Taxing Cue-Triggered Consumption

Nicolaj Verdelin*

Department of Economics, University of Copenhagen & EPRU

May 2006

Abstract

Consumption of a wide array of goods is affected by visceral cues. This is particularly the case for sin goods and not least those with addictive properties. Exposure to a cue is random and increases consumption of the sin good which reduces the disposable income for other goods. This paper shows that when there are no errors in individual decision-making, it is optimal to subsidize cue-triggered consumption goods. The optimal rate of subsidization decreases as cue regulation is tightened. When bounded rationality implies that cue exposure leads to overconsumption of the sin good, it is optimal to tax sin good purchases. The optimal sin tax rate decreases with the degree of cue regulation. This suggests that cue regulation and sin taxes are substitutable instruments, not complementary.

*Department of Economics, University of Copenhagen, Studiestraede 6, DK-1455 Copenhagen K, Denmark. Tel.: +45 3532 3021. Fax: +45 3532 4444. E-mail address: nicolaj.verdelin@econ.ku.dk. I wish to thank my supervisors Henrik Jacobsen Kleven and Claus Thustrup Kreiner as well as Xavier Gabaix for helpful comments and suggestions. All remaining errors are mine.

1 Introduction

In recent years there has been a surge in the interest by the economics profession in the role of temptation and self-control in human decision-making. This research has been motivated, as has Behavioral Economics in general, by the failure of traditional, neoclassical models to accurately capture central aspects of human behavior. Among the range of issues that has been studied are intertemporal consumption allocation (saving), demand for unhealthy foods and addiction. However, the literature has generally failed to converge towards a unifying framework, although the seminal work of David Laibson (1997) has come close to providing a workhorse model for much of the subsequent work on intertemporal choice. In particular, widespread disagreement remains as to the degree of rationality on the part of human beings. Not surprisingly, different assumptions about rationality often lead to starkly contrasting policy implications.

This paper addresses a particular type of temptation-based consumption, that is influenced by so-called *cue-triggered conditioned responses* or *compensatory conditioned responses* (Laibson 2001). In this setting, a visceral cue (the smell of freshly baked bread or seeing somebody smoking) triggers an anticipatory response that affects demand. A compensatory conditioned response occurs when the body prepares itself for future intake upon cue exposure (homeostasis) as is the case with addiction. The temporary, anticipatory emotion that is triggered by exposure to a visceral cue poses a challenge for welfare economics. Should it be given full hedonic value when assessing the well-being of an individual? Or, is it more appropriate to assume that anticipatory emotions do not affect true long-run welfare? This comes back to the central question of the appropriate departure from full rationality.

In the similarly spirited model of Douglas Bernheim and Antonio Rangel (2004), the authors appeal to what they dub a Hedonic Forecasting Mechanism (HFM) which represents a chemical process in the brain that creates an anticipatory emotion which raises the probability of consuming. They base their approach on findings of increased concentrations of dopamine in the brain following exposure to a visceral or environmental cue. It is through this process that the brain forecasts the expected pleasure from consumption, helping the individual to optimize choices. Bernheim and Rangel further argue that cues for addictive goods may inhibit the proper functioning of the HFM.

The present paper provides a simplified and unified framework, encompassing both the case studied by Laibson (2001) and that considered by Bernheim and Rangel (2004), to study optimal tax policy and the role of cue regulation. This leads to new results on optimal paternalism and generates insights into the dependence of policy prescriptions on the specific modelling assumptions made.

Laibson (2001) in his study of addiction follows the tradition of Gary Becker and Kevin Murphy (1988) and assumes that choices are always aligned with true preferences and thus gives full weight to any temporary emotional responses. It is tempting to conclude that with this assumption there is no role for the government to play beyond the correction of externalities. However, as this paper shows, this intuition turns out to be incorrect. The random nature of cue shocks affects the marginal utility of consumption in a stochastic sense which creates a demand for insurance. The government can supply partial insurance by subsidizing the sin good. This result extends to an environment with heterogeneous earnings and taste differences across income groups as long as the government has access to a non-linear income tax. As in cases of optimal insurance design, the size of the subsidy involves a trade-off between the benefits of consumption smoothing and the risk of moral hazard (i.e., increased consumption of the sin good). We find that the optimal subsidy rate decreases with the degree of government cue regulation.

The case for commodity specific insurance is made under the assumption of rational decision-making. In contrast, Bernheim and Rangel (2004) argue that cue-triggered conditioned responses represent a departure from full rationality. If decision utility is misaligned with true utility, standard arguments suggest using sin taxes to correct so-called *internalities*. Importantly, the present paper shows that the optimal corrective tax depends on the degree of regulation of the cue process. As regulation is tightened, the optimal sin tax decreases. When targeted regulation is possible, it improves the trade-off faced by the government. The present paper explores the consequences of using (directed) regulation for optimal sin taxes.

There are two main contributions of this paper. The first is to show that cue regulation and sin taxation are substitutes, not complements irrespective of the degree of individual rationality. The second is to demonstrate that the nature of optimal tax policy depends crucially on the degree of individual rationality. Indeed, as shown below, the two different treatments, proposed in the literature, of the short-term cravings that people experience in the wake of cue exposure lead to opposing policy recommendations. If one follows Laibson (2001) and attaches full hedonic weight to short-term cravings, the optimal policy is a *subsidy* on sin goods. If, instead, one uses the approach of Bernheim and Rangel (2004), the optimum involves a *tax* on sin goods. In this sense, the generality of the result on the substitutability of taxation and cue regulation is in stark contrast to the model-dependent policy prescriptions when it comes to the sign of the optimal sin tax. Naturally, the true degree of rationality is an empirical question but it seems that today there is more widespread acceptance of the notion of errors in individual decision-making. However, for the case of conditioned responses there are more complicated issues involved. Indeed, even if one accepts that demand does not always reveal true preferences, it is not clear that conditioned responses should be given no hedonic weight. The true model may very well be a hybrid of the two polar cases.

The analysis in this paper can be seen as a contribution to the literature on optimal paternalism. While still in its infancy, this literature has shown that the government may well have an important role to play in helping to correct individual errors due to bounded rationality. Ted O'Donoghue and Matthew Rabin (2006) and Jes Winther Hansen (2005) are both examples of this literature. The present paper differs from the existing literature by considering the interdependencies between different policy instruments as well as explicitly contrast the policy prescriptions with and without individual bounded rationality.

The paper proceeds as follows. Section 2 briefly outlines the underlying idea of cue-triggered consumption while section 3 gives an overview of the related literature.

Section 4 presents the model. Section 5 explores the Laibsonian case of full rationality and presents the case for commodity specific insurance. Section 6 considers the case of errors in decision-making and analyzes optimal sin taxes as well as their dependence on the degree of regulation. Section 7 concludes.

2 Cue-Triggered Consumption

The starting points for this paper are the notions of conditioned responses and compensatory conditioned responses known from cognitive neuroscience and psychology. A conditioned response is a "cue-based anticipatory response to a physiological stimulus" and a compensatory conditioned response is a "conditioned response that is homeostatic in nature".[1] Homeostasis is an equilibrating mechanism that stabilizes the physiological system. In more germane terms, a conditioned response is the body's preparatory reaction to a cue such as the smell of freshly baked bread. The cue triggers an anticipation of intake and induces a sense of craving. Compensatory responses arise as preparatory responses to consumption with strong physiological effects such as drug intake. It is aimed at dampening the body's reaction to the anticipated consumption. These insights were formalized by Laibson (2001) by introducing a random cue process and two sets of baseline preferences that are active in the states of exposure and non-exposure, respectively. Building on the original idea of Becker and Murphy (1988), Laibson (2001) introduces a cue-triggered reference-point in the basic utility function. This captures the idea that upon exposure, the marginal utility of consumption increases but the utility derived from consumption is reduced. In terms of observed behavior, this will increase the probability of consuming as well as increase the amount consumed conditional on consuming.

