

OPTIMAL INCOME TAXATION WITH PRESENT BIAS*

Benjamin B. Lockwood[†]

October 2017

Abstract

Work often entails up-front costs in exchange for delayed benefits, and mounting evidence documents present bias over effort in the face of such delays. This paper studies the implications of present bias for the optimal income tax schedule. Optimal tax rates are characterized for a model economy in which workers choose multiple dimensions of labor supply over time. Present bias reduces optimal tax rates, with a larger effect when the elasticity of taxable income is high. Optimal marginal tax rates may be negative at low incomes, providing an alternative, corrective rationale for the Earned Income Tax Credit. (JEL: D04, H21)

*I thank Raj Chetty, Nathaniel Hendren, David Laibson, and Matthew Weinzierl for extensive discussions, and Charles Nathanson, Emmanuel Farhi, Louis Kaplow, Wojciech Kopczuk, Dmitry Taubinsky, Glen Weyl, and seminar participants at Harvard, University of Pennsylvania, Duke, Georgetown, University of Oregon, University of Maryland, Cornell, the National Tax Association and the American Economic Association Annual Meetings for many helpful comments.

[†]University of Pennsylvania and NBER, ben.lockwood@wharton.upenn.edu

1 Introduction

Many countries' tax codes feature work subsidies (sometimes called an Earned Income Tax Credit), which impose negative marginal tax rates on low earners. The US EITC—the nation's largest cash transfer antipoverty program—enjoys strong bipartisan support. Policy makers laud these policies for encouraging work, often while stressing benefits that work confers, such as a sense of meaning, dignity, future advancement, and general happiness.¹

From the viewpoint of standard optimal taxation models, this rhetoric is puzzling. If work has substantial benefits, the logic goes, workers will take that into account when setting their labor supply, so there is no need for a government subsidy. Indeed, a classic result in the optimal taxation literature is that negative marginal tax rates are generally suboptimal (Mirrlees, 1971; Seade, 1982).²

This paper considers an alternative rationale for policies which encourage work: that workers do *not* fully account for the benefits of work when making labor supply decisions. The motivation for this approach is twofold. First, this corrective logic appears more consistent with the arguments policy makers employ to advocate the EITC—both in terms of the sign and magnitude of marginal tax rates, and in terms of comparative statics, such as widespread calls to expand the credit following new evidence that it generated large behavioral responses. Second, a rapidly growing body of empirical evidence points to a specific channel through which workers might underestimate work benefits: present bias over effort.³ In the labor force, work often entails up-front effort costs in exchange for delayed benefits. One must exert search effort before finding a new or better job, for example, and extra effort at work lowers the probability of a future layoff while raising the prospect

¹President Bill Clinton, advocating an EITC expansion before Congress, called the program a “path to independence and dignity.” Congressman Paul Ryan, who proposed an EITC expansion to childless workers, wrote “Truly helping those in need means recapturing a sense of the dignity of work. The trouble is our current welfare system isn't sufficiently designed to encourage work and the dignity it confers.” Senator Marco Rubio, arguing for a similar expansion for childless workers in the form of a direct wage subsidy, wrote “Work, and the sense of purpose and accomplishment that it brings, is essential to human happiness.”

²Saez (2002) argues that an extensive margin of labor supply choice might generate negative marginal tax rates at low incomes. However Jacquet et al. (2013) shows that this result stems from a discretized skill distribution, arguing that a more realistic continuous skill distribution actually generates a discrete workforce participation credit, with positive marginal tax rates at all positive incomes. In any case, the preferred specification in Saez (2002) calls for a modest *positive* marginal rate of 10% on the lowest positive earners (\$0–\$4,000), and much higher rates at higher incomes, suggesting the extensive margin rationale may not generate substantial work subsidies like those under the US EITC.

³For example, experimental subjects prefer to defer effort tasks when choosing between now and a future date, but are indifferent about effort allocation across similarly spaced future dates (Augenblick, Niederle and Sprenger, 2015). See Section 3.1 for a discussion of other empirical evidence.

of a raise or promotion. Present-biased workers may therefore exert less effort than their “long-run” selves desire, generating a corrective rationale for policies which encourage greater effort, such as work subsidies.⁴

The contributions of this paper are both theoretical and empirical. In the theory section, I generalize the benchmark model of optimal income taxation to allow for multiple dimensions of labor supply over time, representing effort along diverse channels such as work quality, on- and off-the-job training, and search effort, in addition to the usual choice of labor hours. I derive necessary and sufficient conditions for the optimal tax in a simplified setting, as well as a necessary (first-order) condition for the optimal income tax under more general conditions with multidimensional heterogeneity and a participation margin.

The theoretical results demonstrate that present bias tends to lower optimal marginal tax rates, consistent with the Pigouvian logic that uninternalized benefits rationalize a corrective subsidy. However the results also highlight a number of more nuanced implications, including two key points for policy design. First, a higher elasticity of taxable income magnifies the optimal corrective component of the subsidy. This contrasts with the standard Pigouvian result that the optimal correction is insensitive to elasticity and is additively separable from the optimal redistributive tax. Yet this finding is consistent with political rhetoric which cites high response elasticities when touting the success of EITC expansions.⁵ Second, initial intuition might suggest the delayed nature of the EITC (which is paid at the end of the tax year) must be highly suboptimal for a policy which combats present bias. Yet Section 2.5 clarifies that a delayed structure does not undermine the corrective force of the EITC; on the contrary, a modest delay is *optimal* under plausible conditions, as it reduces the upward distortion from corrective subsidies along dimensions of labor effort which are chosen at the time of compensation .

Empirically, I review a wide range of evidence to calibrate present bias across the income distribution, which appears to be heavily concentrated at low incomes. Using data on the income distribution and existing tax schedules, I calibrate the skill distribution accounting for present

⁴Although I focus on the channel of present bias in order to calibrate the model based on existing empirical evidence, other biases which generate an undersupply of labor effort yield qualitatively similar policy implications and an analogous optimal tax function.

⁵The Clinton Administration emphasized the increase in single mothers’ labor earnings after the 1993 EITC expansion as evidence of success and reason for further expansion (https://clinton4.nara.gov/WH/New/html/20000112_2.html).

bias, and I simulate the optimal income tax for a range of normative preferences. The quantitative results show that if redistributive motives are modest, the optimal tax system features negative marginal tax rates across a range of low incomes. I further perform an “inverse optimum” exercise, which backs out the marginal social welfare weights (the social value of marginal consumption for households with a given income) which are consistent with the existing EITC, with and without accounting for present bias. Absent present bias, the existing EITC implies a striking unconventional normative feature: welfare weights on the poorest workers are *lower* than at the middle of the income distribution. After accounting for present bias, however, implied welfare weights decline monotonically with income, consistent with conventional normative assumptions.

To preview the quantitative implications of present bias, Figure 1 plots the schedule of optimal marginal tax rates in the baseline model economy with and without accounting for present bias. Results are plotted for a modest degree of inequality aversion (see Section 3 for details), which tends to reduce marginal tax rates. Even in this case, marginal tax rates at low incomes are quite high under the rational optimum, generating a sharp divergence between the simulated optimum and actual marginal tax rates in the US, plotted here for a representative household with two children, including phaseouts of universal benefits.⁶ Stronger redistributive preferences exacerbate this divergence. However after accounting for present bias, optimal marginal tax rates are negative at low incomes, on par with those generated by the EITC.

As is evident from Figure 1, the optimal tax schedule with present bias still diverges from the existing tax schedule in important respects, particularly in the “phase-out” region of the EITC.⁷ Such high marginal tax rates in the phase-out region could be optimal if labor supply elasticities are low for that range of incomes (as suggested, for example, by Chetty et al. (2013)). More broadly, however, the goal of this paper is not to rationalize the precise tax schedule of marginal tax rates (or other policy features) of the US EITC. Instead, this paper is motivated by the widespread prevalence of the *qualitative* policy feature—negative marginal tax rates—combined with the observation that common rhetorical justifications for negative marginal tax rates are inconsistent with conventional

⁶US marginal tax rates are displayed for a single parent with two children residing in Colorado in 2015, computed using \$1000 intervals. Calculations include the phaseout of universally available benefits: SNAP, Medicaid, CHIP, and ACA premium assistance credits. Estimates were computed by Eugene Steuerle and Caleb Quakenbush for congressional testimony.

⁷The US EITC provides a marginal subsidy at low incomes—up to \$13,930 for a household with two children in 2016—followed by a plateau over which the credit is constant, followed by a phase-out region (between \$23,740 and \$50,198 in income), which contributes to the high marginal tax rates across that income range in Figure 1.

models of optimal taxation. If the existing EITC is best understood as a corrective policy, then by formalizing a model of misoptimization, we can hope to improve the policy on those terms, both through the level of marginal tax rates, and through other aspects of policy design.

Relation to the literature

This paper relates to a number of subfields in the optimal taxation literature. The first is optimal taxation with misoptimizing agents. Farhi and Gabaix (2017) give the topic a general treatment, characterizing results in terms of “behavioral sufficient statistics” (misoptimization wedges), without regard to the source of misoptimization. Gerritsen (2016) characterizes optimal income taxation with abstract misoptimization, calibrated using subjective well-being data, again without taking a stand on the reason for misoptimization. This paper takes an alternative (and complementary) approach. By focusing on a specific model of misoptimization, I am able to derive quantitative implications for optimal taxes, with comparative statics and testable hypotheses for labor supply behavior. In this respect, this paper is more similar to Spinnewijn (2015), which calibrates optimal unemployment insurance using data on mistaken beliefs about the probability of reemployment, and to Moser and Olea de Souza e Silva (2017), which characterizes optimal savings policies when agents undersave for retirement.

This paper also contributes to the substantial literature on the optimality of negative marginal tax rates. In the canonical model of redistributive income taxation laid out by Vickrey (1945) and Mirrlees (1971), negative marginal tax rates are suboptimal.⁸ Later analyses explored the sensitivity of this result to positive and normative assumptions in the conventional model. Diamond (1980) and Saez (2002) influentially argued that marginal tax rates could be negative at low incomes if earnings responses are concentrated on the extensive margin—for example, if there are heterogeneous fixed costs of labor force participation. Saez (2002) and Blundell and Shephard (2011) use models with discrete earnings levels to simulate optimal tax rates, finding low (or slightly negative) marginal tax rates on the lowest positive earning type. Jacquet, Lehmann and Van der Linden (2013) refined this insight in a continuous model, showing that extensive margin effects generally call for a positive participation credit (a fixed amount paid to all labor force participants), with positive marginal tax rates at all positive incomes.

⁸This finding has been discussed extensively—see Seade (1977), Seade (1982), Hellwig (2007), and citations therein.

