model would result in time-varying taxes similar to the ones we derive in this paper. Unlike models with durable goods, however, two-part policy instruments to separately correct the externality and the internality (as in Allcott et al. (2014) and Heutel (2015)) are not available in our setting because present consumption and future commitment are determined by the same action.

loss from the electricity sector, even if the second-best optimal tax is implemented. These welfare losses rise sharply if the tax is reduced from the second-best level, if inattention is more severe, or if demand is more elastic. For example, the welfare loss associated with having no tax at all ranges from 181 dollars per household for partially sophisticated agents to 198 dollars per household for naive agents, or 22.8 billion to 24.9 billion dollars.

The key insights from our model are not limited to electricity markets and should be relevant for any price-based policy designed to address market failures, where price/cost inattention and linked intertemporal decision making are commonplace. For example, gasoline consumption 1) emits several types of pollutants; 2) depends on vehicle type and driving habits that persist across many periods; and 3) does not present drivers with explicit, salient prices when making trip-level decisions. Similarly, unhealthy food choices 1) impose health costs on society; 2) are sometimes habit-forming, such that enjoyment in the current period depends on consumption in prior periods; and 3) are generally made at higher frequency than the purchases that make their costs salient. The increasing prevalence of auto-billing and subscription services for a range of purchases will only increase its applicability.

The remainder of the paper is organized as follows. In the next section, we set up the model and derive optimal and second-best taxes. In the subsequent section we derive formulas for the welfare loss from time-invariant taxes. This is followed by a section presenting the results from the numerical simulation. The final section discusses and concludes.

2. Model

We begin with a simple extension of the standard externality model to incorporate time-varying salience. We then extend the model to consider consumption persistence, commitment and agent sophistication in Section 2.2. Utility in each period depends on consumption of a clean numeraire good y_t and a dirty good x_t with price p. Consumption of the dirty good produces social damages φx_t that last only for the current period, where φ is the marginal external cost of damages from consumption of x.[6] For simplicity we model period utility as quasilinear:

$$U(x_t, y_t) = u(x_t) + y_t. \tag{1}$$

We apply a salience factor $\theta_t \in (0, 1]$ to the price and suppose agents optimize as if they face price $\theta_t p$. This reflects common empirical findings that when prices are not perfectly salient, agents' decisions exhibit lower price elasticity than their true preferences (e.g., Chetty et al., 2009; Finkelstein, 2009; Sexton, 2015).[7] In our setting, the price of the dirty good is intermittently salient, so that the price elasticity reflected in observed consumption varies over time while true preferences do not. To fix ideas consider a billing cycle with T periods per cycle; for example, a monthly bill.[8] Denote the first day of the cycle $t = 1$ as the day on which agents receive a salient price signal – a bill. It is our contention that as time passes, the salience of information in the bill decays and reaches some minimum value $\underline{\theta} \in (0, 1)$ at some period $I < T$ within the cycle. At present we assume only that $\theta_1 = 1$ on the first day when the salient price signal or bill arrives, and on subsequent days within the cycle $\underline{\theta} \leq \theta_t \leq \theta_{t-1}$ for $t \leq I$, and $\theta_t = \underline{\theta}$ for $I \leq t \leq T$.[9] In period $T + 1$ the agents receive another salient price signal (a new bill), and the cycle repeats. Agents live for an infinite number of cycles indexed by M.

Consider an arbitrary period i following a price signal or bill in the first billing cycle ($M = 1$). The agent chooses a plan for current and future consumption to maximize the present value of utility across periods within a billing cycle and across cycles, subject to the constraint that the present value of perceived expenditures equals period i wealth A_{i1}:

$$\begin{aligned} V = \max_{\{x_{tM}, y_{tM}\}} & \sum_{t=i}^{T} \beta^{t-i} U(x_{t1}, y_{t1}) + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} U(x_{tM}, y_{tM}) \\ -\lambda & \left[\sum_{t=i}^{T} \beta^{t-i} (\theta_i p x_{t1} + y_{t1}) + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} (\theta_i p x_{tM} + y_{tM}) - A_{i1} \right] \\ & s.t.\ x_{0,1}, \ldots, x_{i-1,1},\ A_{i1}\ given \end{aligned} \tag{2}$$

[6] We do not consider the accumulation of pollution as a stock, but this is an interesting area for future research.

[7] Although it is possible that agents underconsume when prices are not salient, empirical evidence suggests overconsumption is generally prevalent.

[8] This is without loss of generality – the length of the cycle need not be regular or deterministic.

[9] We have found in unreported results that the consumption dynamics, tax, and welfare implications are robust to a variety of alternative assumptions. For example, we bound $\underline{\theta}$ away from zero to reflect that the elasticity of price-insalient goods is generally not zero. However, $\underline{\theta} = 0$ can be handled if $u(x)$ has a bliss point that does not dominate per period income, or if agents develop default behaviors and appliance settings as an adaptation to price insalience, all of which are plausible for residential energy and water use. Defaults and settings result in consumption persistence and commitment devices similar to habit persistence presented in Section 2.2. Similarly, there are many possible smooth or discrete decline functions for θ_t. People may also abruptly forget about the bill in some certain or uncertain time period (similar to the model of internal commitment and self control in Benhabib and Bisin (2005)). Uncertainty over future θ_t may affect current behavior even if a completely insalient state is not realized before the next salient price signal arrives. In that case, the average attention to bills may smoothly decline even if individual attention drops discretely. We impose more structure on the decay of θ_t in our numerical simulations in Section 3.1. Our emphasis is on characterizing the regulator's tax problem, so we do not require a complete cognitive model of attention decay. We have not considered cases where agents may accumulate debt, which is an interesting area for future research although we note that Karlan et al. (2016) allow for debt in their model of inattention to infrequent and large expenses.

where β is the discount factor and subscripts index period within cycle followed by the cycle. Time inconsistency arises as agents solve this problem in each new successive period, which means that earlier plans are revised as θ_i changes. The first order conditions for the plan for present and future consumption as of period i are

$$u'_{tM} = \theta_i p, \quad \text{for all } t \geq i, \text{ and for all } M. \tag{3}$$

In the consumption plan determined by (3), the agent plans as if the current perceived price, $\theta_i p$, is the price in all periods and all cycles. In each period the agent plans to smooth consumption of the dirty good equally across time according to (3), but they revise that plan in each new period as θ_i changes. The actual period i decision is determined by $u'_i = \theta_i p$. As θ_i declines over time, so does the marginal utility of the dirty good. Current and planned consumption, as well as emissions, increase accordingly.

In this simple model without consumption persistence or commitment, there are no adjustment costs from revising the plan each period. There are no dynamic consequences of overconsuming during low salience periods. The agent faces no costs of deviating from her plan in the future, so decisions made during low salience periods do not increase the cost of avoiding externalities in high salience periods. Likewise during high salience periods, a more sophisticated agent who anticipates the decline in θ would prefer to commit to the current plan by raising the cost of deviating in the future. If the agent is sophisticated and anticipates $\theta_t < \theta_i$ in periods $t > i$ when making their period i decision, then they would like to commit themselves to some period t consumption level with $u'_{tM} > \theta_t p$. Without a commitment mechanism, when period t arrives the perceived price is $\theta_t p$ and the agent solves $u'_{tM} = \theta_t p$. The same consumption path arises with or without sophistication. We address these issues in Section 2.2.

2.1. Optimal dynamic corrective taxes

Consumption will deviate from the socially optimal path because of salience internalities and environmental externalities. A system of optimal corrective taxes must address both sources of inefficiency. In this subsection we solve for the dynamic corrective taxes that induce the agent to consume along the socially optimal path.

The social planner chooses a path $\{x_t, y_t\}_{t=1}^{\infty}$ to maximize the intertemporal social welfare function

$$W = \sum_{t=1}^{\infty} \beta^{t-1} \left[U(x_t, y_t) - \varphi x_t\right] - \lambda \left[\sum_{t=1}^{\infty} \beta^{t-1}(px_t + y_t) - A_1\right] \tag{4}$$

subject to A_1 given. The social planner optimizes over all time periods regardless of the billing cycle. We temporarily drop the M subscript and use the t and i subscripts to denote position within a given cycle. The socially optimal path is defined by the first order conditions:

$$u'_t = p + \varphi, \quad \text{for all } t. \tag{5}$$

The agent could then be induced to consume along the optimal path by a dynamic tax that transforms her private first order conditions in each period into the social planner's first order conditions from (5). This would be a tax that solves

$$u'_t = \theta_t(p + \tau(\theta_t)) = (p + \varphi), \quad \text{for all } t. \tag{6}$$

The solution is stated as a proposition.

Proposition 1. *The optimal corrective tax is a time-varying weighted combination of the price internality and pollution externality. The tax grows as θ_t declines following a salient price signal and the weight on the price internality increases:*

$$\tau(\theta_t) = p\frac{1-\theta_t}{\theta_t} + \frac{\varphi}{\theta_t} \geq \varphi, \tag{7}$$

Equation (7) shows that even if there is no externality, a time-varying tax is still necessary to correct the salience internality which leads to private overconsumption. With an externality, however, this tax cannot be linearly separated into a pure internality term and a pure externality term. The presence of θ in the denominator of the externality term shows that declining salience exacerbates both the marginal external cost and the private cost of overconsumption. In the first period when $\theta_1 = 1$ the optimal tax is the standard Pigouvian tax equal to marginal external damages φ. As θ_t declines within a cycle, the tax grows for two reasons. First, having θ_t in the denominator reflects the reduced tax elasticity that comes along with a declining price elasticity, so the tax must grow in order to achieve a given objective. Second, the relative weight on the price internality $(1-\theta_t)$ also increases; as θ_t declines the departure from the first best becomes more affected by the private internality. The tax reaches its maximum at $p\frac{1-\underline{\theta}}{\underline{\theta}} + \frac{\varphi}{\underline{\theta}}$.

Notice that this tax does not depend on whether the agent anticipates changes in price salience. If the agent is sophisticated enough to anticipate declining price salience and they expect to overconsume in the future, they would like to commit to consuming less over time. An optimal dynamic tax would incorporate this pattern of commitment. As noted above, without consumption persistence or commitment there is no mechanism for the agent to control her future internalities and no costs of deviating from a plan made previously. The optimal tax then only incorporates the current value of the salience parameter. We now investigate the role of consumption persistence, commitment, and sophistication.

2.2. Persistence, commitment, and partial sophistication

There are many ways to model commitment, including both internal and external commitment devices. An empirical regularity of the contexts our model addresses (residential electricity, unhealthy food consumption, gasoline consumption) is that consumption decisions persist over time even conditional on prices, income, and other factors. We represent this persistence using Becker-Murphy rational habit formation.

Agents are forward-looking about their consumption persistence. Habit persistence raises the cost of deviating from past consumption decisions. This provides agents a partial commitment mechanism should they expect to make suboptimal decisions in the future. Agents can then stick more closely to the plan chosen during high salience periods.[10] However, the persistence of consumption plans chosen during low salience periods also increases future overconsumption and externalities. The way the agent trades off these effects depends on how sophisticated they are about their own attention decay. In order to highlight the role of commitment when price salience is time-varying, we focus on two types of agents:

- a *naive* agent who dynamically optimizes her habit persistence under perceived prices but does not consider that her attention to prices will decline in the future;
- a *partially sophisticated* agent who dynamically optimizes her habit persistence under perceived prices, anticipating that she will pay less attention to prices in the future.

We first illustrate the dynamics of consumption for each agent type and then discuss the implications of these dynamics for optimal taxation. Naive agents consider how today's decisions affect their future demands, even if they fail to recognize how their price inattention will change in the future. By contrast, partially sophisticated agents expect to pay less attention to prices in the future even if the current price salience is imperfect. We focus on these two definitions of naive and partially sophisticated in order to highlight the role of commitment in models with consumption persistence. When prices are salient, the partially sophisticated agent prefers to commit to a lower consumption path because she anticipates future price inattention and overconsumption. The naive agent does not demand commitment.[11]

Utility in each period now depends on past consumption of the dirty good, x_{t-1} according to

$$U(x_t, x_{t-1}, y_t) = u(x_t - \alpha x_{t-1}) + y_t. \tag{8}$$

The parameter α governs the persistence of consumption decisions across time. Consumption of x in adjacent time periods are "adjacent complements" in the sense that as x_{t-1} increases, the marginal utility of x_t rises in period t and consumption of x_t increases *ceteris paribus*.

2.2.1. Naive agents

The present value of the naive agent's period i consumption plan in the first cycle is given by

$$\begin{aligned} V_N = \max_{\{x_{tM}, y_{tM}\}} & \sum_{t=i}^{T} \beta^{t-i} U(x_{t1}, x_{t-1,1}, y_{t1}) + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} U(x_{tM}, x_{t-1,M}, y_{tM}) \\ & -\lambda \left[\sum_{t=i}^{T} \beta^{t-i} (\theta_i p x_{t1} + y_{t1}) + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} (\theta_i p x_{tM} + y_{tM}) - A_{i1} \right] \\ & \qquad s.t.\ x_{0,1}, \ldots, x_{i-1,1},\ A_{i1}\ given. \end{aligned} \tag{9}$$

From the point of view of period i, the agent perceives the price to be $\theta_i p$ in all periods, and applies this price to their optimal consumption plan for the present and all future periods. The first order conditions for the naive agent's consumption plan are

$$u'_{tM} - \alpha \beta u'_{t+1,M} = \theta_i p, \quad \text{for all } t \geq i, \text{ and for all } M. \tag{10}$$

By recursive substitution we obtain

$$u'_{tM} \approx \theta_i p \sum_{s=0}^{\infty} (\alpha\beta)^s = \frac{\theta_i p}{1 - \alpha\beta}, \quad \text{for all } t \geq i, \text{ and for all } M. \tag{11}$$

[10] This structure is similar to a model of status quo bias with an endogneous status quo, in which past choices affect the current default decision. Introducing durable goods stock accumulation would provide a similar dynamic mechanism. We focus on habits because they are an empirical feature of the polluting goods we are interested in, and because intermittent price salience has important implications for habit formation.

[11] This is the sense in which our concept of sophistication is analogous to models of present bias in which sophisticated agents are aware of their present bias and demand commitment. In present bias models, however, agent sophistication does not vary over time, whereas in our model partially sophisticated agents are only intermittently fully attentive to prices because salience is time-varying. Evidence for sophistication in models of present bias has been documented in Ashraf et al. (2006), Kaur et al. (2010), and Augenblick et al. (2015).