The approach taken by Bernheim and Rangel (2004) is closely related to that of Laibson (2001). They make use of a Hedonic Forecasting Mechanism (HFM) which responds to the brain's production of dopamine. The level of dopamine is used to

[1] These definitions are taken from a Lecture Note by David Laibson.

forecast the pleasure from consumption and is affected by previous experiences with the good in question. This reflects the same anticipatory effect of cue exposure described above. As to decision-making, the agent may find herself in one of two states: A cold mode and a hot mode. In the cold mode, decisions are rationally governed by utility maximization. In the hot mode, the agent consumes the sin good by default. This approach is motivated by evidence that cues related to addictive goods can interfere with the proper functioning of the HFM by increasing the concentration of dopamine. This leads the agent to overestimate the pleasure from consuming the sin good.

The model is best interpreted as one of addiction and formalizes the notions of craving and withdrawal. However, the idea of cue-triggered consumption extends to many other types of consumption, in particular food. In these cases, it is less clear if the utility of consumption is adversely affected by cues[2] but it is clear that behavior differs between the states of exposure and non-exposure, respectively. The discussion in this paper will usually refer to addictive goods such as cigarettes but many of the central points apply to a wider set of consumption goods.

A common trait for many cue-affected consumption goods appears to be that they fall under the general heading of *sin goods*. These are goods where the consumption entails costs, typically in the form of poorer health. This lies at the heart of the before mentioned feature that consuming the good upon exposure usually provides less satisfaction ex post than forecasted when making the choice to consume.

As already hinted, a cue can be interpreted as any kind of object, whether physical or abstract, that triggers an association to consumption of a certain good. This broad interpretation naturally encompasses all kinds of direct advertising as well as environmental cues such as seeing others consume the good or simply seeing the good itself. Cues with a direct associative effect are labelled visceral cues or environmental cues. To give precision to the concept of cue regulation, it is important to distinguish direct measures such as advertising bans from policies such as consumption restrictions where the policy's effect on the flow of cues is but a side-effect. In this paper, the term cue

[2] However, none of the conclusions in this paper depend on this assumption. More on this below.

regulation refers to the former. Among other things, this has the advantage of avoiding the need to model the costs of more intrusive regulation. Instead, the assumption is maintained throughout that public control of the cue flow entails no cost to the consumer.[3] In addition to cue regulation, a policy of banning consumption for specific, easily identifiable groups will be considered.

Naturally, there might be differences across goods as to the importance of different cues and cues might differ across individuals. This is incorporated in the modelling below where a part of the discussion will center on the possibility and advantages of using targeted cue regulation.

The modelling approach taken in this paper is intentionally simple. Thus, intertemporal effects of consumption such as the build-up of an addictive habit are not explicitly modelled. Instead, the focus is on optimal policy and the relation between different policy instruments. The analysis in Bernheim and Rangel (2004) is an example of a full-scale model of addictive consumption. The simplified framework has the advantage of lending itself easily to many different interpretations. Therefore, many of the results would seem to have merit for many different kinds of sin goods with very different characteristics.

3 Related Literature

A growing literature examines the role of temptation and self-control in human decision-making. Laibson (1997) rediscovered the quasi-hyperbolic discounting framework originally suggested by Edmund Phelps and Robert Pollak (1968) in the context of intergenerational preferences. A number of papers has since further explored the (β, δ)-framework. O'Donoghue and Rabin (1999, 2001) analyze intertemporal choices of present-biased agents and contrast 'naïveté' and 'sophistication'. At the same time, other models have been proposed. Laibson (2001) introduced the notion of cue-triggered consumption

[3] All the results presented are robust to the inclusion of an additive cost term. With some added qualifications, they further survive more complicated cost structures that interact with preferences. But as it is empirically unclear how to model potential benefits of cues, these are left out of the analysis.

which is the basis for the dual-self model of Bernheim and Rangel (2004). Laibson did not investigate the policy implications of his model and part of the contribution of this paper is to fill this gap. In contrast, Bernheim and Rangel study a range of instruments and calculate optimal sin tax rates as well as consider conditions under which it may be optimal to subsidize addictive goods as a means of redistribution. A recent contribution to the modelling of interpersonal conflict is the dual-self model of Drew Fudenberg and David Levine (2006).

Faruk Gul and Wolfgang Pesendorfer (2001, 2004) have suggested a different interpretation of temptationist preferences based on an axiomatic, choice-theoretic approach. The assumed full degree of rationality usually implies that there is no corrective role for the government. A notable exception is Gul and Pesendorfer (2006) who find that 'price policies' are genereally inefficient and dominated by bans.

As modelling has progressed, there has been an increasing interest in the policy implications of temptationist preferences as evidenced by the studies by O'Donoghue and Rabin (2001) on optimal sin taxes and asymmetric paternalism. O'Donoghue and Rabin (2006) identify conditions under which optimal sin taxes on non-addictive goods increase with the average degree of irrationality. Hansen (2005) presents an argument for providing public transfers in kind when agents are present-biased. Jonathan Gruber and Botond Köszegi (2001) derive optimal cigarette taxes within a quasi-hyperbolic discounting framework. The focus of these studies is on sin goods whose consumption entails long-run health consequences. Bernheim and Rangel (2004) consider severe addiction.

An important issue in all of public finance is the concern for income distribution. Corrective taxation affects the net income distribution insofar as the consumption of sin goods is correlated with income. Yet, most of the analysis of sin taxes has so far proceeded without considering income heterogeneity.[4] This may be defended on the grounds that the literature on optimal income and commodity taxation (see e.g., Emmanuel Saez 2002) has shown the uselessness of commodity taxation for redistributive

[4] A prominent exception is Bernheim and Rangel (2004).

purposes under fairly general conditions. In contrast, this paper shows that the presence of a non-linear income tax strengthens the case for commodity specific insurance in the Laibsonian case. Intuitively, sin good consumption is correlated with events that are themselves correlated with the social marginal utility of income but not necessarily with income. An interesting avenue that is not explored is the use of sin taxes when the incidence of bounded rationality falls disproportionately on the poor (see Marianne Bertrand et al. (2003) for an expression and some consequences of this view).

4 The Model

The model is very simple and bears some resemblance to those of Laibson (2001) and O'Donoghue and Rabin (2006). In this section, the model is laid out in its most general form allowing for income effects on consumption. This sets the stage for the analysis of the Laibsonian case. The section on bounded rationality will disregard income effects to simplify the analysis.

The model is static. There are N groups of individuals of equal size, indexed by the letter i. Preferences within each group are homogeneous but preferences may differ across groups. Any individual agent allocates consumption across two goods: A composite good x and a sin good z. Labor supply is exogenous and so the income of each member of group i is simply denoted y_i. The preferences of group i can be represented by the following functional form

$$u_i\left(x, z - s \cdot r_i\right), \tag{1}$$

where s is an indicator variable for exposure, taking on the value 1 when exposed and 0 otherwise. Standard assumptions are made about the shape of the utility function.[5] Specifically, u_i is concave, and $u_{ix} = \frac{\partial u_i}{\partial x} > 0$, $u_{iz} = \frac{\partial u_i}{\partial z} > 0$ for $z - s \cdot r_i < \overline{z}_i$ and $u_{iz} \leq 0$ for $z - s \cdot r_i \geq \overline{z}_i$.[6] The possibility of negative marginal utility of consuming the sin good

[5] The form of (1) is not important for the results in this paper. What matters is that the marginal utility of consuming the sin good increases following cue exposure.

