

Optimal Policy Instruments for Externality-Producing Durable Goods under Time Inconsistency

Garth Heutel
Department of Economics
University of North Carolina at Greensboro
PO Box 26165
Greensboro, NC 27402
and NBER
(336) 334 4872 (Phone)
(336) 334 5580 (Fax)
gaheutel@uncg.edu

May 2011

Abstract

When consumers exhibit present bias and are time-inconsistent, the standard solution to market failures caused by externalities—Pigouvian pricing—is suboptimal. I investigate policies aimed at externalities for time-inconsistent consumers. Optimal policy includes an instrument to correct the externality and an instrument to correct the present bias. Either instrument can be an incentive-based policy or a command-and-control policy. Calibrated to the US automobile market, simulation results suggest that the second-best gasoline tax is 18%–30% higher than marginal external damages. The optimal price policy includes a gasoline tax set about equal to marginal external damages and a fuel economy tax that increases the price of an average non-hybrid car by about \$750–\$2200 relative to the price of an average hybrid car.

Keywords: Present Bias; Energy Policy; Gasoline Tax; Quasi-hyperbolic Discounting

JEL Codes: Q58; Q48; D03

I thank Stephen Holland, Chris Ruhm, Jason Shogren, Ken Snowden, Nathan Wozny, and seminar participants for comments, and David Cornejo, Derek Mobley and Alex Smith for valuable research assistance.

A growing body of evidence suggests that consumers regularly and predictably depart from acting in accordance with rational choice theory. In particular, they appear to act with time-inconsistent preferences, or with a present bias: they "underweight" future periods in the present period. This affects decisions over purchases of durable goods with variable operating costs; present bias makes a consumer less likely to spend money upfront to reduce a durable's future operating costs. Many durable goods are energy-intensive and create externalities with consumption, like cars consuming gasoline or appliances consuming electricity. The standard incentive-based solution to the market failure caused by externalities is Pigouvian pricing, but the efficiency of this solution assumes time-consistent preferences.

If consumers are time-inconsistent, does Pigouvian pricing of externalities still lead to a socially optimal outcome? If not, what is the optimal policy? Must it include command-and-control policies instead of or in addition to incentive-based policies? The purpose of this paper is to answer these questions by developing a model of optimal policy for externality-producing durable goods in the presence of time-inconsistent preferences. Surprisingly, very little economic research has yet been undertaken to theoretically derive optimal policy in the presence of behavioral anomalies like time inconsistency, and no paper has answered the questions posed here. Then, I apply the model to the automobile market through simulation and solve for optimal policy.

The question addressed here is policy-relevant for two reasons. First, empirical support for the existence of behavioral anomalies, especially time-inconsistent preferences, is growing.¹ Consumers seem to discount the far future more heavily than the near future, behavior that can be modeled by hyperbolic or quasi-hyperbolic discounting (Laibson 1997). This has been observed in laboratory experiments (Thaler 1981), in individuals' decisions over exercising (Dellavigna and Malmendier 2006) and doing homework assignments (Ariely and Wertenbroch 2002).² It also may be relevant to decisions over energy-efficiency of durable goods. Alcott and Wozny (2010) find that consumers underweight future fuel costs of automobiles at the time of purchase. The well-established "energy paradox" or "energy efficiency gap" finds that

¹ For recent evidence, see Mastrobuoni and Weinberg (2009), Fang and Silverman (2009), Brown et. al. (2009), or Viscusi et. al. (2008).

² The field evidence for this and other types of behavioral anomalies is reviewed in DellaVigna (2009). Andreoni and Sprenger (2010) cite laboratory evidence that fails to find any present bias in preferences, but they suggest that no such bias is expected in laboratory experiments involving money rather than consumption utility. Hastings and Mitchell (2011) combine experimental evidence with data on Chilean households' savings decisions and find that present bias does a better job predicting financial behavior than does financial literacy.

households seem to apply very high discount rates in their decisions over energy-intensive durable goods like air-conditioners (Hausman 1979). Gillingham et. al. (2009) summarize the literature and find implicit discount rates ranging from 25% to 100%.³ This paradox may be explained by present bias.

A second reason the question addressed in this paper is relevant is that environmental and energy policy seems to be moving in a direction towards incentive-based policies, especially tradable permits, and away from command-and-control policies.⁴ This transition has been fueled by arguments from economists that incentive-based policies achieve substantial cost savings compared to command-and-control policies; some empirical evidence has verified this for some policies (Carlson, et al. 2000). If Pigouvian pricing is inefficient under time inconsistency, and if consumers are time-inconsistent, then this push towards these policies may be misguided. More so, if time inconsistency leads to a dominance of some command-and-control policies over Pigouvian pricing, then the push away from command-and-control policies may also be misguided (Shogren and Taylor 2008).

This paper's results are likely to be relevant beyond the domain of environmental policy. Evidence for time-inconsistent preferences appears in a number of consumer decisions, including retirement savings (Laibson, Repetto and Tobachman 1998) and eating (Ruhm 2010). Policies addressing consumer behavior in these areas will not achieve a first-best outcome if they do not account for the consumers' time inconsistency. This paper develops such a framework for policy design for the case of externalities in the presence of time inconsistency; this may serve as a springboard for the analysis of other market failures or policy instruments under time inconsistency. For example, in the case of obesity policy, how does a calorie tax compare to a trans-fat ban?

The theoretical results provide some insight into policy design. First, I show that a Pigouvian tax that only accounts for externalities does not bring about the first-best outcome under time inconsistency. A Pigouvian tax leads to cars that are not fuel-efficient enough and are driven too few miles, compared to the first best. In general, gasoline consumption under time inconsistency can either exceed or fall below the first-best level. Second, the first-best outcome

³ See also Table 1 in Sanstad et. al. (2006). By contrast, Greene's (2010) reading of the econometric literature estimating consumers' valuations of fuel economy finds mixed results, with some studies finding under-valuing and some finding over-valuing of improvements in fuel economy (see his Table 2). Busse et. al. (2009) find that gasoline prices affect the new car market more so than the used car market.

⁴ Command-and-control policies are sometimes referred to as "direct regulatory instruments."

can be attained through a Pigouvian tax and a command-and-control mandate in the initial decision period. This bolsters intuition provided in earlier papers that, for example, fuel economy standards for cars increased efficiency relative to gasoline taxes (Greene 1998). Third, however, I show that the first best can be achieved with an incentive-based policy in the initial decision period rather than a command-and-control policy. Time inconsistency means that future costs are not fully realized by the consumer, but they can be introduced through a price instrument, e.g. a tax on fuel (in)economy. Thus, the common argument that behavioral anomalies give credence to command-and-control mandates over incentive-based mandates is not true in this case; either type of policy can achieve the first best.⁵ Fourth, in contrast to policies that address market failures caused by externalities, under consumer heterogeneity incentive-based policies do not necessarily dominate command-and-control policies. When consumers are time-consistent but heterogeneous in their preferences, a uniform Pigouvian tax on an externality induces the first-best outcome, and a uniform performance standard does not. For time-inconsistent preferences, under heterogeneity, neither a uniform tax nor a uniform performance standard induces the first best. This holds even when consumers are homogeneous with respect to their degree of present bias.

The simulation results suggest that, for the automobile market, the welfare gains from policies that address time inconsistency are substantial, and policies that ignore time inconsistency are substantially different from the optimal policies. The deadweight loss of a policy that addresses externalities from gasoline consumption but does not address time inconsistency ranges from \$160 to \$225 per new vehicle sale, which amounts to an economy-wide deadweight loss of \$1.44 billion to \$2.01 billion annually. Optimal price policy includes a tax that reduces the price differential between the average hybrid car and the average non-hybrid car by \$750 to \$2200. The best tax rate on gasoline is 18% to 30% higher than marginal external damages.

How should the findings of behavioral economics affect energy policy? It is often suggested that behavioral anomalies justify command-and-control policies. Gillingham et. al. (2009) offer as potential policy instruments for behavioral anomalies education, information, and product standards. Allcott and Mullainathan (2010) examine how non-price behavioral interventions ("nudges") affect consumer choices on energy use, and they argue that there are

⁵ See also the discussion of "behavioral feebates" for automobile fuel economy in Allcott and Wozny (2010).

potentially many low-cost instruments available to reduce consumption. For example, simply giving households information in their monthly bills about their relative electricity consumption tends to reduce consumption. Fischer et. al. (2007) argue that strengthening fuel economy standards will be welfare-increasing only if consumers are myopic with short horizons. Yet, no study looks for optimal energy or environmental policy design in the presence of time-inconsistent preferences.

One reason for the lack of much research in optimal policy design under behavioral anomalies is the difficulty of conducting welfare analysis with such anomalies. Standard welfare analysis is based on revealed preference, in which consumers' choices among available bundles gives information about preferences. Under behavioral anomalies, though, choices can be inconsistent (e.g. a consumer prefers A over B in some instance and B over A in another), and thus it is difficult to map them into utility or welfare functions. Several criteria for welfare analysis in the presence of time-inconsistent preferences have been suggested. This paper is agnostic about which welfare criterion to employ in the following sense: I use a particular welfare criterion (the "long-run" criterion, described below) to generate the optimal policy results, then I investigate how robust these results are to two alternate welfare criteria, including those proposed by Bernheim and Rangel (2009).

Some other studies have examined optimal policy in the presence of time-inconsistent preferences. O'Donoghue and Rabin (2006) and Gruber and Koszegi (2001) solve for optimal "sin" taxes on goods that cause future damages (e.g. to health) that are underweighted when consumed because of present bias. The two papers most similar in scope to this paper consider Pigouvian taxation of externalities when individuals exhibit behavioral anomalies. Johannson (1997) considers Pigouvian taxation when individuals exhibit altruism. Intuitively, one might think that when individuals care about the welfare of others, the efficient tax rate on an externality is lower than when individuals are purely self-interested, since their altruism causes them to account for the external damages of others. Johannson (1997) finds that this is not necessarily so. The optimal tax may be higher with altruism than without it because the socially efficient level of the externality may be lower with altruism than without it.

Lofgren (2003) considers Pigouvian taxation when individuals exhibit addiction. She finds that the first-best is achieved with the standard Pigouvian tax: addiction does not affect optimal policy. However, she considers some extensions, including time-inconsistency. Though

her focus is on addiction, some of her results are relevant to the question at hand in this paper. For instance, she finds that under myopic (time-inconsistent) preferences, the optimal tax differs from the Pigouvian tax, a result replicated here in a more complex model of time inconsistency without addiction.

Neither of these two papers, and no paper to my knowledge, directly answers the question of how to optimally address the market failure caused by externalities in an economy where consumers are time-inconsistent.

The next section below presents the base case representative agent model. Section 2 extends the model to multiple heterogeneous agents. Section 3 considers alternate welfare criteria. Section 4 presents simulation results.

I. Representative Agent Model

Consider a representative consumer making a decision over a durable good lasting T periods. The good is purchased in the initial period ($t = 0$). In each subsequent period ($t = 1$ through $t = T$), the consumer chooses the operating intensity of the good. For example, if the good is an automobile, the consumer chooses its fuel economy in the first period (miles per gallon) and chooses how many miles to drive in each subsequent period.

Rational choice theory predicts that a consumer trades off costs and benefits in a time-consistent way. The relative utility weighting of two consecutive time periods will not change over time. Exponential discounting (with a constant discount factor) achieves time consistency. However, consumers exhibit time inconsistency if using quasi-hyperbolic discounting instead of a constant discount factor. Under quasi-hyperbolic discounting, the discount factor applied in the present between any two consecutive future periods is δ , while the discount factor used between the current period and the following period is $\beta\delta$, where $\beta < 1$. The parameter β represents a "present-bias" in preferences, and δ is sometimes called the "long-run" discount factor.

