

NBER WORKING PAPER SERIES

IS ADDICTION “RATIONAL”? THEORY AND EVIDENCE

Jonathan Gruber
Botond Köszegi

Working Paper 7507
<http://www.nber.org/papers/w7507>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2000

We are grateful to Peter Diamond, Michael Grossman, David Laibson, and Sendhil Mullainathan for helpful conversations, to seminar participants at Princeton University, the University of Chicago, rochester University, the University of Texas, Western Michigan University, and the NBER for helpful comments, and to Jim Berry, Jack Gallagher, Kokkeong Puah, choonsiang Tan, and Jon Zinman for research assistance. Gruber gratefully acknowledges funding from the National Science Foundation and the National Institute on Aging. Köszegi gratefully acknowledges support from the NBER’s Fellowship on Aging. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by Jonathan Gruber and Botond Köszegi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Is Addiction “Rational”? Theory and Evidence
Jonathan Gruber and Botond Köszegi
NBER Working Paper No. 7507
January 2000
JEL No. I18, H21

ABSTRACT

A standard model of addictive process is Becker and Murphy’s “rational addiction” model, which has the key empirical prediction that the current consumption of addictive goods should respond to future prices, and the key normative prediction that the optimal government regulation of addictive goods should depend only on their interpersonal externalities. While a variety of previous studies have supported this empirical contention, we demonstrate that these results are very fragile. We propose a new empirical test for the case of cigarettes, using state excise tax increases that have been legislatively enacted but are not yet effective, and monthly data on consumption. We find strong evidence that consumption drops when there are announced future tax increases, providing more robust support for the key empirical prediction of the Becker and Murphy model. But we also propose a new formulation of this model that makes only one change, albeit a major one: the incorporation of the inconsistent preferences which are likely to provide a much better platform for understanding the smoking decision. We find that with these preferences the model continues to yield the predictions for forward-looking behavior that have been tested by others and by ourselves. But it has strikingly different normative implications, as with these preferences optimal government policy should depend as well on the “internalities” imposed by smokers on themselves. We estimate that the optimal tax per pack of cigarettes should be at least one dollar higher under our formulation than in the rational addiction case.

Jonathan Gruber
MIT Department of Economics
E52-355
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
gruberj@mit.edu

Botond Köszegi
MIT
E52-262a
50 Memorial Drive
Cambridge, MA 02142-1347
koszegi@mit.edu

1 Introduction

Many of the habits that pervade everyday life can be properly described as addictive. While the degree of addictiveness varies from activity to activity and person to person, habits such as smoking, drinking, eating, exercise and a host of others often meet the two conditions required for addiction: reinforcement, in that the more you partake of the activity, the more you want to partake; and tolerance, in that the more that you partake of the activity, the lower your future utility given the amount of future consumption (Becker and Murphy, [3]). The importance of addiction for a variety of aspects of consumption behavior has led to a long standing interest in modeling addictive processes. Most of the literature in this area until the mid-1980s modeled addiction as habit formation, capturing the reinforcement aspect of the process through an effect of lagged consumption on the taste for today's consumption of the good.

In a pathbreaking article, Becker and Murphy [3] explored the detailed dynamic behavior of the consumption of addictive goods, and pointed out that many phenomena previously thought to have been irrational are consistent with optimization according to stable preferences. In the Becker and Murphy model, individuals recognize the addictive nature of choices that they make, but may still make them because the gains from the activity exceed any costs through future addiction. That is, in this rational addiction framework, individuals recognize the full price of addictive consumption goods: both the current monetary price, and the cost in terms of future addiction.

This model of rational addiction has subsequently become the standard approach to modeling consumption of goods such as cigarettes. This standard has been reinforced by a sizeable empirical literature, beginning with Becker, Grossman and Murphy (BGM, [4]), which has tested and generally supported the key empirical contention of the Becker and Murphy model: that consumption of addictive goods today will depend not only on past consumption but on future consumption as well. More specifically, this literature has generally assessed whether higher prices next year lead to lower consumption today, as would be expected with forward looking addicts. The fairly consistent findings across a variety of papers that this is the case has led to the acceptance of this framework for modeling addiction.

These past tests, however, run into a number of critical empirical and theoretical problems. On

the empirical side, they rely on the assumption that individuals are appropriately forecasting prices far in advance (as much as one year); as we document below, for cigarettes at least, very few price increases are announced this far in advance. More fundamentally, in many other applications, the fact that the lead of a price variable affects current behavior is taken as the failure of a specification test of the model, not as evidence of forward looking behavior.

Finally, even if forward looking behavior can be demonstrated convincingly, there is a more fundamental theoretical problem: forward looking behavior does not imply time consistency. A key assumption of the rational addiction framework is that individuals are time consistent; their future behavior coincides with their current desires regarding this behavior. But recent developments in behavioral economics have suggested that attention be paid to alternative, time inconsistent models which may be more appropriate for modeling addiction. Importantly, except in extreme cases, these models also imply forward looking consumption decisions. Yet, the implications of these models for government policy are radically different. In particular, while the rational addiction model implies that the optimal tax on addictive bads should depend only on the externalities that their use imposes on society, the time inconsistent alternative suggests a much higher tax that depends also on the "internalities" that use imposes on consumers.

The purpose of our paper is to address both these empirical and theoretical issues with the rational addiction literature, in the context of cigarette consumption. We begin by documenting the problems with previous tests of rational addiction models, and show that sensible changes to the specification completely destroy evidence of adjacent complementarity. We then suggest an alternative test: examining how consumption changes when a tax change is actually announced, but not yet effective. We do so using monthly data on cigarette consumption, as well as sales, matched to information on the enactment and effective dates of all state level cigarette excise tax increases over the recent past. We find, in fact, that in this framework there is evidence for forward looking behavior; cigarette consumption does fall when future price increases are announced but not yet effective. This finding is also robust to the specification tests which prove difficult for previous tests to pass.

We then turn to developing an alternative model which is also consistent with forward looking consumption decisions. We do so by embedding in the Becker-Murphy framework the hyperbolic

discounting preferences pioneered by Laibson [19]. These preferences provide a sensible parameterization which allows us to maintain the optimizing features of the Becker-Murphy framework, while considering time inconsistency in the decision to smoke. We find that this model also generates the prediction that future prices matter for today's consumption; indeed, they matter in ways that are sufficiently similar to the Becker-Murphy model that we are unable to empirically distinguish the two with our data. Yet, we show that this model can deliver very different implications for government policy; our simulations suggest that the optimal tax can be at least a dollar higher for modest time inconsistency in this framework. We also develop a host of additional interesting implications for government policy that arise from this alternative framework for modeling addiction.

Our paper proceeds as follows. In Part II, we review both the Becker-Murphy model and the related empirical literature, and document the failure of the previous empirical framework to deliver a convincing test of forward looking behavior. In Part III, we describe our empirical strategy for testing for forward looking behavior, and we implement this test in Part IV. Part V develops our alternative model of time inconsistent addiction, Part VI solves the model, and Part VII explores the implications of price changes in this framework. Part VIII discusses the implications of the different models for government policy.

2 Previous Work

Models of the consumption of addictive goods have a long tradition. Most of the literature until the mid-1980s focused on the habit formation, or reinforcement, aspect of addictive processes. This aspect leads naturally to the prediction that current consumption of addictive goods will be dependent on the path of past consumption, and a number of articles have demonstrated for goods like cigarettes this backwards-looking intertemporal correlation ¹.

Becker and Murphy [3] presented a novel analysis which greatly advanced the modeling of addictive processes. The key insight of their model was that just because a good is addictive, there is no reason that its consumption can't be analyzed in a standard rationally optimizing framework. Their "rational addicts", in making consumption decisions, recognize that there is a

¹ See Chaloupka and Warner [10] for a superb review of both the theoretical and empirical literatures in this area.

tradeoff with current consumption: while utility rises today from the consumption, long run utility is lower because the individual is building up a stock of the addictive good that has a negative marginal utility. Individuals rationally trade off these factors to consider the appropriate level of consumption of addictive goods.

A key implication of this model is that consumption behavior should exhibit “adjacent complementarity”. Reinforcement arises here through the fact that a larger stock of past consumption raises the marginal utility of current consumption. Thus, the fact that individuals are going to pursue the activity in the future should increase the pursuit today, so as to increase the enjoyment of the activity next period. This insight has led to the central empirical test of the rational addiction model: asking whether consumption today is dependent on consumption tomorrow.

The first paper to carry out this test was Becker, Grossman, and Murphy [4], focusing on cigarette smoking as an addictive behavior. They compile a dataset of cigarette consumption and prices across the U.S. states over the 1955 through 1985 period, and match that to information on cigarette prices across the states. They then estimate models of current consumption as a function of current prices, consumption in the previous year, and consumption in the next year. They recognize the important problem that consumption in both the past and the future are endogenous, and propose an instrumental variables strategy that uses both past and future prices and taxes as instruments. Doing so, they find significant effects of future consumption on current consumption, supporting the forward looking behavior implied by the Becker and Murphy model. This type of test has been carried out by a variety of subsequent studies, on both cigarettes and other substances ².

To inform our discussion of their results, we have replicated the BGM test using data updated through 1997. We have used their same source for consumption, price, and tax information, the (now defunct) Tobacco Institute’s volume The Tax Burden On Tobacco, and have followed their methodology with respect to the construction of the tax and price variables used ³. We do not

² A less than comprehensive list includes Chaloupka [9], Sung, Hu, and Keeler [33], Waters and Sloan [36], Olekalns and Bardsley [29], Grossman and Chaloupka [16], and Grossman, Chaloupka, and Sirtalan [16]. See Chaloupka and Warner [10] for a review.

³ In particular, they construct the average state excise tax rate during the fiscal year over which consumption is measured, using monthly tax data; they incorporate (average) municipal as well as statewide cigarette excise taxes; and they take a weighted average of lagged and current annual price data, incorporating the actual taxes in place over the fiscal year.

include their control variables for income and smuggling, but as we can replicate their basic pattern of results these do not appear critical. All price and tax coefficients are in 1967 dollars, following BGM.

The results of this replication are presented in Table 1. The first column shows BGM's column (1) from Table 3 on page 406 of their 1994 article, with lagged and leaded consumption instrumented by lagged and leaded prices. The second column shows our replication, extending the data through 1997. Our pattern of findings is very similar: significant coefficients on both lagged and lead consumption, with the latter roughly one-third as large as the former.

In the next column, we show the reduced form equivalent of this regression. As would be expected, there is a significant negative coefficient on the future price, consistent with the adjacent complementarity predictions of the Becker and Murphy model. That is, as future prices rise, individuals will know that they will desire lower consumption tomorrow, so they lower their consumption today.

Unfortunately, this test runs into a host of problems. First, conceptually, it is difficult to conceive of individuals being able to forecast well future prices in their state of residence, even if they simply are trying to forecast tax changes. As we document in more detail below, excise tax changes are rarely known one year in advance; only 8 of 160 tax changes over the 1973-96 period were enacted as far as one year in advance. For individuals to forecast prices this far into the future would require a very sophisticated model of expectations.

Second, the dependent variable is sales of cigarettes, not consumption; in particular, this represents sales from wholesalers to retail distributors of cigarettes⁴. If individuals really did anticipate future price changes, then the expected direction of the response is not obvious; to the extent that individuals wish to stockpile cigarettes while they are less expensive, consumption could actually rise in anticipation of price increases.

At an annual frequency, this may not be a major concern, due to cigarette quality deterioration for long periods of storage. But this point interacts with the previous one: if the price change is far in the future, there is unlikely to be stockpiling, but the change is also unlikely to be anticipated;

⁴ While the other criticisms levied here apply to all of the studies in this literature, this one only applies to the subset of studies that use aggregate sales, rather than individual consumption, data.

if the price change is in the near term, then anticipation is more likely, but so is stockpiling.

Third, there may be endogeneity bias to regressing the quantity of cigarettes consumed on their price. This bias is likely to be small, since the primary determinant in within-state specific price changes is changes in excise taxes; existing evidence suggests that excise tax changes are passed through on a slightly more than one-for-one basis (Federal Trade Commission, 1997). But tax changes explain only about 80% of the within state-year variation in prices, so that there is remaining variation in the price that could lead to endogeneity bias in the price-consumption relationship. The true exogenous variation that should be used to identify this model is taxes. However, as column (4) of Table 1 shows, the results are weakened considerably when taxes are used in place of prices in the reduced form specification; the future tax term is significant only at the 10% level. Indeed, when we restrict the analysis to the 1955-85 period used by BGM, the coefficient on future taxes is positive.⁵

Fourth, and perhaps most importantly, this test is unable to distinguish true future price effects from other failures of the fixed effects specification. It is plausible that over such a long time period state effects are not truly fixed. If, for any reason, prices or taxes are slowly rising over time in the states where smoking is falling the most, then this will lead to a finding that future prices are correlated with current consumption. Indeed, in many other applications, the fact that the lead of a price variable affects current behavior is taken as the failure of a specification test of the model, not as evidence of forward looking behavior.⁶ This relationship between price and lagged consumption is consistent with Showalter [32], who documents that an oligopolistic tobacco manufacturer facing a relatively inelastic demand for cigarettes will react to declining consumption by raising price. The same behavior may be true of revenue-maximizing state governments faced with declining cigarette demand, leading to the observed correlation even with taxes.

While this alternative cannot cleanly be distinguished, two simple tests can demonstrate the fragility of the BGM findings to particular specifications of the alternative model. The first is to allow for additional flexibility in the specification by including state-specific time trends. This will control for any slowly moving correlated trends within states in the tastes for smoking and

⁵ The elasticities presented for the tax coefficients are price elasticities, which are evaluated by first estimating models of price as a function of tax, and using the resulting coefficients to estimate pass-through; the coefficients imply pass-through of taxes to prices of 108 to 116%, depending on the specification

⁶ For a discussion of identification issues in a panel data context, see Chamberlain [11].

tax/pricing policy. An alternative test is to move from a fixed effects to a changes specification. If state fixed effects are indeed fixed, then these two alternatives should yield identical answers. But, if there are dependencies of prices on past consumption, then changes can yield very different answers. Neither of these tests is particularly strong, in that they cannot distinguish year-year reverse causality from adjacent complementarity, but they provide reasonable starting points for examining specification fragility. Indeed, a changes specification of the relationship seems a more natural one for testing for adjacent complementarity: it asks whether when prices rise from this year to the next, consumption falls this year.

As we see in the remaining columns of Table 1, however, when we move from a fixed effects specification to either a fixed trends or changes specification, the results for the lead of prices or taxes are greatly weakened. The future price coefficient drops below its standard error, and the future tax coefficient actually becomes positive and, in the last column, statistically significant. These changes are not driven by decreased precision; both the fixed trends and changes specification actually produce smaller standard errors on the lead price/tax term than does the fixed effects specification. Moreover, the coefficient on the contemporaneous price/tax term itself is robust to these alternative specifications, so that the basic demand relationship is not sensitive to functional form.

Overall, the structure of the existing tests of adjacent complementarity appears problematic along a number of dimensions. We next present a test that is designed to remedy these deficiencies, in several ways. First, will rely only on tax increases that have already been announced to identify our anticipatory effect. Second, we will use data on actual cigarette consumption, rather than sales. Third, we will use information on tax changes, not price changes. Finally, we will examine very high frequency changes which are unlikely to suffer from the type of bias that hinders testing for forward looking behavior in annual data; we present tests below to demonstrate the validity of this claim.

3 A New Test of Forward Looking Behavior - Empirical Strategy

Data

We propose an alternative test of forward looking behavior that is consistent in spirit with the test employed by BGM, but improves on the problems noted above. In particular, we have collected from state legislative histories since 1973 the date of the legislative enactment of state excise tax increases, and the date that they were actually effective. By examining cigarette consumption in the intervening period, we can test for forward looking behavior. If individuals are forward looking, but cannot forecast taxes beyond already announced tax increases, then this provides a more appropriate framework for examining adjacent complementarity.

We summarize the information on these tax changes in Table 2, which shows the period of time between the enactment and effective dates of state excise tax increases. Over the full 1973-1996 period, 36 tax changes were enacted and effective in the same month, and 44 in consecutive months. Yet 68 tax changes had at least one month between the enactment and effective dates. The longest gap between enactment and effective dates was 7 months.