[6] A natural benchmark is the case of separability in which $u_i\left(x, z\right) = v_i\left(x\right) + w_i\left(z - s \cdot r_i\right) - c_i\left(z\right)$, where $c_i\left(\cdot\right)$ represents the health costs of consuming the sin good. Apart from the presence of cues, this is the functional form analyzed by O'Donoghue and Rabin (2006).

reflects the fact that the sin good is associated with various costs in the form of e.g., poorer health. Presumably, the marginal costs of sin good consumption are increasing at a non-decreasing rate, implying that when consumption is above some threshold, the marginal cost of consuming exceeds the marginal benefit. Cue exposure has two effects. First, the marginal utility of sin good consumption is increased. This follows from the concavity of u_i. Second, utility drops.[7]

As modelled, cue exposure activates a reference point r_i in the utility function of the exposed agent. This captures the idea, already alluded to, that cue-triggered consumption is associated with a conditioned response which is anticipatory in nature. Following exposure, the agent anticipates consuming the good and this reduces the experienced utility from actually consuming. At the same time, there is a feeling of disappointment if the agent does not consume in that utility is less than what it would be if the agent did not consume and also had not been exposed to a cue. In this sense, the underlying idea clearly bears some resemblence to the concept of 'loss aversion' introduced by Daniel Kahneman and Amos Tversky (1979). More importantly, cue exposure increases the marginal utility of consuming the sin good. This again reflects the anticipatory effect of cue exposure.

The form of (1) is perhaps more easily understood by appealing to the case of smoking. When a smoker is subjected to a cue such as a cigarette commercial, it triggers a craving for nicotine. In this sense, the agent perceives a higher marginal utility of smoking. This response in itself prepares the body for intake and diminishes the derived pleasure from actual consumption. This is captured by the fact that utility is lower for any given consumption vector (x, z) when the cue-triggered reference point is activated $(s = 1)$. At the same time, if the anticipation of nicotine is not satisfied the body responds by generating a feeling of withdrawal as captured by the utility loss following consumption below the expected level (r_i).

The process of cue exposure is random. With probability p_i an agent of type i is

[7] In the original model of Laibson (2001), the utility loss is motivated by appealing to the body's own equilibrating mechanism known as homeostasis.

subjected to a cue. As described above, cue exposure activates a reference point in the utility function for sin good consumption. Due to the effect of cue exposure on the utility function, the model is effectively one of dual selves. Any given agent may find herself in one of two *states*: A state of non-exposure and a state of exposure depending on whether or not the reference point has been triggered by a visceral cue. This is merely one interpretation of the dual-self framework; that the same individual may be in one of two fundamental states of mind. Both the reference point and the probability of exposure is individual-specific as captured by the subscript i. This reflects the fact that people have different sensitivities to cues and different risks of exposure. A heavy smoker most likely has a higher reference point and is also probably more exposed to cues related to smoking. Similarly, children and young adults are perhaps more responsive to cue exposure and also more exposed to advertising. In both cases, this would be captured by high values of r_i and p_i, respectively.

Note the timing of the model. First, cues are transmitted and the consumption choices are made after the resolution of uncertainty i.e., after the state of exposure/non-exposure is known. But at the beginning of time there is individual uncertainty as to the future preferences.

The budget set of a consumer in group i is

$$x_i + qz_i = y_i,$$

where y_i is earned income, q is the price of the sin good and the price of the composite good has been normalized to 1. Individual choices for type i then satisfy

$$\frac{u_{iz}\left(x_i, z_i - s \cdot r_i\right)}{u_{ix}\left(x_i, z_i - s \cdot r_i\right)} = q,$$

where $s = 1$ when the agent has been exposed to a cue and 0 otherwise. From this and the budget constraint we may write the choices as functions of q, y_i and $s \cdot r_i$

$$\begin{aligned} x_i &= x_i\left(q, y_i, s \cdot r_i\right) \\ z_i &= z_i\left(q, y_i, s \cdot r_i\right) \end{aligned}$$

with indirect utility denoted $V_i(q, y_i, s \cdot r_i)$. The first-order conditions imply that

$$\begin{aligned} z_i(q, y_i, r_i) &> z_i(q, y_i) \\ x_i(q, y_i, r_i) = y_i - qz_i(q, y_i, r_i) &< y_i - qz_i(q, y_i) = x_i(q, y_i), \end{aligned} \tag{2}$$

so that a cue-exposed agent will wish to purchase more of the sin good at the expense of other commodities.[8] In addition, the effect of cue exposure is higher the more sensitive is the agent to cues as captured by the reference point r_i. Specifically,

$$\begin{aligned} \frac{\partial x_i}{\partial ri}|_{s=1} &< 0, \ \frac{\partial x_i}{\partial ri}|_{s=0} = 0 \\ \frac{\partial z_i}{\partial r_i}|_{s=1} &> 0, \ \frac{\partial z_i}{\partial r_i}|_{s=0} = 0. \end{aligned}$$

There are two competing interpretations of this set-up. Under one view it considers the agent at a moment in time and there is an instantaneous risk of exposure which will affect the consumption choice. With this interpretation the model is best understood as treating a single out of a large number of independent periods of life as in O'Donoghue and Rabin (2006). Under another view the model is a description of a long-run steady state where an agent is either exposed, which would then be the equivalent of e.g., having build an addictive habit, or the agent is not exposed. Again, in the idiosyncratic state of exposure, consumption of the sin good is higher for any agent than for that same individual in the state of non-exposure.

Note that the model abstracts from any external effects of sin good consumption. While clearly relevant, the consequences of the presence of externalities is well understood and they are left out to focus the analysis on the novel aspects of the model. This is in line with similar work on temptation and self-control such as Bernheim and Rangel (2004) and Gruber and Köszegi (2001).

[8] A more general model would endogenize labor supply. Under weak conditions an agent would smooth the cue shock absorption across all consumption goods and would thus increase labor supply relative to the state of non-exposure. This is potentially important in public finance analysis but is left for future work.

5 The Case of Full Rationality

The model laid out in section 4 involves each agent having two different utility functions depending on the idiosyncratic state of mind. This raises the issue of what is an appropriate welfare criterion. Two competing assumptions have been suggested in the literature. In the model of Laibson (2001), full weight is given to any sense of craving and the agent is assumed to be fully rational as in Becker and Murphy (1988). In this case, utility is truly state-dependent and there is no conflict between choices and self-interest. This approach is the subject of the present section. In the analysis of Bernheim and Rangel (2004), cue exposure is assumed to drive the agent to act irrationally. Thus, there is only one true utility function. In terms of the present model, this assumption is captured by letting true welfare be always represented by the utility function (1)with $s = 0$. Behavior, however, continues to be affected by cue exposure. This case is analyzed in section 6.

The welfare of a non-exposed agent is given by

$$u_i(x, z),$$

while utility for a cue-exposed agent is

$$u_i(x, z - r_i).$$

Individual choices are made according to these state-dependent preferences. It follows from (2) using the Envelope Theorem, the assumption of full rationality and concavity of the individual utility function that

$$\frac{\partial V_i}{\partial y_i} = u_{ix}(x_i(q, y_i, s \cdot r_i)) \Rightarrow \frac{\partial V_i}{\partial y_i}|_{s=1} > \frac{\partial V_i}{\partial y_i}|_{s=0} \tag{3}$$
$$\frac{\partial(\partial V_i/\partial y_i)}{\partial r_i} = u_{ixx}\frac{\partial x_i}{\partial ri} > 0,$$

so the marginal utility of income is higher, conditional on group, for exposed than for non-exposed individuals. In addition, the more cue-sensitive is the individual the greater is the difference in marginal utilities in the two idiosyncratic states.