Quasi-hyperbolic discounting leads to time inconsistency. A consumer at time t will make different future decisions than she will at another time period, even without any changes in information or realizations of uncertainty. Time inconsistency is a specific instance of a behavioral anomaly, an act deviating from predictions of rational choice theory.⁶ Mullainathan

⁶ Behavioral anomalies are termed "behavioral failures" in Shogren and Taylor (2008).

and Thaler (2001) classify behavioral anomalies into three classes; time inconsistency falls under the class "bounded willpower."

To work with a concrete example, let the durable good be an automobile, where the intensity of use is the number of miles driven each period. Consider first the consumer's choice of miles conditional on a particular vehicle with a given fuel efficiency. Let gpm be the fuel economy in gallons of gasoline per mile. Let m_t be the number of miles driven in period t , so that the total fuel consumption for the consumer in period t is $gpm \cdot m_t$. The consumer gets utility (in dollar equivalents) from driving described by a utility function $U(m_t)$, where $U' > 0$ and $U'' < 0$. The cost (in dollars) to the consumer per gallon of fuel is $gas_t + \tau_t$, where gas_t is an exogenous gasoline price and τ_t is a tax set by the government.

The consumer's surplus in period t is $U(m_t) - (gas_t + \tau_t) \cdot gpm \cdot m_t$. The privately optimal number of miles driven in period t conditional on prices and fuel economy is m_t^* , given by the first-order condition $U'(m_t^*) = (gas_t + \tau_t) \cdot gpm$.⁷ This implies that m_t^* is a function of the price of driving one mile: $m_t^* = m^*(gpm \cdot (gas_t + \tau_t))$.

Next consider the consumer's problem in period 0, that is, her decision over the fuel economy of the car (gpm) to maximize total discounted utility. Suppose that the car is not driven in period 0 so that period 0 utility is just the negative of the cost of the car, c . A car with fuel economy gpm costs $c(gpm)$. Assume that $c' < 0$, so that less fuel efficient cars (those with higher gpm) are less expensive, and that $c'' > 0$. The consumer's full problem is thus

$$\max_{gpm, \{m_t\}_{t=1}^T} -c(gpm) + \beta \sum_{t=1}^T \delta^t [U(m_t) - (gas_t + \tau_t) \cdot gpm \cdot m_t]$$

The consumer employs quasi-hyperbolic discounting when $\beta < 1$.⁸ With no uncertainty in gas prices, the consumer can choose m_t for each period at time $t = 0$. Since each choice of m_t is a static problem conditional on gas prices and gpm , the solution to each of those T static problems for m_t can be substituted into the consumer's problem, so that the consumer's problem can be expressed as a choice over just fuel economy:⁹

⁷ Ensure an interior solution by assuming that $U'(m) \rightarrow \infty$ as $m \rightarrow 0$.

⁸ See Laibson (1997). Quasi-hyperbolic discounting is also called (β, δ) discounting or quasi-geometric discounting. I focus on the case of $\beta < 1$ (present bias), although symmetric results arise from $\beta > 1$.

⁹ Here where the purchase decision over the durable good occurs in just the first period, this is equivalent to a multiple-self Nash equilibrium, as in Laibson et. al. (1998).

$$\begin{aligned} \max_{gpm} & -c(gpm) \\ & + \beta \sum_{t=1}^T \delta^t [U(m_t^*(gpm(gas_t + \tau_t))) - (gas_t + \tau_t) \cdot gpm \\ & \cdot m_t^*(gpm(gas_t + \tau_t))] \end{aligned}$$

Continuity and differentiability of c yield a first-order condition. After simplifying through an envelope condition from the consumer's static problem, this becomes

$$-c'(gpm) + \beta \sum_{t=1}^T \delta^t [-(gas_t + \tau_t) \cdot m_t^*(gpm(gas_t + \tau_t))] = 0$$

The first term including the negative sign is positive, and it represents the current-period benefit of a marginal increase in gpm : it is cheaper. The summation is negative, and it represents the discounted cost of a marginal increase in gpm : each future period's utility is lower because the cost of driving is higher. Call the solution to the consumer's problem gpm^* and m_t^* .

Consider next the social planner's problem, which differs from the consumer's problem in two respects. First, suppose that there is an externality associated with the use of fuel.¹⁰ The social planner considers the externality in its social welfare function. The total number of gallons of gasoline used in period t is $m_t \cdot gpm$; let the external damages from gasoline be $d(m_t \cdot gpm)$, where $d(0) = 0$, $d' > 0$, and $d'' \geq 0$.

The consumer's preferences are time inconsistent since $\beta < 1$. The social planner thus encounters a dilemma over deciding what to maximize, since different "selves" of the consumer at different periods have different utility functions. One approach is for the planner to maximize the utility function used by the period-zero self (Krusell, Kuruscu and Smith 2002); this approach might seem unappealing in that it underweights future selves' utilities. An alternate welfare criterion is to maximize a function identical to the initial period consumer's utility function but omitting the present bias, i.e. setting $\beta = 1$. This approach, because it has the planner applying only the long-run discount factor and not the present bias term, is sometimes called the long-run criterion. One interpretation of this criterion is that it represents the preferences of the consumer if she were to decide what to do in the period *before* she had to

¹⁰ Though this is a representative agent framework, the externality can be accommodated by supposing that the consumer does not account for its cost in her decision. (The generalization is that there is a continuum of consumers, all of whom just barely value the miniscule contribution their gasoline use makes to the aggregate externality.)

purchase the car (Gruber and Koszegi 2001). Another interpretation is that the consumer is, in a welfare-relevant way, making a mistake when she applies the present bias term β . That is, the consumer's "decision utility" includes a $\beta \neq 1$ while her "true utility" does not. The social welfare function maximizes her true utility. (True utility is sometimes also called "hedonic utility" or "experienced utility.")

Papers using the long-run criterion to conduct welfare analysis include Carroll et. al. (2009), O'Donoghue and Rabin (2006), and Gruber and Koszegi (2001). A justification for the social planner using a discount rate that differs from the market discount rate is found in Caplin and Leahy (2004). Robson and Samuelson (forthcoming) develop a model based on biological evolution to explain the existence of the discrepancy between decision and true utilities. The long-run criterion, though, requires the paternalistic assumptions that individuals' decisions are not indicative of their true, welfare-maximizing preferences. Gruber and Koszegi (2001), for instance, argue that time-inconsistent preferences demonstrate that people do not act in their best interests. Alternative welfare criteria, including those presented by Bernheim and Rangel (2009) attempt to be less paternalistic. In order to be agnostic about what welfare criterion to employ, I adopt the following strategy. I solve for optimal policy under the long-run criterion, and then I investigate how robust those policy solutions are to alternate welfare criteria. Later, in section 3, I show conditions under which a first-best solution defined according to the long-run criterion is also considered welfare-improving under the alternate criteria.

Under the long-run criterion and accounting for the externality from pollution d , the social planner's problem is

$$\max_{gpm, \{m_t\}_{t=1}^T} -c(gpm) + \sum_{t=1}^T \delta^t [U(m_t) - gas_t \cdot gpm \cdot m_t - d(m_t \cdot gpm)]$$

As with the consumer's problem, each choice of m_t is made in a static setting conditional on gpm . It can be written as a function of gas_t and gpm ; let this be $m_t^{opt} = m_t^{opt}(gas_t, gpm)$.

Then, the social planner's problem can be similarly rewritten:

$$\begin{aligned} & \max_{gpm} -c(gpm) \\ & + \sum_{t=1}^T \delta^t [U(m_t^{opt}(gas_t, gpm)) - gas_t \cdot gpm \cdot m_t^{opt}(gas_t, gpm) \\ & - d(m_t^{opt}(gas_t, gpm) \cdot gpm)] \end{aligned}$$

A first-order condition for the social planner's problem is

$$-c'(gpm) + \sum_{t=1}^T \delta^t [-(gas_t + d'(m_t^{opt}(gas_t, gpm)) \cdot gpm)] \cdot m_t^{opt}(gas_t, gpm) = 0$$

Call the solution to the planner's problem gpm^{opt} and m_t^{opt} . By comparing the first-order conditions of the consumer and the planner, it is apparent that when $\beta = 1$ the first-best outcome will be chosen by the consumer when $\tau_t = d'(m_t^{opt} \cdot gpm^{opt})$ for all $t \in [1, \dots, T]$. This is the Pigouvian tax rate on gasoline; call it τ_t^{pig} . Because there is no deviation between the consumer's decision utility and true utility when $\beta = 1$, she fully accounts for the future variable costs of driving the car when she makes her decision in period zero over fuel economy. If the externality from future driving is internalized through a Pigouvian tax, then her decision is optimal in every period. That is, $m_t^* = m_t^{opt}$ for all $t \in [1, \dots, T]$ and $gpm^* = gpm^{opt}$.

The main results concern the case where $\beta < 1$: the consumer discounts quasi-hyperbolically, but the social planner does not. The first proposition states that no set of gasoline taxes exist, not even the Pigouvian taxes, that lead to the first-best outcome gpm^{opt} and m_t^{opt} , and it describes the direction of the error when using the Pigouvian taxes. Proofs are presented in the Appendix.

Proposition 1: If $\beta < 1$, then there does not exist any set of tax rates $\{\tau_t\}$ for all $t \in [1, \dots, T]$ that lead to the first-best outcome gpm^{opt} and m_t^{opt} . If $\tau_t = \tau_t^{pig}$ for all $t \in [1, \dots, T]$, then $gpm^ > gpm^{opt}$ and $m_t^* < m_t^{opt}$ for all $t \in [1, \dots, T]$.*

Since no gasoline tax exists that will achieve the first best, clearly the Pigouvian tax will not achieve the first best. If τ_t^{pig} is levied, what outcome does it lead to? Intuitively, since the consumer is underweighting the future operating costs of the car, she will pay too little for fuel efficiency in period zero and buy a car with a gpm that is too high. But once that inefficient car is bought, the consumer faces a higher per-mile price of driving compared to the optimal fuel efficiency. So the number of miles driven is fewer than optimal in each period.

Thus, it is not clear how total gasoline consumption (the product of fuel economy and mileage) under the Pigouvian tax compares to the optimal level of gasoline consumption. The present bias in preferences could cause total gasoline consumption and emissions to be greater than or less than the optimal level of gasoline consumption and emissions. Suppose that utility over mileage is iso-elastic with a coefficient of relative risk aversion φ , so that $u(m) = \frac{m^{1-\varphi}}{1-\varphi}$.

The price elasticity of demand for miles driven is $-1/\varphi$ (this is also equal to the price elasticity of demand for gasoline). Under this functional form, present bias ($\beta < 1$) leads to an over-consumption of gasoline if and only if $\varphi > 1$, that is, the absolute value of the price elasticity is less than 1. Your car has too low of a fuel economy because of present bias. If your demand for mileage is price-inelastic, then the decrease in miles driven because of the low fuel economy is small and is not enough to offset the lower fuel economy, and total gasoline consumption increases. Contrariwise, if you are price-elastic, then the decrease in miles driven is large and more than offsets the decreased fuel economy, and total gasoline consumption decreases.

Not just the Pigouvian gasoline tax rates, but no set of gasoline tax rates produces the first-best outcome when $\beta < 1$. Within any single period $t > 0$ the consumer makes no behavioral anomalies, since her decision variable m_t only affects her period- t utility. Given the optimal fuel economy, the optimal miles driven in period t can be achieved only through $\tau_t = \tau_t^{pig}$. But this set of tax rates does not achieve the optimal fuel economy choice in period zero because of the consumer's present bias distorting her period zero decision. If the planner can only tax gasoline consumption in periods $t > 0$, then the optimal decision in period $t = 0$ can never be achieved.