Other work has shown that marginal work subsidies may be justified by normative objectives which differ from those in the conventional model. Most plainly, negative marginal tax rates may be warranted if the tax authority’s goal is to redistribute income *upward*—i.e., if marginal social welfare weights are rising with income. This point was originally made in a discrete context by Stiglitz (1982); a number of recent papers have argued that such weights may arise from multidimensional heterogeneity (Cuff, 2000; Beaudry et al., 2009; Choné and Laroque, 2010).⁹ Fleurbaey and Maniquet (2006) shows how fairness considerations may generate welfare weights which increase with income in equilibrium. Drenik and Perez-Truglia (2014) provides empirical evidence for such views, while noting that such an objective could generate Pareto inefficient policy recommendations. Although the reasoning in this paper is not inconsistent with such normative objectives, these findings demonstrate that negative marginal tax rates may be warranted even under the conventional assumption that policy makers wish to redistribute toward low earners.¹⁰

A third relevant area, sometimes called “positive optimal taxation,” is the growing body of work attempting to use existing policy and rhetoric as a guide to refine models of optimal taxation. These efforts include the “inverse optimum” literature, which inverts the conventional optimal taxation exercise to compute the welfare weights which rationalize existing policy, pioneered by Bourguignon and Spadaro (2012) (see also Hendren (2014), Lockwood and Weinzierl (2016), and the citations therein). Also in this vein are recent efforts to generate models of taxation which more closely reflect the rhetoric used in debates over optimal taxation.¹¹

The remainder of the paper is organized as follows. Section 2 presents a generalized model of optimal taxation with multiple dimensions of labor effort which occur over time, and with present biased workers. Propositions characterize the optimal income tax, including optimal tax rates and

⁹Preference heterogeneity alone is not generally sufficient to generate negative marginal tax rates, however—see simulations in Lockwood and Weinzierl (2015), where optimal tax rates are lower in the presence of such heterogeneity, but remain positive.

¹⁰Two other proposed rationales for negative marginal tax rates have received somewhat less attention. First, negative rates may be justified by “non-welfarist” objectives, for example if the government wishes to minimize poverty (Kanbur, Keen and Tuomala, 1994; Besley and Coate, 1992, 1995) or has preferences directly over the labor and leisure choices of poor individuals (Moffitt, 2006). Second, work may have positive externalities (or poverty negative externalities, e.g., Wane (2001))—indeed, a more recent literature has demonstrated that labor force participation may have positive effects on children’s outcomes (Dahl and Lochner, 2012). This justification is not inconsistent with the present bias rationale in this paper—indeed the distinction between present bias and positive externalities on children is a blurry one. I will focus on shorter term benefits when calibrating present bias in Section 3.1; if work also has positive externalities, I will underestimate the size of optimal work subsidies.

¹¹See Weinzierl (2014a) and Weinzierl (2014b) for an applications to alternative normative frameworks, and Lockwood, Nathanson and Weyl (2017), which considers the importance of using tax policy to usefully allocate talent.

the optimal timing of tax collections. Section 3 presents quantitative results, including calibrations of the optimal income tax (Section 3.2), which draw on recent evidence of present bias over effort to calibrate present bias across the income distribution. Section 3.3 presents an inverse optimum exercise which computes the redistributive preferences implied by the existing EITC, with and without accounting for present bias. Section 4 discusses implications for policy design, and Section 5 concludes.

2 Model

Although the canonical Mirrlees (1971) model contains a single dimension of labor effort (“time worked”), that choice has long been understood to represent a more general setting with a broad array of choices which affect earnings, including training and search effort, some of which may occur well before compensation is received. Feldstein (1999) and Chetty (2009) formalize this generalization, using a model of multidimensional labor choice. I build on that framework to allow for present biased labor workers.

2.1 Individuals

Earnings are a function of a vector of labor supply choices $\{\ell_j\}_{j=1}^J \in \mathbb{R}_+^J$ with strictly convex disutilities $v(\ell) = \sum_j v_j(\ell_j)$ (assumed to be separable, for simplicity), and with associated wages $w = \{w_j\}_{j=1}^J$. Thus taxable earnings are $z = w \cdot \ell = \sum_j w_j \ell_j$. I assume the agent faces a smooth, nonlinear income tax $T(z)$, and that labor supply adjustments are continuous in response to marginal perturbations of the tax function. (An extensive margin of labor force participation is easily added to this setting, though it slightly complicates the resulting optimal tax expressions—see Appendix B for a derivation.) Utility takes the form $u(c) - v(\ell)$, so the agent selects $\ell^* = \arg \max_{\ell} \{u(w \cdot \ell - T(w \cdot \ell)) - v(\ell)\}$.

I modify this conventional model to allow that the labor choices ℓ_j may occur prior to the date of compensation. Let $\tau(j)$ denote the temporal distance between the date of activity ℓ_j and the date of consumption, with $v_j(\ell_j)$ understood as the flow disutility of activity j . A time-consistent

individual with discount factor δ would then prefer a labor supply choice which maximizes

$$U(c, \ell) = - \sum_{j=1}^J \delta^{-\tau(j)} v'_j(\ell_j) + u(c), \quad (1)$$

with $c = w \cdot \ell - T(w \cdot \ell)$. Assuming an interior solution, the optimal labor supply choice ℓ^* favored by a time consistent individual satisfies the set of first-order conditions $v'_j(\ell_j^*)/w_j = \delta^{\tau(j)} u'(w \cdot \ell^* - T(w \cdot \ell^*)) (1 - T'(w \cdot \ell^*))$ for $j = 1, \dots, J$. This preference is not sensitive to the period of evaluation. I adopt the normative assumption that the policy maker agrees with this time-consistent perspective, and seeks to maximize a weighted sum of utility as in equation (1) across individuals.¹² (More on individual heterogeneity below.)

A time-inconsistent individual, on the other hand, may favor different labor supply choices depending on the period of evaluation. I adopt the widespread $\beta - \delta$ representation of time inconsistency, as in Laibson (1997), in which an individual's "time t self" discounts all future periods by β , in addition the exponential discount factor δ . Thus if each individual's time t self can make that period's labor supply choices freely, the individual will choose a "present biased" level of labor supply ℓ^{pb} which satisfies the following set of first-order conditions:

$$\begin{cases} v'_j(\ell_j^{pb})/w_j = \delta^{\tau(j)} u'(w \cdot \ell^{pb} - T(w \cdot \ell^{pb})) (1 - T'(w \cdot \ell^{pb})) & \text{if } \tau(j) = 0 \\ v'_j(\ell_j^{pb})/w_j = \beta \delta^{\tau(j)} u'(\hat{z}_{\tau(j)} - T(\hat{z}_{\tau(j)})) (1 - T'(\hat{z}_{\tau(j)})) & \text{if } \tau(j) > 0 \end{cases} \quad \text{for } j = 1, \dots, J. \quad (2)$$

where $\hat{z}_{\tau(j)}$ denotes the present biased individual's income forecast as assessed (by them) at the time of labor choice j .

The above expressions are written from the perspective of a single individual. Let the notation (i) index heterogeneity in $w(i) = \{w_j(i)\}_{j=1}^J$ and $\beta(i)$. As in Mirrlees (1971), I assume the utility function in (1) is homogeneous, thereby abstracting from difficulties of interpersonal utility

¹²Although the assumption that present bias is a mistake is common, it is not universal—e.g., Bernheim and Rangel (2009) studies welfare analysis without taking a stand on which of the time inconsistent choices is granted normative content. Some suggestive support for the normative assumption in this paper comes from subjective well-being evidence: Meyer and Sullivan (2008) study the change in consumption and labor supply of single mothers following the tax and welfare reforms of the late 1990s (which reduced lump-sum like benefits and inflated work subsidies) and conclude "The significant drop in nonmarket time suggests that utility has fallen for those in the bottom half of the consumption distribution if this nonmarket time is valued at more than \$3 per hour." Yet several papers find that the subjective well-being of single mothers (relative to groups unaffected by these reforms) stayed constant or rose over this period (Ifcher, 2010; Herbst, 2013; Ifcher and Zarghamee, 2014), consistent with an undervaluation of the experienced utility returns to work in the pre-reform period.

comparisons.

2.2 The policy maker’s problem

In keeping with the conventional Mirrleesian approach, I assume the policy maker cannot observe the wages or labor supply choices directly—instead it can only observe compensation $z = w \cdot \ell$.¹³

The policy maker’s problem is to select the tax function $T(z)$ which maximizes “social welfare”—a sum of individual utilities, possibly subjected to a weakly concave transformation $G(U)$ and weighting by Pareto weights $\alpha(i)$:

$$\mathcal{W} = \int \alpha(i)G(U(c(i), \ell(i)))d\mu(i), \tag{3}$$

where $c(i) = z(i) - T(z(i))$. This maximization is performed subject to the policy maker’s budget constraint

$$\int T(z(i))d\mu(i) - E \geq 0, \tag{4}$$

where E denotes an exogenous revenue requirement, and subject to incentive compatibility constraints, which require that $z(i) = w(i) \cdot \ell(i)$, with $\ell(i)$ satisfying the first-order conditions in (2) above.

2.3 Necessary and sufficient conditions for optimal taxes in a simple case

I first focus on an illustrative simple case with unidimensional labor supply and no income effects. In this case, the first-order (necessary) condition for optimal marginal tax rates can be written in terms of model primitives and it is possible to specify sufficient conditions for this formula to represent the optimum.¹⁴ As is well known, such solutions in terms of primitives are generally infeasible except in special cases—later extensions will consider weaker analytic results in more general settings.

¹³Although this assumption may seem counterfactual—the government regularly collects statistics on education levels and hours worked—this information can be understood as a constraint on the activities which can be taxed. The truthful revelation of education and hours may only be incentive compatible because those choices are untaxed, for example, generating a constraint against including them in the tax function.

¹⁴This case is motivated by Diamond (1998), which coincides with this model when $\beta(i) = 1$ for all i .

Specifically, for the purposes of this subsection I impose the following restrictive assumptions:

1. Labor supply is one-dimensional ($J = 1$), and is performed prior to compensation ($\tau(1) > 0$).
2. Utility is quasilinear in consumption ($u(c) = c$), which rules out income effects.
3. Ability w is distributed between $w_{min} \geq 0$ and $w_{max} < \infty$ with full support, and is continuously differentiable with density $f(w)$ and distribution $F(w)$.
4. Constant $\beta(i)$ conditional on ability $w(i)$, with $\beta(i)$ varying continuously with $w(i)$.

Since w and β vary jointly in this case (implying only a single dimension of heterogeneity), I simplify notation to write present bias, Pareto weights, and choice variables as functions of ability w rather than type i . I define $\zeta_\ell(w) = \frac{v'(\ell)}{v''(\ell)\ell}$ to denote the labor supply elasticity, evaluated at the optimal choice of earnings for ability w , which is assumed to be interior, and I define $\zeta_\beta(w) = \beta'(w)w/\beta(w)$, the elasticity of present bias with respect to ability. Then the following proposition characterizes the optimal income tax policy.

Proposition 1. *Under assumptions 1–4, the nonlinear income tax $T(z(w))$ which maximizes (3) subject to (4) and (2) must satisfy the following expression at all points of differentiability:*

$$\frac{T'(z(w))}{1 - T'(z(w))} = \mathcal{A}(w)\mathcal{B}(w) - \mathcal{C}(w), \quad (5)$$

with

$$\mathcal{A}(w) = \frac{1 + 1/\zeta_\ell(w) + \zeta_\beta(w)}{wf(w)} \quad (6)$$

$$\mathcal{B}(w) = \int_{x=w}^{w_{max}} (1 - g(x)) dF(x) \quad (7)$$

$$\mathcal{C}(w) = g(w)(1 - \beta(w)), \quad (8)$$

where

$$g(w) = \frac{\alpha(w)G'(U(c(w), \ell(w)))}{\int_{x=w_{min}}^{w_{max}} \alpha(x)G'(U(c(x), \ell(x)))dF(x)}. \quad (9)$$

This expression is sufficient to characterize the optimum if, under the earnings schedule induced

by this tax, each individual's globally optimal labor supply choice is characterized by the first-order condition $v'(\ell) = w\beta(w)(1 - T'(z(w)))$.