Denote the consumption plan defined by (11) by

$$\hat{x}_{tM} = \hat{x}_{tM}(x_{t-1,M}, p; \theta_i), \quad \hat{y}_{tM} = \hat{y}_{tM}(x_{t-1,M}, p; \theta_i), \tag{12}$$

where this is the plan made during inattentive period i of the current cycle, which the naive agent expects to carry out through all future periods and all future cycles. Note that this depends on the θ_i from the period when the plan was made. In each successive period as θ declines, the sequence of actual decisions is given by

$$\hat{x}_{tM} = \hat{x}_{tM}(x_{t-1,M}, p; \theta_t), \quad \hat{y}_{tM} = \hat{y}_{tM}(x_{t-1,M}, p; \theta_t). \tag{13}$$

2.2.2. Partially sophisticated agents

We suppose that partially sophisticated agents anticipate attention decay, but their sophistication is limited in two ways. First, partially sophisticated agents in our model do not anticipate the beginning of the next cycle when prices become salient again. Second, they do not have perfect foresight of the attention decay function even as they expect themselves to be less attentive in the future. The framework we propose can be easily extended to relax these two assumptions.

Specifically, we assume partially sophisticated agents expect to apply $\underline{\theta} \le \theta_i$ in future decision making periods,[12] so that their anticipated demands are given by

$$\hat{x}_{tM} = \hat{x}_{tM}(x_{t-1,M}, p; \underline{\theta}), \quad \hat{y}_{tM} = \hat{y}_{tM}(x_{t-1,M}, p; \underline{\theta}).$$

Further, by recursively substituting $\hat{x}_{t-1,M}(x_{t-2,M}, p; \underline{\theta})$ we obtain

$$\hat{x}_{tM} = \hat{x}_{tM}(x_{i1}, p; \underline{\theta}), \quad \hat{y}_{tM} = \hat{y}_{tM}(x_{i1}, p; \underline{\theta}), \quad \text{for all } t > i, \text{ and for all } M. \tag{14}$$

The agent acts as if the price is $\theta_i p$ in period i, and expects to make future consumption decisions according to (14). In addition, she considers how the current consumption decision will shift those future demand functions through habit persistence. Let the present value of the partially sophisticated agent's optimal consumption plan be given by

$$\begin{aligned} V_S(\theta_i) = \max_{(x_{i1}, y_{i1})} \; & U(x_{i1}, x_{i-1,1}, y_{i1}) + \sum_{t=i+1}^{T} \beta^{t-i} U(\hat{x}_{t1}(x_{i1}; \underline{\theta}), \hat{x}_{t-1,1}(x_{i1}; \underline{\theta}), \hat{y}_{t1}(x_{i1}; \underline{\theta})) \\ & + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} U(\hat{x}_{tM}(x_{i1}; \underline{\theta}), \hat{x}_{t-1,M}(x_{i1}; \underline{\theta}), \hat{y}_{tM}(x_{i1}; \underline{\theta})) \\ & - \lambda \Bigg[(\theta_i p x_{i1} + y_{i1}) + \sum_{t=i+1}^{T} \beta^{t-i} (\theta_i p \hat{x}_{t1}(x_{i1}; \underline{\theta}) + \hat{y}_{t1}(x_{i1}; \underline{\theta})) \\ & + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} (\theta_i p \hat{x}_{tM}(x_{i1}; \underline{\theta}) + \hat{y}_{tM}(x_{i1}; \underline{\theta})) - A_{i1} \Bigg] \\ & s.t. \; x_{0,1}, \ldots, x_{i-1,1}, A_{i1} \; given \end{aligned} \tag{15}$$

where the $\hat{x}_{tM}(\cdot)$ and $\hat{y}_{tM}(\cdot)$ are given in (14).

The first order condition for x_{i1} is

$$\begin{aligned} \left[u'_{i1} - \alpha\beta u'_{i+1,1}\right] + \Bigg(& \sum_{t=i+1}^{T} \beta^{t-i} \frac{\partial \hat{x}_{t1}}{\partial x_{i1}} \left[u'_{t1} - \alpha\beta u'_{t+1,1} - \theta_i p\right] \\ & + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} \frac{\partial \hat{x}_{tM}}{\partial x_{i1}} \left[u'_{tM} - \alpha\beta u'_{t+1,M} - \theta_i p\right] \Bigg) = \theta_i p \end{aligned} \tag{16}$$

The first term in brackets is the marginal utility across time taking consumption persistence into account. If the agent did not demand commitment to avoid future inattention, this would be equal to the price during period 1 (when prices are fully salient) or the perceived price $\theta_i p$ during period i. This commitment demand is captured by the effects of x_i on each future $\hat{x}_t$. Each term in brackets inside the large parentheses is negative. These terms measure the extent of the deviation from the current plan in

[12] This assumption means that partially sophisticated agents expect attention to fully decay before the next period. In this sense agents are overly cautious about future overconsumption. We could allow them to better forecast their attention decay at the cost of substantial algebraic complication. The results remain qualitatively unchanged.

future periods due to salience effects. Because the agent is overconsuming in the future, the marginal utilities in the future fall below the current perceived price. The partially sophisticated agent anticipates these deviations, albeit imperfectly.

By recursive substitution of (16) we obtain

$$u'_{i1} - \alpha\beta u'_{i+1,1} = p\left(\theta_i + (\theta_i - \underline{\theta})\delta\right) \tag{17}$$

where $\delta = \sum_{t=i+1}^{T} \beta^{t-i} \frac{\partial \hat{x}_{t1}}{\partial x_{i1}} + \sum_{M=2}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-i} \frac{\partial \hat{x}_{tM}}{\partial x_{i1}}$ is the discounted impact of today's consumption decision on all future demand functions.[13] With quasilinear utility δ is a positive constant because terms $\frac{\partial \hat{x}_{tM}}{\partial x_{i1}}$ are constant demand shifters. If there is no consumption persistence (or no commitment mechanism by which to avoid salience effects, i.e., $\alpha = 0$ so $\frac{\partial \hat{x}_{tM}}{\partial x_{i1}} = 0$) then $\delta = 0$.

The partially sophisticated agent uses persistence as a commitment device to reduce future overconsumption, so their demand may be lower than the privately optimal demand would be in the absence of salience effects. This occurs when the current salience θ_i is close to 1, or when the expected decline between current salience and anticipated $\underline{\theta}$ is large, or if δ is large so that current decisions can have a large commitment effect. If the agent anticipates being equally price inattentive in the future so that $\theta_i = \underline{\theta}$, or if there is no commitment so that $\alpha = 0$ and $\delta = 0$, then we have the naive agent's problem. Further if prices are always fully salient, we have a standard habit formation problem with the price on the right hand side of the first order condition.

2.2.3. *Summary of privately optimal consumption paths*

In the absence of internalities from inattention, equating the left hand side of (17) to the price in all time periods will yield the true privately optimal consumption path. By contrast, partially sophisticated agents underconsume relative to the private optimum when θ_i is close to one or far from $\underline{\theta}$ and overconsume as θ_i declines toward $\underline{\theta}$, so that consumption rises throughout the cycle. Comparing the partially sophisticated agent conditions in (17) to those for the naive agent in (10), notice that they differ by the presence of $(\theta_i - \underline{\theta})\delta$. Consumption also rises throughout the cycle for naive agents, but they consume at the true private optimum in the first period and then overconsume in all other periods.

2.3. *Optimal dynamic corrective taxes with sophistication and commitment*

The social planner's problem from (4) is modified only to include habit persistence in household utility. The optimal path in this case is defined by the first order conditions:

$$u'_t - \alpha\beta u'_{t+1} = p + \varphi, \quad \text{for all} t. \tag{18}$$

With habit persistence, a tax at the standard Pigouvian rate of $\tau = \varphi$ would induce households to consume at the social optimum if prices were always perfectly salient. By recursive substitution of (18), we obtain

$$u'_t \approx (p + \varphi) \sum_{s=0}^{\infty} (\alpha\beta)^s = \frac{p + \varphi}{1 - \alpha\beta}. \tag{19}$$

The socially optimal per-period marginal utility is increased (and per-period consumption is restrained) by consideration of both the external costs φ and the recursive impact of today's consumption on future consumption through habit persistence, or the effect of α. Specifically if habit persistence is greater (larger α) then less is consumed in each period along the socially optimal path.

Each agent type could be induced to consume along the optimal path by a dynamic tax that transforms their private first order conditions in each period into the social planner's first order conditions from (19).

For naive agents, the optimal tax solves

$$u'_t = \frac{(p + \tau_N(\theta_t)) \cdot \theta_t}{1 - \alpha\beta} = \frac{p + \varphi}{1 - \alpha\beta}, \tag{20}$$

and for partially sophisticated agents the optimal tax solves

$$u'_t = \frac{(p + \tau_S(\theta_t)) \cdot \left(\theta_t + (\theta_t - \underline{\theta})\delta\right)}{1 - \alpha\beta} = \frac{(p + \varphi)}{1 - \alpha\beta}. \tag{21}$$

[13] Our notion of partial sophistication entails that the agent does not anticipate the salient price signal at the beginning of the next cycle, but optimizes the habit stock over all future periods. Alternative assumptions can be easily accommodated. If the agent anticipates the next salient price signal, then the double summation in (15) can be replaced by a continuation value that depends on x_{i1} and appropriate modifications to the first order conditions and definition of δ. If the agent only optimizes over the remainder of the cycle, then the double summation in equations (15) and (16) can be dropped, and δ will include only the sum from $i + 1$ to T. Optimizing over a fixed number of future periods could be handled similarly. In each of these cases the qualitative results of the paper hold and the only quantitative changes occur through the size of the δ parameter.

The tax that solves equations (20) and (21) is stated in the following proposition.

Proposition 2. *The optimal corrective taxes are time-varying and their magnitudes depend on the forward-looking behavior of the households, i.e., whether households are naive (N) or partially sophisticated (S), as well as how recently a salient price signal was received. Let $j \in \{N, S\}$ denote naive and partially sophisticated agent types, respectively, and let δ_j=0 for naive agents and $\delta_j = \delta$ for partially sophisticated agents. Then*

$$\tau_{tj}(\theta_t) = p\frac{1-(\theta_t+(\theta_t-\underline{\theta})\delta_j)}{\theta_t+(\theta_t-\underline{\theta})\delta_j} + \frac{\varphi}{\theta_t+(\theta_t-\underline{\theta})\delta_j} <> \varphi. \tag{22}$$

Several comments about the optimal dynamic corrective taxes are in order. First, the optimal tax in each period is again a weighted combination of externality correction and price salience internality correction. Now, however, the weights depend on the agent's sophistication level and degree of persistence through δ_j. This can be seen most clearly during the first period (or if otherwise $\theta = 1$) when the taxes are

$$\tau_N(\theta_1 = 1) = \varphi$$
$$\tau_S(\theta_1 = 1) = p\frac{-(1-\underline{\theta})\delta}{1+(1-\underline{\theta})\delta} + \varphi\frac{1}{1+(1-\underline{\theta})\delta} < \varphi.$$

The naive agent's tax is equal to marginal external damage when prices are salient. The partially sophisticated agent tries to avoid future internalities from inattention by reducing consumption when prices are salient, so $\tau_S(\theta_1)$ (and $\tau_S(\theta_i)$ for period i soon enough after period 1) is less than the marginal social damage and can even be a subsidy.

Second, the importance of the inattention parameter varies across time and agent type. For naive agents, the taxes in the initial fully salient period do not depend on the future values of θ; at $t = 1$ naive agents *do not plan* to be inattentive. Therefore, they do not demand commitment, and the optimal tax cannot influence their demand for such commitment.

A smaller θ_t implies a larger tax for both agent types, especially during later periods as price salience declines. During earlier periods when the price is more salient, however, the relationship between θ_t and the optimal tax depends on the sophistication of the agent. A smaller θ_t requires an unambiguously larger tax for naive agents during these earlier periods, but for partially sophisticated agents the gap between the current perceived price and anticipated future inattention ($\theta_t - \underline{\theta}$), and the anticipated effect of today's consumption on all future demand functions (δ), determine the tax. Partially sophisticated agents anticipate greater future overconsumption the smaller is $\underline{\theta}$, and further reduce consumption in earlier periods. This reduced consumption requires a smaller tax during the earlier periods when the price is more salient.

Third, the optimal tax for both agent types increases over the cycle and converges to the same tax as θ_t converges to $\underline{\theta}$. The optimal tax increases because the sensitivity to the tax-inclusive price declines with salience, and because the price internality grows as salience declines. The internalities exacerbate the externality because they lead to excessive consumption of a dirty good. The optimal corrective tax on the partially sophisticated agent is therefore never larger than the tax for the naive agent:

$$\begin{aligned} &\tau_N(\theta_1) = \varphi > \tau_S(\theta_1), \\ &\tau_N(\theta_t) > \varphi > \tau_S(\theta_t), \quad t = 2, \ldots, I \text{ for some } I < T, \\ &\tau_N(\underline{\theta}) = \tau_S(\underline{\theta}) > \varphi, \qquad t = I+1, \ldots, T. \end{aligned} \tag{23}$$

Fourth, the relationship between habit persistence (α) and the optimal tax differs across agent types and across time within the billing cycle. The optimal tax balances correcting for the inattention internality and the pollution externality which are both exacerbated by persistence, but the tax does so differently for the two agent types. The regulator can put the naive agent on the optimal path when the salient price signal is received ($t = 1$) by just correcting the externality, and then keep her on the optimal path by adjusting the tax for inattention in later periods. For the partially sophisticated agent, greater habit persistence implies more effective commitment to a lower consumption path which exacerbates underconsumption in earlier more salient periods. The optimal tax then starts lower and grows more over time when persistence is stronger.