Though no set of gasoline taxes can induce the first best, regulators may be constrained and only have gasoline taxes at their disposal. What might a second-best policy instrument, using only gasoline taxes, look like? Intuitively, one might think that in each period, the second-best τ_t is higher than the Pigouvian tax rate τ_t^{pig} to attempt to overcome the present bias. However, this intuition is not true in general. As discussed above, present bias could cause gasoline consumption to either increase or decrease. It follows that the second-best gasoline tax may exceed the Pigouvian tax or may fall below the Pigouvian tax. Under present bias, the consumer is underweighting future costs of gasoline consumption. The consumer is also underweighting future benefits of gasoline consumption, that is, the utility from driving. If the underweighting of the future benefits dominates the underweighting of the future costs, then the consumer will consume too little gasoline relative to the optimum, and the second-best gasoline tax will be lower than the Pigouvian gasoline tax.¹¹ In this model, a present bias in preferences

¹¹ Suppose that utility is iso-elastic, $c(gpm) = gpm^{-\gamma}$ with $\gamma > 1$, and $d(x) = x^\kappa$ with $\kappa > 1$, and $T = 1$. This allows the consumer's decisions m_t^* and gpm^* to be solved analytically as a function of gas_t and τ_t as well as the functional parameters. Under this parameterization, the Pigouvian tax and the second-best tax can be found. When $\varphi > 1$, then the second-best tax exceeds the Pigouvian tax; the opposite holds when $\varphi < 1$.

does not necessarily increase pollution, and therefore a second-best gasoline tax is not necessarily higher than the Pigouvian tax.¹²

Thus, the regulator needs another policy instrument to achieve the first best. One such instrument is a fuel economy standard.

Proposition 2: If $\beta < 1$, then the first best is achieved by setting $\tau_t = \tau_t^{pig}$ in each period $t > 0$ and setting a fuel economy standard that mandates a maximum gpm of gpm^{opt} .

Proposition 2 may provide some rationale for having gasoline taxes in conjunction with corporate average fuel economy (CAFE) standards for new passenger automobiles. CAFE standards have been in place since the 1978 model year, when they were 18.0 miles per gallon for passenger cars. The 2011 model year standard is 30.2 miles per gallon. The federal gasoline tax is 18.4 cents/gallon and the average state tax rate is 27.2 cents per gallon, as of 2009 Q1. Later in the simulation section, I will compare these values to the values that induce the first best.

With two policy instruments to use, the planner can achieve the first-best outcome. The second instrument, however, need not be a command-and-control standard. Instead, the regulator can set a tax to be paid in period zero based on the car's fuel economy. Call this tax τ_{gpm} .

Proposition 3: If $\beta < 1$, then the first best is achieved by setting $\tau_t = \tau_t^{pig}$ in each period $t > 0$ and setting $\tau_{gpm} = (1-\beta) \cdot \sum_{t=1}^T \delta^t \cdot (gas_t + \tau_t^{pig}) \cdot m_t^{opt}$.

The summation in τ_{gpm} , $\sum_{t=1}^T \delta^t \cdot (gas_t + \tau_t^{pig}) \cdot m_t^{opt}$, is the full discounted benefit of a marginal decrease in gpm . The consumer only accounts for a fraction β of the full benefit, and so the remaining $(1 - \beta)$ is in the tax, bringing about the first-best. The intuition behind the tax on fuel economy τ_{gpm} is analogous to the intuition behind the tax on the externality τ_t^{pig} . With an externality, there is a cost that is not faced by the agent, and a tax that forces the agent to face that cost (a Pigouvian tax) yields the first best. With present bias here characterized as an "internality," there is another cost that is not faced by the agent: part of the future cost of lower fuel economy. This cost is not faced by the agent in period zero because of her present bias. The optimal fuel economy tax τ_{gpm} forces her to face the full cost.

¹² When mileage in each period m is fixed rather than a choice variable (perfectly inelastic demand), then present bias always increases gasoline consumption and the second-best gasoline tax always exceeds the Pigouvian tax.

Behavioral anomalies are often invoked as justification for command-and-control policies over incentive-based policies (Greene 1998).¹³ But just like with externalities, behavioral anomalies can be internalized through price-based incentives. Empirical evidence suggests that consumers do in fact respond to price when making decisions on energy-efficiency investments (Hassett and Metcalf 1995). In this representative agent model, there is no difference between the command-and-control standard and the *gpm* tax. With heterogeneous agents, though, there is reason to suspect that incentive-based policies are cost-effective relative to command-and-control policies. This will be investigated in the following section.¹⁴

Proposition 3 may provide some rationale for the "gas guzzler" tax, a tax paid by the manufacturer on each car that fails to meet a minimum fuel economy threshold. The tax level is based on the car's fuel economy; it ranges from \$1000 for a car with an mpg between 21.5 and 22.5 to \$7700 for a car with an mpg less than 12.5 mpg. However, minivans, pickup trucks and SUVs are not subject to the tax, and thus it only affects low-fuel-economy cars, mainly sports cars. This is a small fraction of total new car sales.

Two market failures require two instruments. Proposition 1 showed that no set of gasoline taxes, without a policy on *gpm*, can achieve the first best. Similarly, no policy on *gpm*, without a policy on gasoline consumption, can achieve the first best.

*Proposition 4: If $\beta < 0$, when $\tau_t = 0$ for all $t > 0$ then no policy on *gpm*, whether a tax τ_{gpm} or an efficiency standard gpm^{max} , can achieve the first best.*

A policy is needed in period $t > 0$ to correct the market failure from the gasoline externality. A policy is needed in period zero to correct the market failure from the behavioral anomaly.

II. Model with Heterogeneous Agents

In the representative agent model, either a performance standard (e.g. a minimum miles-per-gallon requirement) or an incentive-based policy (e.g. a tax on fuel economy) brings about the first-best outcome. One may suspect that under consumer heterogeneity, the incentive-based policy dominates the command-and-control policy; this result is well-known in policies that address externalities. An incentive-based policy for an externality with uniform external costs is

¹³ See also Gillingham et. al. (2009). Their Table 1 (p. 604) lists as potential policy options for behavioral failures relevant to energy efficiency only education, information, and product standards. Pricing is listed as a policy instrument only for market failure, like externalities.

¹⁴ This result is similar to Proposition 3 from Lofgren (2003), where the optimal tax for an addictive, myopic consumer is equal to the Pigouvian tax in the second period but larger than the Pigouvian tax in the initial period. Alcott and Wozny (2010) also consider a tax on automobile fuel economy and label it a "behavioral feebate."

cost-effective. A uniform command-and-control mandate does not provide the flexibility for individuals with different preferences or different abatement costs.

However, this reasoning does not apply to policies that address time-inconsistent preferences. With heterogeneous agents, neither a uniform tax on fuel economy nor a uniform performance standard necessarily brings about the first best outcome. By "uniform" I mean one that does not vary by individual. For heterogeneity in present bias, this result seems obvious. For example, if some consumers exhibit present bias and other do not, then a uniform policy to address time-inconsistent preferences seems like it cannot be optimal. However, the results below do not assume heterogeneity in present bias; all consumers have the same β . Rather, consumers vary only by their instantaneous utility over mileage, U . It is not so obvious that a uniform tax should not induce the first best under this specification of heterogeneity.¹⁵

Consider a model with two consumers, indexed by $i = 1, 2$. The two consumers differ from each other only in their utility function over miles driven; the first consumer's is U_1 , and the second consumer's is U_2 . Both consumers have the same value for the present bias in preferences, β . The social planner maximizes the sum of both true utilities:

$$W = \sum_{i=1}^2 [-c(gpm_i) + \sum_{t=1}^T \delta^t \cdot [U_i(m_{t,i}) - gas_t \cdot gpm_i \cdot m_{t,i} - d(E_t)]]$$

Damages from emissions are again given by $d(E_t)$, where $E_t = \sum_{i=1}^2 m_{t,i} \cdot gpm_i$ is the sum of both consumers' emissions. The first best is given by the solution to the planner's first-order conditions for gpm_i and $m_{t,i}$.

$$\begin{aligned} U'_i(m_{t,i}^{opt}) - (gas_t + 2 \cdot d'(E_t^{opt})) \cdot gpm_i^{opt} &= 0 \\ -c'(gpm_i^{opt}) + \sum_{t=1}^T \delta^t \cdot [-(gas_t + 2 \cdot d'(E_t^{opt})) \cdot m_{t,i}^{opt}] &= 0 \end{aligned}$$

The optimal level of emissions each period E_t^{opt} is given by the optimal fuel economy and miles drive of each consumer. The damages from emissions in each first-order condition are multiplied by the number of consumers (2).

Each consumer maximizes the decision utility function that incorporates quasi-hyperbolic discounting: $-c(gpm_i) + \beta \sum_{t=1}^T \delta^t \cdot [U_i(m_{t,i}) - gas_t \cdot gpm_i \cdot m_{t,i}]$. The consumer may also

¹⁵ Many other sources of heterogeneity are possible, including the time horizon of the automobile T . O'Donoghue and Rabin (2006) consider heterogeneity in β .

face a tax on emissions in each period τ_t , a tax on fuel economy τ_{gpm} , or a restriction on fuel economy gpm_{max} . It can be shown that in the standard case when $\beta = 1$, a uniform tax rate in each period $\tau_t = 2 \cdot d'(E_t^{opt})$ induces the first best; this is the Pigouvian tax τ_t^{pig} .

Proposition 5: In the model with two heterogeneous consumers with $U_1 \neq U_2$, if $\beta < 1$, no combination of a uniform τ_{gpm} and a uniform set of τ_t will necessarily induce the first-best outcome.

Why, with heterogeneous agents, do a uniform emissions tax and a uniform fuel economy tax not achieve the first-best outcome? The externality in this model is a pure public bad; for a given level of emissions, an additional unit of emissions causes the same marginal external damage regardless of who produces it. So, everyone should face the same marginal cost (tax). On the other hand, the marginal cost of the market failure from time-inconsistent preferences is not identical across consumers. It is, in fact, equal to the expression for τ_{gpm} in the proof of Proposition 5. The cost that consumer i fails to face in her decision utility function is a part of her future periods' utility. But, this cost differs between the two consumers since the heterogeneity in utility functions leads to heterogeneity in optimal mileage $m_{t,i}$. The non-uniformity of the optimal tax is analogous to a non-uniform optimal Pigouvian externality tax in a case where damages from emissions are not independent across sources. For instance, if emissions from power plants located close to densely populated areas cause more damage than emissions from power plants far away from populated areas, then the Pigouvian emissions tax rate on the closer power plants is higher than the tax rate on the other plants (Mauzerall, et al. 2005).

Thus, the argument for the dominance of incentive-based policies over command-and-control policies does not apply to policies aimed at addressing the market failure caused by time-inconsistent preferences. Neither policy will attain the first best, although one may induce a second-best outcome with a higher level of social welfare than the other. The ranking of the two policies is unclear in general.

Just as no uniform tax on fuel economy can efficiently address the behavioral market failure, neither can a uniform command-and-control policy.

Proposition 6: In the model with two heterogeneous consumers with $U_1 \neq U_2$, if $\beta < 1$, no combination of a uniform efficiency standard gpm^{max} and a uniform set of τ_t will necessarily induce the first-best outcome.

III. Alternative Welfare Criteria

These results are based on the long-run criterion. This criterion is intuitive: present bias creates an "internality" that is analogous to an externality. Optimal policy involves getting the prices right: forcing the present consumer to pay for the externality and for the internality. The long-run criterion is used frequently in the literature (Carroll, et al. 2009). However, it is controversial, since it abandons the tenet that welfare analysis be guided by revealed preference. The criterion asserts that an individual is not acting in his own best interests and his actions do not maximize his welfare. There is thus a role for paternalistic government intervention.