It is important to note that there is no corrective role for the government: Decision-making is completely aligned with true preferences. So what role may the government beneficially play? Note that cue exposure is a random event which affects the marginal utility of income. The divergence of marginal utilities of income across states creates a desire, ex ante, to transfer resources from the state of non-exposure to that of exposure. This is similar to insurance against an adverse shock. Conceivably, no private insurance company will provide this kind of insurance due to obvious problems of adverse selection under asymmetric information about the state of exposure.[9] This leaves room for the government to beneficially intervene and provide insurance against "affective consumption".

Consider first the case of only one group. In a first best world, the government observes cue exposure and is able to condition lump sum taxes on state of exposure. This is by no means realistic but serves as an illustration of the underlying idea. Let π and π^e denote lump sum taxes for non-exposed and exposed agents, respectively. The government maximizes the expected utility of the representative agent

$$\max_{\pi,\pi^e} pV^e\left(\pi^e\right) + (1-p)\, V_i\left(\pi\right)$$
$$s.t.\;\; p\pi^e + (1-p)\,\pi = 0,$$

where all arguments except the lump sum tax have been supressed for notational convenience and a superscript e denotes exposed agents throughout. The solution calls for the equation of marginal utilities of income across states and by (3) this is equivalent to $\pi_e < 0 < \pi$. This is also the solution that would obtain if people were given access to actuarily fair insurance against the event of cue-exposure. In this setting, the idiosyncratic risk essentially involves spending on 'useless' consumption. Consumption

[9] A caveat is in order. If the agent has access to perfect capital markets, she may self-insure by borrowing against future earnings. This is not possible in the static framework presented here, but when the model is interpreted as representing a single out of a large number of periods, intertemporal consumption smoothing is possible. Even in this case, if individual uncertainty remains, agents will still be better off by pooling idiosyncratic risks. These benefits then have to be weighted against the distortionary effects of taxation. However, under the preferred steady state interpretation of the model, where the state of exposure is in fact a permanent state of addiction, this caveat has no bearing. I am grateful to Claus Thustrup Kreiner for pointing this out.

of the sin good up to the reference point gives no net utility (but avoids a sense of disappointment) but reduces the disposable income of the exposed agent. In the case of cigarettes, alcohol or drug abuse or gambling, the stakes can be quite high.

The case of state dependent lump sum taxes is of no policy relevance but the argument extends to a setting of distortionary commodity taxes. The crucial observation is that sin good consumption is higher for exposed individuals. It follows that a policy of subsidizing sin good consumption will provide commodity specific insurance against the risk of cue exposure.[10] This is formalized in Proposition 1 which also establishes the dependence of the optimal subsidy rate on the degree of cue regulation. The proof is in Appendix A.

Proposition 1. *With an ex ante homogeneous population, it is optimal to subsidize the sin good. Further, the optimal subsidy rate is increasing in p, the probability of cue exposure as long as*

$$\frac{dV^e}{dt} - \frac{dV}{dt} < p\frac{d\left(dV^e/dt\right)}{dy}\frac{d\pi}{dp} + (1-p)\frac{d\left(dV/dt\right)}{dy}\frac{d\pi}{dp},$$

where. $\frac{dV^e}{dt} - \frac{dV}{dt} < 0$ *and* $\frac{d\pi}{dp} > 0$.

To gain intuition for this result, note that a necessary and sufficient condition for the optimal tax on sin goods to be strictly negative is

$$(z^e - z)\left(\frac{\partial V^e}{\partial y} - \frac{\partial V}{\partial y}\right) > 0,$$

when evaluated at $t = 0$ - a condition which is satisfied from above. This condition simply states that a subsidy is optimal whenever there is a positive correlation between consumption of the sin good and the marginal utility of income. This is a well-known condition from the literature on optimal commodity taxation with redistributional considerations (Peter Diamond 1975). The same literature has proceeded to show (Saez

[10] In this two-good economy, a subsidy on the sin good is, of course, equivalent to a tax on the composite good. This merely reflects the general point that the number of government price instruments is always no larger than the number of commodities minus one. More importantly, with endogenous labor supply it may, under certain conditions, be possible to implement the optimum relying only on an income tax.

2002) that commodity taxation is useless under fairly general conditions in the presence of an optimized non-linear income tax. But the analysis presented here is different from standard models in that cue exposure blurs the correlation between the marginal utility of income on the one side and income and ability on the other. A sin good subsidy is able to exploit patterns of consumption to provide partial insurance, something that a system based purely on income taxation can not achieve in this case. Of course, the analysis so far has proceeded under the assumption of homogeneous ability. The case of earnings heterogeneity is treated below.

Note that as long as the sin subsidy is distortionary in the sense that the aggregate compensated elasticity of demand is non-zero, the optimal subsidy is always less than 100% which is the equivalent of complete insurance. In this sense, the model captures the traditional trade-off inherent in optimal insurance design between the risk of moral hazard and the desire to perfectly smooth marginal utilities of consumption. The moral hazard effect is the increase in consumption by both exposees and non-exposees induced by a subsidy. The formal proof is in Appendix A.

An interesting question concerns the dependence of the optimal subsidy rate on the risk of cue exposure. Proposition 1 shows that the optimal subsidy rate increases (i.e., the tax rate decreases) with p. That is, less cue regulation makes it optimal to subsidize the sin good at a higher rate. Intuitively, when regulation is increased the frequency of cue exposure decreases so the risk of experiencing a high marginal utility of income is reduced. The fact that fewer people are exposed implies that the gain from subsidizing the sin good is reduced. At the same time, the cost of subsidizing at a given rate remains unchanged as it is made up by the distortion to sin good consumption for all consumers. As a consequence, the optimal subsidy rate declines as regulation is intensified. This ignores the potential complication introduced by the resence of income effects. Increased regulation improves the net fiscal status of both exposed and non-exposed individuals. This may change the relative marginal utilities of income of the two groups. In addition, the lower frequency of exposure makes it less costly to transfer resources to the exposed state which also tends to increase the subsidy. Finally, the marginal distortionary costs may change as consumption changes due to income effects. This explains the necessary condition listed above. The condition is likely not to be very restrictive. Note that a sufficient condition for the optimal subsidy rate to decrease with regulation is that the marginal welfare effect of taxation increases with income.

The questions addressed in this paper concern only the relation between tax policy and regulation. Proposition 1 is silent on the more fundamental question of optimal regulation - mainly because it requires addressing the issue of the benefits of cues, a topic that is not well analyzed in the literature. Note that in the present model, cues have only adverse effects so the question of optimal regulation is trivial. Thus, the analyses in this and the next section are more appropriately interpreted as the study of how tax policy should respond to incomplete cue regulation. Whether incomplete regulation is indeed part of the optimum would depend on the benefits of cues and the political constraints on the regulatory process.[11]

Importantly, it is never optimal to make use of an outright ban on the consumption of the sin good, whether targeted at certain groups or not. This is because consumption choices are always aligned with true preferences. The case for subsidizing the sin good is not driven by a desire to correct choice errors but is a way to transfer income across idiosyncratic states.