All proofs are provided in Appendix A. If this solution would generate an earnings schedule which decreases over some region of ability, then marginal tax rates are discontinuous— $T(z)$ is kinked—with a range of abilities bunching at the same level of earnings.

The expression in Proposition 1 is a generalization of Diamond (1998), which is identical except for the appearance of $\zeta_\beta(w)$ in $\mathcal{A}(w)$, and the final *corrective term*, $\mathcal{C}(w)$. Intuitively, the role of $\zeta_\beta(w)$ accounts for the fact that if bias is decreasing with ability (so $\beta'(w) > 0$), then high ability types are less tempted to imitate low ability agents by reducing effort, and thus marginal tax rates are optimally higher. The corrective term $\mathcal{C}(w)$, which appears via the Hamiltonian optimization in the proof, can also be derived using economic intuition. Since individuals undersupply effort due to present bias, the optimal tax counteracts this mistake by reducing marginal tax rates (or generating marginal subsidies).

The result in Proposition 1 invites three key observations about the role of present bias which are particularly clear in this setting. The first is a “negative at the bottom” result:

Proposition 2. *If $w_{min} > 0$ and $v'(0) = 0$, in addition to the sufficiency conditions in Proposition 1, then there exists a $w^* > w_{min}$ such that $T'(z(w)) < 0$ for all $w < w^*$.*

This proposition overturns the familiar result in the standard intensive-margin model that optimal marginal tax rates are everywhere nonnegative (Seade, 1982). It is a corollary of the result that in a conventional model without present bias, the optimal marginal tax rate is zero at the bottom under these assumptions (Seade, 1977). Indeed that result is a special case of this proposition—the bottom skill level receives the optimal correction, which is zero in the standard model and negative under present bias.¹⁵ Proposition 2 also shows that this solution is not isomorphic to some (perhaps unconventional) set of welfare weights, as there is no set of finite weights $g(w)$ which result in strictly negative marginal tax rates on the lowest earning type. Importantly, this result may be quite local to the lowest incomes, so it will be important to examine the quantitative behavior of

¹⁵Similar logic would yield an analogous result at the top of the income distribution: when the skill distribution has an upper bound, the famous “zero at the top” result becomes “negative at the top.” The zero at the top result requires the tax authority to know the highest income-earning ability with certainty—an assumption many find implausible, perhaps more so than knowledge of ability at low incomes—and so I do not emphasize that result here.

marginal tax rates across low incomes in the more realistic setting of the numerical simulations in Section 3.

The second observation is that the corrective term $\mathcal{C}(w)$ is proportional to $g(w)$, reflecting the planner’s greater concern for correcting biases among individuals whose consumption is highly valued at the margin. This force leads a more redistributive planner (higher $g(w_{min})$) to favor a larger correction at low incomes. This observation highlights one way in which an internality-correction rationale differs from a conventional (atmospheric) externality logic, where a positive work externality would not call for a larger correction on individuals with higher welfare weights.

The third observation, perhaps least obvious, demonstrates that the size of the optimal corrective subsidy increases with the labor supply elasticity.

Proposition 3. *Consider the effect of a slight increase in present bias ($d\beta < 0$) for some type w^\dagger , which induces a local reduction in the optimal marginal tax rate $dT'_{opt}(z(w^\dagger))$. The magnitude of $|dT'_{opt}(z(w^\dagger))|$ rises with the labor supply elasticity $\zeta_\ell(w^\dagger)$.*

This proposition follows immediately from the expression in Proposition 1, which can be rearranged as

$$T'_{opt} = \left(1 + \frac{1}{\mathcal{A}\mathcal{B} - \mathcal{C}}\right)^{-1}. \quad (10)$$

Differentiating T'_{opt} with respect to $\beta(w^\dagger)$ yields $\frac{dT'_{opt}}{d\beta} = \frac{g}{(\mathcal{A}\mathcal{B} - \mathcal{C} + 1)^2}$, a quantity that becomes larger in magnitude as $\zeta_\ell(w^\dagger)$ rises. Therefore a decrease in β causes a larger corrective reduction in the optimal marginal tax rate when the labor supply elasticity is high.

This result emphasizes the distinction relative to an externality, where the well-known “additivity” result (Sandmo, 1975) implies that the optimal marginal tax rate is the sum of the optimal redistributive tax rate and a Pigouvian correction, which is insensitive to the elasticity of behavior. Intuitively, the distinction in this model comes from the fact that the uninternalized benefit from additional labor effort (in the form of future income) is itself subject to the income tax, and thus the magnitude of that benefit is larger when the marginal tax rate is low, as is the case when labor supply elasticity is high.

2.4 Multidimensional heterogeneity

This section relaxes assumptions 1–4 above. Labor supply is now assumed to be multidimensional, individuals may be heterogeneous across the components of $w(i)$ and $\beta(i)$, and income effects may be nonzero. The optimal income tax which maximizes (3) subject to (4) and (2) in this generalized setting can be characterized by a first-order condition. To write this expression concisely, I define the following notation:

- $g(i) = \frac{\alpha(i)G'(U)u'(c(i))}{\lambda}$: marginal social welfare weights, equal to i 's marginal social value of consumption at the optimum, normalized by the marginal value of public funds λ
- $H(z)$ and $h(z)$: the distribution and density of incomes under the optimal tax.
- $\varepsilon(i) = \frac{dz(i)}{d(1-T')} \frac{1-T'}{z(i)}$: the compensated elasticity of taxable income with respect to the marginal “keep rate” $1 - T'(z(i))$.¹⁶
- $\eta(i) = -\frac{dz(i)}{dT}(1 - T')$: the income effect, representing the change in net earnings due to a small tax credit.
- $\phi(i) = \sum_{\{j|\tau(j)>0\}} w_j \frac{d\ell_j(i)}{d(1-T')} \bigg/ \frac{dz(i)}{d(1-T')}$: the share of earnings response owing to labor supply adjustments prior to compensation.
- $\bar{a}(z^*) = \mathbb{E}[a(i)|z(i) = z^*]$: for any type-dependent statistic $a(i)$, I denote the income-conditional average value using “upper bar” notation.
- $\Sigma_{a,b}^{(z^*)} = \frac{\text{Cov}[a(i),b(i)|z(i)=z^*]}{\bar{a}(z^*)\bar{b}(z^*)}$: the income-conditional covariance between any two type-dependent variables a and b (computed at the optimum), normalized by their income-conditional means. I further define $\Sigma_{a,b,c}^{(z^*)} = \frac{\mathbb{E}[a(i)b(i)c(i)|z(i)=z^*] - \bar{a}(z^*)\bar{b}(z^*)\bar{c}(z^*)}{\bar{a}(z^*)\bar{b}(z^*)\bar{c}(z^*)}$, the extension of this covariance-based definition to three variables (and likewise for more than three).

These parameters are endogenous to the tax system itself, and thus unlike in the previous section, they cannot be written in terms of model primitives. Nevertheless they can be used to characterize a necessary (first-order) condition for the optimal marginal tax rate using the calculus

¹⁶All elasticities and income effects incorporate both the direct change in earnings due to the tax reform, and any additional adjustments due to the change in marginal tax rate as earnings adjust, due to curvature in T . See Jacquet et al. (2013) for a discussion.

of variations approach, and they suggest that the insights highlighted in the previous simple case carry through to this more general setting.

Before proceeding to the optimal income tax characterization, this more general setting requires specifying a present biased individual's forecasted income, denoted by $\hat{z}_{\tau(j)}$ in Equation (2). This may pose a problem if the individual miscalculates their present bias over future labor effort, as they will then also miscalculate future income and the resulting marginal tax rate and marginal utility of consumption. For the purposes of the derivation below, I adopt a simple assumption:

Assumption 1 (Accuracy of income forecasts). *When individuals make labor supply choices ℓ_j^{pb} according to the first-order conditions in (2), they set $\hat{z}_{\tau(j)} = w \cdot \ell_j^{pb}$.*

This assumption abstracts from labor supply misoptimization driven by misunderstandings of projected income. This assumption, which is adopted to simplify the analysis to follow, might seem a reasonable benchmark for several reasons. First, if the individual in question is sophisticated about their present bias, they will correctly anticipate future labor supply decisions. Second, even if the individual is naive, if for all j we have $\tau(j) \in \{0, \tau^*\}$ for some τ^* , meaning all prior effort dimensions are concentrated around the same time (as would arise in a simple two-period model) then Assumption 1 holds automatically, as only the first set of labor supply choices depends on present bias, and thus the individual's assessment of their present bias is irrelevant. Third, the distortion generated by errors in forecasted income are proportional to the curvature of the utility function u and the tax function T . To the extent that the tax function and the utility of consumption are close to linear over the relevant range of possible income variation, forecasting errors will not affect labor supply choice. As a final consideration: there are doubtless many difficulties of forecasting one's income and future tax rate, a thorough study of which would involve measuring and validating income predictions directly, thereby capturing errors arising through present bias as well as any other relevant channels.

Using the notation above, and employing Assumption 1, we have the following proposition characterizing the optimal income tax.

Proposition 4. *The first-order condition for the optimal tax at some income z (where T is twice*

differentiable) can be written

$$\frac{T'(z)}{1 - T'(z)} = \mathcal{A}(z)\mathcal{B}(z) - \mathcal{C}(z), \quad (11)$$

with

$$\mathcal{A}(z) = \frac{1}{\bar{\varepsilon}(z)h(z)z}, \quad (12)$$

$$\mathcal{B}(z) = \int_{s=z}^{\infty} \left[1 - \bar{g}(s) - \bar{\eta}(s) \left(\frac{T'(s)}{1 - T'(s)} + \bar{g}(s)\bar{\phi}(s)(1 - \bar{\beta}(s)) \left(1 + \Sigma_{1-\beta, \eta, g, \phi}^{(s)} \right) \right) \right] dH(s), \quad (13)$$

$$\mathcal{C}(z) = \bar{g}(z)\bar{\phi}(z)(1 - \bar{\beta}(z)) \left(1 + \Sigma_{1-\beta, \varepsilon, g, \phi}^{(z)} \right). \quad (14)$$

As is standard for such first-order conditions, the terms in (11) are endogenous, and therefore the expression is not dispositive with respect to comparative statics. Nevertheless, such expressions are a useful guide to the forces which govern the optimal tax, and the comparative statics suggested by the expression are consistent with the numerical simulations presented in Section 3.

The optimal tax characterization in Proposition 4 clarifies a number of useful insights. First, the parallel structure to Proposition 1 is apparent, suggesting that the basic lessons and intuitions from that simple case extend to the context with multidimensional labor supply and heterogeneity. In particular, as before, present bias tends to reduce marginal tax rates, particular when the marginal social welfare weight or earnings elasticity is high.