We summarize the relationships between the persistence and salience parameters and the optimal taxes across different time periods and agent types in the following corollary to Proposition 2:

Corollary 2.1. *For the naive household:*

- *the tax in all periods is independent of the strength of habit persistence:*

$$\frac{\partial \tau_N(\theta_t)}{\partial \alpha} = 0, \quad \textit{for all } t.$$

- *the greater the inattention (the smaller is θ_t), the larger the tax during inattentive periods, with the absolute magnitude of the effect increasing as θ_t declines:*

$$\frac{\partial \tau_N(\theta_t)}{\partial \theta_t} < 0, \ \frac{\partial^2 \tau_N(\theta_t)}{\partial \theta_t^2} > 0 \ \textit{for all } t.$$

For the partially sophisticated household:

- *stronger habit persistence (larger α and larger δ) reduces the tax when $(\theta_t - \underline{\theta}) > 0$. This effect is stronger the more salience is expected to decline in the future (the greater is $\theta_t - \underline{\theta}$):*

$$\frac{\partial \tau_S(\theta_t)}{\partial \alpha} < 0, \; \frac{\partial^2 \tau_S(\theta_t)}{\partial(\theta_t - \underline{\theta})\partial \alpha} < 0 \text{ for } t = 1, \ldots, I,$$
$$\frac{\partial \tau_S(\theta_t)}{\partial \alpha} = 0 \text{ for } t = I+1, \ldots, T,$$

- *the greater the salience in period t (the larger is θ_t), the smaller the tax in t. The greater the expected future salience as of period t (the larger is $\underline{\theta}$), the larger the tax in t. The more salience is expected to decline in the future salience (the greater is $\theta_t - \underline{\theta}$), the smaller the tax in the current period t:*

$$\frac{\partial \tau_S(\theta_t)}{\partial \theta_t} < 0, \; \frac{\partial \tau_S(\theta_t)}{\partial \underline{\theta}} > 0, \; \frac{\partial \tau_S(\theta_t)}{\partial(\theta_t - \underline{\theta})} < 0 \text{ for all } t.$$

2.4. Time-invariant tax alternatives

For a variety of pragmatic reasons, including administrative costs and political opposition, it may not be feasible to implement a time-varying tax. The most obvious candidate for a time-invariant tax is the standard Pigouvian rate of marginal external damage, φ. This tax would ignore the salience internality, however. We derive a second-best tax that maximizes social welfare subject to the constraint that the tax is time-invariant. This second-best constant tax is a weighted average of the optimal dynamic tax in each period with weights that depend on the behavioral parameters and the internalities associated with each agent type. In the following subsection, we show how to quantify the welfare losses for *any* constant tax relative to the dynamic tax optimum.

Let $\widehat{W}\left(\widehat{\mathbf{x}}_j, \widehat{\boldsymbol{y}}_j\right) = \widehat{W}_j\left(x_0, p; \{\theta_t\}\right)$ be the social welfare function evaluated at the privately optimal consumption paths for $j \in \{N, S\}$. Suppose a small tax τ is added to a previously untaxed environment so that we evaluate $\widehat{W}_j(x_0, p+\tau; \{\theta_t\})$. The second-best optimal constant tax solves $\frac{\partial \widehat{W}}{\partial \tau} = 0$ for each type of agent, the result of which is stated as Proposition 3. One result from Proposition 3 is that for each agent type the second-best tax is an exact-weighted average of the optimal dynamic taxes where the weights depend crucially on the sophistication of the agent and the within-period demand curve slope. These slopes are constant functions of inattention and persistence parameters, the price p, the habit stock at time t, and utility parameters. For notational convenience define

$$\frac{\partial \widehat{x}_{tj}}{\partial \tau} = a_{tj}, \; j \in \{N, S\}.$$

Proposition 3. *Let δ_j=0 for naive agents and δ_j= δ for partially sophisticated agents. The second-best constant tax rate for $j \in \{N, S\}$, is given by the weighted average*

$$\tau_j^\star = \sum_{t=1}^{T} w_{tj} \tau_j(\theta_t), \tag{24}$$

with weights

$$w_{tj} = \frac{\beta^{t-1} a_{tj}(\theta_t + (\theta_t - \underline{\theta})\delta_j)}{\sum_{s=1}^{T} \beta^{s-1} a_{sj}(\theta_s + (\theta_s - \underline{\theta})\delta_j)}, \; t = 1, \ldots, T,$$

and $\tau_j(\theta_t)$ defined in Proposition 2.[14]

A Proof is provided in Appendix A. The weights in $\tau_j^\star$ are the discounted demand curve slopes, adjusted by the salience applied to each period's decision. The second-best tax therefore differs between agent types for two reasons. First, the weights differ between agent types; the demand curves have different slopes because the naive agent plans for persistence but not inattention, and the partially sophisticated agent plans for both persistence and inattention. These different dynamic behaviors create different price-responsiveness within a given period. Second, the optimal dynamic taxes that are being weighted also differ by agent type.

The size of the second-best constant tax for the partially sophisticated agent is more influenced by the persistence parameters than the second best taxes for the naive agent because of the partially sophisticated agent's demand for commitment. If

[14] This tax can also be more compactly expressed as.

$$\tau_j^\star = \frac{\sum_{t=1}^{T} \beta^{t-1} a_{tj}(\theta_t + (\theta_t - \underline{\theta})\delta_j) \cdot \tau_j(\theta_t)}{\sum_{t=1}^{T} \beta^{t-1} a_{tj}(\theta_t + (\theta_t - \underline{\theta})\delta_j)}. \tag{25}$$

consumption is more persistent (if α is large), then the partially sophisticated agent has a more effective commitment device. This makes future demand more responsive to current consumption (i.e., the $\frac{\partial \hat{x}_t}{\partial x_i}$ are large, making δ large), and the agent makes larger corrections in early periods on her own. In this case, the constant tax must be weighted more towards correcting under-consumption in early periods. With a large δ, $\tau_S(\theta_i)$ in earlier periods i become more heavily weighted and also become smaller, so the constant tax $\tau_S^\star$ also becomes smaller. This does not necessarily mean that stronger habit persistence is better for welfare. Intuitively, with stronger habit persistence there is more adjustment in consumption across price salient and insalient time periods. Part of the consumption- and pollution-reducing functions of the second-best constant tax are weakened in order to prevent the partially sophisticated agent from consuming too little when prices are salient. A large δ will also occur if the discount factor is close to one (so that the agent adjusts her habits out of concern for the future and needs less of a tax incentive). Likewise if δ is small the emphasis in the constant tax can be shifted back towards reducing pollution and overconsumption in later periods, and the tax is larger.

If regulators recognize that price salience is a problem but not that it varies over time, they may implement a different, effectively "third-best" constant tax using the received wisdom from static models of tax salience (e.g., Chetty et al., 2009; Chetty, 2009). Such a tax would be third-best because it fails to account for time-varying salience, persistence, commitment, or sophistication, in addition to being constant. In a static model with constant salience, regulators would want to equate marginal utility to the social marginal cost $p + \varphi$ while agents equate marginal utility to the perceived tax-inclusive price $\theta(p + \tau_{stat})$. The τ_{stat} that solves this problem has the same form as the optimal tax in Proposition 1, except that salience is assumed constant:

$$\tau_{stat} = p\frac{1-\theta}{\theta} + \frac{\varphi}{\theta}. \tag{26}$$

This third-best static tax is larger than the second-best constant tax.

2.4.1. Efficiency cost of suboptimal taxes

We now characterize the welfare loss of a suboptimal tax with externalities, internalities, and habit persistence. We focus on the second-best constant tax $\tau_j^\star$, although the formulas apply to any suboptimal tax vector. Focusing on $\tau_j^\star$ puts a lower bound on the efficiency losses from a time-invariant tax. The magnitude of the losses depends on the persistence and salience parameters, as well as the level of sophistication of the agent. Understanding these relationships and the size of potential losses is important not only for welfare analysis of taxes, but also for alternative investment policies. For example, smart electricity grid technologies require large fixed costs but also provide real-time price information and facilitate goal setting or learning, which may alter price salience or habit persistence. These welfare comparisons are therefore important for cost-benefit analysis.

We use the concept of equivalent variation to derive an expression for excess burden. The derivation is presented in Appendix A. Recall that the reference (optimal) dynamic tax takes $I+1$ values in each cycle because $\tau_j(\theta_t) = \tau_j(\underline{\theta})$ for $t = I+1$ to T. These are each compared to the second-best constant $\tau_j^\star$ applied in the same period within the cycle, which we denote $\tau_{(r)j}^\star$ for period r. Let the deviation between the second-best constant tax applied in period r and the optimal period tax in the same period be denoted[15]

$$\Delta\tau_{rj} = \tau_{(r)j}^\star - \tau_j(\theta_r) \text{ for } r = 1, \ldots, I+1.$$

In Appendix A we prove following proposition:

Proposition 4. *The excess burden for agent type j under the second-best tax vector $\underline{\tau}^\star$ instead of the optimal tax vector $\underline{\tau}(\theta)$ is given by*

$$EB_j = -\frac{1}{2}\sum_{M=1}^{\infty}\sum_{t=1}^{T}\beta^{T(M-1)+t-1}\sum_{s=1}^{I+1}\sum_{r=1}^{I+1}\Delta\tau_{sj}\Delta\tau_{rj}\left(\frac{\partial P_{tM}}{\partial x_{t-1,Mj}}\frac{dx_{t-1,Mj}}{d\tau_{sj}} + \frac{\partial P_{tM}}{\partial x_{tMj}}\frac{dx_{tMj}}{d\tau_{sj}} + \frac{\partial P_{tM}}{\partial x_{t+1,Mj}}\frac{dx_{t+1,Mj}}{d\tau_{sj}}\right)\frac{dx_{tMj}}{d\tau_{rj}}. \tag{27}$$

where $P_t(x_{t-1}, x_t, x_{t+1}) = u_t'(x_t - \alpha x_{t-1}) - \alpha\beta u_{t+1}'(x_{t+1} - \alpha x_t)$ is the agent's true intertemporal preferences for x_t, i.e., the welfare-relevant inverse demand at time t. This is the inverse demand curve along the optimal path if prices were always perfectly salient.

This expression is the net present value of a sequence of deadweight loss triangles. The contribution of Proposition 4 is in showing how the size of those triangles is distorted across time by salience and persistence, and in providing a formula for excess burden that is a function of empirical parameters for price salience, habit persistence and demand curves. Because of habit persistence, the base and height of these triangles in any given period are determined by demand rotations and shifts that occur in response to tax changes in other periods. If there were no habits ($\alpha = 0$) and perfect salience ($\theta_t = 1$ for all t)

[15] We also note that any suboptimal time-varying or constant tax could be used to calculate an analogous deviation and applied to the formula in Proposition 4 to calculate its excess burden. This also applies if the regulator is incorrect about agent type. Suppose agents are type j but the regulator implements the second-best constant tax vector for type j' agents, $\underline{\tau}_{j'}^\star$. We can calculate the excess burden by plugging in the differences between the type j' vector of second-best taxes ($\underline{\tau}_{j'}^\star$) and the type j optimal tax vector ($\underline{\tau}_j(\theta)$), i.e., substituting $(\tau_{(s)j'}^\star - \tau_j(\theta_s))(\tau_{(r)j'}^\star - \tau_j(\theta_r))$ for $\Delta\tau_{sj}\Delta\tau_{rj}$ in equation (27). If the regulator is uncertain about agent type rather than simply incorrect, then optimal and second-best taxes are a weighted average of the optimal and second-best tax for each type, weighted by the share of each type in the population. We leave the derivation of these taxes and their excess burden to future work.

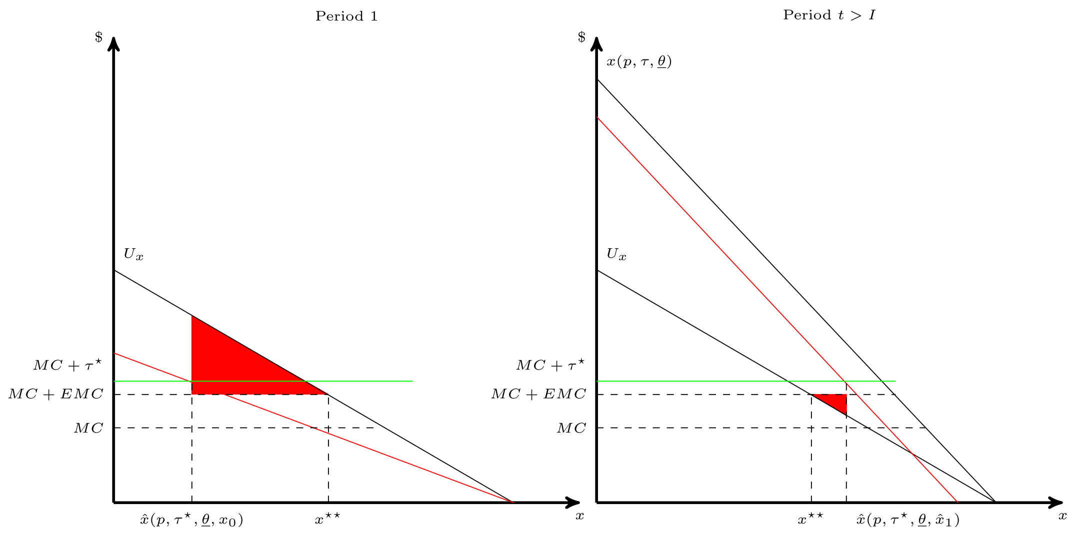

Fig. 1. Excess burden of the second-best constant tax for partially sophisticated agents.

then no such shifts or rotations would occur; in this case the deadweight loss in each period (dropping the M subscript for convenience) would be given by the middle terms in the parentheses in equation (27), which simplify to the standard Harberger (1964) formula:

$$-\frac{1}{2}\Delta\tau_{tj}^2\frac{\partial P_t}{\partial x_{tj}}\left(\frac{\partial x_{tj}}{\partial \tau_{tj}}\right)^2$$

With imperfect salience the ratio of the welfare-relevant demand slope to the decision-relevant demand slope is the expected salience parameter in period t, $\frac{\partial P_t}{\partial x_{tj}}\frac{\partial x_{tj}}{\partial \tau_{tj}} = \theta_t$. If we have inattention without habit persistence, then when prices are not salient these deadweight loss triangles are simply scaled by the salience parameter:

$$-\frac{1}{2}\theta_t\Delta\tau_{tj}\frac{\partial x_{tj}}{\partial \tau_{tj}}\Delta\tau_{tj}$$

This is the static formula derived in Chetty (2009), and shows that price insalience lowers the deadweight loss from taxation because it makes decision-relevant demand less elastic and reduces the quantity distortion. For our model with externalities it may seem counterintuitive that price insalience seems to reduce deadweight loss. This is relative to the socially optimal (rather than privately optimal) consumption path, however, and this is for $\Delta\tau$ of a *given size*. A given deviation from optimal taxes causes a smaller quantity distortion if θ_t is small. It is also important to notice that if θ_t is small then the second best $\Delta\tau$ may be relatively large. This increases the size of the deadweight loss from an imperfect tax. For partially sophisticated agents when prices are not salient the deadweight loss triangles would be $-\frac{1}{2}\left(\theta_t + (\theta_t - \underline{\theta})\delta\right)\Delta\tau_{tj}\frac{\partial x_{tj}}{\partial \tau_{tj}}\Delta\tau_{tj}$ without habit persistence. Notice the expected salience term and the deviation from optimal taxes change over time in this expression. The remainder of the terms in equation (27) capture the additional shifts and rotations from the indirect effects of taxes applied in other periods that further distort the size of the deadweight loss triangles in a given period.

The deadweight loss for partially sophisticated agents is illustrated in Fig. 1 for the first, fully salient, period in the cycle and a later period of the cycle $t > I$ when salience has reached its minimum.