In addition, there were a number of examples of multiple tax increases that were enacted on the same date. The first such change is included in the first 8 rows of the table. The last two rows show that there were 12 second or third changes from such multiple change examples, and that 8 of them were effective more than one year after being enacted. In the empirical work below, we only use the first effective date for such changes.

We have collected two sources of data on cigarette consumption to test for anticipatory responses. The first is the monthly series of tax-paid cigarette sales that underlies the annual data used by BGM and others (as well as in the regressions in Table 1). This was constructed using data from state excise tax collections, as archived at the Tobacco Institute and the North Carolina State University. As with the annual data, these represent withdrawals from wholesale distributors, since this is the point at which the excise tax is paid. We have collected these data from January 1982 through December 1996, with the exception of September and October 1982, for which data were not available. As noted above, however, and as will be documented further below, it is problematic to use data on cigarette sales to test for anticipatory price responses of consumption.

At the same time, data on cigarette consumption for this purpose must meet a difficult criterion, as they must provide state by *month* observations of enough size so as to form reasonable proxies for cigarette consumption. Fortunately, there is a dataset which meets this condition: the Vital Statistics Natality Data. Since 1989, this database has recorded, for every birth in America, whether the mother smokes and how much, as well as the state of residence of the mother. As a result, there are roughly 4 million observations per year on smoking behavior, providing sufficient sample size to measure state by month smoking rates; in our final database, the typical state month cell has 5320 observations. This is clearly not a representative population, but it is a population of particular interest, since maternal smoking and poor subsequent infant health is perhaps the leading externality associated with smoking behavior (Evans, Ringel, and Stech [14]).

We use the full set of 1989-1996 Natality files to measure monthly smoking rates for every state for which the smoking information was collected ⁷. The key question of interest asks women about smoking during pregnancy; we assume that these women are answering with reference to the month of birth. To the extent that they are answering with reference to smoking anytime during pregnancy, we will understate the responsiveness to future increases.

Our dependent variable is the number of cigarettes smoked each day per woman in each state/month cell, which is formed by dividing total cigarettes smoked per day among smoking women by the total number of women in the cell. As Table 3 shows, the weighted (by number of births in a cell) mean of this variable over the 1989-96 period is slightly less than 2. This implies that per capita monthly consumption of (20 cigarette) packs of cigarettes is about 3. This compares to per capital annual packs of cigarettes sold, from our monthly sales data, of 9. The fact that this figure is lower is not surprising, as the women in this sample are less likely to smoke than the typical person, and they smoke less intensively when they do smoke. Overall, as we show in Table 3, the smoking rate for our sample of women is 16.2%, and those who smoke consume on average only about two-thirds of a pack per day. Averaging over all smokers over age 18 for 1989-1996, using data from the Behavioral Risk Factor Surveillance System data, the average smoking rate is 23%, and the average cigarettes smoked per day per smoker is 18.7.

⁷ Smoking data are not available for California, Indiana, and South Dakota in any year; for New York for 1989-93; for Oklahoma for 1989-90; and for Louisiana and Nebraska for 1989.

One concern with these data is that mothers may underreport smoking while pregnant. While we cannot definitively address this concern, it is noticeable that the smoking participation rate in these data is almost exactly the same as that from a National Health Interview Survey supplement in 1991 which provides a retrospective survey of women on their smoking while pregnant.⁸ So there doesn't appear to be any systematic underreporting on birth certificates relative to other surveys. Moreover, underreporting would not lead to a systematic bias to the estimates unless it is somehow correlated with price changes, which seems unlikely.

Empirical Strategy

Our empirical strategy follows that described with reference to Table 1 for examining the impact of current and future prices. We run regressions of the form:

$$SMOK_{sm} = \alpha + \beta * EFFECT_{sm} + \gamma * ENACT_{sm} + \delta * M_m + \phi * S_s + \epsilon \quad (1)$$

where $SMOK$ is the measure of smoking in state s in month m ; $EFFECT$ is the effective tax rate in that state and month; $ENACT$ is the enacted tax rate in that state and month; and M and S are full sets of month (we include dummies for each calendar month in our sample period) and state dummies, respectively. $ENACT$ is the same as $EFFECT$ except when a change has been enacted and not yet effective, so this is our future price variable. For these regressions we exclude both the months in which tax changes are enacted and they are effective, since both of these events can happen at any time during the month, so that the response in that month may be quite muted. In some specifications, we also include a lagged value of the effective rate; we use a twice lagged tax rate, since we are excluding the month of the tax change. All natality data regressions are weighted by the size of the cell to reflect sampling variability in our aggregation strategy.

Note that we use everywhere taxes, and not prices, to test for forward looking behavior, because price data is not available on such a high frequency basis by state. A legitimate question is then whether a finding of anticipatory behavior reflects anticipation by consumers or producers; if cigarette prices in a state are increased in anticipation of tax changes, then demand may be falling through the standard law of demand. This issue is raised by Showalter [32], who finds no evidence of anticipatory pricing at an annual frequency.

⁸ The NHIS supplement data indicate that 20.6% of women smoked at some time during their pregnancy, and 16.6% smoked throughout the pregnancy. The fact that the latter figure so closely matches our data provides further suggestive evidence that women are responding to this question with reference to the month of birth.

While monthly data on cigarette prices are not available, there is quarterly data from the American Chamber of Commerce Research Association (ACCRA) on prices for selected cities of a carton of Winston cigarettes [27]. We have used these data to investigate anticipatory pricing for tax changes where the announcement of the price increase is in a different quarter than the implementation. We found no evidence of sizeable price increases before the tax was actually implemented. For example, the state of Alaska enacted a 70 cent tax increase at the end of May, 1997, which was to be implemented in October. Yet the price of cigarettes rose by only 4 cents in the third quarter before rising by 88 cents in the fourth. Similarly, Michigan at the end of 1993 enacted a 50 cent tax increase which was to be implemented in May. There was only a 5 cent per pack price increase from the end of 1993 through the second quarter of 1994, and then a 50 cent increase in the third quarter. These findings suggest that producers are not increasing prices in anticipation of excise tax increases. This supposition is confirmed below by the fact that we find sales dramatically increasing in anticipation of tax increases, which would not occur if prices had already risen.

A key assumption of our alternative strategy is that over the very high frequency changes that we are observing, state fixed effects can indeed be taken as fixed. That is, whatever is causing the long run correlation between consumption and taxes in the BGM-type specifications will not bias our estimates of the impact of tax increases in the next several months on this month's consumption. This assumption strikes us as reasonable, since state excise tax decisions seem unlikely to be responding to smoking rates in the very recent past, and since we are using a much shorter time frame over which the constant state fixed effects assumption is to be imposed. But we will nevertheless subject our finding to the same set of tests that we applied to the BGM results to demonstrate that it is robust.

4 A New Test - Results

The results of using these two data sources to examine the impact of future tax increases are presented in Table 4. We begin with the packs/capita aggregate sales data. We find from this

data a strong negative effect of the current effective tax rate, with a price elasticity of -0.8 ⁹. This elasticity is somewhat higher than that estimated at an annual frequency in Table 1, and is much higher than that found by BGM. The difference appears to arise from the fact that the tax-induced movement in prices causes larger consumption declines than does the price-induced movement in prices.

In the next row, we include the future tax change term, as well as lagged effective taxes. In fact, the coefficient on the announced rate is actually positive and highly significant. At a monthly frequency, such a positive reaction to future price increases is sensible, as consumers hoard cigarettes at lower prices for future use. This hoarding effect is consistent with the evidence in Keeler et al. (1993), who find cigarette sales rising in the months before a 1989 excise tax increase in California. Indeed, we see this response in our data, for the large increase in the excise tax from 10 to 35 cents in California that was announced in November 1988 and effective in January 1989. In November, cigarette sales were just slightly down from what they had been the previous November, at 6.68 packs per capita. Then, in December, sales jumped to 8.71 packs per capita, before falling back to around 6 over the next few months.

This sizeable hoarding effect could be taken as one type of forward looking behavior by consumers, in that they are stocking up in anticipation of a tax increase. It is not clear how much of the hoarding effect we find is due to consumer vs. retailer behavior, as some state excise tax increases exempt floorstocks held by retailers when the tax changes. For a sample of the 12 largest tax changes in recent years, we found in the legislation or through contact with state taxation officials that in 10 of 12 cases floorstocks were included, so that any hoarding effect would be due to consumers. In either case, this sizeable hoarding effect casts doubt on the usefulness of sales data for testing for anticipatory consumption behavior.

We also find that including the enacted and lagged rate significantly increases the term on the current price, which now implies an elasticity of -1.5 , with a sizeable and significant positive elasticity on the lag as well as the lead. The positive impact of the lagged rate, and the large value

⁹ The magnitudes of these coefficients differ from those in Table 1 for three reasons: we deflate taxes by \$1982 here, rather than the \$1967 used in Table 1 (which lowers the coefficients by a factor of 0.334); we are using monthly data here, rather than annual data (which lowers the coefficients by a factor of 12); and we use the tax rate in dollars, rather than cents (which raises the coefficients by a factor of 100). But the elasticities are of course directly comparable. Once again, we use the coefficients from the regressions in the aggregate annual data of price on taxes to interpret these tax effects as price elasticities.

of the current tax rate, no doubt reflects monthly timing of purchases in our data: if individuals are hoarding in the months before a tax change, then sales will fall most sharply in the month of the change, before rising again somewhat thereafter. This is what we see in the second column of Table 4: a rise in sales in the months before a change, a sharp decline in the month of the change, and then an offsetting increase thereafter.

The next two columns of Table 4 consider the impact of effective and enacted taxes in our natality data. Here, when we just include the effective tax rate term, we find a much smaller impact of taxes, with a price elasticity of just -0.35. This is consistent with the notion that women who are still smoking at the time that they are giving birth may be less sensitive to economic factors such as prices. Nevertheless, the coefficient is highly significant, so that if there is an anticipatory response we should be able to estimate it with this data.

When we include the enacted tax rate as well in the next row of Table 4, the coefficient on the enacted rate is in fact negative and significant. Thus, using this more refined test, we find strong evidence of forward looking behavior, for this population at least. Including the enacted and lagged rate leads in this case to a sizeable fall in the current tax coefficient, with a relatively sizeable (but insignificant) lag term. Thus, the results indicate that the response to a tax next period is equal to the current plus lagged response to a current tax.

One might be tempted to infer something about the degree of rationality from these relative coefficients. Indeed, BGM do some calculations which infer the discount rate by dividing the coefficient on future and lagged consumption. But in fact this is only technically true under one of two conditions: either if the price changes are temporary, which is only true (in an announced sense) for one tax change in our full sample over this period, or the stock of past smoking depreciates fully in one period (since this makes prices further away irrelevant for today's smoking decision). We prove this in appendix B.2. We discuss below the issues involved in using these types of estimates to infer behavioral parameters.

Specification Testing

It is of course important to subject our estimates to the same scrutiny to which we subjected the BGM results. Table 5 therefore includes fixed trends in our model. For the sales data, doing so lowers the impact of current effective rates, with an implied elasticity of -0.62 from the first column

of Table 5. This is very similar to the fixed trends tax elasticity found in Table 1 from the annual sales data. Once again, we find that when the enacted and lagged tax rates are entered, the current tax effect rises, and the enacted and lagged tax rates are both significant and positive.

For the natality data, the results are also similar when fixed trends are added. We find that the impact of the effective tax rate when entered alone (the third column of Table 5) is weakened, with an elasticity now of -0.27. And we once again find that there is a significant negative impact of enacted but not yet effective taxes; the effect is roughly 80% as large as in the fixed effects model. We find here as well that there is virtually no instantaneous impact of effective tax changes on consumption, with most of the effect showing up with a lag. Once again, the current plus lagged effective tax coefficients are roughly equal to the enacted tax coefficient. Thus, unlike the earlier results from BGM (and presumably from other tests using annual data), our findings are relatively robust to the inclusion of fixed trends.

We have also considered differenced models. Here, we cannot simply difference our monthly panel, since it would imply that for tax changes with several months between enactment and effective, the difference in the enacted rate would be zero after the first month. As we discuss below, we do not have sufficient power in our data to estimate month-by-month responses of tax enactment. Therefore, we have pursued an alternative approach of taking the average smoking level over all months between the enactment and effective dates, and taking the difference between this average and the smoking level in the month before the tax change was enacted. We also include in these differenced regressions, as controls, any months where there were no tax changes enacted in that month or in the two months before; we then take the difference between smoking in that month and smoking two months earlier as a control observation. Our true enactment differences are matched to the change in the enacted rate (the change from the old effective rate to the new enacted rate), and our control observations are matched to a tax change of zero. The regression includes a full set of month dummies.

The results of this exercise, for both packs per capita and cigarette consumption from the natality data, are presented in the final row of Table 5. We find very similar results to the fixed trends specification, albeit with slightly larger standard errors.

One additional concern with our finding is that it reflects not anticipatory responses to future

price changes, but rather changes in reporting in the wake of announced tax increases. That is, it is possible that when future price increases are announced, women become more exposed to anti-smoking sentiment and are thus less likely to report that they smoke. While this concern is impossible to address precisely, we have investigated it in two casual ways. First, for the two announced increases in Massachusetts in 1992 and 1996, both of which had roughly two months between the enacted and effective dates, we examined the major newspapers in the Boston area for any evidence of increased anti-smoking counteradvertising. We found no such advertising in the intervening months.

Second, for the differenced models presented in Table 5 that consider how smoking changes when a tax rate is announced, we have included along with the differenced tax rate a dummy for the presence of a tax change. If it is the announcement of the tax change per se that matters, and not the price change, then the inclusion of this dummy should significantly weaken our differences relationship. In fact, however, there is no impact on our coefficient when the dummy is included (although the standard error does rise by about 50%), and the dummy itself is not significant. This test is not definitive, of course, because anti-smoking rhetoric could be proportional to the size of the tax change. But it certainly suggests that it is not the presence of a tax change per se, but rather the future rise in price, that is causing women to reduce their smoking.

Thus, to summarize, we have provided a more robust framework for testing for anticipatory responses by consumers to future changes in the taxation of cigarettes. Even in this more robust framework, we continue to find evidence of adjacent complementarity. This does not, however, necessarily provide support for Becker and Murphy's formulation of the smoking decision, as we document in the remainder of the paper.

5 The frame of the models

The term "rational addiction" obscures the fact that the Becker and Murphy model imposes two assumptions on consumer behavior. The first is that of forward-looking decision-making, which is hard to impugn and which will be a key feature of our alternative models as well. But the second is the assumption that consumers are time consistent. Psychological evidence documents

overwhelmingly that consumers are time inconsistent [1]. In experimental settings, consumers consistently reveal a lower discount rate when making decisions over time intervals further away than for ones closer to the present, raising the specter of intra-personal conflict over decisions that have implications for the future. There is, to date, little non-experimental evidence for time-inconsistency in decisionmaking. But it is important to note that there is *no* evidence, psychological or other, that supports time-consistent preferences over these time inconsistent ones.

The above type of time inconsistency has been recently applied in the context of savings decisions (Laibson [19], Laibson, Repetto, and Tobacman [20], O'Donoghue and Rabin [26]), retirement decisions (Diamond and Kőszegi [13]), and even growth (Barro [2]). Since smoking is a short-term pleasure, and the psychological evidence indicates that time inconsistency is most prevalent with short horizons, this formulation should be especially fruitful in the context of addictive bads such as smoking.

There is also indirect evidence that people's preferences for smoking are time-inconsistent. Two key features distinguish time consistent and time inconsistent agents. The first is the use of *commitment devices* or *self-control techniques*. We distinguish a self-control device from an alternative technology for smoking cessation, quitting aids: whereas quitting aids decrease the disutility from not smoking, self-control devices lower the utility from smoking. Time-consistent decisionmakers might use a quitting aid, but in general they won't use a self-control device—with time consistency, lowering the utility of an undesired alternative is irrelevant for decisionmaking. But for some types of time inconsistent agents (what we label below *sophisticated* agents, who recognize their own time inconsistency), self-control devices are valued as a means of combating one's own time inconsistent tendencies.