The additional insights from the multidimensional context are also apparent from the differences between the terms \mathcal{A} , \mathcal{B} , and \mathcal{C} in Proposition 1 and Proposition 4. First, \mathcal{B} now contains an additional term, proportional to $\bar{\eta}$, capturing both the fiscal externality and behavioral welfare effects from the earning adjustments due to income effects. Intuitively, raising marginal tax rates at z takes resources away from those with higher incomes, who therefore raise their earnings due to income effects. This adjustment is beneficial both through a fiscal externality (the government raises more revenue) and through a present bias correction (working more generates a first-order benefit for the individual in question). Those beneficial effects justify higher marginal tax rates at each z .

A second difference in Proposition 4, relative to Proposition 1, is the presence of the Σ terms,

which incorporate effects of multidimensional heterogeneity.¹⁷ If the statistics β , ε , η , g , and ϕ are mutually orthogonal conditional on income (perhaps a plausible benchmark assumption, until empirical research studies such patterns of heterogeneity more closely) then the Σ terms are zero and can be ignored. Otherwise, they intuitively illustrate the directional effect of income-conditional correlations. For example, if $1 - \beta$ and ε are highly correlated at a particular income, that implies the individuals most responsive to marginal subsidies are those who are most biased (meaning taxes are a more powerful corrective tool) which magnifies the optimal correction term \mathcal{C} .

A final, and potentially important, insight from the more general Proposition 4 is the presence of ϕ . The corrective term \mathcal{C} (and the corrective component of the income effect term in \mathcal{B}) is multiplied by $\bar{\phi}$, the *share* of earnings adjustment which comes along labor supply margins chosen prior to compensation. Intuitively, if work subsidies primarily cause adjustments in labor supply choices contemporaneous with earnings (small ϕ), then corrective subsidies are of little use, since those margins do not benefit from correction. On the other hand, if most adjustments are through labor supply choices made in advance, then subsidies are a powerful corrective policy instrument.

Although the model above and the condition in Proposition 4 does not include an extensive margin (labor force participation) response, Appendix B extends this setting to consider that case.

In terms of empirical implementation, the intensive margin elasticity of taxable income, marginal social welfare weights, and income effects have been discussed at length in the optimal tax literature. However, the present bias parameter β is new to these expressions. That parameter is the topic of thorough study in the behavioral economics literature, both in laboratory experiments and from structural estimation studies. I discuss estimates of β (and how those estimates vary across incomes) in Section 3.

2.5 Extension: optimal tax timing

One question which lies beyond the domain of this reduced two period model is the issue of optimal tax *timing*. The model as described assumes that taxes are levied (and subsidies paid) at the time compensation is received, in contrast to the Earned Income Tax Credit in the US, which remits work subsidies in the form of a tax refund when taxes are filed, typically during the following year.

¹⁷Although individuals are multidimensionally heterogeneous in this setting, through their vector wages $w_j(i)$ and their present bias $\beta(i)$, since the policy maker faces only a single observed outcome (earnings), and has only a single income tax policy instrument, issues of multidimensional screening do not render the problem intractable.

Although a full consideration of optimal tax timing is beyond the scope of this paper, in this section I consider a simple extension to the two period reduced model, in which I allow taxes to be delayed relative to the time income is received.

To abstract from issues of inter-temporal smoothing and saving behavior which are not of central focus, for the purposes of this section I assume consumers live hand-to-mouth (consume all income each period) and have utility that is quasilinear in consumption. As a result, delayed taxes require only a small modification to the model in equations (1)–(4).

Specifically, I now assume that the model has an additional period of consumption, after all labor is performed and compensation is paid, during which some portion $\psi \in [0, 1]$ of taxes are levied (or subsidies paid). This period can be thought of as the date at which tax refunds are paid at the end of the year.¹⁸ The parameter ψ is an additional parameter set by the policy maker, representing the share of taxes withheld from one’s paycheck, or the share of the EITC paid at the end of the year rather than up-front.

For this section I assume that the exponential discount rate δ is equal to 1, to avoid complications of the government paying interest on delayed taxes. (In this case, the precise amount of time delay between consumption and delayed taxes is unimportant.) The individual’s utility function (1) is replaced by

$$U(c, \ell) = - \sum_{j=1}^J v'_j(\ell_j) + w \cdot \ell - (1 - \psi)T(w \cdot \ell) - \psi T(w \cdot \ell). \quad (15)$$

The social welfare function (3) and budget constraint (4) remain unchanged. However, the first-order conditions for labor choice change, since even contemporaneous dimensions of labor effort are perceived to have delayed consequences due to delayed taxes. Therefore (2) is replaced by

$$\begin{cases} v'_j(\ell_j^{pb})/w_j = (1 - \psi(1 - \beta))(1 - T'(w \cdot \ell^{pb})) & \text{if } \tau(j) = 0 \\ v'_j(\ell_j^{pb})/w_j = \beta(1 - T'(w \cdot \ell^{pb})) & \text{if } \tau(j) > 0. \end{cases} \quad (16)$$

The policy maker’s problem is now to select the nonlinear function T and the delay parameter $\psi \in [0, 1]$ that maximizes (3) (where U is given by (15)) subject to (4) and to (16). This modification

¹⁸Of course it is not crucial that no consumption or labor occur during this following period—rather the key feature is that some portion of taxes are levied or paid with a delay.

generates the striking and perhaps surprising result that if workers are present biased, it is optimal to delay taxes and subsidies.

Proposition 5. *The policy (T, ψ) which maximizes (3) and (15) subject to (4) and (16) satisfies the following conditions:*

1. *if $\phi(i) = 1$ for all i , then the optimal tax and overall welfare are not sensitive to the share of delayed taxes ψ ;*
2. *if $\phi(i) < 1$ and $\beta(i) < 1$ for some i with positive earnings, then the optimal policy features $\psi > 0$: a strictly positive share of taxes is delayed at the optimum; and*
3. *if at the optimum $g(i)(1 - \beta(i)) \leq 1$ for all i , then the optimal policy features $\psi = 1$: the optimal tax is fully delayed.*

Part 1 of the proposition states that if all elastic dimensions of effort occur prior to compensation, then it does not matter whether taxes are delayed. This result follows from the insensitivity of advance effort to ψ , and contrasts with the common intuition that the existence of present bias strongly calls for an EITC which is included with one's paycheck. In this case the choice of optimal tax timing may depend on considerations beyond the scope of labor supply correction, such as enforcement, liquidity, or income smoothing and forced savings.

Parts 2 and 3 of the proposition are even more striking. If some share of income elasticity is due to adjustments at the time of compensation, then the optimal tax is delayed (at least partially) at the optimum, and fully delayed if $g(i)(1 - \beta(i)) < 1$ for all individuals. To understand the economic logic of Part 2, consider beginning from the optimal tax when $\psi = 0$, characterized by Proposition 4, and suppose the tax is reformed by slightly raising the share of delayed taxes ψ . Such a reform has no effect on effort exerted in advance, by the logic outlined above. However the reform dampens the effect of taxes on contemporaneous effort. Consider some j' such that $\tau(j') = 0$. If $T'_{opt}(z(i)) > 0$, a delay raises $\ell_j^{pb}(i)$; if $T'_{opt}(z(i)) < 0$, a delay lowers $\ell_j^{pb}(i)$. By the envelope theorem, there is no first-order effect on the individual's utility from this change in labor supply when $\psi = 0$. However there is a beneficial first-order fiscal externality from the adjustment in either case. Therefore the optimal share of delayed taxes is strictly positive. Part 3 states that even when delayed taxes cause an oversupply of contemporaneous effort (relative to the private

optimum), if the welfare-weighted misoptimization $g(i)(1 - \beta(i))$ is weakly less than one, then this private welfare cost is smaller than the beneficial fiscal externality from the labor adjustment.¹⁹

Taken together, the components of Proposition 5 suggest that the logic of present bias does not, in itself, call for an EITC which is included with one’s paycheck. Indeed there may be a positive benefit from a modest delay, since delayed taxes do less to distort contemporaneous dimensions of effort which are not subject to present bias. Note, however, that this delay need not be large to reap these benefits—just sufficiently delayed that the perceived tax benefits are discounted by β at the time of compensation.

3 Numerical analysis

This section considers the quantitative implications of present bias for optimal income taxation. Propositions 1 and 4 highlight the key sufficient statistics which are necessary to evaluate whether a candidate income tax is optimal, which fall into the four categories: (1) behavioral elasticities and income effects which have previously been estimated in the labor supply literature: the compensated elasticity of taxable income with respect to the marginal keep rate ε and the income effect η ; (2) the present bias parameter β , which has been estimated extensively in the behavioral economics literature; and (3) marginal social welfare weights g , which are often exogenously imposed but have also been estimated in the inverse optimum literature.

Although these statistics are sufficient to evaluate the optimality of a status quo income tax, a structural model of behavior is necessary to estimate the optimal tax away from the observed equilibrium. To focus on the implications of present bias for the optimal tax, and to make these simulations comparable to existing results in the optimal taxation literature, I structure these simulations around the simpler setting of Section 2.3 with unidimensional heterogeneity. Specifically, I assume use the following specification for decision utility

$$U_i = -\frac{(z/w(i))^{1+k}}{1+k} + \beta(i)(z - T(z)). \quad (17)$$

Although this specification maps to a single dimension of labor supply choice, it is isomorphic to

¹⁹Under common empirical estimates of β from the behavioral economics literature, and estimates of welfare weights from the inverse optimum literature, the condition in Part 3 is satisfied even at the bottom of the income distribution. These estimates are discussed at more length in the next section.

multiple dimensions of labor supply with a constant elasticity of taxable income and $\phi = 1$ (all earnings responses to tax changes are through labor supply choices made in advance of compensation).

I assume a fixed share of the population is disabled and has $w = 0$. To avoid complications of imperfect screening and optimal disability insurance, I assume that disability status is observable to the tax authority and that the required revenue for disabled individuals is exogenously given.²⁰ I assume the exogenously required income for disabled individuals is \$7,500, equal to average Social Security income in this age group in the Current Population Survey. Thus disability insurance effectively contributes to the government's exogenous revenue requirement. I further assume that present bias β covaries perfectly with ability w , though it may vary across the income distribution. The parameter k controls the labor supply elasticity.²¹

This setup simplifies the model considerably, as it reduces the required structural parameters to the income elasticity parameter k , the distribution w , and the mapping between w and present bias and welfare weights, $\beta(w)$ and $g(w)$. I discuss the calibration of each in turn. I use a baseline value of $1/k = 0.3$, close to the preferred intensive margin elasticity from Chetty (2012) of 0.33, and I also provide results for $1/k = 0.4$ and $1/k = 0.2$. The calibration of β is the primary new challenge for these simulations, and therefore I discuss that calibration in detail in the next subsection. After choosing structural values $\beta(i)$ and k , the structural ability distribution can be calibrated in the usual manner, by inverting the first-order condition for (17) to find the skill distribution which would generate the observed income distribution under the prevailing tax code. Details of the income distribution and status quo tax schedule calibration are discussed in Appendix D.

Finally, I employ a reduced-form representation of declining marginal social welfare weights $g(w)$, which can be interpreted either as declining α Pareto weights, or as declining due to concavity of G (and therefore declining G') at the optimum. I adopt the conventional assumption that weights are declining as consumption increases; the exact patterns for the simulations are discussed below. For the planner's budget constraint, I impose an exogenous government revenue requirement of \$5,000 per capita.