5.1 Introducing Heterogeneity

It is well-known that a social preference for redistribution may affect the use of commodity taxes. This is indeed also the case here if the only instrument available to the government is a sin tax. In this case, it is easily shown that the condition for a small

[11] The analysis so far has focused on the case of continuous consumption. It may well be argued that for many sin goods the extensive margin is more important than the intensive margin. While there is arguably some dispersion in daily cigarette consumption among smokers, the more important margin appears to be the decision of whether to smoke or not. The results in this section are robust to an alternative model where the extensive margin is the only margin of response. In this case, the government's trade-off will be better if the distribution of gains from consumption is more dispersed. This opens the possibility for subsidies that in certain ranges have little or no effect on consumption decisions.

subsidy of the sin good to be welfare improving is

$$\sum_{i \in N} p_i \left(z_i^e - Z\right) \frac{\partial V_i^e}{\partial y} + \sum_{i \in N} \left(1 - p_i\right) \left(z_i - Z\right) \frac{\partial V_i}{\partial y} > 0,$$

where $Z = \frac{1}{N} \sum_i \left(p_i z_i^e + \left(1 - p_i\right) z_i\right)$ is average consumption of the sin good. The two terms give the covariance between consumption of the sin good and the marginal utility of income for exposees and non-exposees, respectively. Thus, on average there must be a positive correlation between consumption of the sin good and the marginal utility of income. But unlike the case of only one agent, differences in marginal utilities of income now derive from both the cue shock and inate heterogeneity in abilities. The condition above reflects the presence of cue shocks by distinguishing the two states such that there are in essence $2N$ types. Importantly, it may be optimal to subsidize the sin good even if there is a negative correlation between ability and sin good consumption in the normal mode (non-exposure). This could happen if exposure increases the consumption of low ability individuals relatively more than it does for high ability individuals. In this model this is equivalent to a greater sensitivy towards cue exposure. If, for instance, there is a small negative correlation between sin good consumption and ability in the normal mode, but the reference point of low ability individuals is sufficiently greater than that of high ability people, the correlation could be reversed for the group of exposed. This case is more likely to lead to an optimal subsidy if the risk of exposure is also greater for low ability groups. However, this shrouds the fact that with heterogeneous earnings the case for a subsidy is no longer driven by an individual demand for commodity specific insurance. The government's instrument is made less effective as it can no longer target its redistribution at the exposed group. Thus, in this model it may be optimal to subsidize the sin good in a representative agent economy but inoptimal in a heteregeneous economy. Note, that even if the government cared not for redistribution based on income, the case for subsidizing the sin good would be weakened in the presence of heterogeneity if the set of instruments is limited to a lump sum tax and a sin tax.[12]

[12] This could be captured by introducing social welfare weights that are higher for people with high income. If appropriately chosen, there would be no desire to redistribute across groups but the case for redistributing within a group would persist.

This is because a sin subsidy would redistribute on balance to those with higher than average consumption of the sin good and this would no longer be perfectly correlated with cue exposure. Rather, insurance would be provided only to those exposees that consume more than average. Only if the effect of cues is so strong that all exposed individuals consume more than the average sin good consumption in the economy is the case for commodity specific insurance unweakened. Effectively, with heterogeneity the price of commodity specific insurance is distorted relative to fairness.

It is interesting what happens if the government has access to a non-linear income tax. As mentioned above, the literature on optimal commodity taxation has demonstrated, starting with Anthony Atkinson and Joseph Stiglitz (1976) and elaborated by Saez (2002), the general uselessness of commodity taxation. This result is derived within a model where there is no argument for a non-zero sin tax in the absence of an exogenous revenue requirement and the presence of homogeneity. This contrasts to the case for commodity specific insurance presented here. It is, however, possible to use the framework of Saez (2002) to examine the robustness of Proposition 1 to the availability of a non-linear income tax. Indeed, access to an income tax revitalises the case for commodity specific insurance relative to when the government has access only to a lump sum instrument. To see this, note that with exogenous labor supply every income group is completely identified by the inate ability of its members.[13] In this case, equation (5) in Saez (2002) effectively reduces to

$$\frac{dW}{dt} = \sum_{i \in N} p_i \left(1 - p_i\right) \left(z_i^e - z_i\right) \left(\frac{\partial V_i}{\partial y} - \frac{\partial V_i^e}{\partial y}\right),$$

which is negative, implying the optimality of a sin good subsidy. Intuitively, with a non-linear income tax, the government can finance the subsidy in a non-uniform way allowing it to provide commodity specific insurance to all income groups.[14]

[13] The assumption of exogenous labor supply is not a requirement for the argument but serves as a simplification. With endogeneous labor supply, exposed individuals work more than non-exposed which increases their income. In this case, it is sufficient for satisfying the Saez condition that sin good demand be non-increasing in ability.

[14] Of course, if there were heterogeneity of preferences conditional on ability and state some members of each group would face a more favorable 'price' for insurance than others and a subsidy may not be optimal.

6 The Case of Bounded Rationality

In this section, the change in behavior upon cue exposure is attributed to bounded rationality as in Bernheim and Rangel (2004). Thus, the short term feeling of temptation that results from exposure to a cue is considered an 'illusion of the brain'. In chemical terms, the brain's release of dopamine as a predictor of the pleasure of intake is disturbed by the presence of the cue. This causes the brain to overestimate the pleasure from consuming. As a consequence, if left to herself, an agent's choices will be suboptimal in the exposed state of mind. Specifically, the agent acts as if true utility were given by (1) and thus were state-dependent. But in fact, utility does not vary with the state.[15] Thus, independently of cue exposure welfare is given by

$$u_i(x, z),$$

while decisions are made in order to maximize (1) after the realization of the idiosyncratic state. It is clear that this leads the agent to overconsume the sin good in the exposed state.[16] This creates scope for government intervention. In the previous section, this was motivated by a demand for insurance that could not conceivably be satisfied by the private market. The case for insurance continues to hold as cue exposure adversely affects the disposable income for goods other than the sin good. But at the same time, there is a need to correct an error in individual choices. With perfect information this trade-off would not exist but when the government cannot condition commodity taxes and transfers on exposure, the concern for income distribution would likely affect the results. However, the subsequent analysis makes the simplifying assumption that utility is quasi-linear, eliminating the demand for insurance.[17] Also, only the case of no cross-sectional variation in earnings ability is considered. In contrast, differences in preference

[15] This can be considered a normalization. It is straigthforward to allow some state-dependence of true preferences but assume that the perceived reference point is higher than the true reference point.

[16] See the chapter by Bernheim and Rangel in Peter Diamond and Hannu Vartiainen (forthcoming) on welfare evaluations with bounded rationality and a discussion specifically within the context of their 2004 model.

[17] Of course, with a concave social welfare function the government would still wish to redistribute to those with low earnings which would be important in a situation of earnings heterogeneity.

intensity for the sin good are allowed.

Assume utility is quasi-linear with a constant marginal utility of consuming the composite good. Further, there is no heterogeneity in earnings. With these assumptions, utility may be written

$$w_i(z) + x$$

Decision utility deviates from welfare by placing weight on the reference point following exposure as described earlier. Thus, consumption choices are made to maximize

$$w_i(z - s \cdot r_i) + x,$$

which deviates from true utility whenever the agent has been subjected to a cue i.e., when $s = 1$. When $s = 0$ there is no conflict between decision utility and true utility.

The government faces a problem of asymmetric information as cue exposure is unobservable and preferences are heterogeneous. Thus, a high consumption of the sin good can be due to either cue exposure or an inately high preference for the good.[18] The government is, therefore, left with the option of affecting behavior only through the use of a distortionary commodity tax. As first shown by O'Donoghue and Rabin (2001) in the context of quasi-hyperbolic discounting, it is optimal for the government to apply at least a small positive tax rate on sin good consumption as long as the demand of exposed individuals is not entirely inelastic. This result easily extends to the present model. Proposition 2 presents the case for an optimal sin tax along with results on the dependence of the optimal tax rate on the frequency (p_i) and impact (r_i) of exposure. The proof is in Appendix B.