As drawn, the deadweight loss in period 1 is larger than in period t, but losses similar to those in period t are experienced over many periods in each cycle. The size of the loss triangles depends on the rotation and shift of the decision-relevant demand curves (the dashed lines). The total excess burden in (27) is increasing in the habit parameter α, holding taxes constant. The intuition is as follows. When the taxes during later, less salient periods are lowered from their optimal dynamic values to meet the second-best constant rate, overconsumption during these periods increases. In terms of Fig. 1, the heavy dotted line in the right panel is lower than the optimal tax-inclusive price for period t so consumption is higher than socially optimal for that period. If α is large, the partially sophisticated agent makes larger adjustments during more salient periods in order to reduce future overconsumption because commitment is more effective. This rotates the period 1 decision-relevant demand curve further downward (the dashed line in the left panel of Fig. 1). The more the demand curve rotates downward, the greater the underconsumption in period 1, and the greater the period 1 deadweight loss, which increases total losses from a suboptimal constant tax.

3. Numerical simulation

One useful feature of our proposed corrective tax structure is that all the parameters in the expressions are estimable or recoverable from recent studies in the literature. For the residential electricity consumption example that is the focus of our simulation, there is a large literature that estimates the demand elasticity, as well as recent estimates of habit persistence and price salience. Marginal social damages of various pollutants from the electricity sector are available from, among others, Graff Zivin, Kotchen, and Mansur (2014). In this section we perform a numerical simulation in order to demonstrate the magnitude of the optimal dynamic and second best corrective taxes, and the excess burden of failing to implement optimal dynamic taxes.

In order to calculate numerical values for the taxes and the excess burden we need to derive specific functional forms for the demand functions for each agent type, and for the slope of the welfare-relevant demand. The complete expressions for these total and partial derivatives, and for the excess burden, are derived in Appendix B. Here we briefly outline the characteristics of the demand curves that are used in these expressions before discussing the calibration and results. We consider a version of the model in which utility is quadratic. Let utility in period s be given by

$$U_s = a(x_s - \alpha x_{s-1}) - \frac{1}{2}b(x_s - \alpha x_{s-1})^2 + y_s.$$

The welfare-relevant inverse demands found by solving the privately optimal plan are given by

$$P(x_{t-1}, x_t, x_{t+1}) = a(1-\beta\alpha) - b(1+\beta\alpha^2)x_t + b\alpha x_{t-1} + b\beta\alpha x_{t+1},$$

We also need to derive expressions for the total derivatives in the excess burden formula $\frac{dx_{tj}}{d\tau_s}$ for each agent type j and period s. These total derivatives require expressions for.

- $\frac{\partial \hat{x}_{tj}}{\partial \tau}$, the slope of the type j demand curve in each period t, and
- $\frac{\partial \hat{x}_{tj}}{\partial x_{rj}}$, the shift in the type j demand curve in period t because of a change in consumption in period r that persists from the past or is anticipated in the future.

In Appendix B, we derive these expressions using the demand curves for each agent type. Quadratic utility produces the following linear (decision-relevant) demand function:

$$x_{tj} = \frac{a(1-\beta\alpha)}{b(1+\beta\alpha^2)} - \frac{(\theta_t + (\theta_t - \underline{\theta})\delta_j)}{b(1+\beta\alpha^2)}p + \frac{\alpha}{1+\beta\alpha^2}x_{t-1} + \beta\frac{\alpha}{1+\beta\alpha^2}x_{t+1}, \tag{28}$$

where $\delta_j = 0$ for naive agents and $\delta_j = \delta$ for partially sophisticated agents.

Equation (28) gives the familiar result that consumption is less responsive to prices when they are not fully salient. If prices are constant, then insalient prices, i.e. periods when θ_t is less than one, imply an upward shift in the intercept of the consumption function, which is consistent with the finding in Gilbert and Graff Zivin (2014).[16] Ignoring the salience factor, this specification is equivalent to the Becker et al. (1994) rational habit demand function. If there is no consumption persistence ($\alpha = 0$) then current consumption also does not depend on past or future consumption.

In a model without consumption persistence or inattention, the response to a price change in the current period is $-1/b$. This is dampened by $1/(1+\beta\alpha^2)$ as agents anticipate the persistence of today's consumption change into the future. For a naive agent the response to current prices is also dampened by salience θ_t. For a partially sophisticated agent, the response to current prices is augmented by $(\theta_t - \underline{\theta})\delta$ as the agent tries to commit to reduced consumption.

3.1. Numerical calibration

We first provide parameters for a central case of the model for U.S. electricity consumption, and then perform sensitivity analysis for key parameters. The parameters for the central case are given in Table 1.

The habit parameter can be calculated from empirical studies on habit formation in energy consumption. Filippini et al. (2018) estimate a coefficient on lagged electricity consumption of 0.422. Using the derivations above, we solve $0.422 = \frac{\alpha}{1+\beta\alpha^2}$ for α, using the root that falls between zero and one.[17] The daily discount factor β was chosen to produce an annual discount factor of 0.9.

The marginal damage parameter is calculated by combining an estimate of the national average marginal CO_2/kwh emissions with an estimate of the marginal social cost of CO_2. Graff Zivin et al. (2014) estimate a national average marginal emissions rate

[16] We assume a constant price in order to isolate the effect of price salience rather than price uncertainty. This is also consistent with our empirical context because electricity prices are usually fixed, regulated rates.

[17] Using Scott (2012) coefficient of 0.787 on lagged gasoline consumption and solving for the real part of the complex roots produces $\alpha = 0.64$, while using Heien and Durham (1991) coefficient of 0.418 on lagged electricity produces $\alpha = 0.54$. Macroeconomic studies using aggregate data (e.g., Fuhrer (2000)) typically estimate the habit parameter closer to 0.8.

Table 1
Calibration parameters.

Parameter		Value
Discount factor	β	0.9997
Habit persistence	α	0.55
Salience	$\underline{\theta}$	0.71
Marginal damage	φ	\$0.03/kwh
Electricity price	p	\$0.129/kwh
Average consumption	$\overline{x}$	29.5 kwh/day
Demand elasticity	$\epsilon = \frac{dx_t}{dp}\frac{p}{\overline{x}}$	−0.1
Cycle length	T	30
Periods with $theta_t > \underline{\theta}$	I	10

of 1.21 lbs of CO_2/kwh. Inflated to 2017 dollars, the U.S. EPA's most recent central estimate of the social cost of carbon is \$49.66 per ton of CO_2 which we converted to pounds to arrive at an estimate of \$0.03/kwh, or about one fourth of the 2017 national average electricity price of \$0.129/kwh (EIA, 2017).[18]

We use the price elasticity of electricity demand, along with the national average daily household consumption and electricity price, in order to calculate an average linear demand slope. We assume that electricity prices are not salient for the majority of each billing month (Gilbert and Graff Zivin (2014)), and so we take our calculated slopes to represent the slope during insalient periods, despite the fact that they "average in" brief periods of salience following payment of a bill. Estimates of the elasticity of electricity demand vary in the literature. We do not wish to provide a full review of the electricity demand elasticity literature here, but as a few recent examples, Alberini and Filippini (2011) estimate short run elasticities between −0.8 and −0.15, Filippini et al. (2018) between −0.12 and −0.27, and Deryugina et al. (2020) of about −0.16. We take the low end of this range and use −0.1 because we have very short run (daily) behavior in mind for our model.

The minimum salience parameter $\underline{\theta}$ is calculated using Sexton (2015) estimate that households that pay their electricity bill using automatic withdrawal consume 4% more energy than houses that receive a bill and take an action to make a payment. We take the automatic withdrawal households as fully inattentive to price, and the active payment households as attentive to price. This likely overstates the minimum level of salience, and we evaluate the sensitivity of our results to this parameter. Using a demand elasticity of −0.1, the implied relative slope of the inattentive and attentive demand curves is 0.71. The astute reader will note that this estimate is larger than those found in the tax salience literature (e.g., Chetty et al. (2009) and Taubinsky and Rees-Jones (2017)), but those studies assume that only the tax portion of the tax-inclusive price is insalient.

Lastly, we assume a cycle length of $T = 30$ to reflect monthly billing. We assume that salience reaches a minimum $\underline{\theta}$ after $I = 10$ periods, and for simplicity we assume the decline to $\underline{\theta}$ is linear such that $\theta_t = 1 - \frac{i-1}{I-1}(1 - \underline{\theta})$. These assumptions are roughly consistent with the findings in Gilbert and Graff Zivin et al. (2014). The magnitudes of our results are similar over different choices for the length of I within a cycle.

3.2. Results

The time path of the optimal dynamic tax vector ($\underline{\tau}_j(\theta)$) for each agent type is shown in Fig. 2, with second-best constant taxes ($\tau_j^\star$) marked on the vertical axis. The magnitude of these taxes incorporating salience, relative to the size of the externality, is noteworthy. The limiting value of the dynamic optimal taxes, which in our setting is equivalent to what the optimal tax would be in a static setting, is more than three times as large as the marginal externality, and the second-best taxes are more than twice as large. The taxes are shown in dollars per unit of electricity, or kilowatthour. For reference, the U.S. average electricity price is \$0.129 per kilowatthour. As stated in Proposition 2, the taxes are rising as price salience declines, and they converge as salience reaches its minimum. The tax for the naive agent begins at the marginal externality of \$0.03 per kilowatthour in the first, fully salient period, and is always at least as large as the tax for the partially sophisticated agent. The tax for the partially sophisticated agent begins below the marginal externality and rises to exceed it as salience declines.

The main results for excess burden are shown in Table 2, including the numerical values of the taxes shown in Fig. 2. The excess burden of a time invariant tax is reported for the second-best constant tax, the static salience tax ignoring dynamic inefficiencies ($\tau_{stat} = 0.095$), and the standard Pigouvian rate of $\varphi = 0.03$. Excess burden is reported in present value dollars over the infinite time horizon. The excess burden of the second-best constant tax is larger for the partially sophisticated agent at \$49.91 per household, compared to \$18.93 per household for the naive agent. With approximately 126 million households in the U.S., this puts the welfare loss associated with moving from the optimal dynamic tax to the optimal, but second-best, constant tax at between \$2.4 billion and \$6.3 billion for the residential electricity sector.

Notice, however, that the excess burden increases significantly as the time-invariant tax moves away from it's second-best optimal level. When taxing at the static optimal rate of \$0.095, which ignores dynamic inefficiencies, the excess burden is \$30.28

[18] EPA (2013) reported the social cost of carbon as \$42 in the year 2020 (valued in 2007 dollars) at a 3% social discount rate. We converted this to 2017 dollars using the CPI.

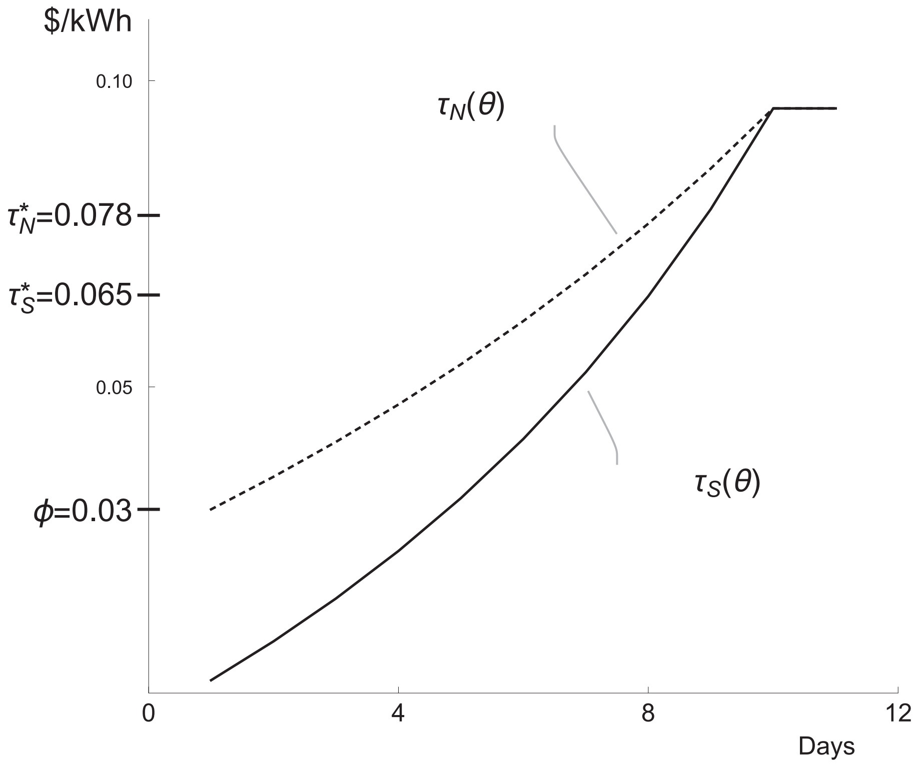

Notes: The dashed line is the optimal dynamic tax for the naive agent, and the solid line is the optimal dynamic tax for the partially sophisticated agent. Bold hash marks on the vertical axis denote second-best constant taxes for each type, as well as the marginal social damage, $\phi = 0.03$. Tax units are in dollars per kilowatthour of electricity; for example, ϕ is three cents per kilowatthour which is about one fourth of the national average electricity price of 12.9 cents per kilowatthour.

Fig. 2. Optimal dynamic and second-best taxes.

or \$90.61 per household if agents are naive or partially sophisticated, respectively. Whe taxing at the Pigouvian rate of \$0.03, which ignores internalities, the excess burden is \$84.84, or \$83.90 per household if agents are naive or partially sophisticated, respectively. Table 2 also reports excess burden figures for implementing no taxes at all, which are not surprisingly larger than any of the tax scenarios discussed. One way to interpret $EB(\varphi)$ in Table 2 is that even when polluting activities are priced at their marginal social cost, approximately \$84 per household, or more than 40 percent of the total welfare loss in $EB(0)$, is lost because of the salience internality.

Fig. 3 illustrates the findings of Corollary 2.1, on the sensitivity of optimal dynamic taxes to persistence (α) and salience ($\underline{\theta}$). The naive agent privately optimizes its habit stock in each period, so the tax only needs to correct for the salience internality and the externality; as a result the habit parameter has no effect on the size of the tax. The partially sophisticated agent also privately optimizes the habit stock, except that during attentive periods the agent is imperfectly committing to a new consumption path through habit persistence and the regulator must optimally address that commitment demand. If commitment (i.e., habit) is stronger, the regulator can reduce the tax during potentially attentive periods.