²⁰Specifically, I assume that 2% of individuals are disabled and unable to work altogether, consistent with the share of respondents in CPS between ages 25 and 55 with positive SSI income.

²¹For comparability to Section 2, $\varepsilon(i) = \frac{1}{k + \frac{1}{1-T}}$, so that at points where the income tax is locally linear, $\varepsilon = 1/k$.

3.1 Calibrating present bias

Table 1 presents estimates of β from several recent papers. The checkmarks highlight two desirable features. First, the column labeled “low income/EITC” identifies papers whose subjects are drawn from a population of low earners, and in some cases EITC recipients, in the US. Such studies are useful for two reasons. First, they mitigate concerns about external validity, as their subjects resemble the population of interest for the simulations in this paper. Second, estimates from low income populations dampens the concern that monetary rewards are problematic for the estimation of present bias, since subjects can use their own funds to replicate (or undo) experimental variation in payoffs. Low income subjects are more likely to face liquidity constraints which prevent such arbitrage, perhaps explaining the substantial measured present bias even over monetary payoffs in those studies.²²

Second, the column labeled “effort” identifies studies which estimate β using intertemporal tradeoffs over effort tasks or labor supply, rather than money. These studies are particularly informative both because they avoid the shortcomings of monetary payoffs, and because the focus of this paper is labor supply, so to the extent that bias varies across domains, these studies identify the parameter of interest. All studies find estimates of β meaningfully (and statistically significantly) below one.

Figure 2 plots these values of β across the incomes in each study and, when possible, plots the relationship between β and income *within* studies. These calculations, the details of which are reported in Appendix C, are in some cases quite rough, and none of these studies was written with the primary goal of estimating present bias across incomes. Indeed, one implication of this paper is that the covariation of present bias with income is worthy of additional empirical research. Nevertheless, as shown in Figure 2, there is a strong positive correlation between β and income, both across studies and (when reported) within individual studies.

This relationship is consistent with a number of possible explanations. First, theory predicts that present-biased individuals endogenously exert less effort and therefore have lower earnings. Second, present bias likely reduces longer term human capital investments, leading to an inverse relationship between the bias wedge and underlying ability. Third, circumstances of material scarcity

²²See Carvalho et al. (2016) for evidence that liquidity constraints generate measurable present bias over monetary payoffs.

might *cause* greater present-bias (Mullainathan and Shafir, 2013). I remain agnostic about the mechanism for the relationship, effectively assuming that the plot in Figure 2 indicates a stable type-specific level of bias as a function of underlying ability.

Figure 2 also displays the best fit line for all the plotted points—I use this relationship for the structural calibration of β across ability in the simulations below, restricting to a maximum value of $\beta = 1$.

3.2 Optimal tax simulations

This section demonstrates the key finding previewed in Figure 1: if redistributive tastes are modest, then negative marginal tax rates on par with those generated by the EITC may be nearly optimal.

For ease of computation and transparency, I encode redistributive preferences directly through declining marginal social welfare weights $g(w)$. For the baseline set of modest redistributive preferences, I select g -weights such that the lowest earners receive a weight 10% more than the median household, and top earners weighted by 40% less than the median, linearly interpolated between percentiles 0, 50, and 100. These weights are substantially less redistributive than those conventionally assumed in the optimal taxation literature (e.g., logarithmic utility over consumption). However there is other evidence that existing tax policies embodies more modest redistributive tastes than is conventionally assumed in that literature (see Hendren (2014) and Lockwood and Weinzierl (2016), for example).

Figure 3 displays the schedule of optimal marginal tax rates for four simulations. The top left panel reproduces the baseline calibration in Figure 1. The top right panel shows optimal taxes with logarithmic redistributive preferences, so that $g(w) = c(w)^{-1}$ at the optimum. The bottom two panels use the baseline set of redistributive weights, with alternatively higher and lower labor supply elasticities.

These simulations highlight some key lessons for tax policy with present biased workers. First, as highlighted in the introduction, negative marginal tax rates like those under the EITC are justified if redistributive preferences are modest. Second, as the log redistributive tastes case illustrates, if redistributive preferences are strong, marginal tax rates remain positive throughout the income distribution, although present bias still tends to reduce tax rates below what is optimal if the population is unbiased. This result may seem surprising in light of the fact that the correction term

in Propositions 1 and 4 are weighted by the marginal social welfare weight, which is higher for lower earners under stronger redistributive preferences. That stronger corrective motive is outweighed, however, by the stronger desire to redistribute *across* low earners under log preferences, reflected by the high level of marginal tax rates under the rational optimum in the log case. These higher marginal tax rates are used to fund a larger lump sum (or basic income), equal to \$21,734, compared to \$4,535 in the less redistributive baseline case. Since the corrective term operates on the term $\frac{T'}{1-T'}$, when marginal tax rates are high (close to one) in the rational case, this fraction is very large, and even a substantial corrective term has little effect.

The third lesson from Figure 3 relates to the labor supply elasticity. A higher elasticity reduces optimal marginal tax rates in the present biased optimum. In fact an elasticity of 0.4, higher than baseline but still well within the range of some empirical estimates, particularly from the macro literature (see Chetty (2012)) generates optimal marginal tax rates as low as -30% .²³ It's worth stressing the divergence between this comparative static and the effect of higher intensive margin elasticities for the the EITC-like subsidies in Saez (2002), where a higher intensive margin labor supply *reduces* the magnitude of marginal tax rates on the poorest workers.

On the other hand, if elasticities are low, as in the lower right panel of Figure 3, then present bias does not generate marginal work subsidies at all. Of course, although these simulations use a constant value for the elasticity, the elasticity may vary with income in practice. Thus if individuals at the bottom of the distribution are particularly elastic, subsidies may be justified even if elasticities are fairly low at higher points in the income distribution.

3.3 Inverse optimum exercise

In this section I adopt an alternative approach, inverting the optimal policy simulation by taking existing policy as given and computing the redistributive preferences with which those policies are consistent. This strategy, implemented by Bourguignon and Spadaro (2012) for European countries and further explored by Hendren (2014) and Lockwood and Weinzierl (2016), provides a reduced-form way to check whether existing policy generates weights which appear reasonable.

²³In this simulation with quasilinear utility and fixed (finite) welfare weights at the bottom, it is possible for consumption to be negative at the optimum for some individuals. I therefore impose an additional constraint on the optimal tax that all incomes must be nonnegative. That constraint binds only in the high elasticity specification, generating high marginal tax rates at the very bottom.

Of course, the definition of “reasonable” is itself controversial, but two features are commonly thought to be sensible requirements in the optimal policy literature: welfare weights are positive at all incomes (Pareto efficiency) and weights are declining with income (redistribution toward lower skilled individuals).

In this section, I show that the implicit welfare weights consistent with the EITC under the usual assumption of perfect optimization exhibit a robust “unreasonable” feature: weights rise substantially with income across low earnings. Taken at face value, the weights suggest that current policy implicitly places greater value on a marginal dollar for middle earners than on a dollar in the hands of the poorest EITC-receiving households (typically working single mothers). I compare these weights to those which arise implicitly from a model which assumes misoptimization due to present bias, calibrated according to Figure 2 as in the previous simulations.

A secondary strength of the inversion approach is that it permits a more detailed representation of the complexities of the actual economy. Since this approach entails only a local inversion of the first-order condition for optimal taxes, it does not require a structural model of earnings responses to non-local tax reforms. As a result, it is possible to incorporate a more detailed calibration of elasticities, including non-constant elasticities of taxable income and positive labor force participation elasticities (see Appendix B for an extension of Proposition 4 to that setting).

In the interest of computing the weights generated by the EITC, I focus on the set of families affected by the credit. Therefore I use a sample different from the one in the benchmark economy of Section 3, though I continue to draw data from the CPS. Specifically, I use the 2015 March CPS, restricted to households with 2 children, and I calibrate marginal tax rates using TAXSIM, as well as the phaseout benefits from CPS. Details of the sample selection and tax rate calibration are discussed in Appendix D.

In addition to the income distribution and schedule of marginal tax rates, the optimal tax condition depends on the compensated elasticity of taxable income, the income effect, and the participation elasticity, as well as the misoptimization wedge, conditional on income. I assume an elasticity of 0.33 at middle and high incomes, the preferred value in Chetty (2012), which lies well in the range of other estimates. For the elasticities at low incomes, I draw from evidence drawn specifically from the EITC-receiving population. Chetty, Friedman and Saez (2013) estimate intensive margin elasticities of 0.31 and 0.14 in the phase-in and phase-out regions of the EITC,

respectively, identified by differences in knowledge of (and, by assumption, responses to) the EITC across geographic regions. Therefore I assume a compensated elasticity of 0.31 for households with incomes below \$15,000 (the approximate center of the EITC plateau) and 0.14 for those with incomes between \$15,000 and \$40,000 (the upper bound of the phase-in region for households with two children).²⁴

For the participation elasticity, Saez (2002) performs calibrations using estimates of 0, 0.1, and 0.5 for the bottom half of the population (and zero for the top half). Chetty et al. (2013) estimates a participation elasticity of 0.19 among EITC recipients—I use this value for the participation elasticity for households with earnings less than \$20,000, and zero participation elasticity at high incomes. Again to avoid a discontinuity in the participation elasticity, I interpolate linearly across incomes from $\rho = 0.19$ at \$20,000 to $\rho = 0$ at \$40,000.

The resulting marginal social welfare weights are plotted in Figure 4, both under the conventional assumption of no misoptimization, and under the assumption that individuals are present biased. Plotted points represent the weight computed locally in \$2500 income bins, while the line plots the smoothed relationship.

The schedule of weights computed without present bias are strikingly unconventional in a key respect: they rise over a substantial range of low incomes, peaking near \$25,000 in annual earnings, before declining with income thereafter. This result indicates that according to the conventional model, policy is designed primarily to benefit EITC-candidate households (households with two children) in the middle of the income distribution, while placing lower welfare weight on very poor EITC households (primarily single working mothers)—a result at odds with conventional normative assumptions. Similarly, the demographic homogeneity within this group of EITC recipients suggests that multidimensional heterogeneity, as in Choné and Laroque (2010), is an unlikely explanation for the pattern.

As shown by the dashed line in Figure 4, this unconventional feature disappears when a calibrated degree of present bias is incorporated into the calculation of welfare weights.²⁵ Weights are substantially higher than 1 at the bottom of the distribution, and decline monotonically with

²⁴Transitions are smoothed by interpolating across a range of \$10,000, centered at the threshold, to avoid sharp discontinuities in welfare weights generated by local jumps in elasticities.

²⁵The monotonic decline of welfare weights with income under present bias is somewhat fragile—some choices of bandwidth and income density estimation produce a slight non-monotonicity at the very bottom.

income.

These results capture the sense in which the EITC is difficult to reconcile with widespread normative assumptions under the conventional model, even allowing for an extensive margin of labor supply. However under a calibrated degree of present bias, the EITC is consistent with conventionally normative assumptions, since the dashed line in Figure 4 does entail a preference for redistribution across low earners, albeit a weak one.