In the relatively small medical literature on self-initiated attempts at quitting smoking, the voluntary use of self-control devices figures prominently. People regularly set up socially managed incentives to refrain from smoking by betting with others, telling them about the decision, and otherwise making it embarrassing to smoke (Prochaska et al [31]). Various punishment and self-control strategies for quitting are also widely studied in controlled experiments on smoking cessation (Miller [22], Murray and Hobbs [23]; see Bernstein [5] for a variety of 'aversive stimulus' techniques), and they are recommended by both academic publications [15] and self-help books [8]. In one study,

for example, subjects tore up a dollar bill for every cigarette they smoked above their given daily limit, and reduced that limit gradually. Presumably, these experiments are incorporating self-control devices because they are seen as the best option for helping individuals quit smoking, as could be the case if individuals were time inconsistent.

A second feature that distinguishes time consistent agents from time inconsistent agents is an inability to actualize predicted or desired future levels of smoking. The former phenomenon is specific to a class of hyperbolic discounters whom we label *naive* below, in that they do not understand that they cannot make consistent plans through time. In fact, unrealized intentions to quit at some future date are a common feature of stated smoker preferences. According to Burns [6], eight of ten smokers in America express a desire to quit their habit. Unfortunately, these desires can be interpreted in a number of ways, and we are not aware of any evidence for adults on their specific predictions or intentions about future smoking behavior. For youths, however, there is clear evidence that they underestimate the future likelihood of smoking. For example, among high school seniors who smoke, 56% say that they won't be smoking 5 years later, but only 31% of them have in fact quit five years hence. Moreover, among those who smoke more than 1 pack/day, the smoking rate five years later among those who stated that they would be smoking (72%) is actually *lower* than the smoking rate among those who stated that they would not be smoking (74%) [28].

Less forceful, but still suggestive, evidence for naive time inconsistency comes from attempted quits. According to Harris [17], 38 of the 46 million smokers in America in 1993 have tried to stop at one point or another, with an average smoker trying to quit once every eight and a half months. Most have tried several times. 54% of serious attempts at quitting fail within one week. These facts do not necessarily contradict a time consistent model which incorporates learning and/or uncertainty, since smokers might experiment with quitting to find out how hard it is or simply 'gamble' in the hope of stumbling on an instance when it's easy. But it seems implausible that smokers learn so slowly or that the situations in which they try quitting are so variable.

Our goal in this section is therefore to take the important insights about forward looking behavior captured in the Becker and Murphy model, and to integrate them with a potentially more realistic description of intertemporal choice in this context. The crucial question we are dealing with is the shape of time discounting. Suppose we are in a T -period decision model. For a

time-consistent agent, discounted utility at time t takes the familiar form

$$\sum_{i=0}^{T-t} \delta^i U_{t+i}, \quad (2)$$

where the U_{t+i} denote the instantaneous utilities. We will contrast this type of discounting with the alternative developed by Laibson [19], quasi-hyperbolic discounting. For quasi-hyperbolic discounters, discounted utility becomes

$$U_t + \beta \sum_{i=1}^{T-t} \delta^i U_{t+i}. \quad (3)$$

β and δ are usually assumed to be between zero and one. The extra discount parameter β is intended to capture the essence of hyperbolic discounting, namely, that the discount factor between consecutive future periods (δ) is larger than between the current period and the next one ($\beta\delta$). At the same time, this formulation still allows one to take advantage of some of the analytical simplicity of the time-consistent model. For a more thorough introduction, see Laibson [19].

Our model marries this intertemporal preference structure with the instantaneous preferences in Becker and Murphy's rational addiction model [3]. We deviate from both Becker and Murphy and Laibson, however, by assuming no savings—some exogenously given income is consumed in each period. The main motivation behind this simplification is to make the sophisticated model tractable. Otherwise, we have to keep track of two state variables (the stock of past smoking and wealth) which evolve differently, and the Euler equation will be considerably more complicated. Although the interaction of savings and addictive behavior is potentially very interesting¹⁰, we are primarily interested in isolating the effect of self-control problems for just the addictive good, which are unlikely to change significantly with the incorporation of a savings decision. In addition, low savings among the low income population that is most likely to smoke renders this assumption relatively innocuous¹¹.

Let a_t and c_t denote, respectively, the consumption of the addictive and the 'ordinary' (non-addictive) goods in period t . Both can take any value on the real line. Furthermore, we denote the

¹⁰ We did solve for the Euler equations with saving, and there are some interesting effects that drop out easily. For example, if wealth makes a person more prone to addiction, she might want to overconsume (consume more than she otherwise would) so as to prevent addiction. Or, if smoking is a substitute for other kinds of consumption, the agent might choose to get addicted on purpose to alleviate her self-control problem in savings. This is the same effect that drives individuals to smoke to lose weight.

¹¹ Of course, there are many addictive goods (cocaine, fine wine, collecting art, etc.) which are very expensive. The no-saving assumption makes our model less applicable to these goods.

period t stock of past consumption by S_t . S_t evolves according to

$$S_{t+1} = (1 - d)(S_t + a_t). \quad (4)$$

d is the depreciation rate of the stock; the higher is d , the less does past behavior influence the stock of accumulated consumption, and thus, indirectly, utility. For notational and arithmetic simplicity, this differs from Becker and Murphy [3], who have $S_{t+1} = (1 - d)S_t + a_t$. As long as depreciation is not full ($d < 1$), their model and our time-consistent case are isomorphic through a simple change of variables.

We assume, as in Becker and Murphy [3], that instantaneous utility is additively separable in these two goods, that is,

$$U_t = U(a_t, c_t, S_t) = v(a_t, S_t) + u(c_t). \quad (5)$$

$v_{as}(a_t, S_t)$ is positive, because consumption of addictive goods generally increases their future marginal utility. Let I_t be period t income and p_t the period t price. We normalize the price of the non-addictive good to be 1.

We consider both agents who discount exponentially (equation 2) and who discount quasi-hyperbolically (equation 3). We will distinguish between two extreme kinds of hyperbolic discounters. *Naive* agents, although they are impatient in the sense that they attach extra value to the current period relative to the future ones, are unaware of their future self-control problem: self t doesn't realize that self $t + 1$ will in turn overvalue period $t + 1$. Thus a naive agent maximizes her intertemporal utility in expression 3, unconscious of the fact that her future selves will change her plans. *Sophisticated* agents, on the other hand, realize their self-control problem: self t knows that self $t + 1$ will want to do something other than what self t would have her do. Therefore the best thing self t can achieve is to make a plan that she will actually follow. Formally, this is modeled as a subgame-perfect equilibrium in a game played by the successive intertemporal selves, the action spaces in our case being the vectors of consumption (a_t, c_t) . See O'Donoghue and Rabin [25] for an excellent discussion of sophistication and naiveté, as well for a few basic behavioral contrasts between the two.

5.1 Time-consistent agents

Standard methods reveal the following Euler-equation for time-consistent agents. The most natural way to think about it is that a small perturbation in consumption in period t that is undone in period $t+1$ doesn't change utility. In contrast to a simple savings problem, however, we also have a $v_s(a_{t+1}, S_{t+1})$ term in the Euler equation, because a change in S_{t+1} affects utility directly, whereas in a savings problem wealth does not.

Lemma 1 *Suppose $u(c_t)$ and $v(a_t, S_t)$ are differentiable. Then, for a time-consistent agent the following Euler-equation holds:*

$$v_a(a_t, S_t) - p_t u'(c_t) = (1 - d)\delta[v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1}) - v_s(a_{t+1}, S_{t+1})] \quad (6)$$

5.2 Naive agents

Naifs solve a very similar maximization problem to time-consistent agents, the only difference being an extra discount factor β between periods t and $t+1$.

Lemma 2 *Suppose $u(c_t)$ and $v(a_t, S_t)$ are differentiable. Then, for naive agents, period t consumption of the addictive and ordinary goods satisfies*

$$v_a(a_t, S_t) - p_t u'(c_t) = (1 - d)\beta\delta[v_a(a_{t+1}^{TC}, S_{t+1}) - p_{t+1}u'(c_{t+1}^{TC}) - v_s(a_{t+1}^{TC}, S_{t+1})] \quad (7)$$

where a_{t+1}^{TC} and c_{t+1}^{TC} denote the amounts the agent is planning to consume in period $t+1$, or, the amounts a time-consistent agent would consume in that period.

5.3 Sophisticated agents

Now we move on to the more difficult problem, the problem for sophisticated agents.

Lemma 3 *Suppose $u(c_t)$ and $v(a_t, S_t)$ are differentiable and that a Markov-perfect subgame-perfect equilibrium with differentiable strategy profiles exists. Then, for each $t \in \{0, \dots, T-1\}$ we have*

$$\begin{aligned} & v_a(a_t, S_t) - p_t u'(c_t) = \\ & = (1 - d)\delta \left[\left(1 + (1 - \beta) \frac{\partial a_{t+1}}{\partial S_{t+1}} \right) (v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1})) - \beta v_s(a_{t+1}, S_{t+1}) \right]. \quad (8) \end{aligned}$$

The proof is in appendix A.

It is worth investigating the sophisticates' first-order condition a little bit further. Rewrite it in the following way:

$$\begin{aligned}
v_a(a_t, S_t) - p_t u'(c_t) &= (1-d)\beta\delta[v_a(a_{t+1}^{TC}, S_{t+1}) - p_{t+1}u'(c_{t+1}^{TC}) - v_s(a_{t+1}^{TC}, S_{t+1})] \\
+ \underbrace{(1-d)\beta\delta[(v_s(a_{t+1}^{TC}, S_{t+1}) - (v_s(a_{t+1}, S_{t+1})))]}_{\text{pessimism effect}} \\
+ \underbrace{(1-d)\beta\delta[(v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1})) - ((v_a(a_{t+1}^{TC}, S_{t+1}) - p_{t+1}u'(c_{t+1}^{TC})))]}_{\text{pessimism effect}} \\
+ \underbrace{(1-d)\delta\left[(1-\beta)\frac{\partial a_{t+1}}{\partial S_{t+1}}(v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1}))\right]}_{\text{incentive effect}} \\
+ \underbrace{(1-d)\delta[(1-\beta)(v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1}))]}_{\text{damage control effect}} \tag{9}
\end{aligned}$$

This formulation allows us to characterize the effects that determine whether sophisticates consume more or less than naifs. The first term on the right-hand side of equation 9 is just the right-hand side of equation 7, so the sign of the rest of the expression determines the relative consumption of sophisticates and naifs.

The first two terms are named the 'pessimism effect'. Assuming that $a_{t+1} > a_{t+1}^{TC}$, that is, that actual consumption of the addictive good is higher than the desired consumption, this term is clearly negative, tending to raise period t consumption relative to that of naifs. This occurs due to the complementarity of consumptions in periods t and $t+1$. In short, the realistic prognosis of sophisticated agents leads them to be pessimistic about future selves' consumption, and since a_t and a_{t+1} are complements, this tends to raise their consumption.

The other two effects operate through self t 's knowledge that she can influence later selves by leaving different amounts of stock. One way this can help is by influencing later selves directly; this is the 'incentive effect,' the third term in the above equation. From self t 's point of view, self $t+1$ consumes too much, not maximizing self t 's discounted utility. Thus a reduction in a_{t+1} leads to a first-order increase in self t 's utility from period $t+1$ on. But self t can give an incentive for self $t+1$ to consume less through leaving a lower S_{t+1} ¹². This tends to lower the consumption of

¹² Assuming, of course, that $\frac{\partial a_{t+1}}{\partial S_{t+1}} > 0$, which should be the case for the good to be appropriately called addictive. We will see below exactly when this condition holds. Also, we are interested in harmful addictions, which implies $v_a(a_{t+1}, S_{t+1}) - p_{t+1}u'(c_{t+1}) > 0$ in all periods other than the last one. See section 6.

sophisticates relative to naifs—naifs think it unnecessary to provide incentives for future selves, as they believe those selves will do the right thing anyway.

We label the last effect, which also tends to lower sophisticates' consumption relative to naifs', 'damage control.' It is a somewhat subtle effect, and the easiest way to understand it is to assume $\frac{\partial a_{t+1}}{\partial S_{t+1}} = 0$, so that there is no incentive effect. However, how badly future selves conduct themselves in the eyes of self t depends not only on what they do, but also on what they should do. That is, if future selves' self-control problem becomes smaller, they make better decisions in the eyes of self t . And this is exactly what is going on here: as we lower S_{t+1} , future self-control problems decrease, lowering the damage later selves do to self t . Thus the term 'damage control'.

These effects might be confusing to distinguish only because they operate on the same space of actions—consumption of a_t . However, that they are conceptually different becomes evident when we have more instruments. Imagine that you have to decide tomorrow whether to play basketball or write your paper. From today's point of view, it is optimal to write the paper, but you know that tomorrow's self won't want to do it. You can do a number of things. First, you can try to set up incentives for tomorrow to write the paper; for example, by making sure you'll have no one to play with. This is the incentive effect. Second, if you think that might not work, you might want to practice your jump shot to at least enjoy the game that you'll inevitably play. This is the equivalent of the pessimism effect. And finally, if you know the commitment won't work, you could assure the outcome won't be so bad, for example by starting your paper today—the damage control effect. A number of hyperbolic discounting models where there is a state variable involved can be understood in these terms (Diamond and Kőszegi [13], Laibson [19]).

The pessimism and incentive effects have been identified by O'Donoghue and Rabin [24] in a related formulation of the consumption decision for addictive goods. In contrast to ours, their setup allows for two consumption choices, hit or not hit; the level of addiction, in turn, can also take on two values: hooked or not hooked, and the agent is hooked if she hit last period. This discreteness assumption generates important differences from our model. As we note below, if the discrete model is extended to include prices, its implications for certain price changes are different from the continuous model we are adapting (and from the empirical evidence we obtain). This feature also makes it a difficult model to use for the optimal tax analysis that we undertake in the

final section of the paper ¹³.

6 Solving the models

Following Becker and Murphy [3], we take $v(a_t, S_t)$ and $u(c_t)$ to be quadratic:

$$v(a_t, S_t) = \alpha_a a_t + \alpha_s S_t + \frac{\alpha_{aa}}{2} a_t^2 + \alpha_{as} a_t S_t + \frac{\alpha_{ss}}{2} S_t^2 \quad (10)$$

$$u(c_t) = \alpha_c c_t + \frac{\alpha_{cc}}{2} c_t^2, \quad (11)$$

where α_a, α_{as} , and α_c are positive and $\alpha_s, \alpha_{aa}, \alpha_{ss}$, and α_{cc} are negative. The key parameter is α_{as} , which measures the effect of past consumption on the marginal utility of current consumption. $\alpha_{as} > 0$ means if you had done more drugs in the past, you will crave them more in the present. This is what can give rise to addictive behavior. The physiological evidence that α_{as} is positive for many goods is overwhelming. For most of this paper, we will take $U(a_t, c_t, S_t)$ to be strictly concave, that is, we suppose its Hessian is negative definite.

In this case, for all three models it is very easy to prove by backward induction that a_t is linear in S_t , that is, $a_t = \lambda_t S_t + \mu_t$, where λ_t and μ_t are constants. The following theorem helps establish that for a general class of parameter values, marginal propensities to addiction are stationary for all three types far from the end of the horizon.

Theorem 1 *Suppose $\beta \geq \frac{1}{2}$ and that $U(a_t, c_t, S_t)$ is strictly concave. Then, $\lim_{j \rightarrow \infty} \lambda_{T-j} = \lambda^{*s}$, where λ^{*s} is given as the unique solution on the interval $(-1, \frac{\alpha_{as}}{-\alpha_{aa} - p^2 \alpha_{cc}})$ of*

$$\lambda^{*s} = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2 \alpha_{cc}}{-\alpha_{aa} - p^2 \alpha_{cc} + \delta(1-d)^2[(1 + (1-\beta)\lambda^{*s})(\alpha_{aa}\lambda^{*s} + \alpha_{as} + p^2\alpha_{cc}\lambda^{*s}) - \beta\alpha_{as}\lambda^{*s} - \beta\alpha_{ss}]}. \quad (12)$$

¹³ Our model also has parallels with another recent paper (Orphanides and Zervos [30]) which examines the role of information in rational addiction. If the addictiveness of a good for the individual is unknown, it might be optimal (even in a rational time-consistent setting) to risk getting addicted. If it turns out the good is addictive, the agent's ex-ante decision looks suboptimal when viewed through the lens of ex-post information, leading to ex-post regret. In this model, addiction-prone individuals who think they will not get addicted are most likely to do so.