Proposition 2. *With quasi-linear utility, inelastic labor supply and no earnings heterogeneity, the optimal sin tax rate,* $t^*(p_i)$*, is positive as long as* $\frac{dz_i^e}{dt}|_{t=0} \neq 0$ *for some* $i \in N$*. Furthermore,*

(a) The optimal sin tax rate is increasing in p_i *for all* $i \in N$ *with* $\frac{dz_i^e}{dt}|_{t=t^*} \neq 0$.

(b) The optimal sin tax rate is increasing in r_i *for all* $i \in N$ *with* $\frac{dz_i^e}{dt}|_{t=t^*} \neq 0$.

[18] This can be due to either large direct gains from consuming or small health costs.

Proposition 2 shows that it is optimal to use a distortionary commodity tax to at least partially correct the 'internality' caused by cue exposure. The intuition underlying Proposition 2 is straightforward and is an example of what the literature has coined 'asymmetric paternalism' (O'Donoghue and Rabin 2003).[19] At a zero tax rate, the true average (for each type) marginal rate of substitution between the sin good and all other goods is less than the price because of the overconsumption of exposees. By raising the tax rate slightly, no harm is done to non-exposees (to a first order) but the welfare of exposees is strictly increased as long as consumption is not entirely price inelastic. Indeed, as is clear from the proof in Appendix B, it is optimal to raise the tax further until the weighted sum across all N types (with weights given by the price responsiveness of sin good consumption for each type) of the difference between the average marginal utility of sin good consumption across states and the marginal rate of transformation equals zero. This will typically involve excessive correction for some types and insufficient correction for others. In any case, taxation interferes with the choices of non-exposees which is at the heart of the trade-off faced by the government.

The lesson from part (a) of Proposition 2 is that there is a close relation between regulation and the optimal sin tax. If the government were to increase regulation of cue emission at a general level, the optimal sin tax rate would decrease. In this sense, sin taxes and cue regulation are substitutes, not complements. In reality, it seems to be more frequently the case that countries with stricter regulation also implement high tax rates. According to the analysis above, this policy is ill-guided. Of course, this is only casual observation and any observed correlation could simply be due to low regulation, low taxation countries having implemented insufficient measures. Importantly, the result shows that regulatory policies aimed only at specific groups should still affect the general sin tax. In this sense will the benefits of targeted regulation be higher if the groups targeted differ vastly from the general population. An example is the case of young people. If the errors of decision-making are large for this group, and they are not

[19]Other terms have been put forth. Richard Thaler and Cass Sunstein (2003) have proposed 'Libertarian Paternalism' while James Choi et al. (2003) suggest 'Benign Paternalism'.

extremely responsive to taxation, there may well be large gains from installing a high degree of targeted regulation and accompany this policy with a decrease in the general sin tax.

Importantly, the scope for using group-specific regulation is dependent not only on a particular group being easily identifiable, which requires its members to have some other observable trait in common, besides unobservable preferences, but also on the technology for regulation. Even if one could identify a certain group, the nature of the group significantly affects the possibilities for implementing targeted regulation. An example where such implementation appears feasible is for the group of young people. Note however, that group-specific regulation could also be interpreted as broad regulation that differentially affects different groups due to differences in preferences. This would strengthen the scope for targeted regulation.

Cue regulation improves the trade-off faced by the government by easing the need to rely on the indiscriminate tax instrument. While the tax instrument has the benefit of only affecting the market price, leaving consumption choices to individual agents, it does force everyone to face the same tax rate regardless of differences in preferences for the sin good. Targeted cue regulation serves the same role as does the sin tax, namely to lower consumption, but does so in a costless way. An example where targeted regulation is both feasible and likely to entail large gains is the case of young people and a wide range of sin goods such as alcohol and cigarettes. Stretching the interpretation of the model somewhat, the indirect utility function may be seen as the expected present discounted value of the future stream of consumption choices with today's choice of sin good consumption entailing long run effects. If for some group the present and future value from consuming the sin good are very uncertain, the utility function is likely to be very concave in the sin good. This will tend to make excessive consumption very damaging. If, at the same time, the sensitivity to cue exposure is high as captured by a high value of r_i, there will be large gains associated with implementing a policy that involves highly intrusive regulation, even extending as far as a complete ban, aimed at the group(s) in question.

So far the discussion has centered on the benefits of using costless cue regulation. However, there may be cases when the use of outright consumption bans for certain groups is desirable. This is the case if for some group the benefits of consuming the good in the non-exposed state are limited and cue exposure leads to significant errors in consumption choices. Depriving certain groups access to consuming the sin good will then allow the government to reduce taxes for the overall population, improving efficiency. This result is in sharp contrast to the suboptimality of bans in the Laibsonian case.

Of course, there may be real-world phenomena not captured by the very stylized model presented here. It may be, for instance, that the population at large fails to understand the full damaging effects of consuming the sin good and that the two signals, regulation and taxation, reinforce each other in the attempt to inform the public of the true costs of consumption. Another important issue is the concern for income distribution. If there is a negative correlation between income and the consumption share of sin goods in total outlays, it may be optimal to lower tax rates for redistributive purposes.[20] Even if this argument is valid, it does not necessarily imply that the tax rate should be decreasing in p. This would depend on the change in the pool of sin good consumers as cue exposure becomes more frequent. One could easily make the argument that as regulation is tightened, the average income and price sensitivity of the consumer pool is reduced, strengthening the case for lower tax rates. This discussion emphasizes the importance of considering heterogeneity in income and a social concern for income distribution.

7 Conclusion

The present paper has explored optimal taxes on cue-triggered consumption goods under two different assumptions about the degree of rationality in individual decision-making.

[20] As is well known from the literature on optimal commodity taxation, the validity of this argument hinges on the income tax system in place. Saez (2002) shows that in the presence of an optimal non-linear income tax, a sufficient condition for the usefullness of commodity taxation is that consumption shares be correlated with ability, not just income.

The policy recommendations from the fully rational Laibsonian case are in stark contrast to those derived within a framework of bounded rationality. In particular, sin goods should be subsidized under fairly general conditions when there are no errors in decision-making. In contrast, when cue exposure increases sin good purchases above the individually optimal amount a corrective sin tax is optimal. In either case, reliance on the price instrument should be reduced as cue regulation is tightened. This implies, for instance, that stricter laws regulating advertising by cigarette companies should be accompanied by lower excise taxes on cigarettes, assuming the tax is initially at its optimal level. Furthermore, outright bans on consumption are never optimal in the Laibsonian case while it may well be an optimal policy when consumers are boundedly rational. Further work on the effects of a social concern for income redistribution and a more precise exploration of the optimal use of different forms of regulation are clearly needed.

References

Atkinson, Anthony B. and Joseph E. Stiglitz (1976). "The design of tax structure: direct versus indirect taxation." *Journal of Public Economics*, 6, 55-75.

Becker, Bary S. and Kevin M. Murphy (1988). "A Theory of Rational Addiction." *Journal of Political Economy*, 96(4), 675-700.

Bernheim, B. Douglas and Antonio Rangel (2004), "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 94(5), 1558-1590.

Bernheim, B. Douglas and Antonio Rangel (forthcoming), "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision-Makers." Forthcoming in Peter A. Diamond and Hannu Vartiainen (eds.), Economic Institutions and Behavioral Economics. Princeton University Press.

Bertrand , Marianne, Sendhil Mullainathan and Eldar Shafir, "A Behavioral-Economics View of Poverty." *American Economic Review Papers & Proceedings*, 94(2), 419-423.