4 Discussion

The analysis from the preceding sections allows the government only a single policy instrument: a nonlinear income tax. This restriction is useful for deriving policy implications that can realistically be implemented, avoiding complicated or impractical history-dependent policies that might arise from a full dynamic model. On the other hand, this constrained environment ignores some possible instruments which might be realistically feasible and useful for tailoring policies more specifically for present biased individuals.

A fully dynamic model of optimal tax timing with present bias is beyond the scope of this paper, and is a topic for additional research. Nevertheless, the results above do provide some suggestive guidance about the effects of various timing structures.

The current EITC is paid in aggregate at the end of the tax year—a structure that has generated mixed reviews. On one hand, spreading the credit across more frequent installments would help smooth consumption across the year and potentially alleviate liquidity constraints. Many EITC recipients carry substantial credit card balances, for example, and more frequent EITC payments would provide additional liquidity and could reduce costly interest expenses. On the other hand, the lump sum nature of the current EITC provides a short-run forced savings mechanism, and recipients often use the large annual payment to invest in durable goods—saving for which might otherwise prove difficult. Indeed, anecdotal evidence suggests EITC recipients do not *want* to receive distributed payments (Halpern-Meehin et al., 2015), a finding consistent with the very low uptake of the “Advance EITC” option, which allowed for more frequent payments. (See Romich and Weisner (2000) for a discussion, and Jones (2010) for experimental evidence of low desire for the Advance EITC.)

Proposition 5 in Section 2.5 suggests that it is beneficial to levy taxes with a lag—thus the delayed nature of the present EITC is not necessarily detrimental for combatting present bias, as initial intuition might suggest. However Section 2.5 recommends a rather different structure than the current very lumpy schedule of EITC payments. Rather than paying subsidies at the end of each year, subsidies should be smoothed with a small, constant delay—perhaps on the order of one month.

4.1 Limitations

This paper abstracts from several important complications. Here I discuss several simplifying assumptions and their implications.

No human capital acquisition. Although I allow for delayed benefits from labor effort, I assume these benefits are separable from subsequent labor supply decisions. This effectively rules out considerations of human capital accumulation. I make this assumption for a number of reasons. The first is pragmatic: the complexities of adequately accounting for human capital acquisition in optimal taxation are substantial, even without incorporating misoptimization (Stantcheva, 2014). For the sake of simplicity and transparency, and to generate results comparable to the existing literature exploring the (sub)optimality of negative marginal tax rates, I restrict consideration to a static distribution of ability. Second, the model in this paper still allows for a reduced-form relationship between human capital and present bias by calibrating bias conditional on income. Indeed, results in Section 3.1 suggest that bias is concentrated among those with lowest ability, as would be expected if some component of ability variation is due to human capital in which present biased individuals underinvest. Third, there is some empirical evidence that individuals who randomly receive work subsidies do not experience *persistent* increases in income relative to those who do not (Card and Hyslop, 2005, 2009)—a finding inconsistent with the notion that such subsidies raise human capital via on-the-job training effects.

To the extent that human capital effects are important for the design for optimal work subsidies, the directional implications are ambiguous. On one hand, some human capital is surely acquired on the job, raising the delayed benefits of work and likely inflating the size of optimal subsidies. On the other hand, human capital acquisition might at times be a *substitute* for work with even more

delayed benefits—for example, one may need to forego work to attend college. In that case work subsidies could exacerbate bias by discouraging human capital acquisition. This latter possibility may generate a rationale for exempting college-age individuals from EITC eligibility.

Perfectly competitive labor markets. In keeping with much of the optimal taxation literature, I assume that workers are employed in a perfectly competitive labor market, and that labor demand is infinitely elastic. This assumption has been questioned by Rothstein (2010), who argues that the incidence of work subsidies falls partly on employers. Finitely elastic labor demand undermines the argument for an EITC relative to guaranteed minimum income with high marginal tax rates, since the latter regime tends to reduce labor supply, raising wages and total transfers from employers to employees. Also in this vein, Kroft, Kucko, Lehmann and Schmieder (2015) incorporate endogenous wages and unemployment (not all job seekers find jobs) using a sufficient statistics approach. Their empirical results favor a negative income tax (rather than an EITC with negative marginal tax rates at low incomes) in a discrete model in the style of Saez (2002).

Fixed present bias. I assume throughout the model and calibrations that although present bias may vary with ability, it is fixed within individuals. In practice, biases may be mutable. One possibility, for example, is that exposure to the costs of present bias might lead individuals to improve their self control—a channel which would undermine the optimality of corrective subsidies that dampen such exposure. Another possibility, explored by Mullainathan et al. (2012) and Mullainathan and Shafir (2013), is that the conditions of poverty exacerbate behavioral biases. This possibility could be accommodated by writing bias as a decreasing function of consumption—a modification which would favor raising degree of overall redistribution. As with the case of human capital accumulation, this model would predict that temporary work subsidies should have persistent impacts on labor supply, inconsistent with the findings by Card and Hyslop cited above. Additionally, recent work by Carvalho et al. (2016) suggests that although liquidity constraints exacerbate measured present bias over monetary payments, they do not affect present bias over labor effort—consistent with a stable degree of structural bias. Still, optimal taxation with endogenous present bias is a promising avenue for further exploration.

5 Conclusion

As the study of optimal taxation begins to account for imperfect rationality and behavioral biases, a critical challenge is to quantify misoptimization accurately. This paper focuses on a particularly robust and well calibrated source of misoptimization—present bias—which, recent evidence suggests, generates substantial effort distortions.

A model of optimal taxation with present bias generates new theoretical implications, including a “negative at the bottom” result and a surprising implication for optimal tax timing—if workers are present biased and face multiple dimensions of labor choice, then it is beneficial to pay work subsidies with a delay.

A compilation of existing estimates of present bias provides strong evidence of such bias; more suggestive evidence indicates bias is heavily concentrated at low incomes. Although estimates of misoptimization will surely continue to improve, the consistency of results across methodologies provides some hope that misoptimization can be estimated with sufficient precision to provide clear guidance for policy design.

The implications for optimal tax policy depend on one’s view of optimal redistribution. If redistributive tastes are modest, like those in the baseline calibrations in this paper, then the negative marginal tax rates generated by the EITC may be close to optimal. In that case the policy implications of these results are clear: the existing EITC need not be reduced or greatly reformed—indeed, social welfare would rise if the EITC were extended to workers without children, and made more salient. On the other hand, if redistributive preferences are strong, in line with the welfare weights often assumed by optimal tax theorists (such as logarithmic utility of consumption) then negative marginal tax rates at low incomes appear to be suboptimal even in the context of present biased workers. As a descriptive matter, the existing EITC schedule appears inconsistent with the conventional normative assumption of decreasing marginal social welfare weights in the standard model; a calibrated model of present bias resolves this inconsistency and gives rise to declining weights.

References

- Augenblick, Ned, Niederle, Muriel and Sprenger, Charles.** (2015). ‘Working over time: Dynamic inconsistency in real effort tasks’, *The Quarterly Journal of Economics* p. qjv020.
- Augenblick, Ned and Rabin, Matthew.** (2015). ‘An Experiment on Time Preference and Misprediction in Unpleasant Tasks’, *Working paper, Haas School of Business and Harvard University*.
- Beaudry, Paul, Blackorby, Charles and Szalay, Dezsö.** (2009). ‘Taxes and Employment Subsidies in Optimal Redistribution Programs’, *The American Economic Review* 99(1), 216.
- Bernheim, B Douglas and Rangel, Antonio.** (2009). ‘Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics’, *The Quarterly Journal of Economics* 124(1), 51–104.
- Besley, Timothy and Coate, Stephen.** (1992). ‘Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs’, *American Economic Review* 82(1), 249–261.
- Besley, Timothy and Coate, Stephen.** (1995). ‘The design of income maintenance programmes’, *The Review of Economic Studies* 62(2), 187–221.
- Blundell, Richard and Shephard, Andrew.** (2011). ‘Employment, hours of work and the optimal taxation of low-income families’, *The Review of Economic Studies* p. rdr034.
- Bourguignon, François and Spadaro, Amedeo.** (2012). ‘Tax–benefit revealed social preferences’, *The Journal of Economic Inequality* 10(1), 75–108.
- Card, David and Hyslop, Dean R.** (2005). ‘Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers’, *Econometrica* 73(6), 1723–1770.
- Card, David and Hyslop, Dean R.** (2009). ‘The dynamic effects of an earnings subsidy for long-term welfare recipients: Evidence from the self sufficiency project applicant experiment’, *Journal of Econometrics* 153(1), 1–20.

- Carvalho, Leandro S, Meier, Stephan and Wang, Stephanie W.** (2016). ‘Poverty and economic decision-making: Evidence from changes in financial resources at payday’, *The American Economic Review* 106(2), 260–284.
- Chetty, Raj.** (2009). ‘Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance’, *American Economic Journal: Economic Policy* 1(2), 31–52.
- Chetty, Raj.** (2012). ‘Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply’, *Econometrica* 80(3), 969–1018.
- Chetty, Raj, Friedman, John N and Saez, Emmanuel.** (2013). ‘Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings’, *American Economic Review* 103(7), 2683–2721.
- Choné, Philippe and Laroque, Guy.** (2010). ‘Negative marginal tax rates and heterogeneity’, *The American Economic Review* 100(5), 2532–2547.
- Clinton, Bill.** (1993), ‘Address to Joint Session of Congress’.
URL: <http://www.presidency.ucsb.edu/ws/?pid=47232>
- Cuff, Katherine.** (2000). ‘Optimality of workfare with heterogeneous preferences’, *Canadian Journal of Economics/Revue canadienne d’économique* 33(1), 149–174.
- Dahl, Gordon B and Lochner, Lance.** (2012). ‘The impact of family income on child achievement: Evidence from the earned income tax credit’, *The American Economic Review* 102(5), 1927–1956.
- DellaVigna, Stefano, Lindner, Attila, Reizer, Balázs and Schmieder, Johannes F.** (2015). ‘Reference-Dependent Job Search: Evidence from Hungary’, *Unpublished Working Paper*.
- DellaVigna, Stefano and Paserman, M Daniele.** (2005). ‘Job search and impatience’, *Journal of Labor Economics* 23(3), 527–588.

- Diamond, Peter.** (1980). ‘Income taxation with fixed hours of work’, *Journal of Public Economics* 13(1), 101–110.
- Diamond, Peter A.** (1998). ‘Optimal income taxation: an example with a U-shaped pattern of optimal marginal tax rates’, *American Economic Review* 88(1), 83–95.
- Drenik, Andres and Perez-Truglia, Ricardo.** (2014). ‘Sympathy for the Diligent and the Demand for Workfare’, *Available at SSRN 1599368* .
- Fang, Hanming and Silverman, Dan.** (2009). ‘Time-inconsistency and welfare program participation: Evidence from the NLSY’, *International Economic Review* 50(4), 1043–1077.
- Farhi, Emmanuel and Gabaix, Xavier.** (2017). ‘Optimal Taxation with Behavioral Agents’, *NBER Working Paper No. 21524* .
- Feldstein, Martin.** (1999). ‘Tax avoidance and the deadweight loss of the income tax’, *Review of Economics and Statistics* 81(4), 674–680.
- Fleurbaey, Marc and Maniquet, Francois.** (2006). ‘Fair income tax’, *The Review of Economic Studies* 73(1), 55–83.
- Gerritsen, Aart.** (2016). ‘Optimal taxation when people do not maximize well-being’, *Journal of Public Economics* 144, 122–139.
- Halpern-Meehin, Sarah, Edin, Kathryn, Tach, Laura and Sykes, Jennifer.** (2015), *It’s Not Like I’m Poor*, University of California Press.
- Hellwig, Martin F.** (2007). ‘A contribution to the theory of optimal utilitarian income taxation’, *Journal of Public Economics* 91(7), 1449–1477.
- Hendren, Nathaniel.** (2014). ‘The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function’, *NBER Working Paper No. 20351* .
- Herbst, Chris M.** (2013). ‘Welfare reform and the subjective well-being of single mothers’, *Journal of Population Economics* 26(1), 203–238.
- Ifcher, John.** (2010). ‘The happiness of single mothers after welfare reform’, *The BE Journal of Economic Analysis & Policy* 11(1).