Naive hyperbolic discounters also have an informational problem that influences their consumption decisions: they are overoptimistic about their future self-control problem. Descriptively, this sounds like the Orphanides and Zervos [30] effect, but there are two major differences. First, naifs' misperception problem is not about the addictiveness of the good—they are perfectly clear on how consumption in one period affects utility later. Rather, they don't understand how this intertemporal complementarity will play out in their future behavior. This is a much more general informational problem that carries across contexts. Second, it's not merely overoptimism that leads naifs to consume too much—we've seen that sophisticates, who are realistic about their self-control problems, also consume too much, maybe even more than naifs.

Furthermore, $\lambda^{*s} > 0$ (that is, there is adjacent complementarity) if and only if

$$\alpha_{as} > \frac{\beta\delta(1-d)^2}{1-\delta(1-d)^2}(-\alpha_{ss}). \quad (13)$$

The proof is technical and left to the appendix.

Notice that one of the assumptions of the theorem is $\beta \geq \frac{1}{2}$. If this is not the case, it seems possible (though we conjecture it won't usually be the case) that the agent exhibits wild cyclical behavior characterized by periodic binges and brutal cuts. However, most of the psychological literature points to a β above one-half, at least for the time period we consider most relevant for time inconsistency in smoking decisions, a few weeks or few months. For small rewards, a weekly discount rate of 10 to 30 percent seems reasonable (Kirby and Herrnstein [18]); this implies that β is at least 0.7 (and even higher if δ is less than one over this period as well). Thaler [34] finds monthly discount rates on the order of 20 to 30 percent, and three-monthly discount rates of up to 50 percent. The evidence reviewed by Ainslie [1] indicates that yearly discount rates are about 40 percent, and $\beta = 0.6$ is the estimate used by Laibson [19].

Setting $\beta = 1$ in the above expression gives the implicit expression for the marginal propensity to respond to the stock for time-consistent agents, λ^{*TC} . And the naifs' first-order condition 7 can then be used to show that

$$\lambda^{*n} = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2\alpha_{cc}}{-\alpha_{aa} - p^2\alpha_{cc} + \beta\delta(1-d)^2[\alpha_{aa}\lambda^{*TC} + \alpha_{as} + p^2\alpha_{cc}\lambda^{*TC} - \alpha_{as}\lambda^{*TC} - \alpha_{ss}]}. \quad (14)$$

Theorem 2 $\lambda^{*n} \geq \lambda^{*s} \geq \lambda^{*TC}$.

Once again, the proof is left to the appendix.

This theorem states that a given increase in the stock of past consumption increases today's consumption most for naifs, least for time-consistent agents, and somewhere in-between for sophisticates.

It is not surprising that $\lambda^{*n}, \lambda^{*s} \geq \lambda^{*TC}$: for large S it is bad to leave a higher stock, and neither naifs nor sophisticates take this sufficiently into account, consuming more than time-consistent agents with the same long-run preferences.

However, the fact that $\lambda^{*n} \geq \lambda^{*s}$ has potentially interesting and broad implications. Since S is a 'bad' (it measures the level of addiction), in a very loose sense this result is saying that the

worse their situation gets, the worse naifs do relative to sophisticates. Naifs, in a sense, believe even when they are very hooked that things are not so bad, since their future self will just do things right. Thus, they are enslaved to their current craving, with no controls. On the other hand, sophisticates' knowledge that things *are* bad put a restraint on their craving.

7 Implications of price changes

For this section, we will assume that $\alpha_{cc} = 0$, thereby eliminating income effects, which are probably very small for small price changes in many addictive goods. We also assume constant income, $I_t = I$. We will assume for much of what follows that all three models exhibit adjacent complementarity. We want the problem to be well-behaved, that is, for μ_{T-j} to converge for all three types as $j \rightarrow \infty$, so that consumption rules are approximately stationary far from the end of the horizon. For naifs and time-consistent agents, this is simple, and a precise proof is contained in appendix B.1.

For sophisticates, the incentive and damage control effects complicate the analysis. For small β , the model can exhibit some 'violent' characteristics with respect to price changes. For example, when the price increases permanently, the current self knows that this will act as a deterrent for the future self, decreasing her need for incentives and damage control. This could lead her to increase consumption drastically. Since a drastic increase in consumption in response to a current price increase sounds implausible, we make sufficient assumptions in appendix B.1 to rule out this possibility ¹⁴

Under these conditions, it is easy to derive responses to different price changes for the three types. We do so in appendix B.1 and summarize some responses to permanent price increases in table 7. From this table, it is clear that as long as the good is sufficiently addictive, all three types respond to a future price increase by decreasing consumption. In particular, for all three types, the knowledge (or expectation) that future selves will decrease their consumption decreases the marginal utility of consumption today due to the complementarity of intertemporal consumption levels (the 'make quitting easier' effect) ¹⁵. Therefore, Becker, Grossman, and Murphy's [4] test

¹⁴ After doing this, there are still phenomena operating through the incentive effect that look more reasonable. For example, the incentive effect might explain some yuppie binges: that before a big project or a new job, many normally restrained people get wasted on alcohol or high on drugs, only because they know that their job is important enough for them not to keep up with the habit permanently.

¹⁵ An effect going the other way is the usual substitution effect: one wants to shift consumption toward times when

cannot distinguish the rational addiction model from alternatives such as ours.

In principle, the price responses to changes at different points in the future can be used to back out the parameters β and δ using the formulas in table 7, thereby allowing us to assess the degree of time inconsistency in smoking decisions. Moreover, a less ambitious test of time inconsistency per se is also feasible. Operating under the null hypothesis of time consistency, we can back out d and λ^{*TC} as described in appendix C.1. Then, the ratio of the response to current price changes and to those occurring in one period allows us to estimate δ . If in reality the agent is a hyperbolic discounter, this method underestimates her long-term discount factor. The underestimation, in turn, is revealed by comparing the responses to price changes two periods versus one period ahead, since the former depends more on the long-term discount factor. This approach depends neither on whether the agent is naive or sophisticated, nor really on the exact form of discounting she uses (as long as it is hyperbolic). Section C of the appendix describes the above test precisely and proves formally that the methodology actually works for both kinds of agents.

In practice, however, both of these tests have proved difficult to implement. The basic problem is apparent from our discussion of differenced models: even grouping together all of the months between enactment and effective dates, we obtain an estimate of the price response which is only 2.3 times its standard error. The result is that it is impossible to break down this period into the smaller windows required to carry out these tests; the estimates for windows of different lengths (e.g. price change in one vs. two periods ahead) are simply too imprecise to permit comparison to each other. Future work with more precise data can perhaps implement these suggested tests to assess the shape of discounting.

It is also worth noting that if the discrete model of O'Donoghue and Rabin [24] is extended to include prices, it turns out that none of the types will ever quit in response to a future price change. This is because quitting is a one-time decision (not a smooth decline in consumption as in the continuous model), so that one might as well wait until the price change to quit—quitting earlier wouldn't be any easier ¹⁶.

it's cheaper. For time-consistent agents, the former effect dominates iff the good is addictive, $\lambda^{*TC} > 0$. For both types of quasi-hyperbolic discounters, one needs stronger addictivity for the former to dominate.

¹⁶ This complete lack of response to future price changes in the O'Donoghue and Rabin model is due to their assumption of discreteness *and* full depreciation, $d = 1$. If $d = 1$ is relaxed while retaining the discreteness of consumption choices, the argument that it would be just as hard to quit today as tomorrow is invalidated. However,

8 Optimal government policy

A key implication of the rational addiction framework for modeling addiction is that government regulatory policy towards addictive goods should depend only on their interpersonal externalities. Just as the government has no cause, absent market failures, for interfering with revealed preference in the realm of non-addictive goods, there is no reason to take addictiveness per se as a call to government action, if individuals are pursuing these activities "rationally". It is this framework that underlies the well known efforts of Manning et al. [21] and others to formulate optimal taxation of cigarettes and alcohol as a function of the size of their external costs. These estimates, which are frequently cited and influential in debates over excise taxation, suggest that the optimal tax rate for cigarettes in particular is fairly low, since the net external costs of smoking are small.

But models with time-inconsistent agents extend the role of government policy by breaking down revealed preference concepts of consumer choice. The argument that people act in their best interests, so – barring well-known qualifications – the government should leave them alone, is immediately invalidated in our setting. Therefore, though our models are explicitly of the no-externality type¹⁷, a benevolent social planner would want to intervene in this economy.

Of course, the question arises why we consider only government interventions to combat self-control problems. If a sophisticated agent had access to an effective private self-control device, she would take advantage of it, reducing the value of a government intervention. However, we find it unlikely that fully effective self-control devices can be found in this context. Market-provided self-control mechanisms are probably undercut by the market mechanism itself: although firms have a financial incentive to provide self-control to agents, other firms have a financial incentive to break it down. For example, if a firm developed a self-control shot that causes pain when the consumer smokes, another firm has an incentive to develop a drug that relieves these effects for agents who temporarily want to get rid of their commitment. Other problems arise in contracting setups. If there are ex post gains to be made, the future self might want to renegotiate today's contract. But

it is not invalidated completely: old smokers, who have reached their steady-state level of stock, won't increase it any more if they smoke today, so quitting tomorrow will be just as hard. Thus, there is an extra wrinkle relative to the predictions of the continuous model: only *new* smokers (those who have taken up the habit recently) should respond to future price changes.

¹⁷ At least in the inter-personal sense. One might look at the intra-personal conflicts that are generated by a hyperbolic model as intra-personal externalities.

even if there are none ¹⁸, there is an ex post incentive to cheat on the contract: smoking is hard to verify in court. This leaves us with privately provided self-control mechanisms like betting with others or becoming involved in situations where it is very difficult to smoke, but these mechanisms are likely to run into similar enforcement problems to those discussed above.

8.1 Setup

As in any model where different socially relevant actors have different tastes, a discussion of optimal government policy must start with the setup of the social welfare function. In the context of hyperbolic discounting, these actors are not separate individuals, but different intertemporal incarnations of the same individual. The question of social welfare maximization in such a situation has largely been ignored, so we face the difficult problem of specifying the social preferences to be used for our purposes.

We will consider two approaches. In the first, we take the agent's long-run preferences as those relevant for social welfare maximization. Clearly, if the representative agent were to vote in a tax change today that is instituted starting tomorrow, these are the preferences she would use in choosing the new tax rate. Alternatively, we can just maximize utility according to the preferences of the self when the tax change is instituted, thus using a quasi-hyperbolically discounted welfare function. The latter approach might seem odd in the hyperbolic context, since it completely disregards future selves' preferences in the maximization. This 'dictatorship of the present,' however, is equally unwarranted in the standard time-consistent framework (Caplin and Leahy [7]). Even there, there is no compelling reason to think that different selves should have identical preferences over streams of consumption ¹⁹. Since 'dictatorship of the present' is commonly assumed in economics, we will consider it as well. In addition, if we make sufficient assumptions in the exponential model to make 'dictatorship of the present' normatively appealing, the natural extension of the preferences to hyperbolic discounting leads to an exponentially discounted welfare function, giving back our

¹⁸ For example, the agent could post a bond, which she loses if she smokes.

¹⁹ Even if a lifetime of cigarette smoking is optimal from a 20-year-old's point of view, a 40-year-old might dislike the resulting stock of past smoking, which is bad for her health. A social planner should take this into account. The only case when the 20-year-old's preferences can in general be taken to be identical to the social optimum is when the 40-year-old derives utility from past smoking (for example, through memory) and, moreover, discounts *her own* utility just as the 20-year-old does.

first approach ²⁰. Both of these approaches lead to an optimal tax of zero in the limiting case of $\beta = 1$, so our optimal tax is purely a 'self-control tax' (and, in the naive case, a misperception tax), conceptually distinct from an 'intertemporal redistribution tax' that could be warranted by Caplin and Leahy's arguments.

We first consider the case of a representative consumer with a very long life and a social planner restricted to a tax on the addictive good that is invariant over time ²¹. Notice that our assumption of no income effects ($\alpha_{cc} = 0$) guarantees that the social planner has no influence on λ^* , but she can essentially choose μ^* by setting the appropriate tax. Since in general λ^{*s} and λ^{*n} are greater than λ^{*TC} , the government is restricted to a second-best policy. For each $i = n, s$, the social planner solves

$$\begin{aligned} \max_{\mu} \sum_{t=0}^{\infty} \delta^t \left[v(\lambda^{*i} S_t + \mu, S_t) + \alpha_c (I_t - p(\lambda^{*i} S_t + \mu)) \right] \\ \text{s.t. } S_0, S_{t+1} = (1-d)(S_t + \lambda^{*i} S_t + \mu) \end{aligned} \quad (17)$$

Since the price p in the objective function is the pre-tax price, this formulation implicitly assumes that the tax receipts are lump-sum redistributed in each period.

The above is a quadratic in μ with a negative prime coefficient ²², so the first-order condition gives the optimal tax. One can think of the derivative of 17 as a sum of the derivatives with respect to each period's consumption. Then the first-order condition for the optimal choice of μ is

$$0 = \sum_{t=0}^{\infty} \delta^t \left(v_a(a_t, S_t) - p\alpha_c + \sum_{k=1}^{\infty} \delta^k (1-d)^k (1 + \lambda^{*i})^{k-1} [(v_a(a_{t+k}, S_{t+k}) - p\alpha_c) \lambda^{*i} + v_s(a_{t+k}, S_{t+k})] \right). \quad (18)$$

²⁰ Dictatorship of the present is normatively appealing only when all selves have the same preferences over streams of consumption:

$$\sum_{i=0}^T \delta^i U(a_i, c_i, S_i) \quad (15)$$

Then, the most natural way to introduce hyperbolic discounting is to assume that self t puts a weight $\frac{1}{\beta}$ on consumption in period t :

$$\sum_{i \neq t} \delta^i U(a_i, c_i, S_i) + \frac{1}{\beta} \delta^t U(a_t, c_t, S_t) \quad (16)$$

Summing this social welfare function across all selves, we get back the previous utility function. That is, an additive social welfare function is equivalent to one that takes the agent's long-run preferences as those relevant for welfare analysis.

²¹ In the language used to describe the policies, we will assume throughout that the addiction in question is harmful in the no-tax setting—that is, consuming more today reduces future discounted utility. Because of the quadraticity of the utility functions, this is not a necessary consequence of our model: for any a_t , there is a region where utility is increasing in S_t . The results for these beneficial addictions should be symmetric to those below.

²² Otherwise the maximum utility would be infinite, a non-starter.

8.2 Sophisticates

First consider sophisticates. Sophisticates solve

$$0 = v_a(a_t, S_t) - (p + \tau)\alpha_c \quad (19)$$

$$+ \beta \sum_{k=1}^{\infty} \delta^k (1-d)^k \prod_{j=1}^{k-1} \left(1 + \frac{\partial a_{t+j}}{\partial S_{t+j}} \right) \left[(v_a(a_{t+k}, S_{t+k}) - (p + \tau)\alpha_c) \frac{\partial a_{t+k}}{\partial S_{t+k}} + v_s(a_{t+k}, S_{t+k}) \right],$$

where τ is the unit tax on the consumption of the addictive good. Combining this with 18 and rearranging gives

$$(1 - \beta) \sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - (p + \tau)\alpha_c) = \beta \frac{1}{1 - \delta} \frac{1 - \delta(1 - d)}{1 - \delta(1 - d)(1 + \lambda^{*s})} \tau \alpha_c. \quad (20)$$

It is easy to show that the optimal tax is positive: the derivative of 17 at $\mu = \mu^{*s}$ can be written in the form

$$\sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - p\alpha_c + \delta(1 - d)V_S^s(S_{t+1})), \quad (21)$$

where $V^s(S_t)$ stands for the exponentially discounted utility from leaving stock S_t and consuming according to the sophisticated consumption function from then on. A hyperbolic discounting agent solves $v_a(a_t, S_t) - p\alpha_c + \beta\delta(1 - d)V_S^s(S_{t+1}) = 0$, and since by assumption $V_S^s(S_{t+1})$ is negative, the above derivative is negative for $\beta < 1$. Therefore the optimal μ is lower than μ^{*s} , and consequently the optimal tax is greater than zero.