Because of the strong assumptions behind the long-run criterion, in this section I explore how robust the above results are to two other welfare criteria. The first alternate welfare criterion models the decision of the individual over time as an intrapersonal game, where each "self" at a period of time is a distinct player. Welfare analysis considers Pareto optima or Pareto improvements among the various selves of the game (Bhattacharya and Lakdawalla 2004). The second criterion is based on the recent work by Bernheim and Rangel (2009). Under both alternative criteria, I present conditions under which the optimal allocation, as defined under the long-run criterion, is also welfare-improving under the alternate criteria compared to the standard, Pigouvian solution that ignores present bias.

Multiself Pareto Optima

Consider an intrapersonal game between the different selves. Employ a Nash equilibrium; every self takes the strategies of all other selves as given and chooses a best response. In the representative agent model above, the solution that I find is identical to a Nash equilibrium of an intrapersonal game. Each self $t > 0$ responds only to the fuel economy of the car that it inherits. Self $t = 0$ chooses gpm to maximize its utility, given that each future self will optimize over m_t in each future period.

Welfare analysis for intrapersonal games can be done by evaluating multiself Pareto improvements. If an outcome can be altered by a planner such that each self is at least as well off and at least one self is better off, then this is an unambiguous welfare improvement, according to this criterion.

Is the optimal policy outcome under the long-run criterion a multiself Pareto optimum? Yes, because the planner is maximizing a weighted sum of the individual selves' utilities.

Likewise, the solution to the individual's maximization problem is also a multiself Pareto optimum, because the period zero self is also maximizing a weighted sum of the individual selves' utilities. Given that the individual's solution is a Pareto optimum, it does not appear that a welfare enhancing social policy, i.e. a Pareto improvement, is possible.

However, the representative agent model does not explicitly model the return of revenue from the emissions tax collection. Revenues are returned lump-sum. With only one representative agent, the tax payments just equal the lump-sum return, so there is no change in utility from the tax payments per se. But, the tax on gasoline consumption induces lower mileage, increasing the consumer's welfare because of the reduction in the externality. The tax is therefore welfare-increasing. Likewise, with a tax in period zero on gpm , the tax payments are returned to the period zero self.

More generally, though, the planner may be able to achieve a multiself Pareto improvement by reallocating tax revenues across time. Because the gpm policy may lower the period zero self's utility to increase all of the other selves' utilities, some of the tax revenues from periods $t = 1, \dots, T$ could be returned to the period zero self to make him at least as well off.

In particular, is it possible for a planner to reallocate from the Pigouvian policy to an alternate policy and create a Pareto improvement? Define the "Pigouvian policy" to be the policy that sets the Pigouvian tax rate in each period $t > 0$ and no tax on gpm in period zero, and that returns all tax revenue in each period lump sum. That is, defining s_t to be the lump sum payment to the individual in period t , $s_t = \tau_t \cdot gpm \cdot m_t$ for all $t > 0$. Define the outcomes under the Pigouvian policy as gpm^0 and m_t^0 for $t > 0$. Define the "optimal policy" as the Pigouvian tax combined with an optimal tax on gpm , τ_{gpm} , as defined in Proposition 5. The optimal policy also includes subsidies s_t that need not be returned in full in each period. Rather, the subsidies can be reallocated across time so long as an overall budget constraint is met: $\sum_{t=0}^T s_t \leq gpm^{opt} \cdot \tau_{gpm} + \sum_{t=1}^T gpm^{opt} \cdot m_t^{opt} \cdot \tau_t^{pig}$. The left-hand-side of this budget constraint is total lump sum payments over the $T + 1$ periods, and the right-hand-side is the total tax revenues (the gpm tax in the zero period and the Pigouvian tax in each remaining period). Define the outcomes under the optimal policy as gpm^{opt} and m_t^{opt} for $t > 0$. Because they are lump sum payments, the s_t values do not affect these outcomes.

Given this structure, the following proposition presents a condition under which the optimal policy represents a multiself Pareto improvement over the Pigouvian policy.

Proposition 7: The optimal policy is a multiself Pareto improvement over the Pigouvian policy as long as the following condition is met: $\sum_{t=1}^T [u(m_t^{opt}) - u(m_t^0) + (d(gpm^0 \cdot m_t^0) - d(gpm^{opt} \cdot m_t^{opt})) + gas_t \cdot (gpm^0 \cdot m_t^0 - gpm^{opt} \cdot m_t^{opt})] > -c(gpm^0) + c(gpm^{opt})$.

If that condition holds, then the planner can improve upon the Pigouvian solution by implementing the optimal *gpm* tax and a set of intertemporal transfers. The optimal *gpm* tax alone increases the planner's maximand (the "true utility") but may decrease the period zero self's maximand (the "decision utility"). However, the lump sum payments transfer resources from future selves to the period zero self and end up increasing the period zero self's maximand.

Does the condition in Proposition 7 hold? It depends on functional forms and parameter values. In the calibrated simulation results presented in the next section, the inequality holds. Thus, the first-best results found for that market under the long-run criterion can also attain a multiself Pareto optimum.

Bernheim and Rangel Criteria

A second alternative welfare criterion is introduced in Bernheim and Rangel (2009) (hereafter BR). They develop a choice-based welfare economics that accommodates nonstandard behavioral models. In particular, consumers can exhibit choice behavior that violates the standard model of well-defined choices. In some situation, an individual can choose x over y , and in other situations the individual, facing the same budget constraint, chooses y over x . The only difference between the two situations consists of features that we do not think ought to be relevant to a social planner; these features are called ancillary conditions. Examples of ancillary conditions include the order in which options are presented, or the assignment of a default option. In the case of time-inconsistent preferences, the ancillary condition of interest is the point of time in which the decision is made.

Under such a framework, BR define weak and strict revealed preference relations analogously to those relations defined in the standard model (without ancillary conditions). So, if a bundle x is strictly unambiguously chosen over y , that means that y is never chosen in any budget-ancillary condition combination (termed a generalized choice set) where x is available. If some bundle x in a generalized choice set has no bundles that are strictly unambiguously chosen over it, then it is said to be an individual welfare optimum. This criterion thus respects

the choices that individuals make, regardless of the presence of ancillary conditions rendering such choices seemingly inconsistent under the standard model. A welfare improvement is moving from one bundle to another where the second bundle is strictly unambiguously chosen over the first, that is, under no ancillary conditions will it not be chosen.

The BR framework allows for analysis of behavioral anomalies broader than just time-inconsistent preferences. However, they directly apply their framework to the case of $\beta\delta$ preferences and present a theorem to describe when one bundle will be strictly unambiguously chosen over another bundle given such preferences. Let a bundle x be defined by a vector of consumption scalars x_t from $t = 1$ to T , and likewise for bundle y . Their Theorem 4 states that x is strictly unambiguously chosen over y if and only if $\sum_{t=1}^T (\beta\delta)^{t-1} u(x_t) > u(y_1) + \beta \sum_{t=2}^T \delta^{t-1} u(y_t)$. The right-hand-side of this inequality is the first period's decision utility from bundle y . The left-hand-side is the utility that would be received from bundle x under a time-consistent discount factor $\beta\delta$.

Asking the analogous question from the above subsection on multiself Pareto optima: Is the optimal policy strictly unambiguously chosen over the Pigouvian policy? As when considering a multiself Pareto optimum, the answer is no: the optimal policy by definition must make the first period decision utility lower, since the Pigouvian policy maximizes the first period decision utility.¹⁶ But, as before, we can consider a system of intertemporal intrapersonal transfers that could make the inequality hold. Define the transfer to the individual in period t as s_t ; the government's budget constraint is identical to that in the previous subsection. Suppose also that the Pigouvian policy includes no intertemporal transfers; each period's tax revenue is immediately returned to the consumer. Then, the BR inequality describing when the optimal policy is strictly unambiguously chosen over the Pigouvian policy is:

¹⁶ That is, the first-period decision utility under x is greater than the first-period decision utility under y is a necessary but not sufficient condition for the inequality in the text to hold (see BR, p. 70). This result, and in fact all of their Theorem 4, depends on the assumption that utility in each period is non-negative. The model here can be accommodated to that assumption with a suitable constant additive term in each period's utility.

$$\begin{aligned}
& -c(gpm^{opt}) + s_0 - gpm^{opt} \cdot \tau_{gpm}^{opt} \\
& + \sum_{t=1}^T (\beta\delta)^t \cdot [u(m_t^{opt}) - (gas_t + \tau_t^{pig}) \cdot gpm^{opt} \cdot m_t^{opt} - d(gpm^{opt} \cdot m_t^{opt}) \\
& + s_t] > -c(gpm^*) + \beta \cdot \sum_{t=1}^T \delta^t \cdot [u(m_t^*) - gas_t \cdot gpm^* \cdot m_t^* - g(gpm^* \cdot m_t^*)]
\end{aligned}$$

If a sequence of transfers s_t that satisfies the intertemporal budget constraint can be devised that satisfies this inequality, then such a sequence of transfers, combined with the optimal policy, would be strictly unambiguously chosen over the Pigouvian policy.

As with the inequality in Proposition 7, whether or not this inequality holds depends on functional forms and parameters of the model. Thus whether or not the optimal policy under the long-run criterion is robust to these alternate welfare criteria is indeterminate in general. In the simulations that I describe in the following section, both of these inequalities are satisfied. Therefore, the optimal policy results for this particular application are robust to these alternate criteria.

IV. Numerical Simulation

I now turn to a model to be solved computationally, in an attempt to find the magnitude of the effects described in the analytical models above. The model is calibrated to consider consumer decisions over automobile purchases and gasoline consumption. The numerical model adds a number of elements that are absent from the analytic model. In particular, it allows more broadly for heterogeneity among agents (there need not be just one or two types of agents). It also allows a scrapping decision and thus makes endogenous the lifetime of a car, T .

Calibration

The consumer's utility over miles driven in a single period is $u(m) = C \frac{m^{1-\varphi}}{1-\varphi}$, where $1/\varphi$ is the price elasticity of demand for miles driven (also equal to the price elasticity of demand for gasoline) and C is a constant. This is the short-run price elasticity, since it is the response in miles traveled to a change in price holding constant the fuel economy of the car. This price elasticity has been estimated; values vary but it is generally found to be significantly less than one (Hughes, Knittel and Sperling 2008), (Espey 1998), (Espey 1996). I use the preferred value

of -0.34 , from a meta-analysis (Brons, et al. 2008) of short-run gasoline demand elasticities. This implies $\varphi = 2.941$. Later in the simulation results, I explore what this value of the short-run elasticity implies about the long-run elasticity.

To calibrate the scale parameter in utility C , I use data on mileage of passenger cars among US households from the 2001 National Household Travel Survey (NHTS).¹⁷ The form of the utility function predicts that a household's optimal mileage, m^* , is a function of the total price of a mile driven in the following way: $m^* = \left(\frac{\text{price}}{c}\right)^{-1/\varphi}$. The total price of a mile driven is the price of a gallon of gasoline time the car's fuel economy in gallons per mile.

Heterogeneity is incorporated into the model by allowing for different types of consumers through different values for C in the utility function. In particular, I calibrate C separately for drivers of four different vehicle types identifiable in the NHTS: cars (including station wagons), vans (mini/cargo/passenger), SUVs, and pickup trucks. For each vehicle type I calculate the total price of a mile driven based on the national average tax-inclusive gasoline price (\$2.72)¹⁸ and the average fuel economy for vehicle type. Total annual miles traveled is reported in the NHTS; for cars the mean value is 11681. Given a price elasticity of -0.34 , this pins down the value of C . Under this specification of heterogeneity, all consumers within one of the four consumer groups are identical, and consumers are unable to choose between groups. Instead, within a vehicle type consumers are able to choose a fuel economy level at a cost, $c(\text{gpm})$, specific to that vehicle type. To the extent that policy can move individuals into different types of vehicle rather than just different fuel economy levels within a vehicle type, these simulation results may misspecify optimal policy.¹⁹

The cost of a vehicle as a function of fuel economy, $c(\text{gpm})$, is calibrated separately for each vehicle type from manufacturer's suggested retail prices (MSRPs) of models with varying fuel economy. MSRP data are from Automotive News, and fuel economy data are from the Environmental Protection Agency.²⁰ In theory, one would regress price on fuel economy for

¹⁷ Available here: <http://nhts.ornl.gov/download.shtml>.