For both agent types, the tax can be lower during inattentive periods if inattention is less severe (i.e., if $\underline{\theta}$ is closer to one). During attentive periods the relationship between $\underline{\theta}$ and the optimal tax differs between agent types. The naive agent does not anticipate inattention in the future, so $\underline{\theta}$ only affects the tax during inattentive periods. The partially sophisticated agent underconsumes in early periods because of (imperfect) commitment to reduce future overconsumption. Because of early underconsumption, the tax starts below marginal damage in the first, perfectly salient period and rises to exceed marginal damage as the probability of being inattentive rises. The tax in early periods can therefore be larger (closer to marginal damage) if future inattention is less severe, as this alleviates underconsumption during the early periods.

Fig. 4 illustrates the sensitivity of excess burden to changes in the habit persistence and salience parameters. Darker shading corresponds to larger excess burden. The scale of the subfigures differs for each agent type because the excess burden increases so rapidly as the salience parameter declines, but at different rates for each agent type; it is difficult to represent sensitivity to $\underline{\theta}$ on the same scale. For both agent types, the major source of variation in excess burden is the salience parameter, because this is the main source of the internality. However, there is a slight decline in excess burden for the naive agent as persistence

Table 2
Excess burden and tax results for the base case.

Taxes	Naive	Partially Sophisticated
Excess burden of time-invariant taxes in present value $		
$EB(\tau_j^\star)$	18.93	49.91
$EB(\tau_{stat})$	30.28	90.61
$EB(\varphi)$	84.84	83.90
Excess burden when there is no tax		
$EB(0)$	198.4	181.4
Second-best constant taxes in $ per kwh		
$\tau_j^\star$	0.078	0.065
Optimal dynamic taxes in $ per kwh		
$\tau(\theta_1)$	0.030	0.002
$\tau(\theta_2)$	0.035	0.008
$\tau(\theta_3)$	0.041	0.015
$\tau(\theta_4)$	0.047	0.023
$\tau(\theta_5)$	0.054	0.032
$\tau(\theta_6)$	0.061	0.041
$\tau(\theta_7)$	0.068	0.052
$\tau(\theta_8)$	0.077	0.065
$\tau(\theta_9)$	0.086	0.079
$\tau(\theta_{10})$	0.095	0.095
$\tau(\theta_{I+1})$	0.095	0.095

Notes: This table reports the optimal dynamic taxes, second-best constant taxes, and excess burden of several time-invariant taxes using parameters from Table 1. The taxes are reported in dollars per unit of electricity (kilowatthour) while the excess burden values are reported in net present value dollars per household over an infinite horizon. The marginal externality is $\varphi = \$0.03$ per unit of electricity (kilowatthour) and the optimal tax for a static model is $\tau_{stat} = \$0.095$ per kilowatthour. For comparison, the U.S. average residential electricity price is $0.129 per kilowatthour.

increases. In Fig. 5 and Fig. 6 of Appendix B, we report similar results for the excess burden of the static salience tax and the Pigouvian tax.

Lastly, Table 3 reports the sensitivity of excess burden to the elasticity of demand. Specifically, the table shows that the excess burden for each of the time-invariant taxes and each of the agent types grows linearly as demand becomes more elastic. The welfare losses for a good that has unit elastic demand are ten times larger than the welfare losses in our residential electricity example assuming a demand elasticity of −0.1. Considering the variety of types of goods to which our model applies (e.g., gasoline, unhealthy food, etc.), the economy-wide losses from imperfect taxation and time-varying salience are likely to be non-trivial.

4. Conclusion

Economies across the globe are becoming increasingly cashless and many payment systems have become automated, driving a temporal wedge between consumption and payment and generally making the costs of consumption intermittently salient. Since this inconsistent price salience alters demand elasticities, it is a particular concern for goods that generate externalities and the price-based policies deployed to address them. This paper develops a simple model of consumer behavior when prices of a good that creates some social harm are intermittently salient. In this context, optimal corrective taxes are dynamic and heavily influenced by the persistence of consumption decisions across time and the level of sophistication of agents about salience. Consumption persistence is an empirical feature of intertemporal consumption for which there is a large volume of evidence across contexts such as habit formation, durable goods adoption, and status quo bias, to name a few. In this setting the forward-looking capacity of the agents shapes the nature of dynamic externalities and internalities, and is therefore crucial for optimal taxation and welfare analysis.

We derive an optimal dynamic tax schedule for two types of agents: a *naive* agent who does not anticipate future inattention but is forward looking about persistence; and a *partially sophisticated* agent who anticipates future inattention and plans for consumption persistence. This comparison allows us to illustrate the role of commitment on the part of partially sophisticated agents. In the naive case, the degree of persistence is irrelevant to the optimal tax schedule because there is no internality from persistence and also no demand for commitment. Lower salience still results in larger optimal taxes, however. In the partially sophisticated case, the agent demands commitment to avoid future suboptimal decisions when prices are salient, and consumption persistence provides a mechanism to partially achieve that commitment. Greater persistence therefore lowers the

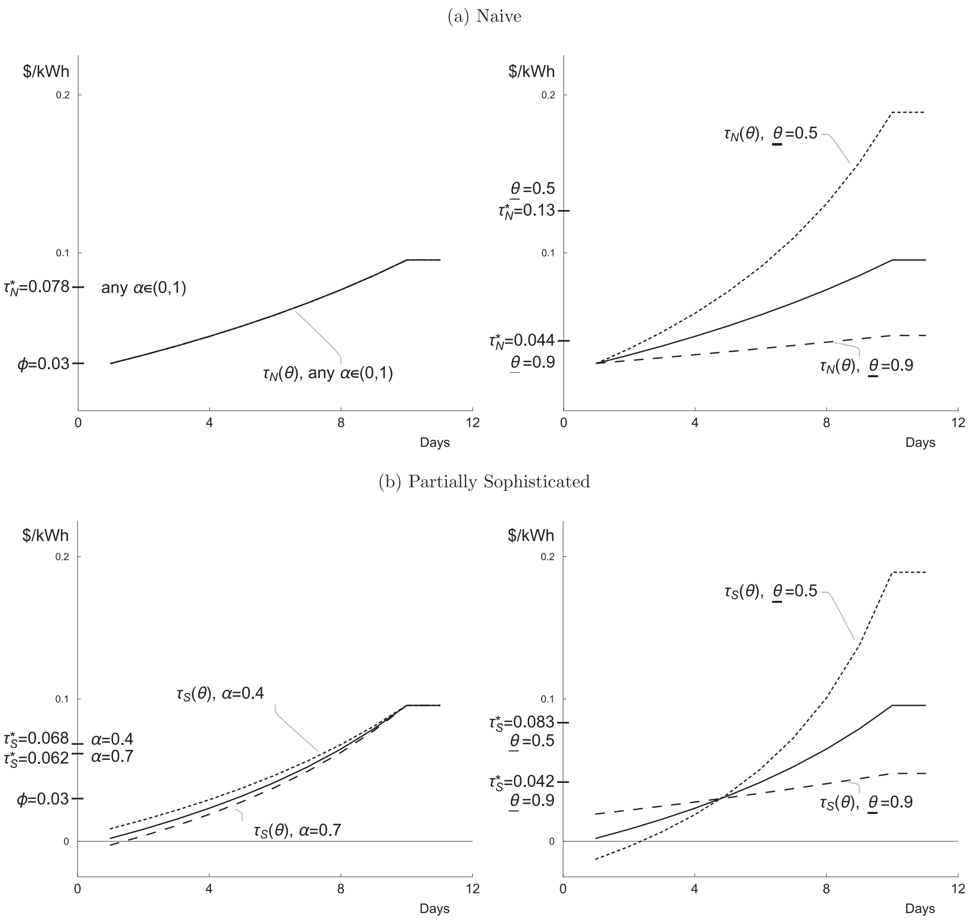

Notes: The solid black line on each graph is the optimal dynamic tax in the base case, with $\alpha = 0.55$ and $\underline{\theta} = 0.71$. Bold hash marks on the vertical axes denote second-best constant taxes in each case, as well as the marginal social damage, $\phi = 0.03$. Tax units are in dollars per kilowatthours of electricity; for example, ϕ is three cents per kilowatthour which is about one fourth of the national average electricity price of 12.9 cents per kilowatthour. For naive agents, optimal taxes are invariant to α but are larger with lower $\underline{\theta}$. For partially sophisticated agents, optimal taxes are lower during more salient periods if α is larger, whereas a lower $\underline{\theta}$'s causes the time path of optimal taxes to rotate down in the earliest periods and up in later periods.

Fig. 3. Sensitivity of optimal taxes to persistence α (left panels) and salience $\underline{\theta}$ (right panels).

taxes required to put the agent on the optimal consumption path, while lower salience results in greater variation in the tax schedule: lower taxes when agents are paying attention to prices, and higher optimal taxes when they are not.

Since time-varying taxes can be bureaucratically expensive and politically difficult to implement, we also explore the consumption and welfare implications under a variety of time-invariant tax strategies. In particular, we derive a second-best constant tax that maximizes welfare among time-invariant taxes. We also consider the welfare consequences of ignoring internalities and taxing at the Pigouvian rate of marginal external cost, as well as the welfare consequences of ignoring dynamics and imposing a tax that would be optimal in a static model with inattention and externalities.

The key insights from our model are then placed into context through a simulation exercise based on data from the U.S. residential electricity market. This market is attractive for a number of reasons, not least of which because there is well-documented persistence in consumption over time within households, price inattention, intermittent billing, and socially costly environmental externalities. Our base case results suggest that optimal taxes are significantly larger than those corresponding to the simple Pigouvian case. The magnitude of this difference is driven by the size of the internalities associated with inattention and persistence that vary by agent type. The optimal tax is more than double the Pigouvian rate. We also find that even the optimal time-invariant tax generates welfare losses that range from 2.4 to 6.3 billion dollars when aggregated across the US electricity market.

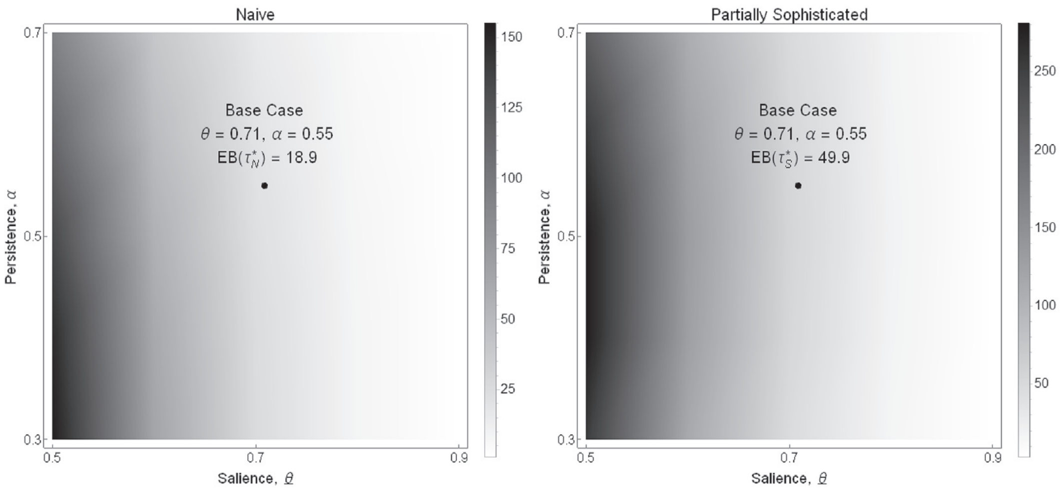

Notes: Excess burden is quantified in per household net present value dollars over an infinite time horizon, with darker shading indicating larger excess burden. Note that the scale is different for each agent type. This is necessary because the variation in excess burden across the parameter space and agent types is nonlinear. Salience is the main factor in the size of the excess burden of a constant tax for both agent types. However, excess burden does vary with persistence; excess burden is larger when persistence is lower for naive agents, whereas excess burden is concave in persistence for partially sophisticated agents.

Fig. 4. Excess burden of second-best constant tax ($\tau = \tau^\star$), Sensitivity to persistence α and salience $\underline{\theta}$.

Table 3
The sensitivity of excess burden to demand elasticity.

	Naive			Partially Sophisticated		
ϵ	$EB_N(\tau_N^\star)$	$EB_N(\varphi)$	$EB_N(\tau_{stat})$	$EB_S(\tau_S^\star)$	$EB_S(\varphi)$	$EB_S(\tau_{stat})$
−0.1	18.9	84.8	30.3	49.9	83.9	90.6
−0.2	37.9	170	60.6	99.8	168	181
−0.3	56.8	255	90.8	150	252	272
−0.4	75.7	339	121	200	336	362
−0.5	94.6	424	151	250	419	453
−0.6	114	509	182	299	503	544
−0.7	132	594	212	349	587	634
−0.8	151	679	242	399	671	725
−0.9	170	764	272	449	755	815
−1.0	189	848	303	499	839	906
−1.1	208	933	333	549	923	997
−1.2	227	1020	363	599	1010	1090
−1.3	246	1100	394	649	1090	1180
−1.4	265	1190	424	699	1170	1270
−1.5	284	1270	454	749	1260	1360

Notes: This table reports the excess burden of time invariant taxes as a function of the elasticity of demand, ϵ. Three time-invariant taxes are evaluated for each agent type: the second best constant tax ($\tau_j^\star$), the Pigouvian tax which ignores internalities (φ), and the static salience tax (τ_{stat}). In each case, excess burden scales linearly with the elasticity of demand.

The sensitivity of these results to alternative parameter assumptions is also explored. Since the excess burden from second-best tax strategies scales with the elasticity of demand, the welfare losses in other empirical contexts will very much depend on the nature of demand for the good in question. Goods with a unit elasticity of demand, for example, would incur welfare losses ten times that of our electricity example. The degree of price salience also plays a significant role in determining the welfare losses from second best time-invariant tax policies. A modest thirty percent change in the salience parameter generates an excess burden that is five to ten times larger than in our base case, underscoring the magnitude of heterogenous impacts across different consumption contexts.

Our study is not without its limitations. Our model of inattention is intentionally stylized. We do not allow price inattention to be determined by past decisions, the magnitude of relative prices, or other environmental cues that may determine attention in a particular context. While our conceptualization of consumption persistence is general, it abstracts from important aspects of persistence that could arise due to complementary durable good purchases. Efforts to endogenize inattention and further formalize the persistence relationship constitute the next logical steps in advancing the insights from our model. On the empirical side, our simulation is hampered by the lack of data and empirical tests that could be used to distinguish naive from partially sophisticated agents. Understanding which agent type is most prevalent in a particular setting, or better yet how to target each of them through careful mechanism design, is critical for the design of good policy and an area ripe for future research.