There are several important additional implications for this case of sophisticates. First, since the optimal tax is positive for $\beta < 1$, the left-hand side of the equation 20 is positive. But $v_a(a_t, S_t) - (p + \tau)\alpha_c > 0$ means that the addiction is harmful—higher consumption lowers utility from future periods. Therefore, at least in an average sense, the optimal tax is not so large so as to make the addiction harmless on the margin. The reason is that the tax is there to correct a marginal self-control problem. If there was no self-control problem (on average), there would be nothing to correct—the agent's different intertemporal selves wouldn't disagree, so the losses to consuming more would be second-order. But then, the selves would be consuming too little, since their private costs are higher than the social costs due to the tax.

Second, we can also determine the dependence of the optimal tax on S_0 . For any t , the total derivative of $v_a(a_t, S_t)$ with respect to S_t is $\alpha_{aa}\lambda^{*s} + \alpha_{as}$. This is greater than zero because $\lambda^{*s} < \lambda_T$.

Therefore the derivative of the left-hand side of equation 20 with respect to S_0 is positive, and so the optimal tax is increasing in the level of initial addiction.

The reason for an optimal tax that is increasing in S_0 follows from the nature of the tax. It is solely a 'self-control tax,' or, in other words, a tax that is intended to aid in overcoming the agent's self-control problem. As such, it has to increase as the self-control problem becomes more serious. And in our model the marginal harm done by smoking more, and so the self-control problem, increases with S_0 .

Third, as a variation on this theme, one can ask what the government would do if it had access to a time-variant tax, for example one that can be set differently in period 0 and the rest of the periods. If the agent has already reached her steady-state consumption level when the first tax is instituted, then the period-0 tax will be higher than the long-term tax. This occurs exactly because the tax breaks the agent's habit, so $S_1 < S_0$. For distributional reasons, the traditional prescription is exactly the opposite. According to this view, a high up-front tax for cigarettes is undesirable because it hurts addicted consumers too much, and these consumers tend to be disproportionately poor. One should instead wait for these people to quit and then raise taxes high to make sure they don't get readdicted. But this distributional concern should be dealt with separately from taxation that is used as a self-control device; high up-front taxes perhaps combined with redistribution of the resources to low income smokers would be a preferred strategy to phasing in the tax.

More generally, if taxes vary period to period, the first-order condition for the optimal choice of τ_t , as a variant of condition 20, is

$$\begin{aligned} \tau_t \alpha_c &= -(1 - \beta) \left(\sum_{k=1}^{\infty} \delta^k (1 - d)^k (1 + \lambda^{*s})^{k-1} [(v_a(a_{t+k}, S_{t+k}) - p\alpha_c)\lambda^{*s} + v_s(a_{t+k}, S_{t+k})] \right) \\ &\quad - \beta \delta (1 - d) \lambda^{*s} \sum_{i=0}^{\infty} \delta^i (1 - d)^i (1 + \lambda^{*s})^i \tau_{t+1+i} \alpha_c. \end{aligned} \quad (22)$$

From this formulation, it is clear that if the good is sufficiently addictive, taxes in different periods of time are *substitutes*, a consequence of the intertemporal complementarity in consumption levels.

Thus, if we think the taxes are too low in certain periods of life, due to addictivity taxes should be higher than otherwise in earlier periods. Similarly, if the model is written over space, if we can't regulate smoking in the home, we should overregulate it in other settings such as restaurants or bars.

We now return to the form of the social welfare function. It is important to note that the exact form of the social welfare function is not as crucial as one might think, at least not to make the point that the optimal tax should be positive. It is easy to prove that a small positive tax is Pareto-improving—it increases the discounted utility of each intertemporal incarnation of the agent. To see this, note first that a small decrease in a self's consumption causes a second-order loss to her discounted utility, while a decrease in future selves' consumption gives a first-order gain. In addition, future selves gain through the fact that they receive a lower stock of consumption S .

Of course, to get an actual tax, we still have to cardinalize how we weigh off different selves' utilities. We remark on one alternative specification to the above, that of maximizing self 0's discounted utility subject to the constraint that the tax is constant across periods. In this formulation, the optimal tax is smaller than with an exponentially discounted social welfare function. The reason is simply that with exponential discounting the social planner wants to correct every self's self-control problem, including self 0's, whereas with hyperbolic discounting self 0 wants to respect her own preferences and just correct future selves' behavior. In other words, self 0 wants the government to leave her alone, and then tax the future selves so as to maximize the exponentially discounted sum of utilities. Since the tax is restricted to be the same across periods, self 0 chooses something between the two solutions.

8.3 Naifs

Let us move on to naifs. Rewrite equation 18 as

$$0 = \sum_{t=0}^{\infty} \delta^t \left(v_a(a_t, S_t) - (p + \tau)\alpha_c + \delta(1 - d)V_S^n(S_{t+1}) + \left(1 + \sum_{k=1}^{\infty} \delta^k (1 - d)^k (1 + \lambda^{*n})^{k-1} \lambda^{*n} \right) \tau \alpha_c \right), \quad (23)$$

where $V^n(S_t)$ is the exponentially discounted value function that results from naif consumption at prices $p + \tau$. Further rewrite this as

$$\begin{aligned} 0 &= \sum_{t=0}^{\infty} \delta^t \left(v_a(a_t, S_t) - (p + \tau)\alpha_c + \delta(1 - d)V_S^{TC}(S_{t+1}) + \delta(1 - d)(V_S^n(S_{t+1}) - V_S^{TC}(S_{t+1})) \right) + \\ &+ \sum_{t=0}^{\infty} \delta^t \left(\left(1 + \sum_{k=1}^{\infty} \delta^k (1 - d)^k (1 + \lambda^{*n})^{k-1} \lambda^{*n} \right) \tau \alpha_c \right). \end{aligned} \quad (24)$$

Knowing that naifs solve $v_a(a_t, S_t) - (p + \tau)\alpha_c + \beta\delta(1 - d)V_S^{TC}(S_{t+1}) = 0$, we can finally put the first-order condition for naifs in the form

$$(1 - \beta) \sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - (p + \tau)\alpha_c) = \beta \frac{1}{1 - \delta} \frac{1 - \delta(1 - d)}{1 - \delta(1 - d)(1 + \lambda^{*n})} \tau\alpha_c + \beta\delta(1 - d) \sum_{t=0}^{\infty} \delta^t (V_S^n(S_{t+1}) - V_S^{TC}(S_{t+1})). \quad (25)$$

The naifs' first-order condition looks similar to that of sophisticates, equation 20, with the last term on the right-hand side being the major difference. This term creeps in because naifs perceive the wrong value function—they think a time-consistent agent's value function applies to them, whereas the naif's does. For example, if the marginal utility of leaving more S is smaller in reality than naifs think, that tends to increase the optimal tax.

The extra term generates important differences from the sophisticated case. First, it is not in general true that if the addiction is harmful, then the optimal tax is greater than zero: it could be the case that even though the addiction is harmful, naifs think it is even more harmful than it is ($0 > V_S^n(S_{t+1}) > V_S^{TC}(S_{t+1})$), so they consume too little. However, the possibility of a subsidy on a harmful good is more of a theoretical curiosity than a real recipe. Since $V^n(S)$ is dominated by $V^{TC}(S)$, for large S we have $V_S^n(S) < V_S^{TC}(S)$, so the optimal tax is greater than zero²³.

Second, contrary to the optimal tax for sophisticates, it seems possible that $\sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - (p + \tau)\alpha_c) < 0$, i.e. that the good appears beneficial on the average. The reason is that here the tax not only corrects a self-control problem, but also a misperception problem—the agent is wrong in predicting her future behavior. This is a very important qualitative difference in terms of optimal taxation. Whereas in the sophisticated case taxation that eliminates all harmful consumption can never be justified, even if the good is very addictive and people have severe self-control problems (low β 's), it might be the best policy for the naive case²⁴. To put it in more plain terms, a sort of 'cautious' paternalism is recommended for parts of the population that realize they have a self-control problem, while a more 'short-leashed' policy should apply to those who don't.

Third, once again, we can easily show that the optimal tax is increasing in S_0 . For this, simply note that the left-hand side of equation 25 is increasing, while the right-hand side is decreasing in

²³ Formally, this can be proven in a similar way to the sophisticated case.

²⁴ To be more precise, naifs will *think* that consuming the good is beneficial, whereas in reality it isn't. Thus, naifs might say, 'There is no harm in smoking this one cigarette, so why don't you let me?', and they would be right—if they really consumed according to their plans. Ultimately, they are not right because they don't think they'll get addicted, and they will.

S_0 (we know $V_{SS}^n(S_t) < V_{SS}^{TC}(S_t)$ because $V^{TC}(S_t)$ dominates $V^n(S_t)$ for any t .) For naifs, both their self-control problem and their overoptimism (their self-deceptive view of the value function) get more problematic as S_0 increases, so the optimal tax has to increase. Also, the optimal time profile of the taxes would include a decreasing burden.

These results have particularly interesting implications for the age profile of government intervention. Our finding that the optimal tax rises with the stock suggests taxes that rise with age; in our model, older smokers are doing more harm to themselves than are their younger counterparts, and therefore have a greater value for the self-control imparted by taxation. Over time, this age profile could be flattened as individuals are broken of their addiction²⁵. On the other hand, if youths are naive and adults are sophisticated (or time-consistent), banning smoking for children and not banning but taxing it for adults might be more appropriate, even if children are aware of the negative health consequences of smoking²⁶. Finally, if there are political or other constraints on addressing smoking among adults, then the substitutability of taxes over time provides another rationale for higher taxes or more regulation of smoking for youths.

8.4 A calibration exercise

In this section, we attempt to calibrate our model and calculate an actual optimal tax for sophisticated agents. To do so, we will assume that the disutility associated with smoking is linear. Let h_S denote the money equivalent of the per-period future marginal utility of an extra cigarette (so it should be negative.) Theoretically, there are two possible interpretations of this externality. First, h_S could be the pure utility effect of stock, $v_s(a_{t+k}, S_{t+k})$. Alternatively, since a higher stock leads the future self to reoptimize behavior, we can think of the externality taking into account this change. In that case, $h_S = (v_a(a_{t+k}, S_{t+k}) - p\alpha_c)\lambda^{*s} + v_s(a_{t+k}, S_{t+k})$. We will calculate the optimal tax according to the first interpretation, since we believe that the health data we will take advantage of in this section corresponds more closely to that one. It is also the interpretation that leads to a lower optimal tax.

²⁵ Moreover, since the harm from a higher stock fades as individuals near death, optimal taxes may fall towards the end of the life-cycle.

²⁶ Indeed, children appear not only to know but to actually overestimate the dangers of smoking (Viscusi [35]). Interestingly, they also grossly underestimate their own future propensity to smoke. This pair of beliefs is consistent with what a naive hyperbolic discounter might think.

Starting from equation 20 for the optimal tax, substituting on the left-hand side using the sophisticates' first-order condition expressed in terms of stock, and then setting $h_s = v_s(a_t, S_t)$ for each t gives

$$(1 - \beta) \frac{\delta(1 - d)}{1 - \delta(1 - d)(1 + (1 - \beta)\lambda^{*s})} (-h_s) = \frac{1 - \delta(1 - d)}{1 - \delta(1 - d)(1 + \lambda^{*s})} \tau, \quad (26)$$

which reduces to

$$\tau = \frac{1 - \delta(1 - d)(1 + \lambda^{*s})}{1 - \delta(1 - d)(1 + (1 - \beta)\lambda^{*s})} (1 - \beta)(-H_S). \quad (27)$$

Notice that—contrary to a Pigouvian intuition—the optimal tax is smaller than $1 - \beta$ times the marginal externality of the stock. The reason is the incentive effect. Even under optimal taxes, a sophisticated agent feels a need to exert control on the future selves by consuming less. This effect helps the government, and therefore it is not necessary to tax the full marginal externality.

One difficulty with estimating the optimal tax is parameterizing h_s (and H_S). Clearly, there is a lot of disutility associated with smoking that is hard to quantify. Unpleasant consequences in this category include coughing, increased vulnerability to various diseases, and the inability to enjoy sports and food to the greatest possible degree. We will ignore all these, and assume that the only disutility from smoking is in the increased chance of early death. In particular, we assume that a cigarette costs a person seven minutes of life [21], and that the valuation of one year of life is \$100,000 [12]. At these figures, the cost in terms of life years lost per pack of cigarettes is \$27.

Another serious difficulty lies in choosing the right period length for our purposes. The quasi-hyperbolic discounting model is only a theoretical device meant to capture the essence of hyperbolic discounting, and is not designed for actual policy simulations. In particular, it implicitly limits the scope of intrapersonal conflict to exactly one period—there is no self-control problem *within* a period or regarding future periods. The appropriate period length should be chosen taking these limitations into account.

Since our empirical analysis was done in terms of a monthly time period, we will continue to work with this time frame, and assume that $\beta = 0.9$ and $\delta = 1$ over this time period. Our choice of β is intended to parameterize a modest time inconsistency problem—most of the psychological evidence indicates that monthly discount rates are substantially higher than 10 percent. Physiological and empirical evidence suggests that λ^{*s} is fairly high for smoking. Evidence is less clear on the

depreciation rate.

Since the tax is quite sensitive to d and λ^{*s} , so table 6 shows the optimal tax for a few values. The implied tax rates are still very high except for a combination of low d and high λ^{*s} . We haven't been able to pin down d and λ^{*s} empirically, so we can't identify the relevant cell in table 6. However, there is information in our data that can be exploited to rule out at least some combinations of d and λ^{*s} : namely, the speed of convergence to a new steady state after an enacted and immediately effective price change. A combination of low d and high λ^{*s} would imply that this convergence is very slow—even long after the price change, agents should be reducing their consumption from period to period. In particular, the change in smoking from month $N + 1$ to month $2N + 1$ over the change in smoking from month 1 to month $N + 1$ will equal $(1 - d)^N(1 + \lambda^{*s})^N$.

In our sample the convergence to a new steady state seems relatively fast, ruling out a combination of high λ^{*s} and low d . For example, the consumption change from months 7 to 13 is only about 0.08 of the consumption change from months 1 to 7. With $d = 0.5$ and $\lambda^{*s} = 0.9$, this figure should be 0.74; even with $d = 0.5$ and $\lambda^{*s} = 0.7$, the figure should be 0.38. This drop is compatible with $d = 0.6$ and $\lambda^{*s} = 0.7$. All exercises that we tried in this vein ruled out $d = 0.5$ and $\lambda^{*s} = 0.9$, while most also rule out $d = 0.5$ and $\lambda^{*s} = 0.7$. We can therefore safely put a lower bound on the optimal tax at over a dollar per pack, and a reasonable number is perhaps closer to \$1.50.

Moreover, this estimate is likely to be a lower bound; if we took into account the full internality, not only the fatal health consequences of smoking, the tax would likely be much higher. And, as noted earlier, we have chosen a degree of time inconsistency which is considerably lower than that used in the previous literature.

On the other hand, if we use the second social welfare function, which aggregates the preferences of today's hyperbolic agents, the optimal tax would be somewhat lower. But, for high δ , it will be quite close, and for $\delta = 1$ it will be the same. Our assumption of $\delta = 1$ amounts to assuming an infinite horizon, and in this case the optimal tax is not different for the two social welfare functions. For $\delta < 1$, the difference is large for small β 's—for example, $\beta = 0$ gives a tax of zero for hyperbolic preferences, which is not the case with exponential ones. The reason is that for a small β self 0 doesn't care about correcting other selves' self-control problem, but she does want the government to leave her alone. With the time period we have chosen, δ should be very close to 1, and β is not

Table 6: Optimal taxes for various values of d and λ^{*s}

	$d = 0.5$	$d = 0.6$	$d = 0.7$	$d = 0.8$	$d = 0.9$
$\lambda^{*s} = 0.9$	\$0.29	\$1.13	\$1.70	\$2.11	\$2.42
$\lambda^{*s} = 0.7$	\$0.86	\$1.49	\$1.92	\$2.24	\$2.47
$\lambda^{*s} = 0.5$	\$1.40	\$1.83	\$2.14	\$2.36	\$2.53

too small, so this social welfare function should lead to essentially the same optimal tax.