¹⁸ Available weekly at http://www.eia.doe.gov/petroleum/data_publications/wrgp/mogas_home_page.html. The value used here is taken from July 19, 2010.

¹⁹ The direction of bias or misspecification is unclear. For example, suppose that optimal policy actually involves moving some consumers from SUVs to cars. A higher gasoline tax than the one found in this model may be necessary to achieve that vehicle type switch. Or, with the option available to switch, the tax may need not be as high as in this case where there is no option to switch.

²⁰ Available at <http://www.autonews.com/section/prices> and <http://www.fueleconomy.gov/feg/download.shtml>, respectively.

each vehicle class to find $c(gpm)$. However, even after controlling for observables like size and horsepower, the regression invariably finds a *negative* correlation between price and fuel economy; $c(gpm)$ is increasing. This is because more fuel-efficient cars tend to be smaller and come with fewer features; luxury cars are less fuel-efficient. The data are not rich enough to control for these features that affect price and are correlated with fuel economy.

Instead of a regression, I calibrate the cost function based on two data points: the average cost and fuel economy for a non-hybrid car and for a hybrid car. Hybrid versions of cars are nearly equivalent save for the increased fuel economy (e.g. several models come in a hybrid and non-hybrid version). The price differential should thus represent the additional cost of higher fuel economy only. For cars, for example, the mean price of non-hybrid models is \$28,932 and the mean fuel economy in miles per gallon is 23.1; the mean price for hybrid models is \$34,515 and the mean fuel economy in miles per gallon is 35.2. Between these two points on $c(gpm)$, an infinite number of functions could be fit. I thus consider only quadratic functions that must be decreasing in gpm and convex, to satisfy the second-order condition in the analytical model's maximization problem. These restrictions do not pin down a particular function, so in sensitivity analysis I examine the impacts of varying $c(gpm)$.

External damages from gasoline consumption are assumed to be linear, so that marginal external damages are constant ($d(m \cdot gpm) = d \cdot m \cdot gpm$). This constant is taken from a recent assessment of the optimal gasoline tax in the US (Parry 2009). This paper finds that the optimal tax rate (i.e., marginal external damages) is \$1.23 per gallon. Of these damages, the majority comes from congestion externalities (52 cents) and accident externalities (41 cents). Climate change externalities account for nine cents, and other pollutants account for 12 cents. The remaining 10 cents is from oil dependence externalities. Importantly, the two dominant categories comprising the optimal tax (congestion and accidents) are externalities tied not to gasoline consumption but rather to mileage. I include these in the base case simulations but examine alternate specifications in sensitivity analysis.

The discount factors are taken from Laibson et. al. (2007), who estimate the parameters of a quasi-hyperbolic discounting model using a structural model and data on savings and consumption choices of US households. Their benchmark estimates imply $\beta = 0.7$ and $\delta = 0.96$ for annual periodicity. Brown et. al. (2009) find in laboratory experiments a value of β ranging from 0.6 to 0.7. These values are comparable to estimates in Allcott and Wozny (2010), who

examine the US automobile market and find that consumers value \$1 worth of expected discounted gasoline expenditures only 61 cents (though Allcott and Wozny (2010) are agnostic about whether the underweighting is due to present bias). By contrast, Fang and Silverman (2009) find a much lower β of 0.35 from data on welfare program participation.

The consumer's utility function must account for the fact that older cars are less preferable to newer cars, otherwise a consumer would never replace her car. I add a negative term to the utility function linearly increasing in vehicle age ($= -D \cdot v$, where v is the vehicle's age). The coefficient D is calibrated via simulation so that the predicted average vehicle age matches the average vehicle age in the NHTS. For cars, the average age is 8.98 years. The resulting calibrated value of the coefficient of the age disutility term is 322; the dollar equivalent utility of a car is reduced by \$322 for each year of its age.

In the analytic model, a consumer choosing to buy a car in period t pays for it in period t . However, most new car purchases (70%) in the US are financed. As of 2007, the average down payment on a new car loan is 10% of the total price, the average loan length is six years, and the average interest rate is 7%, according to Edmunds.com. The financing option is very relevant to consumers' decisions under present bias. If the entire cost must be paid up front, then present-biased consumers will probably be less likely to buy durables. If only a small fraction of the cost is paid up front, then present-biased consumer will probably be more likely to buy durables. Thus, present bias is often cited as a reason why consumers buy "too much" of items that are financed (e.g. new cars) and "too little" of items that are not financed (e.g. new appliances). In the simulation, I assume that all new car purchases are financed at the average rate terms (10% down, six years, 7% interest rate). Depending on the values of the consumer's present bias (β) and long-run discount factor (δ), this could increase or decrease the apparent cost of durables to the consumer.

Table 1 summarizes the parameters and functional forms of the model.

Simulation

The model is solved for a finitely lived agent through backwards induction, as in other papers that model quasi-hyperbolic discounting (Laibson, Repetto and Tobacman 2007), (Laibson, Repetto and Tobachman 1998). The representative consumer lives 90 years but does

not begin owning a car until age 21. In each period, the agent decides how many miles to drive his current car and decides whether or not to scrap the car and replace it with a new one.

The decision over miles traveled, contingent on type of car, can be solved in the same manner as in the analytic model. Then, given optimal miles traveled, the consumer's decision over scrapping and replacement is solved computationally. Because of the fixed time horizon, this can be solved through backwards induction. That is, in the last period the agent chooses the optimal miles traveled for each possible vehicle and then decides the optimal scrapping decision. (Of course, the agent will never choose to buy a new car in the final period since he will not be around to use it in the following period.) Then, given the agent's optimal behavior in the final period, the agent in the second-to-last period chooses a course of action. This continues until the first period. The agent at each period is discounting quasi-hyperbolically and knows that his future selves optimize with a time-inconsistent discount function.²¹ The solution is identical to an intrapersonal game, in which the agent is a distinct player at each time period, and solutions are restricted to Markov equilibria (Strotz 1955-1956). Software implementing the solution method is available on the author's website.

Given a representative agent's solution, policy options can be analyzed by evaluating the agent's true utility at his solution. The true utility includes the external damages from gasoline consumption and does not include the present bias factor β . Then, optimal or second-best policy can be found by maximizing the value of true utility over the policy variable, for example the gasoline tax. The first-best solution can also be found for comparison, by evaluating the agent's problem without present bias ($\beta = 1$) and where the gasoline tax is set to equal marginal external damages. I run simulations based on 100 representative consumers, whose initial vehicle allocation (ages and fuel economy) is chosen to represent the household-owned vehicle fleet described in the NHTS.

Results

I present simulation results from three different specifications of the model. In the first specification, I omit both heterogeneity and endogenous vehicle lifetimes. Without these two features, the model is identical to the theoretical model. Thus the results of that model can be

²¹ Thus it is said that agents are *sophisticated* about their time-inconsistency. In contrast, a *naïve* agent would act now as if his future selves would be consistent, although in the future they would not (O'Donoghue and Rabin 1999).

quantified and verified computationally. I omit heterogeneity by including only cars and not the three other vehicle types; I omit the endogenous lifetime of vehicles by fixing the lifetime at 18 years (about twice the mean car age in the data). As in the theoretical model, consumers choose in the initial period what vehicle to purchase and then have no option to scrap or replace it.²²

Table 2 presents summary statistics from this first specification under various policy alternatives. The first column presents statistics from the first-best outcome, which occurs when the agent is time-consistent and the gasoline tax equals marginal external damages. The row for "policy instrument" is not relevant to the first-best outcome. The first statistic is the deadweight loss of the policy, equal to the discounted value of true utility evaluated at the solution minus the value for the first-best outcome. The units of deadweight loss are dollars, and the values are per vehicle. The remaining statistics are the mean annual mileage, the mean annual gas consumption, and the mean fuel economy in gallons per mile.

The first alternate policy simulation is presented in column 2 of Table 2. In that simulation (as in all simulations subsequent to column 1) the agent exhibits present bias. The gasoline tax is set at the Pigouvian level of marginal external damages (\$1.23/gallon). Comparing deadweight loss in column 2 to column 1, present bias reduces the level of true utility, making the agent worse off. The value of deadweight loss is \$226 per vehicle, which is the discounted sum of deadweight loss over the vehicle's entire lifetime (18 years). The total number of new passenger vehicle sales in the US in 2009 was 8.9 million,²³ so the total annual deadweight loss from the new vehicles purchased in that year is about \$2.01 billion. The average mileage is lower, and the average *gpm* is higher, under the Pigouvian tax than in the first best. This, and the fact that the outcome under the Pigouvian tax does not achieve the first best, confirm Proposition 1.

The Pigouvian tax does not bring about the first-best outcome, and Proposition 1 claims that no gasoline tax, without any other policy instrument, can bring about the first-best outcome. But what is the best outcome that can be obtained with only a gasoline tax; that is, what is the second-best gasoline tax? I solve for the second-best gasoline tax computationally. For each value of a gasoline tax, I calculate the value of true utility, and then I choose the tax rate that

²² Under this specification, because of the assumption of a fixed vehicle lifetime, the model can be solved analytically rather than through backwards induction.

²³ Available from the Bureau of Transportation Statistics:
http://www.bts.gov/publications/national_transportation_statistics/html/table_01_17.html.

maximizes true utility. The resulting value of the second-best gasoline tax, shown in column 3, is \$1.60, 30% higher than marginal external damages. This suggests that ignoring present bias in preferences can lead to policy prescriptions that are significantly different than optimal levels. Compare outcomes in column 3 to those in column 2. Under the higher, second-best tax of column 3, mileage is lower, gas consumption is lower, and the average fuel economy is higher (*gpm* is lower). Deadweight loss is lower with the second-best tax, but still not zero (the first best is not achieved). This verifies Proposition 1.

Columns 4 and 5 model policies that achieve the first best using two instruments. Proposition 2 shows that the first best is achieved with the Pigouvian gasoline tax and a fuel economy standard requiring a maximum *gpm* equal to its optimal level. Column 4 enacts these policies, and the outcomes are identical to those in column 1. Proposition 3 shows that the first best is achieved with the Pigouvian gasoline tax and a tax on *gpm*. The optimal *gpm* tax equals 148440. This tax rate is multiplied by a vehicle's *gpm*. For the average non-hybrid car, with a *mpg* of 23.11, this tax payment is \$6423. For the average hybrid car, with a *mpg* of 35.18, this tax payment is \$4219. Thus the optimal fuel economy tax increases the price of a non-hybrid relative to a hybrid by \$2200. Before the fuel economy tax, the relative price difference between the two cars is \$5500.

Lastly, column 6 considers a different second-best policy that, like in column 3, has only one instrument. In column 6, the sole instrument is a fuel economy standard; the gasoline tax rate is set to zero. As Proposition 4 predicts, the first best is not achieved. In fact, without any gas tax, this policy is the worst of all policies, based on deadweight loss. Although the fuel economy is actually higher (lower *gpm*) than the first best, mileage and gas consumption are both much higher, since gasoline is so much cheaper.