- Ifcher, John and Zarghamee, Homa.** (2014). ‘The Happiness of Single Mothers: Evidence from the General Social Survey’, *Journal of Happiness Studies* 15(5), 1219–1238.
- Jacquet, Laurence, Lehmann, Etienne and Van der Linden, Bruno.** (2013). ‘Optimal redistributive taxation with both extensive and intensive responses’, *Journal of Economic Theory* 148(5), 1770–1805.
- Jones, Damon.** (2010). ‘Information, Preferences, and Public Benefit Participation: Experimental Evidence from the Advance EITC and 401 (k) Savings’, *American Economic Journal: Applied Economics* pp. 147–163.
- Kanbur, Ravi, Keen, Michael and Tuomala, Matti.** (1994). ‘Optimal non-linear income taxation for the alleviation of income-poverty’, *European Economic Review* 38(8), 1613–1632.
- Kaur, Supreet, Kremer, Michael and Mullainathan, Sendhil.** (2015). ‘Self-control at work’, *Journal of Political Economy* 123(6), 1227–1277.
- Kroft, Kory, Kucko, Kavan J, Lehmann, Etienne and Schmieder, Johannes F.** (2015). ‘Optimal Income Taxation with Unemployment and Wage Responses: A Sufficient Statistics Approach’, *NBER Working Paper No. 21757*.
- Laibson, David.** (1997). ‘Golden eggs and hyperbolic discounting’, *The Quarterly Journal of Economics* 112(2), 443–478.
- Laibson, David, Mextad, Peter, Repetto, Andrea and Tobacman, Jeremy.** (2017). ‘Estimating discount functions with consumption choices over the lifecycle’, *Working paper*.
- Lockwood, Benjamin B, Nathanson, Charles G and Weyl, E Glen.** (2017). ‘Taxation and the Allocation of Talent’, *Journal of Political Economy* 125(5), 1635–1682.
- Lockwood, Benjamin B. and Weinzierl, Matthew.** (2015). ‘De Gustibus non est Taxandum: Heterogeneity in preferences and optimal redistribution’, *Journal of Public Economics* 124(0), 74–80.

- Lockwood, Benjamin B. and Weinzierl, Matthew.** (2016). ‘Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions’, *Journal of Monetary Economics* 77, 30–47.
- Meier, Stephan and Sprenger, Charles D.** (2015). ‘Temporal stability of time preferences’, *Review of Economics and Statistics* 97(2), 273–286.
- Meyer, Bruce D and Sullivan, James X.** (2008). ‘Changes in the consumption, income, and well-being of single mother headed families’, *The American Economic Review* pp. 2221–2241.
- Mirrlees, James A.** (1971). ‘An exploration in the theory of optimum income taxation’, *The Review of Economic Studies* 38(2), 175–208.
- Moffitt, Robert.** (2006). ‘Welfare work Requirements with Paternalistic Government Preferences’, *The Economic Journal* 116(515), F441–F458.
- Moser, Christian and Olea de Souza e Silva, Pedro.** (2017). ‘Optimal Paternalistic Savings Policies’, *Columbia Business School Research Paper No. 17-51* .
- Mullainathan, Sendhil, Schwartzstein, Joshua and Congdon, William J.** (2012). ‘A Reduced-Form Approach to Behavioral Public Finance’, *Annual Review of Economics* 4, 17–1.
- Mullainathan, Sendhil and Shafir, Eldar.** (2013), *Scarcity: Why having too little means so much*, Macmillan.
- Paserman, M Daniele.** (2008). ‘Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation’, *The Economic Journal* 118(531), 1418–1452.
- Romich, Jennifer L and Weisner, Thomas.** (2000). ‘How families view and use the EITC: Advance payment versus lump sum delivery’, *National Tax Journal* pp. 1245–1265.
- Rothstein, Jesse.** (2010). ‘Is the EITC as good as an NIT? Conditional cash transfers and tax incidence’, *American Economic Journal: Economic Policy* pp. 177–208.
- Rubio, Marco.** (2015), *American Dreams: Restoring Economic Opportunity for Everyone*, Sentinel.

- Ruggles, Steven, Genadek, Katie, Goeken, Ronald, Grover, Josiah and Sobek, Matthew.** (2015), ‘Integrated Public Use Microdata Series: Version 6.0 [dataset]’.
- Ryan, Paul.** (2014), *The Way Forward: Renewing the American Idea*, Hachette UK.
- Saez, Emmanuel.** (2001). ‘Using elasticities to derive optimal income tax rates’, *The Review of Economic Studies* 68(1), 205–229.
- Saez, Emmanuel.** (2002). ‘Optimal income transfer programs: intensive versus extensive labor supply responses’, *The Quarterly Journal of Economics* 117(3), 1039–1073.
- Sandmo, Agnar.** (1975). ‘Optimal taxation in the presence of externalities’, *The Swedish Journal of Economics* pp. 86–98.
- Seade, Jesus.** (1982). ‘On the sign of the optimum marginal income tax’, *The Review of Economic Studies* 49(4), 637–643.
- Seade, Jesus K.** (1977). ‘On the shape of optimal tax schedules’, *Journal of public Economics* 7(2), 203–235.
- Spinnewijn, Johannes.** (2015). ‘Unemployed but optimistic: optimal insurance design with biased beliefs’, *Journal of the European Economic Association* 13(1), 130–167.
- Stantcheva, Stefanie.** (2014). ‘Optimal Taxation and Human Capital Policies over the Life Cycle’, *NBER Working Paper No. 21207*.
- Steuerle, C. Eugene.** (2015). ‘Congressional Testimony: Marginal Tax Rates and 21st Century Social Welfare Reform’.
- Stiglitz, Joseph E.** (1982). ‘Self-selection and Pareto efficient taxation’, *Journal of Public Economics* 17(2), 213–240.
- Vickrey, William.** (1945). ‘Measuring marginal utility by reactions to risk’, *Econometrica: Journal of the Econometric Society* pp. 319–333.
- Wane, Waly.** (2001). ‘The optimal income tax when poverty is a public ‘bad’’, *Journal of Public Economics* 82(2), 271–299.

Weinzierl, Matthew. (2014a). ‘Revisiting the Classical View of Benefit-Based Taxation’, *NBER Working Paper No. 20735* (14-101).

Weinzierl, Matthew C. (2014b). ‘The Promise of Positive Optimal Taxation: Normative Diversity and a Role for Equal Sacrifice’, *Journal of Public Economics* 118, 128–142.

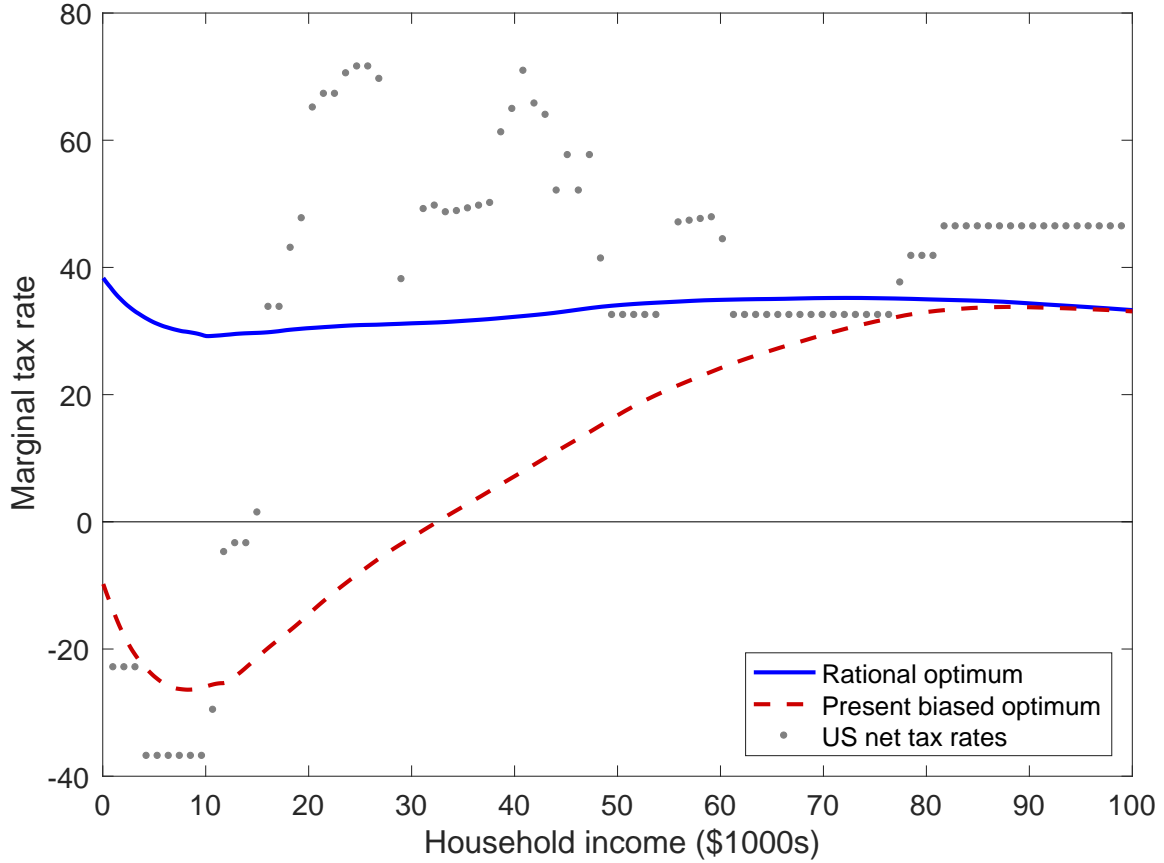
Tables and Figures

Table 1: Empirical Evidence of Present Bias

| | implied β | low income/EITC | effort |
|--|-----------------|-----------------|--------|
| Augenblick, Niederle and Sprenger (2015) | 0.89 | | ✓ |
| Augenblick and Rabin (2015) | 0.83 | | ✓ |
| Kaur, Kremer and Mullainathan (2015) | 0.71 | | ✓ |
| Meier and Sprenger (2015) | 0.69 | ✓ | |
| Laibson, Repetto and Tobacman (2017) | 0.50 | ✓ | |
| Paserman (2008) | 0.65 | ✓ | |
| Fang and Silverman (2009) | 0.34 | ✓ | ✓ |
| DellaVigna et al. (2015) | 0.58 | ✓ | |


This table reports several estimates of the present bias β parameter conditional on income. See Appendix C for details.