Appendix A. Model Propositions

Proof of Proposition 3

Proof. By plugging the privately chosen consumption path for each agent[19] into the social welfare function and differentiating with respect to τ, we obtain

$$\frac{\partial \widehat{W}_j}{\partial \tau} = 0 = \sum_{M=1}^{\infty} \beta^{T(M-1)} \left[\sum_{t=1}^{T} \beta^{t-1} a_{tMj} \left[u'_{tM} - \alpha \beta u'_{t+1,M} - p - \varphi \right] \right] = 0$$

In the long run, behavior is repeated within each cycle, and the expression becomes

$$\frac{\partial \widehat{W}_j}{\partial \tau} = 0 = \frac{1}{1-\beta^T} \left[\sum_{t=1}^{T} \beta^{t-1} a_{tj} \left[u'_t - \alpha \beta u'_{t+1} - p - \varphi \right] \right]$$

If $j = N$, the first order conditions in (10) (modified to include the introduction of a small tax) can be used to simplify $\frac{\partial \widehat{W}_N}{\partial \tau}$ to

$$0 = \sum_{t=1}^{T} \beta^{t-1} a_{tN} \left[(p+\tau)\theta_t - p - \varphi \right]$$

Solving for τ and plugging in the formulas for $\tau_N(\theta_t)$ gives the result.

Similarly if $j = S$, the first order conditions in (10) and (17) (again modified to include the introduction of a small tax) can be used to simplify $\frac{\partial \widehat{W}_S}{\partial \tau}$ to

$$0 = \sum_{t=1}^{T} \beta^{t-1} a_{tS} \left[(p+\tau)(\theta_t + (\theta_t - \underline{\theta})\delta) - p - \varphi \right]$$

Solving for τ and plugging in the formulas for $\tau_S(\theta_t)$ gives the result. ■

Derivation of excess burden and Proof of Proposition 4

The loss from imperfect taxation is the net present value of the wealth society would be willing to forgoe in order to avoid the imperfect tax instrument (relative to the optimum), net of any changes in tax revenue between the two policies. Because the reference (optimal) tax level varies and takes $I+1$ values in each cycle, we will represent society's expenditure function and indirect welfare function as depending on a tax-inclusive price vector $\underline{\tau}(\theta)$ with $I+1$ elements. In the case of the second-best constant tax, all $I+1$ elements of the vector, denoted $\underline{\tau}^\star$, will be identical because the tax-inclusive price is a constant. Temporarily suppressing the agent type subscripts N and S, the general expression for the excess burden of deviating from the $(I+1)$-vector of optimal taxes $\underline{\tau}(\theta)$ to $\underline{\tau}^\star$ is

$$\begin{aligned} EB(\underline{\tau}^\star) = A_1 - e\left(p+\tau(\theta_1),\dots,p+\tau(\theta_{I+1}), \widehat{W}(p+\tau^\star,\dots,p+\tau^\star, A_1, x_0; \{\theta_t\})\right) \\ - \left(R(\tau^\star_{(1)},\dots,\tau^\star_{(I+1)}, A_1) - R(\tau(\theta_1),\dots,\tau(\theta_{I+1}), A_1)\right) \end{aligned} \tag{29}$$

Although $\tau^\star$ is a constant, in order to compare its application in different time periods with $\underline{\tau}(\theta) = (\tau(\theta_1),\dots,\tau(\theta_{I+1}))$, we denote $\underline{\tau}^\star = \left(\tau^\star_{(1)},\dots,\tau^\star_{(I+1)}\right)$ to be the vector containing $\tau^\star$ applied in the analogous time periods, with $\tau(\theta_{I+1})$ and $\tau^\star_{(I+1)}$ applied in $t = I+1,\dots,T$. The function $e(\cdot)$ is society's expenditure function. It is the amount of wealth at the optimal tax vector that would achieve the level of welfare obtained under the second-best constant tax.[20] The difference between $e(\cdot)$ and the net present value of income is the amount of wealth society would be willing to forgoe to retain the optimal tax structure. The last term in parentheses measures any gain or loss in tax revenue from moving from the optimal dynamic taxes to the alternative set of taxes.[21]

Following Auerbach (1985), we derive a more convenient expression for $EB(\underline{\tau}^\star)$ using a second-order Taylor expansion. We

[19] The privately chosen consumption path is implicit in (10) if $j = N$, and given in (13) and the solution to (16) if $j = S$.

[20] Note that we could use the same procedure to calculate the excess burden of any vector of sub-optimal taxes, including any other time-invariant tax option.

[21] For example.

$$R(\tau(\theta_1),\dots,\tau(\theta_{I+1}), A_1) = \sum_{M=1}^{\infty} \beta^{T(M-1)} \sum_{t=1}^{T} \beta^{t-1} \tau(\theta_t) \widehat{x}_{tM}(\tau(\theta_1),\dots,\tau(\theta_{I+1}), A_1, x_0; \{\theta_t\}).$$

calculate the expansion around the optimal dynamic tax vector $\underline{\tau}(\theta)$.

$$EB(\underline{\tau}^{\star}) \approx \frac{\partial EB}{\partial \underline{\tau}'}(\underline{\tau}(\theta)) \cdot \left(\underline{\tau}^{\star} - \underline{\tau}(\theta)\right) + \frac{1}{2}\left(\underline{\tau}^{\star} - \underline{\tau}(\theta)\right)' \cdot \frac{\partial^2 EB}{\partial \underline{\tau} \partial \underline{\tau}'}(\underline{\tau}(\theta)) \cdot \left(\underline{\tau}^{\star} - \underline{\tau}(\theta)\right) \tag{30}$$

The envelope theorem guarantees that the first-order terms in the Taylor expansion will be zero when evaluated at the optimal tax vector. The second order terms require taking derivatives of marginal utilities in each period, evaluated at the optimal level of demand (i.e., at demand when $\underline{\tau} = \underline{\tau}(\theta)$). The time inseparability of habit persistence requires that these derivatives be taken with respect to tax changes in the same and all other periods. Proposition 4 operationalizes the excess burden in equation (30) for our model, but first we introduce some simplifying notation.

In order to capture the marginal utilities underlying equation (30), let $P_t(x_{t-1}, x_t, x_{t+1}) = u_t'(x_t - \alpha x_{t-1}) - \alpha\beta u_{t+1}'(x_{t+1} - \alpha x_t)$ represent the agent's true intertemporal preferences for x_t, i.e., the welfare-relevant inverse demand at time t. This is the inverse demand curve along the optimal path if the agent was perfectly attentive. The decision-relevant demands will deviate from this welfare-relevant demand because of inattention.

For an optimizing agent, past and anticipated future quantities shift current demand. The assumption of rational habit formation means that both agent types adjust current demand in response to anticipated future tax changes because they expect the persistence of current consumption to affect their future tax burden. The partially sophisticated agent also anticipates being less attentive to the future tax-inclusive price; in order to avoid paying additional future taxes due to inattention, the agent further reduces consumption today. A future tax increase therefore results in an inward rotation in the early period demand curves for partially sophisticated agents. That inward rotation causes reduced consumption to persist through lower habit in future periods, causing future period demand curves to shift inward. We represent the accumulation of these effects from a tax change in some arbitrary period i on demand in some period t for agent type j by the total derivative

$$\frac{dx_{tMj}}{d\tau_{(i)j}}.$$

This total derivative captures direct effects of $\tau_{(i)}$ on x_t as well as indirect effects through x's chosen in different time periods that affect x_t through the habit stock. We can decompose this total derivative into direct and indirect effects depending on the relative position of t and i, and how relatively salient prices are at t and i within the billing cycle. The total derivative is used in the expression for excess burden in Proposition 4, and its decomposition may be helpful for understanding the way tax changes affect excess burden across time because of consumption persistence.

Consider the effect on demand in some arbitrary period and cycle tM of a tax change applied in period $i \le I$ of each cycle, holding the tax in every other period of every cycle constant. This effect will include a direct effect if $i = t$, as well as indirect effects that persist from changes in demand in the $i'th$ period of previous cycles, and indirect effects from anticipated demand changes in the $i'th$ period of future cycles:

$$\frac{dx_{tMj}}{d\tau_{(i)j}} = \begin{cases} \underbrace{\sum_{m=1}^{M-1} \frac{\partial x_{tMj}}{\partial x_{imj}} \frac{\partial x_{imj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects persisting from earlier} \\ \text{periods } i \text{ in previous cycles}}} + \underbrace{\frac{\partial x_{tMj}}{\partial \tau_{(i)j}}}_{\text{Direct effect}} + \underbrace{\sum_{n=M+1}^{\infty} \frac{\partial x_{tMj}}{\partial x_{inj}} \frac{\partial x_{inj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects anticipated from} \\ \text{changes in later cycles}}} & \text{for } t = i, \\ \underbrace{\sum_{m=1}^{M} \frac{\partial x_{tMj}}{\partial x_{imj}} \frac{\partial x_{imj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects persisting from earlier} \\ \text{periods } i \text{ in current and previous cycles}}} + \underbrace{\sum_{n=M+1}^{\infty} \frac{\partial x_{tMj}}{\partial x_{inj}} \frac{\partial x_{inj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects anticipated from} \\ \text{changes in later cycles}}} & \text{for } t > i, \\ \underbrace{\sum_{m=1}^{M-1} \frac{\partial x_{tMj}}{\partial x_{imj}} \frac{\partial x_{imj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects persisting from earlier} \\ \text{periods } i \text{ in previous cycles}}} + \underbrace{\sum_{n=M}^{\infty} \frac{\partial x_{tMj}}{\partial x_{inj}} \frac{\partial x_{inj}}{\partial \tau_{(i)j}}}_{\substack{\text{Indirect effects anticipated from later} \\ \text{periods } i \text{ in current and future cycles}}} & \text{for } t < i. \end{cases} \tag{31}$$

The $I + 1$'th element of the tax vector is applied in each period from $I + 1$ to T, so the direct effect on demand in each of those periods has an indirect effect on each other period in the current and all other cycles. If the period t of interest is after period I in a particular cycle, the total effect of a change in the $I + 1$'th element of the tax vector is

$$\frac{dx_{tMj}}{d\tau_{(I+1)j}} = \underbrace{\sum_{m=1}^{M-1} \left(\frac{\partial x_{tMj}}{\partial x_{I+1,mj}} \frac{\partial x_{I+1,mj}}{\partial \tau_{(I+1)j}} + \ldots + \frac{\partial x_{tMj}}{\partial x_{Tmj}} \frac{\partial x_{Tmj}}{\partial \tau_{(I+1)j}} \right)}_{\substack{\text{Indirect effects from insalient} \\ \text{periods in previous cycles}}}$$

$$+ \underbrace{\sum_{t'=I+1}^{t-1} \frac{\partial x_{tMj}}{\partial x_{t'Mj}} \frac{\partial x_{t'Mj}}{\partial \tau_{(I+1)j}}}_{\substack{\text{Indirect effects from past} \\ \text{insalient periods in current cycle}}} + \underbrace{\frac{\partial x_{tj}}{\partial \tau_{(t)j}}}_{\text{Direct effect}} + \underbrace{\sum_{t''=t+1}^{T} \frac{\partial x_{tMj}}{\partial x_{t''Mj}} \frac{\partial x_{t''Mj}}{\partial \tau_{(I+1)j}}}_{\substack{\text{Indirect effects from later} \\ \text{insalient periods in current cycle}}}$$

$$+ \underbrace{\sum_{n=M+1}^{\infty} \left(\frac{\partial x_{tMj}}{\partial x_{I+1,nj}} \frac{\partial x_{I+1,nj}}{\partial \tau_{(I+1)j}} + \ldots + \frac{\partial x_{tMj}}{\partial x_{Tnj}} \frac{\partial x_{Tnj}}{\partial \tau_{(I+1)j}} \right)}_{\substack{\text{Indirect effects from insalient} \\ \text{periods in future cycles}}}, \quad \text{for } t > I, \tag{32}$$

If the period t of interest is before period $I+1$ in a particular cycle, the total effect is similar except that there is no direct effect, and all indirect effects from the current cycle occur in the future. In this case the total effect of a change in the $I+1$'th element of the tax vector is

$$\frac{dx_{tMj}}{d\tau_{(I+1)j}} = \underbrace{\sum_{m=1}^{M-1} \left(\frac{\partial x_{tMj}}{\partial x_{I+1,mj}} \frac{\partial x_{I+1,mj}}{\partial \tau_{(I+1)j}} + \ldots + \frac{\partial x_{tMj}}{\partial x_{Tmj}} \frac{\partial x_{Tmj}}{\partial \tau_{(I+1)j}} \right)}_{\substack{\text{Indirect effects from insalient} \\ \text{periods in previous cycles}}}$$

$$+ \underbrace{\sum_{n=M}^{\infty} \left(\frac{\partial x_{tMj}}{\partial x_{I+1,nj}} \frac{\partial x_{I+1,nj}}{\partial \tau_{(I+1)j}} + \ldots + \frac{\partial x_{tMj}}{\partial x_{Tnj}} \frac{\partial x_{Tnj}}{\partial \tau_{(I+1)j}} \right)}_{\substack{\text{Indirect effects from insalient periods} \\ \text{later in current and future cycles}}}, \quad \text{for } t \le I. \tag{33}$$

Finally, let the deviation between the second-best constant tax applied in period r and the optimal period r tax be denoted

$$\Delta \tau_{rj} = \tau^{\star}_{(r)j} - \tau_{rj}(\theta).$$

Proof of Proposition 4

Proof. Write the derivative of excess burden with respect to the r'th element of the tax vector as

$$\begin{aligned} \frac{\partial EB_j}{\partial \tau_r} &= P_1(x_0, x_1, x_2) \frac{dx_{1j}}{d\tau_r} + \beta P_2(x_1, x_2, x_3) \frac{dx_{2j}}{d\tau_r} + \ldots \\ &= \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} P_{tM}(x_{t-1,Mj}, x_{tMj}, x_{t+1,Mj}) \frac{dx_{tMj}}{d\tau_r}. \end{aligned}$$

Taking second- and cross-partials of this expression gives

$$\frac{\partial^2 EB_j}{\partial \tau_s \partial \tau_r} = \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} \left(\frac{\partial P_{tM}}{\partial x_{t-1,Mj}} \frac{dx_{t-1,Mj}}{d\tau_{sj}} + \frac{\partial P_{tM}}{\partial x_{tMj}} \frac{dx_{tMj}}{d\tau_{sj}} + \frac{\partial P_{tM}}{\partial x_{t+1,Mj}} \frac{dx_{t+1,Mj}}{d\tau_{sj}} \right) \frac{dx_{tMj}}{d\tau_{rj}}.$$

Because the first order terms in the Taylor approximation to excess burden are zero at the optimal tax, we can rewrite (30) as

$$EB(\underline{\tau}^{\star}) \approx -\frac{1}{2} \left(\underline{\tau}^{\star} - \underline{\tau}(\theta) \right)' \begin{bmatrix} \frac{\partial^2 EB}{\partial \tau_1^{\star 2}} & \cdots & \frac{\partial^2 EB}{\partial \tau_1^{\star} \partial \tau_{I+1}^{\star}} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 EB}{\partial \tau_1^{\star} \partial \tau_{I+1}^{\star}} & \cdots & \frac{\partial^2 EB}{\partial \tau_{I+1}^{\star 2}} \end{bmatrix} \left(\underline{\tau}^{\star} - \underline{\tau}(\theta) \right) \tag{34}$$

Expanding equation (34) and plugging in the second- and cross-partial derivatives produces equation (27). ■

Equation (27) is useful for quantifying the social losses more generally if some tax vector $\underline{\widetilde{\tau}}$ is implemented that differs from $\underline{\tau}_j(\theta)$. For example, we can evaluate the excess burden from any time-invariant tax. If the regulator uses the conventional wisdom of taxing at marginal external damage $\varphi_{(s)}$ in each period s then the excess burden is simply

$$EB_{j,\varphi} = -\frac{1}{2} \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} \sum_{s=1}^{I+1} \sum_{r=1}^{I+1} (\varphi_{(s)} - \tau_{sj}(\theta))(\varphi_{(r)} - \tau_{rj}(\theta))$$

$$\cdot \left(\frac{\partial P_{tM}}{\partial x_{t-1,Mj}} \frac{dx_{t-1,Mj}}{d\tau_s} + \frac{\partial P_{tM}}{\partial x_{tMj}} \frac{dx_{tMj}}{d\tau_s} + \frac{\partial P_{tM}}{\partial x_{t+1,Mj}} \frac{dx_{t+1,Mj}}{d\tau_s} \right) \frac{dx_{tMj}}{d\tau_r}. \tag{35}$$

The losses from implementing alternative taxes can be calculated in a similar manner. For example, if the regulator uses some other tax vector we could substitute $\underset{\sim}{\tilde{\tau}}$ in for $\varphi_{(s)}$ in equation (35).