All of the results in Table 2 perfectly corroborate the propositions of the basic theoretical model with one representative agent. This is unsurprising, since the specification in Table 2 is no different from that in the theoretical model. Next, I expand on the theoretical model by considering heterogeneity in consumers and vehicle types. Table 3 presents results from a specification including the four consumer and vehicle types described earlier. The assumption of a fixed vehicle lifetime (of 18 years) is maintained. Each summary statistic presented in the table is an average of the statistic for the four consumer types, weighted by the market share of the vehicle types from the NHTS.

Comparing columns 1 and 2 again verifies the predictions of Proposition 1. The Pigouvian tax fails to achieve the first-best outcome and instead leads to vehicles being too fuel-inefficient and mileage too low compared to the first best. The second-best gasoline tax is solved for in this specification, and results are presented in column 3. The second-best gasoline tax rate is again higher than marginal external damages, here by 18%.

As in Table 2, in Table 3, columns 4 and 5 compare two sets of policies that each have two policy instruments. Column 4 considers a gasoline tax combined with a fuel economy standard, and column 5 considers a gasoline tax combined with a fuel economy tax. Note that all policies are uniform across all consumer types. In the homogeneous specification in Table 2, each such set of policies attained the first best, as verified by Propositions 2 and 3. However, in the specification with heterogeneity, Propositions 5 and 6 show that no such set of policies attains the first best, so long as the fuel economy tax rate and the fuel economy standard are uniform across consumers. The policies in columns 4 and 5 are thus both second-best policies. The value of both policy instruments in each column is found numerically by maximizing true utility. Note that the level of the gasoline tax is not equal to marginal external damages in either column. In column 4, the gasoline tax is about equal to the second-best gasoline tax in column 3. In column 5, the gasoline tax is just slightly greater than the Pigouvian tax. Although both policies in columns 4 and 5 are second-best policies, note that the outcomes under the gasoline tax and fuel economy tax are remarkably closer to the first-best outcomes than are the outcomes under the gasoline tax and the fuel economy standard. In fact, from Table 3 it appears that the policy in column 5 achieves the first best. However, this is due to rounding; the deadweight loss is about 25 cents.

The theoretical model was only able to show that neither policy achieved the first best. The numerical simulation shows that the policy that includes a fuel economy tax dominates the policy that includes a fuel economy standard. Although the policy that includes a fuel economy tax does not achieve the first best, it comes remarkably close. Why does this policy dominate the policy with a fuel economy standard? Note that a uniform gasoline tax (equal to marginal external damages) plus a fuel economy standard unique to each consumer type would achieve the first best, as would the same uniform gasoline tax plus a fuel economy tax unique to each consumer type. The non-uniform policy values can be found from the theoretical model. The optimal fuel economy standard varies across consumer types from a minimum of .0335 gpm to a

maximum of .0558 gpm, a range that equals 67% of the minimum value. By contrast, the optimal fuel economy tax varies across consumer types from a minimum of 148440 dollars per gpm to a maximum of 171870 dollars per gpm, a range that equals just 16% of the minimum value. By this measure, the heterogeneity in the optimal standard exceeds the heterogeneity in the optimal tax. Therefore, a uniform standard gets it "more wrong" than does a uniform tax. This ranking is dependent on the calibration and should not be expected to hold for any parameter values, in contrast to the well-known result that Pigouvian taxes dominate uniform command-and-control standards under heterogeneity without time inconsistency.

The last column in Table 3 presents the policy simulation that finds the second-best fuel economy standard when the gasoline tax is fixed at zero. As in Table 2, this policy is the worst of all presented. Notice that the second-best fuel economy standard ($gpm_{max} = .0558$) has the same value as the second-best standard when the gasoline tax is allowed to be non-zero (column 4). In fact, this value is also almost equal to the first-best fuel economy for the most fuel-inefficient vehicle type, the pickup truck. Because the cost of higher fuel economy for pickups is so high, any standard that forces those consumers to increase fuel economy creates costs that outweigh any benefits. Another reason why the fuel economy tax of column 5 does so much better than the fuel economy standard of column 4 is that with the standard there is not much "room to move."

The last set of results using the preferred parameter values, presented in Table 4, adds in the endogenous choice of scrapping, and thus uses the backward induction solution method described earlier. It can be shown numerically as a verification that without present bias ($\beta = 1$), true utility is maximized with a gasoline tax equal to marginal external damages. Table 4 presents the same summary statistics as the earlier tables and adds a row displaying the mean age of vehicles. In the earlier specifications, this statistic was invariant to policy because of the absence of a scrapping choice. In Table 4, the deadweight loss values are the discounted values over an individual's lifetime (age 21 years to age 90 years), rather than a vehicle's lifetime.

In column 2, the average car age is about one quarter of a year older, the average annual mileage is 125 miles lower, and the average annual gasoline consumption is about 10 gallons higher, compared to column 1. All of these changes decrease utility, with a per-person discounted deadweight loss of \$4200 over the person's lifetime (this is the weighted average deadweight loss over the four consumer types). The mean fuel economy is higher in gpm ,

meaning that the average car is less fuel efficient. These results conform to the theoretical predictions of Proposition 1: the chosen level of gpm is higher than the optimal level and the chosen level of m is lower than the optimal level. The level of gasoline consumption is higher under column 2; this holds because gasoline demand is price inelastic.

The second-best gasoline tax, shown in column 3, is \$1.44. This is about 17% higher than marginal external damages. Deadweight loss decreases from column 2 to column 3. The average car age decreases, but not quite to the optimal level of 8.59 years. Annual mileage actually decreases, pushing it farther away from the optimal level than it was under the Pigouvian tax. Gasoline consumption is less than optimal. Finally, cars are more fuel efficient than under the Pigouvian policy, but not at the optimal level of fuel efficiency.

The last two columns of Table 4 examine policies that combine a gasoline tax with a policy on fuel economy in period zero. Column 4 presents the results for the best combination of a gasoline tax and a uniform fuel economy standard, while in column 5 the policy combines a gas tax with a fuel economy tax. In neither column is the first best achieved. The optimal level of the fuel economy tax in column 5 is \$50663 times the car's gpm . For the average non-hybrid car, with a fuel economy of 23.1 miles per gallon, this equals \$2193. For the average hybrid car, with a fuel economy of 35.2 miles per gallon, this equals \$1440. The optimal fuel economy tax thus makes the average non-hybrid car about \$750 costlier relative to the average hybrid car.

Comparing deadweight loss across policies, it can be seen that none of these policies achieves the first best. The policy option with the lowest deadweight loss is the one from column 5, with a tax on gasoline and a tax on fuel economy. Both policy options that contain two separate policies (columns 4 and 5) achieve lower deadweight loss than either policy option that contains just one policy (columns 2 and 3). As with the results in Table 3, combining a gasoline tax with a fuel economy tax is closer to the first best than combining a gasoline tax with a fuel economy standard. However, in Table 4 the policy combining a gasoline tax and a fuel economy tax results in a substantially different outcome than the first best.

Finally, in Table 5 I examine how sensitive these results are to the calibrated parameter values. I investigate five alternative parameter values. For each alternative parameter value, I present the summary statistics from the first best and from the Pigouvian tax policy. All simulations assume the fixed vehicle lifetime, heterogeneous agents specification as in Table 3.

Each pair of columns in Table 5 is analogous to columns 1 and 2 of Table 3. Deadweight loss in Table 5 is the total discounted value over a vehicle's lifetime.

The first two columns of Table 5 consider less present bias in preferences, i.e. a β closer to 1. (Clearly, when $\beta = 1$, all of the standard results hold and the Pigouvian solution equals the first best.) Unsurprisingly, with a β closer to 1, the Pigouvian outcome is closer to the first-best outcome than when $\beta = 0.7$ as in the base case. When $\beta = 0.95$, the deadweight loss from employing only a Pigouvian tax is just \$5 per vehicle (\$45 million annual economy-wide deadweight loss, based on the number of new vehicles sold in the US).

I next investigate φ , the negative inverse of the short-run price elasticity of gasoline demand. This is a short-run elasticity since it does not account for the agent changing his automobile choice in response to changes in gas prices. The long-run elasticity, which allows for choice of fuel economy, can be calculated from this model.²⁴ In the base case where the short-run price elasticity is -0.34 , the long-run price elasticity is -0.40 . This is lower than the preferred estimate of -0.84 from Brons et. al. (2008), but about equal to the median value found in the meta-analysis in Espey (1998) of -0.43 .

In columns 3-6 of Table 5, I alter φ and simulate the outcomes. I also change the value of C , the constant in the utility function, to match the mean mileage in the data for each of the four vehicle types. When $\varphi = 1$, gasoline demand is unitarily elastic. As suggested in earlier discussion, in this instance present bias does not change gasoline consumption: the effects from a lower fuel economy and from lower mileage just offset. Furthermore, when $\varphi = 0.9$, gasoline demand is elastic, so present bias reduces gasoline consumption. Although not shown in Table 5, the second-best gasoline tax in this instance is lower than the Pigouvian tax.

Columns 7 and 8 are from simulations where the marginal external damages from emissions are lower than in the base case. In particular, I use the same source for estimates of damages (Parry 2009), but include only those externalities based directly on gasoline consumption (climate change, other pollutants, oil dependence) and not those based instead on mileage (accidents and congestion). The marginal damages are 31 cents per gallons. This reduces the deadweight loss from not considering the behavioral failure to just over \$100 per

²⁴ The implicit function theorem can be performed on the agent's first-order condition for the choice of fuel economy to find $dgpm/dgas$, which can be plugged into the formula for price elasticity of demand. While the utility function is defined such that the short-run elasticity is constant at $-1/\varphi$, the long-run elasticity is a function the optimal mileage, fuel economy, the cost function c , as well as the short-run elasticity.

vehicle. Optimal mileage and gasoline consumption is of course higher, since marginal damages are lower.

Lastly, columns 9 and 10 consider an alternate specification of the cost functions of fuel economy c . In the base case results, these functions (there is a unique function for each vehicle type) are based on comparing average costs of hybrid and non-hybrid vehicles, fit to a quadratic. But this does not pin down a function. Thus, the simulations in columns 9 and 10 use the same two points and fit a decreasing function through them that is less convex than is the function specified in the base case. As a result, the deadweight loss from using the Pigouvian tax is more than twice the deadweight loss with the base case cost function, suggesting that on that dimension the base case results are conservative.

V. Conclusion

Growing support is arising from the field of behavioral economics for the claim that consumers regularly exhibit time-inconsistent preferences and make decisions under a present bias. Little is known about how this phenomenon impacts optimal policy design or interacts with market failures. This paper examines how policies addressing externalities perform under time-inconsistent preferences. The paper's theoretical model finds that policymakers need to acknowledge time inconsistency to achieve socially optimal outcomes. The numerical simulation suggests that ignoring time inconsistency can yield policy prescriptions that substantially differ from those that would bring about the first-best outcome.

The intuitive results from the model come from a particular and perhaps controversial specification of welfare-maximization: the long-run criterion. In order to be agnostic about what welfare criterion ought to be used, I compare the results from the long-run criterion to those under alternative welfare criteria, and I find conditions under which the main results are robust to the alternatives. The numerical simulations satisfy these robustness conditions.

The simulation results are based on several assumptions, any of which could be relaxed to provide even more sensitivity analyses. For instance, producer behavior is not modeled; manufacturers also respond to price policies, and this response could affect market prices and quantities. More heterogeneity could be added in many places: more types of consumers, regional or temporal variance in gasoline prices or in external gasoline or mileage damages, more types and features of vehicles. Any of these extensions would no doubt capture more features of the market. But, the purpose of this simulation is not to pin down optimal policy

point estimates, but rather to provide an idea of the magnitude of the effects of time-inconsistency on policy prescriptions.