Choi, James J., David I. Laibson, Brigitte C. Madrian and Andrew Metrick. (2003), "Optimal Defaults." *American Economic Review Papers & Proceedings*, 93(2), 180-185.

Diamond, Peter A. (1975), "A many-person Ramsey rule." *Journal of Public Economics*, 4, 335-342.

Fudenberg, Drew and David K. Levine (2006), "A Dual Self Model of Impulse Control." Harvard Institute of Economic Research Discussion Paper No. 2112.

Gruber, Jonathan and Botond Köszegi (2001), "Is Addiction "Rational"? Theory and Evidence." *Quarterly Journal of Economics*, 116(4), 1261-1303.

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and Self-Control." *Econometrica*, 69, 1403-1436.

Gul, Faruk and Wolfgang Pesendorfer (2004), "Self-control and the Theory of Consumption." *Econometrica*, 72, 119-158.

Gul, Faruk and Wolfgang Pesendorfer (2006), "Harmful Addiction." *Review of Economic Studies*, forthcoming.

Hansen, Jes Winther (2005), "Present-Biased Individuals, Optimal Paternalism and Transfers in Kind." EPRU Working Paper 05-11.

Laibson, David I. (2001), "A Cue Theory of Consumption." *Quarterly Journal of Economics*, 116(1), 81-120.

Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory." *Econometrica*, 47(2), 263-292.

O'Donoghue Ted and Matthew Rabin (1999), "Doing It Now or Later." *American Economic Review*, 89(1), 103-124.

O'Donoghue Ted and Matthew Rabin (2001), "Choice and Procrastination." *Quarterly Journal of Economics*, 116(1), 121-160.

O'Donoghue Ted and Matthew Rabin (2003), "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review Papers & Proceedings*, 92(2), 186-191.

O'Donoghue Ted and Matthew Rabin (2006), "Optimal Sin Taxes." *Journal of Public Economics* (forthcoming).

Phelps, Edmund S. and Robert A. Pollak (1968), "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35, 185-199.

Saez, Emmanuel (2002), "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes." *Journal of Public Economics*, 83, 217-230.

Thaler, Richard H. and Cass R. Sunstein (2003), "Libertarian Paternalism." *American Economic Review Papers & Proceedings*, 93(2), 175-179.

Appendix A

Proof of Proposition 1. The social welfare function is simply expected utility for the representative agent

$$W = pV^e(t,\pi) + V(t,\pi)$$

with $\pi + tZ(t,\pi) = 0$ where $Z(t,\pi) = pz^e(t,\pi) + (1-p)z(t,\pi)$ is aggregate demand for the sin good. Using the budget constraint to express π as a function of t we find

$$\frac{d\pi}{dt} = -Z - t\left(\frac{\partial Z}{\partial q} - \frac{\partial Z}{\partial y}\frac{d\pi}{dt}\right) \Longleftrightarrow \frac{d\pi}{dt} = -\frac{Z + t\frac{\partial Z}{\partial q}}{1 - t\frac{\partial Z}{\partial y}} = -\frac{Z\left(1 + \frac{t}{q}\varepsilon_z\right)}{1 - t\frac{\partial Z}{\partial y}},$$

where $\varepsilon_z = \frac{q}{Z}\frac{\partial Z}{\partial q}$ is the aggregate elasticity of sin good demand. The derivative of social welfare w.r.t. t becomes

$$\begin{aligned}
\frac{dW}{dt} &= p\left(\frac{\partial V^e}{\partial q} - \frac{\partial V^e}{\partial y}\frac{d\pi}{dt}\right) + (1-p)\left(\frac{\partial V}{\partial q} - \frac{\partial V}{\partial y}\frac{d\pi}{dt}\right) \\
&= p\left(-z^e - \frac{d\pi}{dt}\right)\frac{\partial V^e}{\partial y} + (1-p)\left(-z - \frac{d\pi}{dt}\right)\frac{\partial V}{\partial y} \\
&= p\left(\frac{Z\left(1 + \frac{t}{q}\varepsilon_z\right)}{1 - t\frac{\partial Z}{\partial y}} - z^e\right)\frac{\partial V^e}{\partial y} + (1-p)\left(\frac{Z\left(1 + \frac{t}{q}\varepsilon_z\right)}{1 - t\frac{\partial Z}{\partial y}} - z\right)\frac{\partial V}{\partial y},
\end{aligned}$$

where use is made of Roy's Identity. Evaluated at $t = 0$ this gives

$$\begin{aligned}
\frac{dW}{dt}|_{t=0} &= p(Z - z^e)\frac{\partial V^e}{\partial y} + (1-p)(Z - z)\frac{\partial V}{\partial y} \\
&= -p(1-p)(z^e - z)\frac{\partial V^e}{\partial y} + p(1-p)(z^e - z)\frac{\partial V}{\partial y}.
\end{aligned}$$

This is negative if and only if

$$\left(z^e(0) - z(0)\right)\left(\frac{\partial V^e}{\partial y} - \frac{\partial V}{\partial y}\right) > 0.$$

Returning to the derivative above, the FOC is

$$p\left(\frac{Z\left(1+\frac{t}{q}\varepsilon_z\right)}{1-t\frac{\partial Z}{\partial y}} - z^e\right)\frac{\partial V^e}{\partial y} + (1-p)\left(\frac{Z\left(1+\frac{t}{q}\varepsilon_z\right)}{1-t\frac{\partial Z}{\partial y}} - z\right)\frac{\partial V}{\partial y} = 0.$$

Complete insurance would equalize marginal utilities across states and a 100% subsidy would achieve this. Using the Slutsky Equation, this policy is only optimal if

$$\begin{aligned}
-Z + p\frac{Z\left(1+\frac{t}{q}\varepsilon_z\right)}{1-t\frac{\partial Z}{\partial y}} + (1-p)\frac{Z\left(1+\frac{t}{q}\varepsilon_z\right)}{1-t\frac{\partial Z}{\partial y}} &= 0 \\
&\Updownarrow \\
t\left(Z\frac{\partial Z}{\partial y} + \frac{\partial Z}{\partial q}\right) &= 0 \\
&\Updownarrow \\
t\left(Z\frac{\partial Z}{\partial y} + \frac{\partial Z^c}{\partial q} - Z\frac{\partial Z}{\partial y}\right) &= 0 \\
&\Updownarrow \\
\frac{\partial Z^c}{\partial q} &= 0,
\end{aligned}$$

as the expenditure function allows perfect aggregation.

Turning to the dependence of t on the frequency of cue exposure, note that the FOC for the choice of sin tax rate is

$$\frac{dW}{dt} = p\frac{dV^e}{dt} + (1-p)\frac{dV}{dt} = 0$$

The derivative of the FOC w.r.t. p evaluated at $t = t^*$ is

$$\begin{aligned}
\frac{d\left(dW/dt\right)}{dp}\Big|_{t=t^*} &= \frac{dV^e}{dt} - \frac{dV}{dt} + p\frac{d\left(dV^e/dt\right)}{dp} + (1-p)\frac{d\left(dV/dt\right)}{dp} \\
&= \frac{dV^e}{dt} - \frac{dV}{dt} - p\frac{d\left(dV^e/dt\right)}{dy}\frac{d\pi}{dp} - (1-p)\frac{d\left(dV/dt\right)}{dy}\frac{d\pi}{dp},
\end{aligned}$$

where

$$\frac{dV^e}{dt} = p\left(-z^e - \frac{d\pi}{dt}\right)\frac{\partial V^e}{\partial y}$$
$$\frac{dV}{dt} = (1-p)\left(-z - \frac{d\pi}{dt}\right)\frac{\partial V}{\partial y}.$$

The signs of these expressions are determined by the signs of the terms in parentheses. From above

$$\text{sign}\left\{-z^e - \frac{d\pi}{dt}\right\} = \text{sign}\left\{Z + t^*\frac{\partial Z}{\partial q} - \left(1 - t^*\frac{\partial Z}{\partial y}\right)z^e\right\}$$

which gives (using the Slutsky Equation)

$$Z - z^e + t^*\frac{\partial Z}{\partial q} + t^*\frac{\partial Z}{\partial y}z^e = (Z - z^e)\left(1 - t^*\frac{\partial Z}{\partial y}\right) + t^*\frac{\partial Z^c}{\partial q} < 0,$$

as $z^e > Z$. A similar argument shows that $\frac{dV}{dt} > 0$.