Appendix B. Simulation

In this appendix we derive an explicit expression for excess burden in the case of quadratic utility for use in our numerical simulation. In order to apply a specific functional form, it is convenient to rewrite the expression for excess burden from Proposition 4 as

$$EB_j = -\frac{1}{2} \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} \sum_{s=1}^{I+1} \sum_{r=1}^{I+1} \Delta\tau_{sj}\Delta\tau_{rj} \cdot u''_{tM} \left(\frac{dx_{tMj}}{d\tau_{sj}} - \alpha \frac{dx_{t-1,Mj}}{d\tau_{sj}} \right) \left(\frac{dx_{tMj}}{d\tau_{rj}} - \alpha \frac{dx_{t-1,Mj}}{d\tau_{rj}} \right). \tag{36}$$

To see that this is equivalent to the expression in Proposition 4, rewrite the derivative of the excess burden with respect to the r'th element of the tax vector as

$$\begin{aligned}
\frac{\partial EB_j}{\partial \tau_r} &= P_1(x_0, x_1, x_2) \frac{dx_{1j}}{d\tau_r} + \beta P_2(x_1, x_2, x_3) \frac{dx_{2j}}{d\tau_r} + \dots \\
&= (u'_1 - \beta\alpha u'_2) \frac{dx_{1j}}{d\tau_r} + \beta(u'_2 - \beta\alpha u'_3) \frac{dx_{2j}}{d\tau_r} + \dots \\
&= u'_1 \left(\frac{dx_{1j}}{d\tau_r} - \alpha \cdot 0 \right) + \beta u'_2 \left(\frac{dx_{2j}}{d\tau_r} - \alpha \frac{dx_{1j}}{d\tau_r} \right) + \dots \\
&= \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} u'_{tM} \left(\frac{dx_{tMj}}{d\tau_{rj}} - \alpha \frac{dx_{t-1,Mj}}{d\tau_{rj}} \right).
\end{aligned}$$

Taking second- and cross-partials of this expression gives

$$\frac{\partial^2 EB_j}{\partial \tau_s \partial \tau_r} = \sum_{M=1}^{\infty} \sum_{t=1}^{T} \beta^{T(M-1)+t-1} u''_{tM} \left(\frac{dx_{tMj}}{d\tau_{sj}} - \alpha \frac{dx_{t-1,Mj}}{d\tau_{sj}} \right) \left(\frac{dx_{tMj}}{d\tau_{rj}} - \alpha \frac{dx_{t-1,Mj}}{d\tau_{rj}} \right).$$

In the quadratic case, conveniently $u''_t = -b$. We now need to derive expressions for $\left(\frac{dx_{tMj}}{d\tau_s} - \alpha \frac{dx_{t-1,Mj}}{d\tau_s} \right)$ for each j and s. As mentioned in the text, these expressions will utilize.

- $\frac{\partial \hat{x}_{tj}}{\partial \tau}$, the slope of the type j demand curve in each period t, and
- $\frac{\partial \hat{x}_{tj}}{\partial x_{rj}}$, the shift in the type j demand curve in period t because of a change in consumption in period r that persists from the past or is anticipated in the future.

As discussed in the text, agents have the following linear demand:

$$x_{tMj} = \frac{a(1-\beta\alpha)}{b(1+\beta\alpha^2)} - \frac{(\theta_t + (\theta_t - \underline{\theta})\delta_j)}{b(1+\beta\alpha^2)} p + \frac{\alpha}{1+\beta\alpha^2} x_{t-1,Mj} + \beta \frac{\alpha}{1+\beta\alpha^2} x_{t+1,Mj}, \tag{37}$$

where $\delta_j = 0$ for naive agents and $\delta_j = \delta$ for partially sophisticated agents.

For notational convenience, define the response to past consumption as

$$\gamma = \frac{\alpha}{1+\beta\alpha^2}$$

and the absolute value of the slope of current period demand as

$$d'_{tMj} = \left| \frac{\partial x_{tMj}}{\partial \tau} \right| = \frac{(\theta_t + (\theta_t - \underline{\theta})\delta_j)}{b(1+\beta\alpha^2)}. \tag{38}$$

Current consumption is increasing in expected future consumption as well as past consumption. If the naive agent is attentive to prices in period t - but not sophisticated about inattention and salience - then they take into account how period t consumption affects period $t+1$ decisions, but not that period $t+1$ decisions will be suboptimal. In the partially sophisticated case the agent also considers expected impacts on future demand through δ. Recall that

$$\delta = \sum_{t=r+1}^{\infty} \beta^{t-r} \frac{\partial \widehat{x}_t}{\partial x_r}$$

As before, $\frac{\partial \widetilde{x}_{t+1}}{\partial x_t} > 0$ is a constant, and more specifically in the quadratic case,

$$\frac{\partial \widetilde{x}_{t+1}}{\partial x_t} = \gamma, \quad \frac{\partial \widetilde{x}_{t+2}}{\partial x_t} = \gamma^2, \quad \dots \text{ etc.},$$

whereas

$$\frac{\partial \widetilde{x}_t}{\partial x_{t+1}} = \beta\gamma, \quad \frac{\partial \widetilde{x}_t}{\partial x_{t+2}} = \beta^2\gamma^2, \quad \dots \text{ etc..}$$

We can therefore rewrite

$$\delta = \sum_{t=r+1}^{\infty} \beta^{t-r}\gamma^{t-r} = \frac{\beta\gamma}{1-\beta\gamma}$$

From equation (28) we obtain

$$\frac{dx_{tj}}{d\tau_r} = \begin{cases} -d'_{rj}\left(\gamma^{t-r}\sum_{m=1}^{M}\gamma^{T(m-1)} + (\beta\gamma)^{r-t}\sum_{m=M+1}^{\infty}(\beta\gamma)^{T(m-M)}\right) \text{ if } r \le I,\ r \le t \\ -d'_{rj}\left(\gamma^{t-r}\sum_{m=1}^{M-1}\gamma^{T\cdot m} + (\beta\gamma)^{r-t}\sum_{m=M}^{\infty}(\beta\gamma)^{T(m-M)}\right) \qquad \text{if } \ r \le I,\ r \ge t. \end{cases} \tag{39}$$

which can be simplified to

$$\frac{dx_{tj}}{d\tau_r} = \begin{cases} -d'_{rj}\left(\frac{\gamma^{t-r}(1-\gamma^{T\cdot M})}{1-\gamma^T} + \frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^T}\right) \qquad \text{if } r \le I,\ r \le t, \\ -d'_{rj}\left(\frac{\gamma^{t-r+T}(1-\gamma^{T(M-1)})}{1-\gamma^T} + \frac{(\beta\gamma)^{r-t}}{1-(\beta\gamma)^T}\right) \text{ if } r \le I,\ r \ge t. \end{cases} \tag{40}$$

Also from equation (28) we can obtain

$$\frac{dx_{tj}}{d\tau_{(I+1)}} = \begin{cases} -\frac{\underline{\theta}}{b(1+\beta\alpha^2)}\left(\frac{\gamma^t(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma} + \frac{(\beta\gamma)^{T+I+1-t}}{1-(\beta\gamma)^T}\frac{1-(\beta\gamma)^{T-I}}{1-\beta\gamma}\right. \\ \qquad \left. + \frac{\gamma(1-\gamma^{t-I-1})}{1-\gamma} + \frac{1-(\beta\gamma)^{T-t+1}}{1-\beta\gamma}\right) \qquad \text{if } I < t, \\ -\frac{\underline{\theta}}{b(1+\beta\alpha^2)}\left(\frac{\gamma^t(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma} + \frac{(\beta\gamma)^{I+1-t}}{1-(\beta\gamma)^T}\frac{1-(\beta\gamma)^{T-I}}{1-\beta\gamma}\right) \qquad \text{if } I \ge t. \end{cases} \tag{41}$$

From these expressions algebra shows that for $r \le I$,

$$\frac{dx_{tj}}{d\tau_r} - \alpha\frac{dx_{t-1,j}}{d\tau_r} = \begin{cases} -d'_{rj}\left(\frac{\gamma^{1-r+T}(1-\gamma^{T(M-1)})}{1-\gamma^T} + \frac{(\beta\gamma)^{r-1}}{1-(\beta\gamma)^T}\right) \qquad \text{if } \ r \ge 1, \ \ t = 1, \\ -d'_{rj}\left((\gamma-\alpha)\frac{\gamma^{t-1-r+T}(1-\gamma^{T\cdot M})}{1-\gamma^T} + (1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t}}{1-(\beta\gamma)^T}\right) \text{ if } r \ge t > 1, \\ -d'_{rj}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^T} + (1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^T}\right) \qquad \text{if } r < t, \ \ r \ge 1. \end{cases} \tag{42}$$

and also that

$$\frac{dx_{tj}}{d\tau_{(I+1)}} - \alpha\frac{dx_{t-1,j}}{d\tau_{(I+1)}} = \begin{cases} -\frac{\theta}{b(1+\beta\alpha^2)}\left((1-\alpha)\left(\frac{1}{1-\beta\gamma}+\frac{\gamma}{1-\gamma}\right)+\frac{\alpha-\gamma}{1-\gamma}\gamma^{t-I-1}\right. \\ \qquad \left. +(\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{T+1-t}((\beta\gamma)^I-1)}{1-(\beta\gamma)^T}\right) & \text{if } I+1<t, \\ -\frac{\theta}{b(1+\beta\alpha^2)}\left((\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{I+1-t}(1-(\beta\gamma)^{T-I})}{1-(\beta\gamma)^T}\right) & \text{if } I+1\geq t. \end{cases} \tag{43}$$

We can now express the excess burden for both agent types. The expression can be written

$$EB_j(\underline{\tau}^\star) \approx -\frac{1}{2}\left(\Delta\underline{\tau}_j\right)'\cdot\begin{bmatrix} & \vdots & \\ \dots & \frac{\partial^2 EB_j}{\partial\tau_s\partial\tau_r} & \dots \\ & \vdots & \end{bmatrix}\cdot\left(\Delta\underline{\tau}_j\right) \tag{44}$$

where

$$\begin{aligned} \frac{\partial^2 EB_j}{\partial\tau_r^2} = -b\left(-d'_{rj}\right)^2\sum_{M=1}^{\infty}\beta^{T(M-1)}\Bigg[&\left(\frac{\gamma^{1-r+T}(1-\gamma^{T(M-1)})}{1-\gamma^T}+\frac{(\beta\gamma)^{r-1}}{1-(\beta\gamma)^T}\right)^2 \\ &+\sum_{t=2}^{r}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r+T}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t}}{1-(\beta\gamma)^T}\right)^2 \\ &+\sum_{t=r+1}^{T}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^T}\right)^2\Bigg] \text{ if } r\leq I, \end{aligned} \tag{45}$$

$$\begin{aligned} &\frac{\partial^2 EB_j}{\partial\tau_{(I+1)}^2} = -b\left(-\frac{\theta}{b(1+\beta\alpha^2)}\right)^2\times \\ &\Bigg\{\sum_{M=1}^{\infty}\beta^{T(M-1)}\Bigg[\sum_{t=1}^{I+1}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{I+1-t}(1-(\beta\gamma)^{T-I})}{1-(\beta\gamma)^T}\right)^2 \\ &+\sum_{t=I+2}^{T}\beta^{t-1}\left((1-\alpha)\left(\frac{1}{1-\beta\gamma}+\frac{\gamma}{1-\gamma}\right)+\frac{\alpha-\gamma}{1-\gamma}\gamma^{t-I-1}+(\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^T}\frac{1-\gamma^{T-I}}{1-\gamma}\right. \\ &\left.+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{T+1-t}((\beta\gamma)^I-1)}{1-(\beta\gamma)^T}\right)^2\Bigg]\Bigg\}, \end{aligned} \tag{46}$$

and for $r<s<I+1$:

$$\begin{aligned} &\frac{\partial^2 EB_j}{\partial\tau_r\partial\tau_s} = -b(-d'_{sj})(-d'_{rj})\times \\ &\Bigg\{\sum_{M=1}^{\infty}\beta^{T(M-1)}\Bigg[\left(\frac{\gamma^{1-r+T}(1-\gamma^{T(M-1)})}{1-\gamma^T}+\frac{(\beta\gamma)^{r-1}}{1-(\beta\gamma)^T}\right)\left(\frac{\gamma^{1-s+T}(1-\gamma^{T(M-1)})}{1-\gamma^T}+\frac{(\beta\gamma)^{s-1}}{1-(\beta\gamma)^T}\right) \\ &+\sum_{t=2}^{r}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r+T}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t}}{1-(\beta\gamma)^T}\right)\times \\ &\left((\gamma-\alpha)\frac{\gamma^{t-1-s+T}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{s-t}}{1-(\beta\gamma)^T}\right) \\ &+\sum_{t=r+1}^{s}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^T}\right)\times \\ &\left((\gamma-\alpha)\frac{\gamma^{t-1-s+T}(1-\gamma^{T\cdot M})}{1-\gamma^T}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{s-t}}{1-(\beta\gamma)^T}\right) \end{aligned}$$