The theory provided a rather specific example of a market failure: a durable good that creates externalities. The simulation was even more specific: automobiles. The theoretical model is applicable to other externality-producing durable goods, like home appliances home energy efficiency investments. Furthermore, the framework here may be applicable elsewhere. For example, time inconsistency is often attributed as relevant to the rise in obesity (Ruhm 2010). The framework developed in this paper could be used to analyze policy options like taxes on unhealthy foods, limitations on the availability of certain foods, or subsidies to gym memberships. These results could similarly be extended with applications to retirement savings or addictive behavior.

References

- Allcott, Hunt, and Nathan Wozny. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Working Paper*, November 2010.
- Allcott, Hunt, and Sendhil Mullainathan. "Behavior and Energy Policy." *Science* 327 (March 2010): 1204-1205.
- Andreoni, James, and Charles Sprenger. "Estimating Time Preferences from Convex Budgets." *NBER Working Paper #16347*, September 2010.
- Ariely, Dan, and Klaus Wertenbroch. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Science* 13, no. 3 (May 2002): 219-224.
- Bernheim, B. Douglas, and Antonio Rangel. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124, no. 1 (February 2009): 51-104.
- Bhattacharya, Jay, and Darius Lakdawalla. "Time-Inconsistency and Welfare." *NBER Working Paper #10345*, March 2004.
- Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld. "A meta-analysis of the price elasticity of gasoline demand. A SUR approach." *Energy Economics* 30, no. 5 (September 2008): 2105-2122.
- Brown, Alexander, Zhikang Eric Chua, and Colin Camerer. "Learning and Visceral Temptation in Dynamic Saving Experiments." *Quarterly Journal of Economics* 124, no. 1 (February 2009): 197-231.
- Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. "Pain at the Pump: The Differential Effect of Gasoline Prices on New and Used Automobile Markets." *NBER Working Paper #15590*, December 2009.
- Caplin, Andrew, and John Leahy. "The Social Discount Rate." *Journal of Political Economy* 112, no. 6 (December 2004): 1257-1268.

- Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen Palmer. "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy* 108, no. 6 (December 2000): 1292-1326.
- Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics* 124, no. 4 (November 2009): 1639-1674.
- Congdon, William, Jeffrey Kling, and Sendhil Mullainathan. "Behavioral Economics and Tax Policy." *NBER Working Paper 15328*, 2009.
- Costa, Dora, and Matthew Kahn. "Energy Conservation "Nudges" and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment." *NBER Working Paper 15939*, 2010.
- DellaVigna, Stefano. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47, no. 2 (June 2009): 315-372.
- Dellavigna, Stefano, and Ulrike Malmendier. "Paying Not to Go to the Gym." *American Economic Review* 96, no. 3 (June 2006): 694-719.
- Espey, Molly. "Explaining the variation in elasticity estimates of gasoline demand in the United States: A meta-analysis." *Energy Journal* 17, no. 3 (1996): 49-60.
- Espey, Molly. "Gasoline Demand Revisited: An International Meta-analysis of Elasticities." *Energy Economics* 20, no. 3 (June 1998): 273-295.
- Fang, Hanming, and Dan Silverman. "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY." *International Economic Review* 50, no. 4 (November 2009): 1043-1077.
- Fischer, Carolyn, Winston Harrington, and Ian Parry. "Should Automobile Fuel Economy Standards be Tightened?" *The Energy Journal* 28, no. 4 (2007): 1-29.
- Gillingham, Kenneth, Richard Newell, and Karen Palmer. "Energy Efficiency Economics and Policy." *Annual Review of Resource Economics* 1 (October 2009): 597-619.
- Greene, David. "Why CAFE Worked." *Energy Policy* 26, no. 8 (July 1998): 595-613.
- Greene, David. "Why the Market for New Passenger Cars Generally Undervalues Fuel Economy." *Joint Transport Research Centre Discussion Paper No. 2010-6*, January 2010.
- Gruber, Jonathan, and Botond Koszegi. "Is Addiction "Rational"? Theory and Evidence." *The Quarterly Journal of Economics* 116, no. 4 (November 2001): 1261-1303.
- Hassett, Kevin, and Gilbert Metcalf. "Energy Tax Credits and Residential Conservation Investment: Evidence from Panel Data." *Journal of Public Economics* 57, no. 2 (June 1995): 210-217.
- Hastings, Justine, and Olivia Mitchell. "How Financial Literacy and Impatience Shape Retirement Wealth and Investment Behaviors." *NBER Working Paper #16740*, January 2011.
- Hausman, Jerry. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics* 10, no. 1 (Spring 1979): 33-54.

- Hughes, Jonathan, Christopher Knittel, and Daniel Sperling. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29, no. 1 (2008): 113-134.
- Johansson, Olof. "Optimal Pigovian Taxes under Altruism." *Land Economics* 73, no. 3 (August 1997): 297-308.
- Krusell, Per, Burhanettin Kuruscu, and Anthony Smith. "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting." *Journal of Economic Theory* 105, no. 1 (July 2002): 42-72.
- Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112, no. 2 (May 1997): 443-477.
- Laibson, David, Andrea Repetto, and Jeremy Tobachman. "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity*, no. 1 (1998): 91-172.
- Laibson, David, Andrea Repetto, and Jeremy Tobachman. "Estimating Discount Functions with Consumption Choices over the Lifecycle." *NBER Working Paper #13314*, 2007.
- Li, Shanjun, Christopher Timmins, and Roger Von Haefen. "How Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy* 1, no. 2 (August 2009): 113-137.
- Lofgren, Asa. "The Effect of Addiction on Environmental Taxation in a First and Second-Best World." *Working Papers in Economics no. 91, Department of Economics, Goteborg University*, 2003.
- Mastrobuoni, Giovanni, and Matthew Weinberg. "Heterogeneity in Intra-monthly Consumption Patterns, Self-Control, and Savings at Retirement." *American Economic Journal: Economic Policy* 1, no. 2 (August 2009): 163-189.
- Mauzerall, Denise, Babar Sultan, Namsoung Kim, and David Bradford. "NO_x Emissions from Large Point Sources: Variability in Ozone Production, Resulting Health Damages and Economic Costs." *Atmospheric Environment* 39, no. 16 (May 2005): 2851-2866.
- Metcalf, Gilbert, and Kevin Hassett. "Measuring the Energy Savings from Home Improvement Investments: Evidence from Monthly Billing Data." *Review of Economics and Statistics* 81, no. 3 (August 1999): 516-528.
- Mullainathan, Sendhil, and Richard Thaler. "Behavioral Economics." In *International Encyclopedia of the Social and Behavioral Sciences*, edited by Neil Smelser and Paul Bates. 2001.
- O'Donoghue, Ted, and Matthew Rabin. "Doing It Now or Later." *American Economic Review* 89, no. 1 (March 1999): 103-124.
- O'Donoghue, Ted, and Matthew Rabin. "Optimal Sin Taxes." *Journal of Public Economics* 90, no. 10-11 (November 2006): 1825-1849.
- Parry, Ian. "How Much Should Highway Fuels be Taxed?" *RFF Discussion Paper 09-52*, December 2009.
- Robson, Arthur, and Larry Samuelson. "The Evolution of Decision and Experienced Utilities." *Theoretical Economics*, forthcoming.
- Ruhm, Chris. "Understanding Overeating and Obesity." *NBER Working Paper #16149*, 2010.

- Sanstad, Alan, W. Michael Hanemann, and Maximillian Auffhammer. *End-use Energy Efficiency in a "Post-Carbon" California Economy: Policy Issues and Research Frontiers*. California Climate Change Center, UC Berkeley, 2006.
- Shogren, Jason, and Laura Taylor. "On Behavioral-Environmental Economics." *Review of Environmental Economics and Policy* 2, no. 1 (Winter 2008): 26-44.
- Strotz, R. H. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23, no. 3 (1955-1956): 165-180.
- Thaler, Richard. "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters* 8, no. 3 (1981): 201-207.
- Viscusi, Kip, Joel Huber, and Jason Bell. "Estimating Discount Rates for Environmental Quality from Utility-Based Choice Experiments." *Journal of Risk and Uncertainty* 37, no. 2-3 (December 2008): 199-220.

Table 1: Calibration of Numerical Model			
Function or Parameter	Description	Value	Source
φ	Inverse of short-run price elasticity of gasoline demand	2.941	(Brons, et al. 2008)
C	Utility function scale parameter	Varies by vehicle type	Calibrated from mileage data in NHTS
$c(gpm)$	Cost of car as function of fuel economy	Varies by vehicle type	Calibrated from EPA fuel economy data and MSRP data.
d	External damages from gasoline consumption	\$1.23/gallon	(Parry 2009)
β	Present bias discount factor (annual)	0.7	(Laibson, Repetto and Tobacman, Estimating Discount Functions with Consumption Choices over the Lifecycle 2007)
δ	Long-run discount factor (annual)	0.96	
D	Consumer disutility from vehicle age	Varies by vehicle type	Calibrated to match average vehicle age in NHTS

	(1)	(2)	(3)	(4)	(5)	(6)
	First-best	Pigouvian gasoline tax	Second-best gasoline tax	Optimal gasoline tax and fuel economy standard	Optimal gasoline tax and fuel economy tax	Second-best fuel economy standard
Policy instrument(s)	N/A	$\tau = 1.23$ $\tau_{gpm} = 0$	$\tau = 1.60$ $\tau_{gpm} = 0$	$\tau = 1.23$ $gpm_{max} = .0335$	$\tau = 1.23$ $\tau_{gpm} = 148440$	$\tau = 0$ $gpm_{max} = .0333$
Deadweight Loss (\$)	0	226	192	0	0	525
Mean mileage	11417	11070	10743	11417	11417	13309
Mean gas consumption (gallons)	382.49	405.61	388.46	382.49	382.49	442.82
Mean fuel economy (gpm)	.0335	.0366	.0362	.0335	.0335	.0333

Notes: Deadweight loss is the total discounted value, per new car, over the lifetime of the car ($T = 18$ years). Gasoline taxes τ are in dollars per gallon.

Table 3: Summary Statistics from Simulation with Heterogeneous Consumers, fixed Vehicle Lifetime						
	(1)	(2)	(3)	(4)	(5)	(6)
	First-best	Pigouvian gasoline tax	Second-best gasoline tax	Second-best gasoline tax and fuel economy standard	Second-best gasoline tax and fuel economy tax	Second-best fuel economy standard
Policy instrument(s)	N/A	$\tau = 1.23$ $\tau_{gpm} = 0$	$\tau = 1.4454$ $\tau_{gpm} = 0$	$\tau = 1.4315$ $gpm_{max} = .0558$	$\tau = 1.2304$ $\tau_{gpm} = 150421$	$\tau = 0$ $gpm_{max} = .0558$
Deadweight Loss (\$)	0	162	147	138	0	1071
Mean mileage	11800	11565	11348	11371	11800	13329
Mean gas consumption (gallons)	489.30	505.88	494.01	493.70	489.30	596.70
Mean fuel economy (gpm)	.0412	.0434	.0432	.0431	.0412	.0444

Notes: Deadweight loss is the total discounted value, per new car, over the lifetime of the car ($T = 18$ years), averaged over the four vehicle types, weighted by their market shares. Gasoline taxes τ are in dollars per gallon.

Table 4: Summary Statistics from Numerical Simulations					
	(1)	(2)	(3)	(4)	(5)
	First-best	Pigouvian gasoline tax	Second-best gasoline tax	Second-best gasoline tax and fuel economy standard	Second-best gasoline tax and fuel economy tax
Policy instrument(s)	N/A	$\tau = 1.23$	$\tau = 1.44$	$\tau = 1.26,$ $gpm_{max} = .0569$	$\tau = 1.29,$ $\tau_{gpm} = 50663$
Deadweight loss (\$)	0	4202.7	3850.4	3798.5	3653.8
Mean vehicle age (years)	8.5941	8.8155	8.7870	8.8154	9.1228
Mean mileage	11654	11529	11322	11528	11576
Mean gas consumption (gallons)	499.89	509.33	497.28	505.38	497.82
Mean fuel economy (gpm)	0.0426	0.0439	0.0436	0.0436	0.0428

Notes: Deadweight loss is the total discounted value, per consumer, over his or her entire lifetime, averaged over the four consumer types, weighted by their market shares. Gasoline taxes τ are in dollars per gallon.