We will not attempt to calculate the optimal tax for naifs, because that would involve making assumptions about the degree of underestimation of the self-control problem. However, if there is such underestimation, the tax for naifs is going to be higher, and much higher if the underestimation is serious.

9 Conclusions

The theoretical and empirical insights of Becker and Murphy and of Becker, Grossman, and Murphy, as well as subsequent work in the vein of their pioneering efforts, have greatly advanced the modeling of addictive processes by economists. These are important and timely advances, as policy makers are becoming increasingly interested in regulating addictive behaviors. In the case of cigarettes, recent years have seen increased state taxation, regulation of smoking in public places, and a spate of court cases brought by the states and now the Justice Department against the industry.

We have attempted in this paper to make two contributions to the literature on addiction. First, we have suggested a more convincing framework for testing the central hypothesis of the rational addiction model, that individuals are forward looking with respect to their decisions to consume addictive goods. We find that announced but not yet effective tax increases lead to both increased sales and decreased consumption of cigarettes, which is very consistent with forward looking behavior by consumers. The use of announced price changes, and the robustness of our finding to specification checks, provides the strongest evidence to date for adjacent complementarity.

Second, we have noted that the rational addiction framework embeds another important as-

sumption besides forward-looking consumption behavior: time consistency. This assumption is at odds with virtually all laboratory experiments and with a variety of casual real-world evidence on smoking decisions. When we change the Becker and Murphy model to incorporate time inconsistent preferences, we obtain predictions for price changes which are very similar to what are delivered by their model. But we obtain radically different implications for policy. Instead of the standard result that the optimal tax on cigarettes depends only on their associated externalities, we find that there are substantial "internalities" as well which justify government intervention. For very modest parameterization of these internalities, and ignoring any costs other than those associated with the excess mortality of smoking, we find that there are sizable additional taxes suggested by these internalities, on the order of \$1 per pack or more. We also find that differences in the stock of past smoking and the nature of time inconsistency across populations may call for very different types of regulatory interventions, for example by age of smoker. In particular, if younger smokers are either more "naive" than older smokers, or if there are political pressures that make it impossible to tax appropriately smoking at older ages, it would justify tougher regulatory interventions for youth smoking.

This result should not be surprising. The key feature of smoking, particularly in contrast to other "addictive bads" such as drinking, is that its internal effects dwarf its external costs: the vast majority of harm done by a smoker is to him or herself. At standard values of the value of a life/year, we estimate above that a pack of cigarettes costs \$27 in terms of lost life expectancy. If even a small share of these internal costs are to be considered by government policy makers, the resulting justification for intervention easily outweighs any externalities associated with smoking.

Of course, we have not proven time inconsistency in smoking decisions; we were unable to design and implement a test that could effectively distinguish quasi-hyperbolic and exponential preferences in this context. At the same time, the fact that there is no empirical support or even laboratory support for exponential discounting in this or related contexts suggests that alternative models of the type that we have derived be taken seriously. The important general point for thinking about government policy in this context is that, when standard public finance analyses suggest that the tax on addictive bads is simply equal to their external costs, those analyses are implicitly embracing a rational addiction model. Given the enormous magnitude of the internal costs to smoking, however,

alternative models such as ours must be considered seriously in designing regulatory policy towards addictive goods.

A Proofs of theorems

We start with a proof of lemma 3.

Lemma 3 *Suppose $u(c_t)$ and $v(a_t, S_t)$ are differentiable and that a Markov-perfect subgame-perfect equilibrium with differentiable strategy profiles exists. Then, for each $t \in \{0, \dots, T-1\}$ we have*

$$\begin{aligned} & v_a(a_t, S_t) - p_t u'(c_t) = \\ & = (1-d)\delta \left[\left(1 + (1-\beta) \frac{\partial a_{t+1}}{\partial S_{t+1}} \right) (v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1})) - \beta v_s(a_{t+1}, S_{t+1}) \right]. \end{aligned} \quad (28)$$

Proof: Since instantaneous utilities and future selves' strategies are differentiable, self t 's discounted utility is differentiable in a_t ²⁷. Furthermore, since an equilibrium exists, self t 's maximization problem must have a solution. Then, as self t 's consumption of the addictive product is unrestricted, the derivative of her discounted utility at a_t is zero. Therefore,

$$\begin{aligned} 0 & = v_a(a_t, S_t) - p_t u'(c_t) \\ & + \beta \delta \left[(v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1})) (1-d) \frac{\partial a_{t+1}}{\partial S_{t+1}} + v_s(a_{t+1}, S_{t+1}) (1-d) \right] \\ & + \beta \delta^2 (1-d)^2 \left(1 + \frac{\partial a_{t+1}}{\partial S_{t+1}} \right) \left[(v_a(a_{t+2}, S_{t+2}) - p_{t+2} u'(c_{t+2})) \frac{\partial a_{t+2}}{\partial S_{t+2}} + v_s(a_{t+2}, S_{t+2}) \right] \\ & + \dots \\ & = v_a(a_t, S_t) - p_t u'(c_t) \\ & + \beta \sum_{i=1}^{T-t} \delta^i (1-d)^i \prod_{j=1}^{i-1} \left(1 + \frac{\partial a_{t+j}}{\partial S_{t+j}} \right) \left[(v_a(a_{t+i}, S_{t+i}) - p_{t+i} u'(c_{t+i})) \frac{\partial a_{t+i}}{\partial S_{t+i}} + v_s(a_{t+i}, S_{t+i}) \right]. \end{aligned} \quad (29)$$

The complicated second term on the right-hand side comes from the following consideration. It is trivial to prove by induction that the derivative of S_{t+i} with respect to a_t is $(1-d)^i \prod_{j=1}^{i-1} \left(1 + \frac{\partial a_{t+j}}{\partial S_{t+j}} \right)$. Now this has two effects on future instantaneous utility. First, it affects utility directly—the stock of past consumption is assumed to affect current utility. That's the $v_s(a_{t+i}, S_{t+i})$ term. Second, it affects utility through changing the consumption of self $t+i$. That's the $(v_a(a_{t+i}, S_{t+i}) - p_{t+i} u'(c_{t+i})) \frac{\partial a_{t+i}}{\partial S_{t+i}}$ term.

We can write the same optimality condition for self $t+1$:

$$\begin{aligned} 0 & = v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1}) \\ & + \beta \sum_{i=1}^{T-t-1} \delta^i (1-d)^i \prod_{j=1}^{i-1} \left(1 + \frac{\partial a_{t+1+j}}{\partial S_{t+1+j}} \right) \left[(v_a(a_{t+1+i}, S_{t+1+i}) - p_{t+1+i} u'(c_{t+1+i})) \frac{\partial a_{t+1+i}}{\partial S_{t+1+i}} \right] \\ & + \beta \sum_{i=1}^{T-t-1} \delta^i (1-d)^i \prod_{j=1}^{i-1} \left(1 + \frac{\partial a_{t+1+j}}{\partial S_{t+1+j}} \right) v_s(a_{t+1+i}, S_{t+1+i}). \end{aligned} \quad (30)$$

Multiplying equation 30 by $\delta(1-d) \left(1 + \frac{\partial a_{t+1}}{\partial S_{t+1}} \right)$ and subtracting it from equation 30 we get

$$\begin{aligned} 0 & = v_a(a_t, S_t) - p_t u'(c_t) \\ & + \beta \delta \left[(v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1})) (1-d) \frac{\partial a_{t+1}}{\partial S_{t+1}} + v_s(a_{t+1}, S_{t+1}) (1-d) \right] \\ & - \delta(1-d) \left(1 + \frac{\partial a_{t+1}}{\partial S_{t+1}} \right) (v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1})). \end{aligned} \quad (31)$$

Rearranging this gives the desired Euler equation. \square

²⁷Notice that self t really only has one choice variable, because a choice of a_t ties down c_t due to the no-saving assumption.

We pick up the discussion from the observation that for each t and each type of agent, $a_t = \lambda_t S_t + \mu_t$ for some constants λ_t and μ_t ²⁸. Then

$$S_{t+1} = (1-d)(S_t + a_t) = (1-d)(S_t + \lambda_t S_t + \mu_t) \quad (32)$$

$$a_{t+1} = \lambda_{t+1} S_{t+1} + \mu_{t+1} = \lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1} \quad (33)$$

Plugging this into the sophisticates' first-order condition, equation 28, and assuming $p_t = p$ in each period:

$$\begin{aligned} & \alpha_a + \alpha_{aa}(\lambda_t S_t + \mu_t) + \alpha_{as} S_t - p[\alpha_c + \alpha_{cc}(I_t - p(\lambda_t S_t + \mu_t))] = \\ & = (1-d)\delta[(1 + (1-\beta)\lambda_{t+1})[\alpha_a + \alpha_{aa}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) \\ & + \alpha_{as}(1-d)(S_t + \lambda_t S_t + \mu_t) - p(\alpha_c + \alpha_{cc}(I_{t+1} - p(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1})))] \\ & - \beta(\alpha_s + \alpha_{as}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) + \alpha_{ss}(1-d)(S_t + \lambda_t S_t + \mu_t))] \end{aligned} \quad (34)$$

The above has to be true for all S_t , so the coefficient of S_t in the expression has to be zero. After 'some' manipulation, this implies

$$\lambda_t = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2 \alpha_{cc}}{-\alpha_{aa} - p^2 \alpha_{cc} + \delta(1-d)^2[(1 + (1-\beta)\lambda_{t+1})(\alpha_{aa}\lambda_{t+1} + \alpha_{as} + p^2 \alpha_{cc}\lambda_{t+1}) - \beta\alpha_{as}\lambda_{t+1} - \beta\alpha_{ss}]} \quad (35)$$

This is a backward recursion for the λ 's in the different periods. It looks quite scary, but can be understood with some effort. That is what theorem 1 does.

Theorem 1 Suppose $\beta \geq \frac{1}{2}$ and that $U(a_t, c_t, S_t)$ is strictly concave. Then the backward recursion 35 converges, that is, $\lim_{j \rightarrow \infty} \lambda_{T-j} = \lambda^*$, where λ^* is given as the unique solution on the interval $(-1, \frac{\alpha_{as}}{-\alpha_{aa} - p^2 \alpha_{cc}})$ of

$$\lambda^* = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2 \alpha_{cc}}{-\alpha_{aa} - p^2 \alpha_{cc} + \delta(1-d)^2[(1 + (1-\beta)\lambda^*)(\alpha_{aa}\lambda^* + \alpha_{as} + p^2 \alpha_{cc}\lambda^*) - \beta\alpha_{as}\lambda^* - \beta\alpha_{ss}]} \quad (36)$$

Furthermore, $\lambda^* > 0$ if and only if

$$\alpha_{as} > \frac{\beta\delta(1-d)^2}{1 - \delta(1-d)^2}(-\alpha_{ss}). \quad (37)$$

Proof. Define the function $f_s(\lambda)$ according to equation 35. We will prove that $f_s(\lambda_T) < \lambda_T = \frac{\alpha_{as}}{-\alpha_{aa} - p^2 \alpha_{cc}}$, $f_s(-1) > -1$, and that f_s is continuous and increasing on $(-1, \lambda_T)$. This is sufficient to establish that λ_{T-i} converges. Then clearly $\lambda^* > 0$ iff $f_s(0) > 0$, which is equivalent to 13.

First, notice that the second term of $f_s(\lambda)$ is the reciprocal of a quadratic with a negative coefficient on λ^2 . Then if this term is positive for two points on the real line, it is also positive in-between these two points. Moreover, it is easy to show that on the interval where this term is positive, f_s is strictly convex²⁹. Therefore, it is sufficient to show that $f_s(-1) > -1$, $\lambda_T > f_s(\lambda_T) > -1$, and that $f'_s(-1) \geq 0$. The first two ensure that we are on the continuous and strictly convex section of f_s , and the last one (together with convexity) ensures that f_s is increasing on $(-1, \lambda_T)$.

The rest is just carrying out the above. We have

$$f_s(-1) = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2 \alpha_{cc}}{-\alpha_{aa} - p^2 \alpha_{cc} + \beta\delta(1-d)^2[-\alpha_{aa} + 2\alpha_{as} - p^2 \alpha_{cc} - \alpha_{ss}]} > -1 \quad (38)$$

as both the numerator and the denominator are positive in the second term. Proceeding,

$$\begin{aligned} f_s(\lambda_T) & = f_s\left(\frac{\alpha_{as}}{-\alpha_{aa} - p^2 \alpha_{cc}}\right) = -1 + \frac{\alpha_{as} - \alpha_{aa} - p^2 \alpha_{cc}}{-\alpha_{aa} - p^2 \alpha_{cc} + \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]} \\ & = \frac{\alpha_{as} - \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]}{-\alpha_{aa} - p^2 \alpha_{cc} + \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]} \end{aligned} \quad (39)$$

²⁸To be more precise, for the solution to the first-order condition to give a maximum, we need strict concavity at each stage. But we know that if a function $C(a_t, S_t)$ is strictly concave and continuously differentiable, then $C_a(a_t, S_t) = 0$ gives the global maximum for a fixed S_t , and this maximum is strictly concave in S_t . This consideration implies for time-consistent agents that the previous period's problem is also strictly concave. Then for the sophisticated problem to be strictly concave, notice that her value function starting in the next period is quadratic, and dominated by the time-consistent agent's quadratic value function.

²⁹The second derivative of the reciprocal of a quadratic q is $-\frac{q^2 q'' - q(q')^2}{q^4}$, which is positive as long as q is positive and concave.

This being $< \lambda_T = \frac{\alpha_{as}}{-\alpha_{aa} - p^2\alpha_{cc}}$ is equivalent to $-\alpha_{as}\lambda_T - \alpha_{ss} > 0$. But the latter can be rewritten as $\alpha_{as}^2 < \alpha_{ss}(\alpha_{aa} + p^2\alpha_{cc})$, and since owing to the concavity of $U(a_t, c_t, S_t)$ we have $\alpha_{as}^2 < \alpha_{ss}\alpha_{aa}$, this inequality holds. $-\alpha_{as}\lambda_T - \alpha_{ss} > 0$ also implies that the second term is positive, so that $f(\lambda_T) > -1$.

Finally,

$$f'_s(\lambda) = -\frac{(\alpha_{as} - \alpha_{aa} - p^2\alpha_{cc})(2(1-\beta)\lambda(\alpha_{aa} + p^2\alpha_{cc}) + \alpha_{aa} + p^2\alpha_{cc} + (1-2\beta)\alpha_{as})}{[-\alpha_{aa} - p^2\alpha_{cc} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as} + p^2\alpha_{cc}\lambda) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]]^2}, \quad (40)$$

which gives

$$f'_s(-1) = -(1-2\beta) \left[\frac{\alpha_{as} - \alpha_{aa} - p^2\alpha_{cc}}{-\alpha_{aa} - p^2\alpha_{cc} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as} + p^2\alpha_{cc}\lambda) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]} \right]^2 \geq 0. \quad (41)$$

This completes the proof. \square

Note that the above proof doesn't work if $U(a_t, c_t, S_t)$ is not concave, possibly leading to 'wild' behavior on the part of the agent. In particular, in that case one can't prove, and it's not in general true, that $f_s(\lambda_T) < \lambda_T$. Also, it is not the case that the backward program that just looks at first-order conditions at each stage finds a maximum for every period. But in as much as it does, we can say the following. Carefully looking at the graph of $f_s(\lambda)$, it seems possible that λ_{T-} first increases, then jumps to around -1 , then starts increasing again, restarting the cycle. Behaviorally, this means that the agent goes through periods of addiction followed by brutal 'cold turkey' types of quits, a phenomenon described by Becker and Murphy [3].

Theorem 2 $\lambda^{*n} \geq \lambda^{*s} \geq \lambda^{*TC}$.