$$+\sum_{t=s+1}^{T}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^{T}}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^{T}}\right)\times$$
$$\left((\gamma-\alpha)\frac{\gamma^{t-1-s}(1-\gamma^{T\cdot M})}{1-\gamma^{T}}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{s-t+T}}{1-(\beta\gamma)^{T}}\right)\Bigg]\Bigg\}, \tag{47}$$

and finally

$$\frac{\partial^2 EB_j}{\partial\tau_r\partial\tau_{(I+1)}}=-b(-d'_{rj})\left(-\frac{\underline{\theta}}{b(1+\beta\alpha^2)}\right)\times$$
$$\Bigg\{\sum_{M=1}^{\infty}\beta^{T(M-1)}\Bigg[\left(\frac{\gamma^{1-r+T}(1-\gamma^{T(M-1)})}{1-\gamma^{T}}+\frac{(\beta\gamma)^{r-1}}{1-(\beta\gamma)^{T}}\right)\times$$
$$\left((\gamma-\alpha)\frac{(1-\gamma^{T(M-1)})}{1-\gamma^{T}}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{I}(1-(\beta\gamma)^{T-I})}{1-(\beta\gamma)^{T}}\right)$$
$$+\sum_{t=2}^{r}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r+T}(1-\gamma^{T\cdot M})}{1-\gamma^{T}}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t}}{1-(\beta\gamma)^{T}}\right)\times$$
$$\left((\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^{T}}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{I+1-t}(1-(\beta\gamma)^{T-I})}{1-(\beta\gamma)^{T}}\right)$$
$$+\sum_{t=r+1}^{I+1}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^{T}}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^{T}}\right)\times$$
$$\left((\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^{T}}\frac{1-\gamma^{T-I}}{1-\gamma}+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{I+1-t}(1-(\beta\gamma)^{T-I})}{1-(\beta\gamma)^{T}}\right)$$
$$+\sum_{t=I+2}^{T}\beta^{t-1}\left((\gamma-\alpha)\frac{\gamma^{t-1-r}(1-\gamma^{T\cdot M})}{1-\gamma^{T}}+(1-\alpha\beta\gamma)\frac{(\beta\gamma)^{r-t+T}}{1-(\beta\gamma)^{T}}\right)\times$$
$$\left((1-\alpha)\left(\frac{1}{1-\beta\gamma}+\frac{\gamma}{1-\gamma}\right)+\frac{\alpha-\gamma}{1-\gamma}\gamma^{t-I-1}+(\gamma-\alpha)\frac{\gamma^{t-1}(1-\gamma^{T(M-1)})}{1-\gamma^{T}}\frac{1-\gamma^{T-I}}{1-\gamma}\right.$$
$$\left.+\frac{1-\alpha\beta\gamma}{1-\beta\gamma}\frac{(\beta\gamma)^{T+1-t}((\beta\gamma)^{I}-1)}{1-(\beta\gamma)^{T}}\right)\Bigg]\Bigg\}. \tag{48}$$

We now have all the expressions we need to parameterize and calculate the optimal taxes, second-best taxes, and excess burden for each household type.

Appendix Figures

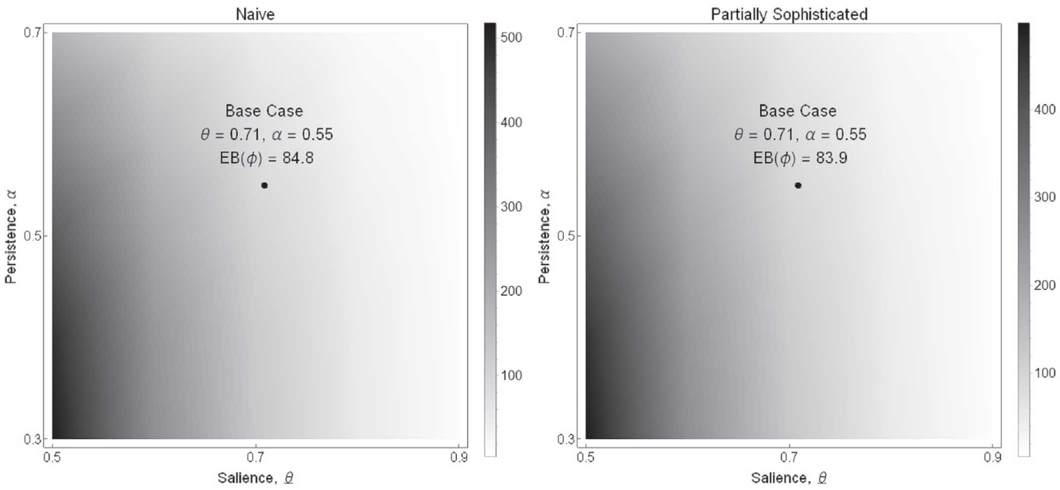

Notes: Excess burden is quantified in per household net present value dollars over an infinite time horizon, with darker shading indicating larger excess burden. Note that the scale is different for each agent type. This is necessary because the variation in excess burden across the parameter space and agent types is highly nonlinear. For both agent types, salience is the main factor in the size of the excess burden of a constant tax. However, excess burden does vary somewhat with persistence; the excess burden of taxing at the standard Pigouvian rate is larger for when persistence is lower.

Fig. 5 Excess burden of standard Pigouvian tax ($\tau = \varphi$), Sensitivity to persistence α and salience $\underline{\theta}$.

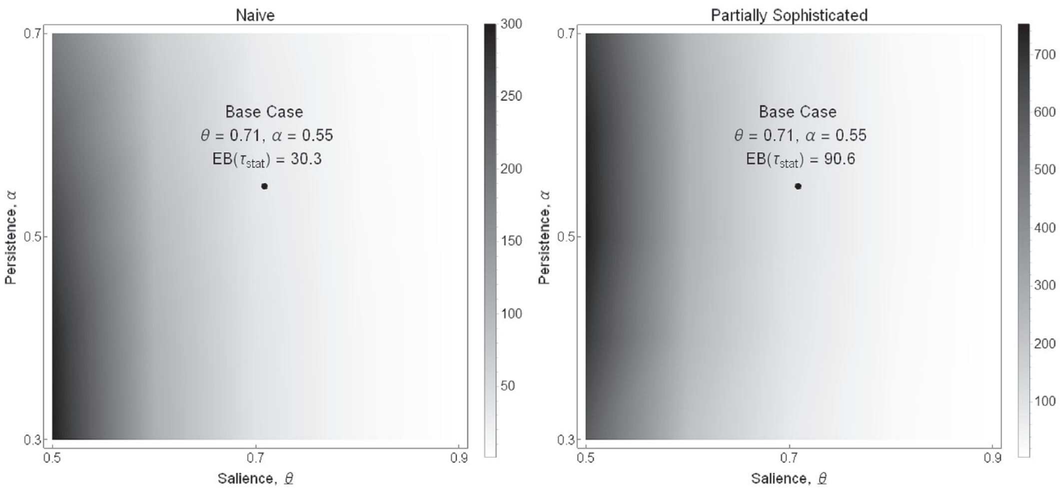

Notes: Excess burden is quantified in per household net present value dollars over an infinite time horizon, with darker shading indicating larger excess burden. Note that the scale is different for each agent type. This is necessary because the variation in excess burden across the parameter space and agent types is highly nonlinear. For both agent types, salience is the main factor in the size of the excess burden of a constant tax. However, excess burden does vary somewhat with persistence; excess burden is larger when persistence is lower for naive agents, whereas excess burden is concave in persistence for partially sophisticated agents.

Fig. 6 Excess burden of static salience tax ($\tau = \tau_{stat}$), Sensitivity to persistence α and salience $\underline{\theta}$.

References

Alberini, A., Filippini, M., 2011. Response of residential electricity demand to price: the effect of measurement error. Energy Econ. 33 (5), 889–895.
Allcott, H., Taubinsky, D., 2015. Evaluating behaviorally motivated policy: experimental evidence from the lightbulb market. Am. Econ. Rev. 105 (8), 2501–2538.
Allcott, H., Mullainathan, S., Taubinsky, D., 2014. Energy policy with externalities and internalities. J. Publ. Econ. 112, 72–88.
Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. Q. J. Econ. 121 (2), 635–672.
Auerbach, A.J., 1985. The theory of excess burden and optimal taxation. In: Handbook of Public Economics, vol. 1. Elsevier, pp. 61–127.
Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over time: dynamic inconsistency in real effort tasks. Q. J. Econ. 130 (3), 1067–1115.
Becker, G.S., Murphy, K.M., 1988. A theory of rational addiction. J. Polit. Econ. 96 (4), 675–700.
Becker, G.S., Grossman, M., Murphy, K.M., 1994. An empirical analysis of cigarette addiction. Am. Econ. Rev. 84 (3), 396–418.
Benhabib, J., Bisin, A., 2005. Modeling internal commitment mechanisms and self-control: a neuroeconomics approach to consumption–saving decisions. Game. Econ. Behav. 52 (2), 460–492.
Chetty, R., 2009. The Simple Economics of Salience and Taxation. No. w15246. National Bureau of Economic Research.
Chetty, R., Looney, A., Kroft, K., 2009. Salience and taxation: theory and evidence. Am. Econ. Rev. 99 (4), 1145–1177.
Costa, F., Gerard, F., 2018. Hysteresis and the Welfare Effect of Corrective Policies: Theory and Evidence from an Energy-Saving Program. No. w24608. National Bureau of Economic Research.
DellaVigna, S., 2009. Psychology and economics: evidence from the field. J. Econ. Lit. 47 (2), 315–372.
Deryugina, T., MacKay, A., Reif, J., 2020. The long-run dynamics of electricity demand: evidence from municipal aggregation. Am. Econ. J. Appl. Econ. 12 (1), 86–114.
U.S. Energy Information Administration, 2017. Electric Power Annual 2017. U.S. Department of Energy.
Farhi, E., Gabaix, X., 2020. Optimal taxation with behavioral agents. Am. Econ. Rev. 110 (1), 298–336.
Feinberg, R.A., 1986. Credit cards as spending facilitating stimuli: a conditioning interpretation. J. Consum. Res. 13 (3), 348–356.
Filippini, M., Hirl, B., Masiero, G., 2018. Habits and rational behaviour in residential electricity demand. Resour. Energy Econ. 52, 137–152.
Finkelstein, A., 2009. E-ztax: tax salience and tax rates. Q. J. Econ. 124 (3), 969–1010.
Flavin, M., Nakagawa, S., 2008. A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence. Am. Econ. Rev. 98 (1), 474–495.
Fuhrer, J., 2000. Habit formation in consumption and its implications for monetary-policy models. Am. Econ. Rev. 90 (3), 367–390.
Gilbert, B., Graff Zivin, J., 2014. Dynamic salience with intermittent billing: evidence from smart electricity meters. J. Econ. Behav. Organ. 107 (A), 176–190.
Gilbert, B., Graff Zivin, J., 2018. Dynamic Corrective Taxes with Time-Varying Salience. No. w25014. National Bureau of Economic Research.
Graff Zivin, J., Kotchen, M., Mansur, E., 2014. Spatial and temporal heterogeneity of marginal emissions: implications for electric cars and other electricity-shifting policies. J. Econ. Behav. Organ. 107 (A), 248–268.
Grubb, M.D., Osborne, M., 2015. Cellular service demand: biased beliefs, learning, and bill shock. Am. Econ. Rev. 105 (1), 234–271.
Gruber, J., Koszegi, B., 2004. Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. J. Publ. Econ. 88 (9–10), 1959–1987.
Harberger, A.C., 1964. The measurement of waste. Am. Econ. Rev. 54 (3), 58–76.
Heien, D., Durham, C., 1991. A test of the habit formation hypothesis using household data. Rev. Econ. Stat. 189–199.
Heutel, G., 2015. Optimal policy instruments for externality-producing durable goods under present bias. J. Environ. Econ. Manag. 72, 54–70.
Hirschman, E.C., 1979. Differences in consumer purchase behavior by credit card payment system. J. Consum. Res. 6 (1), 58–66.
Karlan, D., McConnell, M., Mullainathan, S., Zinman, J., 2016. Getting to the top of mind: how reminders increase saving. Manag. Sci. 62 (12), 3393–3411.
Kaur, S., Kremer, M., Mullainathan, S., 2010. Self-control and the development of work arrangements. Am. Econ. Rev. 100 (2), 624–628.
Landry, P., 2019. Bad habits and the endogenous timing of urges. Rev. Econ. Stud. 1–22 0.
Leth-Petersen, S., 2007. Habit formation and consumption of energy for heating: evidence from a panel of Danish households. Energy J. 35–54.
O'Donoghue, T., Rabin, M., 1999. Doing it now or later. Am. Econ. Rev. 89 (1), 103–124.
O'Donoghue, T., Rabin, M., 2006. Optimal sin taxes. J. Publ. Econ. 90 (10–11), 1825–1849.
Pigou, A.C., 1920. The Economics of Welfare. Macmillan, London.
Prelec, D., Simester, D., 2001. Always leave home without it: a further investigation of the credit-card effect on willingness to pay. Market. Lett. 12 (1), 5–12.
Samuelson, W., Zeckhauser, R., 1988. Status quo bias in decision making. J. Risk Uncertain. 1 (1), 7–59.
Scott, K.R., 2012. Rational habits in gasoline demand. Energy Econ. 34 (5), 1713–1723.
Sexauer, B., 1977. The role of habits and stocks in consumer expenditure. Q. J. Econ. 127–142.
Sexton, S., 2015. Automatic bill payment and salience effects: evidence from electricity consumption. Rev. Econ. Stat. 97 (2), 229–241.
Taubinsky, D., Rees-Jones, A., 2017. Attention variation and welfare: theory and evidence from a tax salience experiment. Rev. Econ. Stud. 85 (4), 2462–2496.
Thomas, M., Desai, K.K., Seenivasan, S., 2011. How credit card payments increase unhealthy food purchases: visceral regulation of vices. J. Consum. Res. 38 (1), 126–139.
Thunstrom, L., 2010. Preference heterogeneity and habit persistence: the case of breakfast cereal consumption. J. Agric. Econ. 61 (1), 76–96.
U.S. Environmental Protection Agency, 2013. Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866. (May 2013, Revised August 2016).