Table 5: Sensitivity Analysis										
	$\beta = 0.95$		$\varphi = 1$		$\varphi = 0.9$		$d = 0.31$		c less convex	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	First-best	Pigouvian gasoline tax	First-best	Pigouvian gasoline tax	First-best	Pigouvian gasoline tax	First-best	Pigouvian gasoline tax	First-best	Pigouvian gasoline tax
Deadweight Loss (\$)	0	5	0	150	0	150	0	104	0	348
Mean mileage	11800	11758	10782	10150	10598	9909	12958	12762	11894	11362
Mean gas consumption (gallons)	489.30	492.16	448.80	448.80	441.65	439.07	555.56	569.98	488.80	525.31
Mean fuel economy (gpm)	.0412	.0416	.0417	.0440	.0418	.0441	.0426	.0443	.0411	.0459

Notes: Deadweight loss is the total discounted value, per new car, over the lifetime of the car ($T = 18$ years), averaged over the four vehicle types, weighted by their market shares.

Appendix

Proofs of Propositions

Proposition 1: To prove the first statement, suppose the contradiction: there exists a set of tax rates $\{\tau_t^{opt\beta}\}$ that lead to $m_t^* = m_t^{opt}$ for all $t > 0$ and $gpm^* = gpm^{opt}$. The consumer's first-order condition for choice of m_t in period t is $U'(m_t^*) = (gas_t + \tau_t^{opt\beta}) \cdot gpm^*$, or $U'(m_t^{opt}) = (gas_t + \tau_t^{opt\beta}) \cdot gpm^{opt}$. Since U'' is strictly negative, the last equation can only be satisfied when $\tau_t^{opt\beta} = \tau_t^{pig}$. When $\beta = 1$, then τ_t^{pig} necessarily induces the optimal solution. When $\beta < 1$, the optimal solution does not change, since the planner does not consider the quasi-hyperbolic discount factor β . But, the consumer's decision and first-order condition differ from the $\beta = 1$ case. Thus, it does not equal the planner's solution.

To prove the second statement, note that the consumer's choice of gpm^* is given by her first-order condition; call this equation F . The implicit function theorem can be used to show how gpm^* varies with β

$$\frac{dgpm^*}{d\beta} = \frac{-dF/d\beta}{dF/dgpm} = \frac{\sum_{t=1}^T \delta^t [(gas_t + \tau_t^{pig}) \cdot m_t^*]}{-c''(gpm) + \beta \sum_{t=1}^T \delta^t [-(gas_t + \tau_t^{pig}) \cdot m_t^*]}$$

The numerator is positive. The denominator is negative from the second-order condition of the consumer's optimization problem. Thus, $dgpm^*/d\beta < 0$. Since $gpm^* = gpm^{opt}$ when $\beta = 1$, it follows that $gpm^* > gpm^{opt}$ when $\beta < 1$.

The consumer's choice of m_t^* in each period is a function of the total price of a mile of driving, $(gas_t + \tau_t^{pig}) \cdot gpm$, from the first-order condition $U'(m_t^*) = (gas_t + \tau_t^{pig}) \cdot gpm$. Since $U'' < 0$, $dm_t^*/dgpm < 0$. When $gpm = gpm^{opt}$, $m_t^* = m_t^{opt}$. But when $\beta < 1$, $gpm^* > gpm^{opt}$, so $m_t^* < m_t^{opt}$ at each period $t > 0$.

Proposition 2: Consider the consumer's problem with the added constraint that $gpm \leq gpm^{opt}$.

$$\max_{gpm, \{m_t\}_{t=1}^T} -c(gpm) + \beta \sum_{t=1}^T \delta^t [U(m_t) - (gas_t + \tau_t^{pig}) \cdot gpm \cdot m_t]$$

Subject to

$$\begin{aligned} U'(m_t) - (gas_t + \tau_t^{pig}) \cdot gpm &= 0 \forall t > 0 \\ gpm &\leq gpm^{opt} \end{aligned}$$

Consider this problem's Lagrangian, where the constraint from the period t choice of m_t has a multiplier λ_t and the inequality constraint on gpm has a multiplier μ . The first-order condition with respect to m_t is

$$\beta \delta^t [U'(m_t^*) - (gas_t + \tau_t^{pig}) \cdot gpm^*] + \lambda_t \cdot U''(m_t^*) = 0$$

The term in brackets is zero from the first-order condition from the static choice of m_t . Since U'' is strictly negative, $\lambda_t = 0$ for all $t > 0$. Then, the first-order condition for gpm is

$$-c'(gpm^*) + \beta \sum_{t=1}^T \delta^t [-(gas_t + \tau_t^{pig}) \cdot m_t^*] - \mu = 0$$

If $\mu = 0$, then this condition mimics the consumer's first-order condition in the problem without the fuel economy constraint. Proposition 1 shows that in that case $gpm^* > gpm^{opt}$. This violates the constraint in this problem. Hence, $\mu > 0$ and $gpm^* = gpm^{opt}$. Then, from the first-order condition for each decision over m_t , each period's choice over m_t results in m_t^{opt} .

Proposition 3: With a tax on gpm , the consumer's problem is

$$\max_{gpm, \{m_t\}_{t=1}^T} -c(gpm) - gpm \cdot \tau_{gpm} + \beta \sum_{t=1}^T \delta^t [U(m_t) - (gas_t + \tau_t^{pig}) \cdot gpm \cdot m_t]$$

Subject to

$$U'(m_t) - (gas_t + \tau_t^{pig}) \cdot gpm = 0 \forall t > 0$$

From the same argument as in the proof to Proposition 2, the Lagrangian multiplier λ_t on the period t constraint equals zero, so the first-order condition over gpm is

$$-c'(gpm^*) - \tau_{gpm} + \beta \sum_{t=1}^T \delta^t [-(gas_t + \tau_t^{pig}) \cdot m_t^*] = 0$$

With the value of τ_{gpm} as given, this first-order condition can be written as

$$-c'(gpm^*) + \sum_{t=1}^T \delta^t [-(gas_t + \tau_t^{pig}) \cdot (\beta \cdot m_t^* + (1 - \beta) \cdot m_t^{opt})] = 0$$

When $gpm^* = gpm^{opt}$, then $m_t^* = m_t^{opt}$ for all $t > 0$ since $\tau_t = \tau_t^{pig}$. Then, plugging $m_t^* = m_t^{opt}$ in the first-order condition above makes it equal to the planner's first-order condition, and thus gpm^{opt} is a solution by definition. So gpm^{opt} and m_t^{opt} solve the consumer's problem, and by the second-order condition this is a unique solution.

Proposition 4: Suppose that some policy exists that induces the first-best outcomes gpm^{opt} and m_t^{opt} . The consumer's first-order condition for the choice of m_t must be satisfied: $U'(m_t) = (gas_t + \tau_t) \cdot gpm$. At the first best this is only satisfied with $\tau_t = \tau_t^{pig} \neq 0$.

Proposition 5: Suppose that such a set of policies exist that induce the first best. Consumer i 's first order conditions are

$$U'_i(m_{t,i}^*) - (gas_t + \tau_t) \cdot gpm_i^* = 0$$

$$-c'(gpm_i^*) + \beta \sum_{t=1}^T \delta^t \cdot [-(gas_t + \tau_t) \cdot m_{t,i}^*] = 0$$

Under the supposition, $m_{t,i}^* = m_{t,i}^{opt}$ and $gpm_i^* = gpm_i^{opt}$. Comparing the consumer's condition for $m_{t,i}$ with the planner's, it must be true that $\tau_t = \tau_t^{pig}$ in each period t . Then, subtracting consumer i 's condition for gpm_i from the planner's condition for gpm_i gives

$$\tau_{gpm} = (1 - \beta) \sum_{t=1}^T \delta^t \cdot (gas_t + \tau_t^{pig}) \cdot m_{t,i}^{opt}$$

This right hand side is not independent of i , except in the special case where $m_{t,1}^{opt} = m_{t,2}^{opt}$ for each t . Thus, τ_{gpm} is not uniform, a contradiction.

Proposition 6: Suppose that such a set of policies exists that induces the first best. Consumer i 's problem is to maximize utility subject to the emissions tax τ_t and the restriction on gpm . The first-order conditions, where the inequality constraint's multiplier is μ , is

$$-c'(gpm_i^*) + \beta \sum_{t=1}^T \delta^t \cdot [-(gas_t + \tau_t) \cdot m_{t,i}^*] - \mu = 0$$

$$U'_i(m_{t,i}^*) - (gas_t + \tau_t) \cdot gpm_i^* = 0$$

Comparing the second equation, the consumer's first-order condition for her choice of $m_{t,i}$, with the planner's equivalent equation, it follows that $\tau_t = \tau_t^{pig}$, as in the last proof. Then, subtracting the consumer's first-order condition for choice of gpm_i with the planner's condition yields

$$\mu = (1 - \beta) \sum_{t=1}^T \delta^t \cdot (gas_t + \tau_t^{pig}) \cdot m_{t,i}^{opt}$$

That is, the shadow value of the constraint equals the price that would induce the first best (τ_{gpm} from the last proof). This is strictly positive, so the constraint binds, and $gpm_i^* = gpm^{max}$. However, this value is not independent of i , and thus the optimal policy value for gpm^{max} cannot be uniform.

Proposition 7: I will show that the optimal policy Pareto improves the Pigouvian policy by setting the lump sum payments s_t such that each self in period $t > 0$ is just as well off in the optimal policy as in the Pigouvian policy, and (if the condition holds) the period zero self is strictly better off. In the Pigouvian policy, the single-period utility for any period $t > 0$ is $u(m_t^0) - gas_t \cdot gpm^0 \cdot m_t^0 - d(gpm^0 \cdot m_t^0)$. The tax payment and lump sum payment just cancel each other out. In the optimal policy, the single-period utility for any period $t > 0$ is $u(m_t^{opt}) - (gas_t + \tau_t^{pig}) \cdot gpm^{opt} \cdot m_t^{opt} + s_t - d(gpm^{opt} \cdot m_t^{opt})$. Choose lump sum payments s_t in these periods such that each single-period utility value is equal under both policies. This implies $s_t = u(m_t^0) - u(m_t^{opt}) - (d(gpm^0 \cdot m_t^0) - d(gpm^{opt} \cdot m_t^{opt})) - gas_t \cdot (gpm^0 \cdot m_t^0 - gpm^{opt} \cdot m_t^{opt}) + \tau_t^{pig} \cdot gpm^{opt} \cdot m_t^{opt}$. Then, the government's budget constraint defines the zero period lump sum payment s_0 : $s_0 = gpm^{opt} \cdot \tau_{gpm} + \sum_{t=1}^T gpm^{opt} \cdot \tau_t^{pig} \cdot m_t^{opt} - \sum_{t=1}^T s_t$. The single-period utility in the zero period under the Pigouvian policy is $-c(gpm^0)$, and under the optimal policy it is $-c(gpm^{opt}) + s_0 - gpm^{opt} \cdot \tau_{gpm}$. For the single-period utility in period zero to be strictly higher under the optimal policy than under the Pigouvian policy, it must be that $s_0 > -c(gpm^0) + c(gpm^{opt}) + gpm^{opt} \cdot \tau_{gpm}$. Substituting in the expressions for s_0 and s_t from the expressions above yields the condition in the proposition.