Now consider $\frac{d\pi}{dp}$. We find

$$\frac{d\pi}{dp}|_{t=t^*} = -t^*\frac{dZ}{dp} = t^*\left(\frac{\partial Z}{\partial y}\frac{d\pi}{dp} - z^e + z\right) \Longrightarrow \frac{d\pi}{dp} = -\frac{t^*\left(z^e - z\right)}{1 - t^*\frac{\partial Z}{\partial y}} > 0$$

It follows that

$$\frac{d\left(dW/dt\right)}{dp}|_{t=t^*} < 0 \Longleftrightarrow \frac{dV^e}{dt} - \frac{dV}{dt} < p\frac{d\left(dV^e/dt\right)}{dy}\frac{d\pi}{dp} + (1-p)\frac{d\left(dV/dt\right)}{dy}\frac{d\pi}{dp}$$

Q.E.D.

Appendix B

Proof of Proposition 2. The government solves

$$\max_{t,\pi} W = \frac{1}{N}\sum_{i\in N}\left[p_i V_i^e\left(t,\pi\right) + \left(1-p_i\right)V_i\left(t,\pi\right)\right]$$
$$s.t. \quad \pi + tZ\left(t\right) = 0,$$

where

$$Z = \frac{1}{N}\sum_{i\in N}\left[p_i z_i^e\left(t\right) + \left(1-p_i\right)z_i\right]$$

is aggregate demand for the sin good. By the quasi-linearity of individual utility and the utilitarian shape of the social welfare function, income distribution is of no importance. Indeed,

$$
\begin{aligned}
& \frac{1}{N}\sum_{i\in N}\left[p_i V_i^e(t,\pi)+(1-p_i)V_i(t,\pi)\right] \\
= & \frac{1}{N}\sum_{i\in N}\left[p_i\left(w_i\left(z_i^e(t)\right)+y_i-\pi-(q+t)z_i^e(t)\right)+(1-p_i)\left(w_i\left(z_i(t)\right)+y_i-\pi-(q+t)z_i(t)\right)\right] \\
= & \pi-t\frac{1}{N}\sum_{i\in N}\left[p_i z_i^e(t)+(1-p_i)z_i\right]+\frac{1}{N}\sum_{i\in N}\begin{bmatrix} p_i\left(w_i\left(z_i^e(t)\right)+y_i-qz_i^e(t)\right) \\ +(1-p_i)\left(w_i\left(z_i(t)\right)+y_i-qz_i(t)\right)\end{bmatrix} \\
= & \frac{1}{N}\sum_{i\in N}\left[p_i\left(w_i\left(z_i^e(t)\right)+y_i-qz_i^e(t)\right)+(1-p_i)\left(w_i\left(z_i(t)\right)+y_i-qz_i(t)\right)\right],
\end{aligned}
$$

from the balanced budget condition. The government then sets t to maximize the above expression. The FOC becomes

$$
\begin{aligned}
\sum_{i\in N}\left[p_i\left(w_{iz}\left(z_i^e(t)\right)-q\right)\frac{dz_i^e(t)}{dt}+(1-p_i)\left(w_{iz}\left(z_i(t)\right)-q\right)\frac{dz_i(t)}{dt}\right] &= 0 \\
&\Updownarrow \\
\sum_{i\in N}\left[p_i w_{iz}\left(z_i^e(t)\right)+(1-p_i)w_{iz}\left(z_i(t)\right)-q\right]\frac{dz_i(t)}{dt} &= 0
\end{aligned}
$$

Evaluated at $t=0$ this becomes

$$
\begin{aligned}
& \sum_{i\in N}\left[p_i\left(w_{iz}\left(z_i^e(0)\right)-q\right)\frac{dz_i^e(0)}{dt}+(1-p_i)\left(w_{iz}\left(z_i(0)\right)-q\right)z_i(t)\frac{dz_i(0)}{dt}\right] \\
= & \sum_{i\in N}p_i\left(w_{iz}\left(z_i^e(0)\right)-q\right)\frac{dz_i^e(0)}{dt} \\
\geq & \ 0,
\end{aligned}
$$

as $\frac{dz_i^e(0)}{dt}\leq 0$ and from the individual FOC, $w_{iz}\left(z_i^e(0)-r_i\right)-q=0\Longrightarrow w_{iz}\left(z_i^e(0)\right)-q<0$. The expression holds with strict inequality whenever $\frac{dz_i^e}{dt}|_{t=0}\neq 0$ for at least one i. It follows by continuity that a small tax is welfare improving relative to no tax. From the concavity of the individual indirect utility functions follows the concavity of social welfare, so the optimal sin tax must be positive.

(a) Differentiating the FOC w.r.t. p_i, holding $t = t^*$ constant, gives

$$\begin{aligned}
& \frac{d\left(dW/dt\right)}{dp_i}\Big|_{t=t^*} \\
= \; & \left(w_{iz}\left(z_i^e\left(t^*\right)\right) - q\right)\frac{dz_i^e\left(t^*\right)}{dt} - \left(w_{iz}\left(z_i\left(t^*\right)\right) - q\right)\frac{dz_i\left(t^*\right)}{dt} \\
& + \left[w_{izz}\left(z_i^e\left(t^*\right)\right) - w_{izz}\left(z_i\left(t^*\right)\right)\right]\frac{dz_i\left(t^*\right)}{dt}\frac{dz_i^e\left(t^*\right)}{dp_i} \\
& + \left[p_i\left(w_{iz}\left(z_i^e\left(t^*\right)\right) - q\right) + \left(1 - p_i\right)\left(w_{iz}\left(z_i\left(t^*\right)\right) - q\right)\right]\frac{d\left(dz_i\left(t^*\right)/dt\right)}{dp_i} \\
= \; & \left(w_{iz}\left(z_i^e\left(t^*\right)\right) - w_{iz}\left(z_i\left(t^*\right)\right)\right)\frac{dz_i\left(t^*\right)}{dt} \\
\geq \; & 0,
\end{aligned}$$

as $\frac{dz_i^e(t)}{dt} = \frac{dz_i(t)}{dt}$ from the relation $z_i^e = z_i + r_i$ and $\frac{dz_i^e(t^*)}{dp_i} = 0$ as there are no income effect on sin good consumption. The inequality is strict whenever $\frac{dz_i(t^*)}{dt} \neq 0$. It follows that the optimal sin tax rate increases with p_i.

(b) Taking the derivative of the FOC w.r.t. r_i at $t = t^*$ gives

$$p_i w_{izz}\left(z_i^e\left(t^*\right)\right)\frac{dz_i^e\left(t^*\right)}{dt} > 0,$$

as $\frac{dz_i^e(t)}{dr_i} = 1$ and $\frac{d(dz_i(t^*)/dt)}{dr_i} = 0$. Q.E.D.