Proof. To prove $\lambda^{*s} \geq \lambda^{*TC}$, define f_{TC} as f_s except with $\beta = 1$. The difference between the denominators of the second terms of f_s and f_{TC} , ignoring the positive multiplicative constant $\delta(1-d)^2$, is

$$(1-\beta)[\lambda^2(\alpha_{aa} + p^2\alpha_{cc}) + 2\alpha_{as}\lambda + \alpha_{ss}] = (1-\beta)(\lambda \ 1 \ p\lambda) \begin{pmatrix} \alpha_{aa} & \alpha_{as} & 0 \\ \alpha_{as} & \alpha_{ss} & 0 \\ 0 & 0 & \alpha_{cc} \end{pmatrix} \begin{pmatrix} \lambda \\ 1 \\ p\lambda \end{pmatrix} \leq 0 \quad (42)$$

by the concavity of $U(a_t, c_t, S_t)$. This implies that for any $\lambda \in (-1, \lambda_T)$, $f_s(\lambda) > f_{TC}(\lambda)$, and thus $\lambda^{*s} \geq \lambda^{*TC}$.

To prove the inequality $\lambda^{*n} \geq \lambda^{*s}$, let us introduce relevant value functions for sophisticates and time-consistent agents:

$$V^s(S_{t+1}) = \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^s, c_{t+i}^s, S_{t+i}^s) \\ V^{TC}(S_{t+1}) = \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^{TC}, c_{t+i}^{TC}, S_{t+i}^{TC}), \quad (43)$$

where the superscripts refer to the two different agents. Since the utility function is quadratic and strategies are linear, both of these are quadratic in S_{t+1} . Furthermore, it is clear that $V^{TC}(S_{t+1}) \geq V^s(S_{t+1})$ for all S_{t+1} , so $V^{TC}(S_{t+1})$ has at least as large a prime coefficient as $V^s(S_{t+1})$. In addition, these coefficients are negative³⁰.

Agent i ($i = n, s$) solves

$$\max_{a_t} v(a_t, S_t) + u(I_t - pa_t) + \beta V^i((1-d)(S_t + a_t)), \quad (44)$$

leading to the first-order condition

$$v_a(a_t, S_t) - pu'(c_t) + \beta(1-d)V_S^i(S_{t+1}) = 0. \quad (45)$$

Differentiating this totally with respect to S_t gives

$$\frac{\partial a_t}{\partial S_t} = \frac{v_{as}(a_t, S_t) + \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})}{-v_{aa}(a_t, S_t) - p^2 u''(c_t) - \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})} \\ = -1 + \frac{v_{as}(a_t, S_t) - v_{aa}(a_t, S_t) - p^2 u''(c_t)}{-v_{aa}(a_t, S_t) - p^2 u''(c_t) - \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})}. \quad (46)$$

Since v_{as} , $-u''$, $-v_{aa}$, and $-V_{SS}^i$ are positive constants, and $-V_{SS}^s \geq -V_{SS}^{TC}$, the above implies $\lambda^{*n} \geq \lambda^{*s}$. \square

³⁰The easiest way to see this is to check that $V_S^{TC}(S_{t+1}) = \delta v_s(a_{t+1}^{TC}, S_{t+1}^{TC})$ is negative for a sufficiently high S_{t+1} . Alternatively, refer to an earlier footnote on the strict concavity of value functions.

B The responsiveness to price and its implications

B.1 Deriving price responsiveness

Condition 34 implies for sophisticates

$$\begin{aligned} \mu_t^s &= \text{constant} - \frac{p_t \alpha_c - p_{t+1} \alpha_c (1-d) \delta (1 + (1-\beta) \lambda_{t+1})}{-\alpha_{aa} + \delta (1-d)^2 [(1 + (1-\beta) \lambda_{t+1}) [\alpha_{aa} \lambda_{t+1} + \alpha_{as}] - \beta \alpha_{as} \lambda_{t+1} - \beta \alpha_{ss}]} \\ &+ \frac{(1-d) \delta [\beta \alpha_{as} - (1 + (1-\beta) \lambda_{t+1}) \alpha_{aa}]}{-\alpha_{aa} + \delta (1-d)^2 [(1 + (1-\beta) \lambda_{t+1}) [\alpha_{aa} \lambda_{t+1} + \alpha_{as}] - \beta \alpha_{as} \lambda_{t+1} - \beta \alpha_{ss}]} \mu_{t+1}^s \end{aligned} \quad (47)$$

The limit of the coefficient of μ_{t+1}^s in the expression exists and is equal to

$$\frac{(1-d) \delta [\beta \alpha_{as} - (1 + (1-\beta) \lambda^{*s}) \alpha_{aa}]}{-\alpha_{aa} + \delta (1-d)^2 [(1 + (1-\beta) \lambda^{*s}) [\alpha_{aa} \lambda^{*s} + \alpha_{as}] - \beta \alpha_{as} \lambda^{*s} - \beta \alpha_{ss}]}, \quad (48)$$

which is clearly positive. It is easy to see that in general μ_{T-j}^s converges if and only the above limit is less than 1.

Lemma 4 μ_{T-j}^{TC} converges.

Proof. Suppose not. Then one can easily choose parameters so that μ_{T-j}^{TC} diverges, i.e. $|\mu_{T-j}^{TC}| \rightarrow \infty$. Now consider any S . There are real numbers M and N such that for a small enough t , $M < V^{t, TC}(S) < N$. The lower bound comes from the consideration that one can just consume $\frac{d}{1-d} S$ in each period (the steady-state consumption corresponding to S), giving a discounted utility $\frac{1}{1-d} v(\frac{d}{1-d} S, S) + \alpha_c (I - p \frac{d}{1-d} S)$ in the limit. The upper bound comes from the fact that $U(a_t, c_t, S_t)$ is strictly concave quadratic, and thus has a global maximum. (Note that this also implies that N can be chosen independently of S .)

Now since $|\mu_{T-j}^{TC}| \rightarrow \infty$, also $|V_S^{T-j, TC}(S)| = |v_S(a_{T-j}, S)| \rightarrow \infty$. But since $\lim_{j \rightarrow \infty} V_{SS}^{T-j, TC}(S) = \alpha_{ss} + \alpha_{as} \lambda^{*TC}$, this contradicts that the value function is bounded from above. \square

From the naifs' first-order condition

$$\begin{aligned} \mu_t^n &= \text{constant} - \frac{p_t \alpha_c - p_{t+1} \alpha_c (1-d) \beta \delta}{-\alpha_{aa} + \beta \delta (1-d)^2 [\alpha_{aa} \lambda_{t+1}^{TC} + \alpha_{as} - \alpha_{as} \lambda_{t+1}^{TC} - \alpha_{ss}]} \\ &+ \frac{(1-d) \beta \delta (\alpha_{as} - \alpha_{aa})}{-\alpha_{aa} + \beta \delta (1-d)^2 [\alpha_{aa} \lambda_{t+1}^{TC} + \alpha_{as} - \alpha_{as} \lambda_{t+1}^{TC} - \alpha_{ss}]} \mu_{t+1}^{TC}, \end{aligned} \quad (49)$$

and since μ_{T-j}^{TC} and λ_{T-j}^{TC} converge, so does μ_{T-i}^n . Thus we have established the following theorem:

Theorem 1 *If prices are constant over time, μ_{T-j}^{TC} and μ_{T-j}^n converge.*

Call the limits μ^{*TC} and μ^{*n} . Unfortunately, for sophisticates we don't necessarily get convergence. To see this, rewrite 48 as

$$\begin{aligned} &\frac{(1-d) \delta [\alpha_{as} - \alpha_{aa}]}{-\alpha_{aa} + \delta (1-d)^2 [(1 + (1-\beta) \lambda^{*s}) [\alpha_{aa} \lambda^{*s} + \alpha_{as}] - \beta \alpha_{as} \lambda^{*s} - \beta \alpha_{ss}]} \\ &- \frac{(1-d) \delta (1-\beta) (\alpha_{aa} \lambda^{*s} + \alpha_{as})}{-\alpha_{aa} + \delta (1-d)^2 [(1 + (1-\beta) \lambda^{*s}) [\alpha_{aa} \lambda^{*s} + \alpha_{as}] - \beta \alpha_{as} \lambda^{*s} - \beta \alpha_{ss}]} \end{aligned} \quad (50)$$

As $\beta \rightarrow 0$, this approaches $(1-d) \delta (1 + \lambda_T)$, which can easily be greater than 1³¹. If μ_{T-j}^s diverges, the model exhibits certain 'violent' characteristics (for example, there would be huge consumption reactions to even minimal price changes, possibly even in the 'wrong' way), which we don't want to deal with. Therefore we make a sufficient assumption for 48 to be less than 1. We assume that $(1-d) \delta (1 + \lambda^{*s}) < 1$. By the continuity of λ^{*s} with respect to β , this is true if β is sufficiently close to 1.

We will use the following lemma extensively to study consumption responses to price changes.

Lemma 5 *For a recursion of the form $x_{j+1} = l + kx_j$ with $0 < k < 1$ we have*

$$\frac{d}{dl} \lim_{j \rightarrow \infty} x_j = \frac{1}{1-k}. \quad (51)$$

³¹Even though we have not proved that λ_{T-j}^s converges for small β 's, it probably does for most parameter values. Even if it doesn't, it is clear that agents who ignore the future should have a $\lambda_t = \lambda_T$.

Proof. An easy geometric argument. \square

We will consider permanent price decreases of Δp that start either last period, this period, next period, or two periods into the future. A direct application of lemma 5 yields a consumption response to a present permanent decrease in price for sophisticates of

$$\frac{\Delta p \alpha_c}{m^s} \frac{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}{1-k}, \quad (52)$$

where

$$m_s = -\alpha_{aa} + \delta(1-d)^2[(1 + (1-\beta)\lambda_{t+1})[\alpha_{aa}\lambda_{t+1} + \alpha_{as}] - \beta\alpha_{as}\lambda_{t+1} - \beta\alpha_{ss}],$$

$$k = \frac{(1-d)\delta[\beta\alpha_{as} - (1 + (1-\beta)\lambda_{t+1})\alpha_{aa}]}{-\alpha_{aa} + \delta(1-d)^2[(1 + (1-\beta)\lambda_{t+1})[\alpha_{aa}\lambda_{t+1} + \alpha_{as}] - \beta\alpha_{as}\lambda_{t+1} - \beta\alpha_{ss}]}.$$
 (53)

If the price change is next period, μ_{t+1} changes by the same amount as if it was this period, so the change in μ_t is also the same except for a term $\frac{\Delta p \alpha_c}{m^s}$:

$$\frac{\Delta p \alpha_c}{m^s} \frac{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}{1-k} - \frac{\Delta p \alpha_c}{m^s}. \quad (54)$$

For time-consistent agents, the expressions are similar, but they reduce nicely because in that case $k = (1-d)\delta(1 + \lambda^{*TC})$. The most convenient forms are given in table 7.

The problem is somewhat different for naifs. From equation 49

$$\mu^{*n} = \text{constant} - \frac{p\alpha_c(1 - (1-d)\beta\delta)}{m^n} + (1-d)\beta\delta(1 + \lambda^{*n})\mu^{*TC}, \quad (55)$$

where

$$m^n = -\alpha_{aa} + \beta\delta(1-d)^2[\alpha_{aa}\lambda^{*TC} + \alpha_{as} - \alpha_{as}\lambda^{*TC} - \alpha_{ss}]. \quad (56)$$

With a price change, the last term changes due to a change in μ^{*TC} , and the second-to-last directly due to the price change. Taking the change in μ^{*TC} from table 7, the response to a present permanent decrease in price is

$$(1-d)\beta\delta(1 + \lambda^{*n}) \frac{(1 - (1-d)\delta)}{1 - (1-d)\delta(1 + \lambda^{*TC})} \frac{\Delta p \alpha_c}{m^{TC}} + (1 - \beta\delta(1-d)) \frac{\Delta p \alpha_c}{m^n}. \quad (57)$$

Noting that $\frac{1}{m^{TC}} = \frac{1 + \lambda^{*TC}}{\alpha_{as} - \alpha_{aa}}$ and $\frac{1}{m^n} = \frac{1 + \lambda^{*n}}{\alpha_{as} - \alpha_{aa}}$, this becomes

$$\begin{aligned} & \frac{\Delta p \alpha_c(1 + \lambda^{*n})}{\alpha_{as} - \alpha_{aa}} \left[\frac{(1-d)\beta\delta(1 - (1-d)\delta)(1 + \lambda^{*TC})}{1 - (1-d)\delta(1 + \lambda^{*TC})} + 1 - \beta\delta(1-d) \right] \\ &= \frac{\Delta p \alpha_c(1 + \lambda^{*n})}{\alpha_{as} - \alpha_{aa}} \left[\frac{(1-d)\beta\delta\lambda^{*TC} + 1 - (1-d)\delta(1 + \lambda^{*TC})}{1 - (1-d)\delta(1 + \lambda^{*TC})} \right]. \end{aligned} \quad (58)$$

If the price change is next period, the consumption response is the same except for the $\frac{\Delta p \alpha_c}{m^n}$ term. The expression then simplifies to

$$\frac{\Delta p \alpha_c(1 + \lambda^{*n})}{\alpha_{as} - \alpha_{aa}} \frac{(1-d)\beta\delta\lambda^{*TC}}{1 - (1-d)\delta(1 + \lambda^{*TC})}. \quad (59)$$

B.2 Estimating δ from price responses

As we have mentioned, estimating δ from the information available to us is not as easy as dividing the responses to future and past price changes, as Becker, Grossman, and Murphy [4] claim. In our setup, that ratio is easily derived from table 7:

$$\frac{\delta\lambda^{*TC}}{(1 + \lambda^{*TC})(1 - (1-d)\delta)}, \quad (60)$$

For this to be equal to δ , one would need $1 - (1-d)\delta(1 + \lambda^{*TC}) = 0$, which is not possible by lemma 4. In general, to recover δ , one needs to estimate d and λ^{*TC} , as we do in appendix C.

The BGM [4] method of calculating δ only works in general if the price changes in question are *temporary*³². If they are permanent, it works if $d = 1$, since this makes prices further away irrelevant for today's smoking decision—in two periods, the effects of today's consumption will have disappeared.

³²From expression 47, the response to a future temporary price increase is $(1-d)\delta\lambda^{*TC}$ times the response to an equivalent current temporary price increase, while the response to a past increase is $(1-d)\lambda^{*TC}$ times as much.

Table 7: Summary of price responses

	this period	next period
time-consistent	$-\frac{(1-d)\delta}{1-(1-d)\delta(1+\lambda^{*TC})} \frac{\Delta p\alpha_c}{m^{TC}}$	$-\frac{\Delta p\alpha_c}{m^{TC}} \frac{(1-d)\delta\lambda^{*TC}}{1-(1-d)\delta(1+\lambda^{*TC})}$
naive	$-\frac{\Delta p\alpha_c(1+\lambda^{*n})}{\alpha_{as}-\alpha_{aa}} \left[\frac{(1-d)\beta\delta\lambda^{*TC} + 1 - (1-d)\delta(1+\lambda^{*TC})}{1-(1-d)\delta(1+\lambda^{*TC})} \right]$	$-\frac{\Delta p\alpha_c(1+\lambda^{*n})}{\alpha_{as}-\alpha_{aa}} \frac{(1-d)\beta\delta\lambda^{*TC}}{1-(1-d)\delta(1+\lambda^{*TC})}$
sophisticated	$-\frac{\Delta p\alpha_c}{m^s} \frac{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-k}$	$-\frac{\Delta p\alpha_c}{m^s} \frac{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-k} + \frac{\Delta p\alpha_c}{m^s}$

	two periods ahead
naive	$-\frac{\Delta p\alpha_c(1+\lambda^{*n})}{\alpha_{as}-\alpha_{aa}} \left[\frac{\beta\delta^2(1-d)^2\lambda^{*TC}(1+\lambda^{*TC})}{1-(1-d)\delta(1+\lambda^{*TC})} \right]$
sophisticated	$-k \left[\frac{\Delta p\alpha_c}{m^s} \frac{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-k} - \frac{\Delta p\alpha_c}{m^s} \right]$

C A test of time consistency

C.1 Recovering d and λ^{*i}

For all three types, the response to a past price change is $1 + (1-d)\lambda^{*i}$ times the response to a current price change, since the constant μ^{*i} changes by the same amount but there is a change in S due to the change in last period's consumption. Therefore, looking at consumption responses to present vs. past price changes, one can recover $(1-d)\lambda^{*i}$.

In addition, it is easy to show that steady-state consumption for type i is

$$a^i = \frac{\frac{d}{1-d}}{\frac{d}{1-d} - \lambda^{*i}} \mu^{*i}. \quad (61)$$

Thus comparing the steady-state change in consumption with the instantaneous one identifies $\frac{\frac{d}{1-d}}{\frac{d}{1-d} - \lambda^{*i}}$. This, along with the earlier estimate of $(1-d)\lambda^{*i}$, let us get at d and λ^{*i} .

C.2 The test

We operate under the null hypothesis of time consistency. Using the above estimates for d and $lstartc$, we can estimate δ from the ratio of responses to price changes today versus one period ahead. Let the estimates be denoted by hats. Now, from table 7, the null hypothesis of time consistency implies that the ratio of responsiveness to price changes two periods versus one period ahead is equal to $(1-d)\delta(1+\lambda^{*TC})$. So the null hypothesis would be that this ratio is equal to $(1-\hat{d})\hat{\delta}(1+\lambda^{*TC})$. We prove that the alternative hypothesis of quasi-hyperbolic discounting implies that the ratio is greater than $(1-\hat{d})\hat{\delta}(1+\lambda^{*TC})$.

Naifs: Clearly, the above method yields the correct d and λ^{*n} . From table 7 we get that when $\hat{\delta}$ is calculated assuming time consistency, we have

$$\frac{(1-d)\hat{\delta}\lambda^{*n}}{1-(1-d)\hat{\delta}} = \frac{(1-d)\beta\delta\lambda^{*TC}}{1-(1-d)\delta(1+(1-\beta)\lambda^{*TC})}. \quad (62)$$

We first prove that $\hat{\delta} < \delta$. Since $\lambda^{*n} > \lambda^{*TC}$ and the left-hand side of equation 62 is increasing in δ , for this it suffices to show that

$$\frac{\beta}{1-(1-d)\delta(1+(1-\beta)\lambda^{*TC})} < \frac{1}{1-(1-d)\delta}. \quad (63)$$

Using that $(1-d)\delta(1+\lambda^{*TC}) < 1$ (which follows from lemma 4), one can expand this inequality, and after brief manipulation it in turn reduces to $(1-d)\delta(1+\lambda^{*TC}) < 1$. Therefore $\hat{\delta} < \delta$. The following lemma provides the result on which the proposed empirical test relies.

Lemma 6 $(1-d)\hat{\delta}(1+\lambda^{*n}) < (1-d)\delta(1+\lambda^{*TC})$.

Proof. Start from equation 62. Cross-multiplying and rearranging gives

$$(1-d)\hat{\delta}(1+\lambda^{*n}) = 1 - \frac{(1-(1-d)\hat{\delta})(1-(1-d)\delta(1+\lambda^{*TC}))}{1-(1-d)\delta(1+(1-\beta)\lambda^{*TC})}. \quad (64)$$

Since $\hat{\delta} < \delta$ and $\lambda^{*TC} > 0$, the above is less than

$$1 - (1 - (1 - d)\delta(1 + \lambda^{*TC})) = (1 - d)\delta(1 + \lambda^{*TC}). \quad (65)$$

□

Sophisticates: We can move straight to the lemma that establishes the validity of the proposed empirical test. The lemma in this case requires an extra assumption—the same assumption that we needed to ensure that the sophisticated model is not too ‘wild.’

Lemma 7 *If $(1-d)\delta(1+\lambda^{*s}) < 1$, then $k > (1-d)\hat{\delta}(1+\lambda^{*s})$.*

Proof. Start from the equivalent expression to the naive case, pulled from table 7:

$$\frac{(1-d)\hat{\delta}\lambda^{*s}}{1-(1-d)\hat{\delta}} = \frac{k - (1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})} \quad (66)$$

Multiplying out and rearranging gives

$$(1-d)\hat{\delta}(1+\lambda^{*s}) = \frac{k - (1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})} - (1-d)\hat{\delta} \frac{k-1}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})}. \quad (67)$$

This being less than k is equivalent to

$$k - (1-d)\delta(1+(1-\beta)\lambda^{*s}) - (1-d)\hat{\delta}k + (1-d)\hat{\delta} < k(1 - (1-d)\delta(1+(1-\beta)\lambda^{*s})), \quad (68)$$

which reduces to

$$\delta(1+(1-\beta)\lambda^{*s}) > \hat{\delta} \quad (69)$$

Using equation 66 and noting that its left-hand side is increasing in $\hat{\delta}$, this is equivalent to

$$\frac{(1-d)\delta(1+(1-\beta)\lambda^{*s})\lambda^{*s}}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})} > \frac{k - (1-d)\delta(1+(1-\beta)\lambda^{*s})}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})} \quad (70)$$

The above easily becomes

$$(1-d)\delta(1+(1-\beta)\lambda^{*s})(1+\lambda^{*s}) > k, \quad (71)$$

But easily $k = (1-d)\delta(1+\lambda^{*s}) - (1-d)\delta(1-\beta)\frac{\alpha_{aa}\lambda^{*s} + \alpha_{aa}}{\alpha_{aa} - \alpha_{aa}} < (1-d)\delta(1+\lambda^{*s})$, completing the proof. □

References

- [1] George Ainslie. *Picoeconomics*. Cambridge University Press, 1992.
- [2] Robert Barro. Ramsey meets laibson in the neoclassical growth model. *Quarterly Journal of Economics*. forthcoming.
- [3] G. S. Becker and K. M. Murphy. A theory of rational addiction. *Journal of Political Economy*, 1988.
- [4] G.S. Becker, M. Grossman, and K. M. Murphy. An empirical analysis of cigarette addiction. *American Economic Review*, 1994.
- [5] Douglas A. Bernstein. The modification of smoking behavior: An evaluative review. In William A. Hunt, editor, *Learning Mechanisms in Smoking*, pages 3–41. Aldine Publishing Company, Chicago, 1970.
- [6] J. Burns. Looking to the future. In J. Burns, editor, *Special Report: Business and Health*, pages 21–22. Medical Economics Publishing, 1992.
- [7] Andrew Caplin and John Leahy. The social discount rate. Mimeo, 1999.
- [8] CDC. You can quit smoking. Tobacco Information and Prevention Source. Webpage: www.cdc.gov/tobacco.
- [9] F. J. Chaloupka. Rational addictive behavior and cigarette smoking. *Journal of Political Economy*, 1991.
- [10] Frank J. Chaloupka and Kenneth E. Warner. The economics of smoking. In Joseph Newhouse and Anthony Culyer, editors, *The Handbook of Health Economics*. 1998. forthcoming.

- [11] Gary Chamberlain. Panel data. In Zvi Griliches and Michael D. Intriligator, editors, *Handbook of Econometrics*, volume II, pages 1248–1318. Elsevier Science, New York, 1984.
- [12] David Cutler and Elizabeth Richardson. Measuring the health of the u.s. population. *Brookings Papers on Economic Activity*, pages 217–282, 1997.
- [13] Peter Diamond and Botond Kőszegi. Hyperbolic discounting and retirement. Mimeo, 1998.
- [14] William Evans, Jeanne Ringel, and Diana Stech. Tobacco taxes and public policy to discourage smoking. In James Poterba, editor, *Tax Policy and the Economy*, pages 1–56. Cambridge, MA: MIT Press, 1999.
- [15] J. Grabowski and S. M. Hall. Tobacco use, treatment strategies, and pharmacological adjuncts: An overview. In J. Grabowski and S. M. Hall, editors, *Pharmacological Adjuncts in Smoking Cessation*, pages 1–13. National Institute on Drug Abuse Monograph 53, 1985.
- [16] Michael Grossman, Frank J. Chaloupka, and Ismail Sirtalan. An empirical analysis of alcohol addiction: Results from monitoring the future panels. *Economic Inquiry*, 1998.
- [17] Jeffrey E. Harris. *Deadly Choices: Coping with Health Risks in Everyday Life*. New York: Basic Books, 1993.
- [18] Kris N. Kirby and R. J. Herrnstein. Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, 6(2):83–89, 1995.
- [19] David Laibson. Hyperbolic discount functions, undersaving, and savings policy. *National Bureau of Economic Research Working Paper*, 1996.
- [20] David I. Laibson, Andrea Repetto, and Jeremy Tobacman. Self-control and saving for retirement. In *Brookings Papers on Economic Activity*, pages 91–196. The Brookings Institution, 1998.
- [21] Willard G. Manning, E. B. Keeler, J. P. Newhouse, E. M. Sloss, and J. Wasserman. *The Costs of Poor Health Habits*. Harvard University Press, 1991.
- [22] J. L. Miller. Self-control in the elimination of cigarette smoking: Case histories using a changing criterion design. Master's thesis, Western Michigan University, 1978.
- [23] R. G. Murray and S. A. Hobbs. Effects of self-reinforcement and self-punishment in smoking reduction: Implications for broad-spectrum behavioral approaches. *Addictive Behaviors*, 6(1):63–67, 1981.
- [24] Ted O'Donoghue and Matthew Rabin. Addiction and self-control. Mimeo, 1997.
- [25] Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999.
- [26] Ted O'Donoghue and Matthew Rabin. Procrastination in preparing for retirement. In Henry J. Aaron, editor, *Behavioral Dimensions of Retirement Economics*, pages 125–156. Brookings Institution Press, 1999.
- [27] American Chamber of Commerce Research Association. Accra cost of living index, various years.
- [28] U.S. Department of Health and Human Services. *Preventing Tobacco Use Among Young People: A Report of the Surgeon General*. National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health, 1994.
- [29] Nilss Olekalns and Peter Bardsley. Rational addiction to caffeine: An analysis of coffee consumption. *Journal of Political Economy*, 1996.
- [30] Athanasios Orphanides and David Zervos. Rational addiction with learning and regret. *Journal of Political Economy*, 1995.
- [31] James O. Prochaska, Pamela Crimi, Duane Lapsanski, Leslie Martel, and Peter Reid. Self-change processes, self-efficacy and self-concept in relapse and maintenance of cessation and smoking. *Psychological Reports*, 51:983–990, 1982.
- [32] Mark H. Showalter. Firm behavior in a market with addiction. forthcoming.
- [33] H-Y Sung, T-W Hu, and T-E Keeler. Cigarette taxation and demand: an empirical model. *Contemporary Economic Policy*, 1994.
- [34] Richard Thaler. Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8:201–207, 1981.
- [35] W. Kip Viscusi. Public perception of smoking risks. *International Conference on the Social Costs of Tobacco*, 1998.
- [36] Teresa M. Waters and Frank A. Sloan. Why do people drink? tests of the rational addiction model. *Applied Economics*, 1995.

Table 1: Issues with Becker-Grossman-Murphy Empirical Approach

	BGM '94	Replication	Red. Form Price	Red. Form Tax	Fixed Trnds Price	Fixed Trnds Tax	Changes Price	Changes Tax
C_{t+1}	0.418 (0.047)	0.502 (0.049)						
C_{t+1}	0.135 (0.055)	0.183 (0.055)						
P_t	-1.388 (0.155) [-0.407]	-1.644 (0.163) [-0.429]						
P_{t-1}			-1.839 (0.349) [-0.582]		-1.413 (0.205) [-0.369]		-0.278 (0.109) [-0.072]	
P_t			-2.231 (0.273) [-0.480]		-1.679 (0.254) [-0.438]		-1.639 (0.107) [-0.427]	
P_{t+1}			-0.965 (0.259) [-0.252]		-0.218 (0.200) [-0.057]		-0.032 (0.107) [-0.008]	
τ_{t-1}				-2.268 (0.387) [-0.514]		-1.497 (0.280) [-0.361]		-0.205 (0.143) [-0.049]
τ_t				-2.956 (0.513) [-0.670]		-2.831 (0.360) [-0.683]		-2.894 (0.141) [-0.685]
τ_{t+1}				-0.630 (0.367) [-0.143]		0.420 (0.267) [0.101]		0.282 (0.139) [0.067]
No. Obs.	1,415	1,976	1,976	1,976	1,976	1,976	1,932	1,932

Notes: Regressions use 1955 to 1997 annual packs/capita, price, and tax information. Standard errors in parentheses; implied elasticities in brackets. In columns (1) and (2), consumption is instrumented by lagged and leaded price. First four columns include fixed effects for state and year; the next two also include state-specific trends; final two columns remove the state fixed effects and estimate a differenced model.

Table 2: Length of Time Between Enactment and Effective Dates of Excise Tax Increases

	1973-1996 (Full dataset)	1982-1996 (Packs/Capita Available)	1989-1996 (Natality Data Available)
Same Month	36	27	14
Consecutive Months	44	38	18
1 Month Between	23	19	11
2 Months Between	27	21	12
3 Months Between	9	6	2
4 Months Between	6	5	2
5 Months Between	2	2	
7 Months Between	1	1	1
Multiple Changes, < 1 year Between	4	4	3
Multiple Changes, > 1 Year Between	8	8	5

Notes: Source is authors' tabulations of data on state excise tax enactment and effective dates. Each row shows number of tax changes with the noted length of time between enactment and effective dates. Last two rows refer to tax events where multiple future changes are enacted at the enactment date.

Table 3: Means of the Analysis Samples

Packs/Capita Sample	
Packs/Capita per Month	9.01 (2.47)
Effective Tax Rate	0.171 (0.082)
Number of Observations	8,885
Natality Sample	
Cigs/Day	1.92 (0.72)
Smoking Rate	0.162 (0.050)
Cigs/Day if Smoke	12.24 (1.19)
Effective Tax Rate	0.187 (0.096)
Number of Births Per State/Month Cell	5319 (4945)
Number of Observations	4,446

Notes: Tabulations of monthly packs/capita and Natality cigarette consumption data. Standard deviations in parentheses.

Table 4: Effect of Tax Announcement on Smoking - Fixed Effects Models

	Aggregate Sales Data		Natality Consumption Data	
Effective Rate	-7.998 (0.306) [-0.803]	-14.13 (1.664) [-1.502]	-0.660 (0.050) [-0.347]	-0.215 (0.226) [-0.113]
Enacted Rate		2.307 (1.066) [0.232]		-0.344 (0.142) [-0.181]
Effective Rate (-2)		4.001 (1.330) [0.402]		-0.118 (0.180) [-0.062]
Number of Obs	8678	8675	4342	4341

Notes: Coefficients from regression of consumption on listed variables, as well as a full set of dummies for state of residence and calendar month of data. Standard errors in parentheses; implied price elasticities in square brackets. Regressions for aggregate packs/capita data in first two columns; regressions for natalty cigarette consumption data in second two columns. All regressions exclude month of enactment of tax increase and month that it is effective.

Table 5: Effect of Tax Announcement on Smoking - Alternative Models

	Aggregate Sales Data		Natality Consumption Data	
	Fixed Trend Model			
Effective Rate	-5.834 (0.400) [-0.624]	-12.86 (1.618) [-1.374]	-0.490 (0.057) [-0.274]	-0.035 (0.177) [-0.020]
Enacted Rate		2.195 (1.061) [0.235]		-0.271 (0.117) [-0.152]
Effective Rate (-2)		5.181 (1.290) [0.554]		-0.238 (0.142) [-0.133]
Number of Obs	8678	8675	4342	4341
	Differences			
Enacted Rate	2.486 (1.608) [0.266]		-0.299 (0.130) [-0.168]	
Number of Obs	8383		4182	

Notes: Coefficients from regression of consumption on listed variables, as well as a full set of dummies for state of residence and calendar month of data. First panel presents fixed trends regressions, which include as well a full set of state-specific time trends; second panel presents differences regressions, which include just month dummies. Standard errors in parentheses; implied price elasticities in square brackets. Regressions for aggregate packs/capita data in first two columns; regressions for natalty cigarette consumption data in second two columns. All regressions exclude month of enactment of tax increase and month that it is effective.