Figure 1: Simulated optimal marginal tax rates with and without present bias.



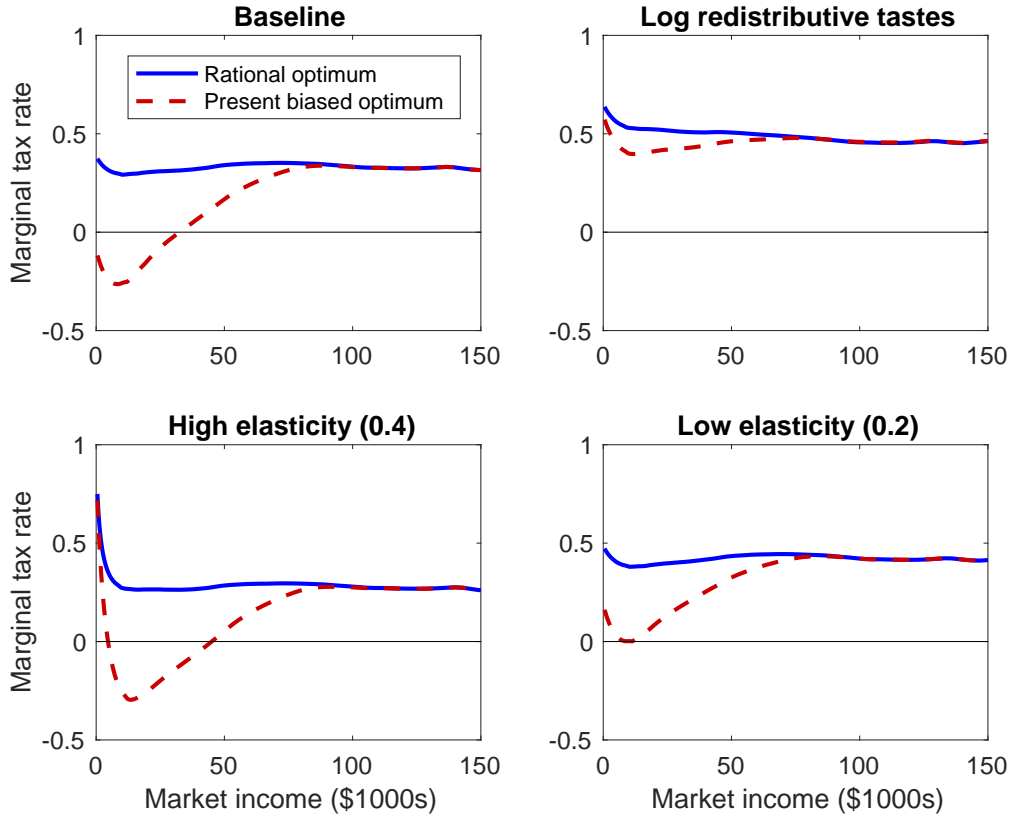
This figure displays simulated optimal tax rates under modest redistributive preferences, with and without accounting for present bias. Net marginal tax rates, including the phaseout of universal benefits, for a representative EITC-receiving household are plotted in gray (see footnote 6 for details). The details of the model economy are discussed in Section 3, where the relationship between redistributive preferences and marginal tax rates is discussed at length.

Figure 2: Estimated relationship between income and present bias parameter β .




This figure plots estimates of β across income from several papers. The dotted “best fit line” is used in simulations for schedule of present bias across the skill distribution. See Appendix C for details.

Figure 3: Optimal income tax simulations.



This figure displays optimal simulated income taxes under four different specifications, varying both redistributive preferences and labor supply elasticities. The baseline uses a labor supply elasticity of 0.3 and modest redistributive preferences (bottom earners have a welfare weight 10% above the median, top earners have a welfare weight of 40% below the median, linearly interpolated). The second panel uses logarithmic redistributive preferences (see text for details), while the bottom two panels use the same modest redistributive preferences with alternative higher and lower values of labor supply elasticities.

Figure 4: Inverse Optimum Illustration



This figure plots the welfare weights implicit in US policy under the conventional assumption of perfect optimization, and under calibrated present bias. Weights are computed by generating a smoothed income density based on 5th-order polynomial regression, then computing weights locally within \$2500 income bins. Lines are generated using kernel regression with a bandwidth of \$10,000.

Appendix (for online publication)

Appendix A Proofs

Proof of Proposition 1

Proof. Let $\mathcal{V}(w)$ denote rescaled utility function as perceived by type w 's self at the time labor supply is chosen:

$$\mathcal{V}(w) := w\ell(w) - T(w\ell(w)) - \frac{v(\ell(w))}{\beta(w)}. \quad (18)$$

(Any exponential discount factor δ can be absorbed by rescaling v .) The individual's optimization implies

$$1 - T'(w\ell(w)) = \frac{v'(\ell(w))}{\beta(w)w}, \quad (19)$$

and so we have

$$\mathcal{V}'(w) = \frac{\ell(w)v'(\ell(w))}{\beta(w)w} + \frac{v(\ell(w))\beta'(w)}{\beta(w)^2}. \quad (20)$$

Normative utility is equal to

$$\mathcal{V}(w) + \left(\frac{1 - \beta(w)}{\beta(w)} \right) v(\ell(w)), \quad (21)$$

and, in a modified version of the standard optimal control setup, we can take $\mathcal{V}(w)$ as the state variable and $\ell(w)$ as the control variable, writing the problem as

$$\max \int_{w_{min}}^{w_{max}} \alpha(w) G \left(\mathcal{V}(w) + \left(\frac{1 - \beta(w)}{\beta(w)} \right) v(\ell(w)) \right) f(w) dw \quad (22)$$

subject to the (appropriately rewritten) budget constraint with required exogenous expenditures E :

$$\int_{w_{min}}^{w_{max}} \left(\mathcal{V}(w) + \frac{v(\ell(w))}{\beta(w)} \right) f(w) dw \leq \int_{w_{min}}^{w_{max}} w\ell(w) f(w) dw - E. \quad (23)$$

and

$$\mathcal{V}'(w) = \frac{\ell(w)v'(\ell(w))}{\beta(w)w} + \frac{v(\ell(w))\beta'(w)}{\beta(w)^2}. \quad (24)$$

Letting λ denote the multiplier on the budget constraint in (23), and letting $m(w)$ denote the

multipliers on the constraint in (24), the Hamiltonian for this problem is

$$\mathcal{H} = \left[\alpha(w)G \left(\mathcal{V}(w) + \left(\frac{1 - \beta(w)}{\beta(w)} \right) v(\ell(w)) \right) - \lambda \left(\mathcal{V}(w) + \frac{v(\ell(w))}{\beta(w)} - w\ell(w) \right) \right] f(w) + m(w) \left(\frac{\ell(w)v'(\ell(w))}{\beta(w)w} + \frac{v(\ell(w))\beta'(w)}{\beta(w)^2} \right). \quad (25)$$

The usual solution technique requires

$$m'(w) = -\frac{\partial \mathcal{H}}{\partial \mathcal{V}} = (\lambda - \alpha(w)G') f(w). \quad (26)$$

Maximizing \mathcal{H} with respect to $\ell(w)$, we have

$$\begin{aligned} & \left(-\alpha(w)G' \times \left(\frac{1 - \beta(w)}{\beta(w)} \right) v'(\ell(w)) + \lambda \left(\frac{v'(\ell(w))}{\beta(w)} - w \right) \right) f(w) \\ & = m(w) \left(\frac{v'(\ell(w)) + \ell(w)v''(\ell(w))}{\beta(w)w} + \frac{v'(\ell(w))\beta'(w)}{\beta(w)^2} \right). \end{aligned} \quad (27)$$

Using the fact that $m(w_{max}) = 0$ (no distortion at the top) we have

$$m(w) = \int_{w_{max}}^w m'(w)dw = - \int_w^{w_{max}} m'(w)dw = \int_w^{w_{max}} (\alpha(w)G' - \lambda) f(w)dw. \quad (28)$$

Substituting into (27) and rearranging yields

$$\frac{T'}{1 - T'} = \frac{1}{f(w)} \int_{x=w}^{w_{max}} \left(1 - \frac{\alpha(x)G'}{\lambda} \right) f(x)dx \left(\frac{1 + \frac{\ell(w)v''(\ell(w))}{v'(\ell(w))}}{w} + \frac{\beta'(w)}{\beta(w)} \right) - \left(\frac{\alpha(w)G'}{\lambda} \right) (1 - \beta(w)). \quad (29)$$

Then substituting in the expressions (from the text) for the elasticities of labor supply and present bias yields the expression in Proposition 1.

The expression for λ is derived by noting that the shadow value of public funds must equal the social welfare generated by a uniform marginal increase in consumption. \square

Proof of Proposition 2

Proof. By assumption, all individuals' earnings satisfy the first-order condition $v'(\ell) = w\beta(w)(1 - T'(z(w)))$ at the optimum. Consider the limit

$$\lim_{w \downarrow w_{min}} \left\{ \frac{1 + 1/\zeta_\ell(w) + \zeta_\beta(w)}{wf(w)} \int_{x=w}^{w_{max}} (1 - g(x)) dF(x) \right\} = \frac{1 + 1/\zeta_\ell(w_{min}) + \zeta_\beta(w_{min})}{w_{min}} \left(\lim_{w \downarrow w_{min}} \frac{1}{f(w)} \int_w^{w_{max}} (1 - g(x)) f(x)dx \right), \quad (30)$$

where equality follows because all constituent functions are continuous in w . If $\lim_{w \downarrow w_{min}} f(w) > 0$ (the limit as w approaches w_{min} from above) then the limit evaluates to zero. If instead

$\lim_{w \downarrow w_{min}} f(w) = 0$, the limit in parentheses can be evaluated, using l'Hôpital's rule, as

$$\lim_{w \downarrow w_{min}} \frac{(g(w) - 1) f(w)}{f'(w)}, \quad (31)$$

which in turn evaluates to zero. Therefore

$$\lim_{w \downarrow w_{min}} \left\{ \frac{T'(z(w))}{1 - T'(z(w))} \right\} = -g(w_{min}) (1 - \beta(w_{min})), \quad (32)$$

which, by assumption that $\beta(w)$ is bounded below 1, implies the marginal tax rate on the lowest earners is negative and bounded away from zero. By continuity of $\frac{T'(z(w))}{1 - T'(z(w))}$ in w , there is thus a range of w sufficiently close to w_{min} which face negative marginal tax rates. \square

Proof of Proposition 4

Proof. Consider a reform to the optimal income tax which slightly raises the marginal tax rate by $d\tau$ in a narrow range of width ϵ around some income level z^* , where the optimal tax is assumed to be continuous and twice differentiable. The first-order effects of this reform can be decomposed into a mechanical effect dM (through raised revenue and a reduction in welfare), a *local* behavioral effect dL through the behavioral responses of individuals who earn z^* , and an *inframarginal* effect dI through the behavioral responses of individuals who earn more than z^* . Each of these terms represents the derivative of social welfare (through the channel in question) with respect to $d\tau$, taking the limit as $\epsilon \rightarrow 0$, normalized by the value of public funds. At the optimum, the sum of these effects must equal zero. (I normalize the size of each effect by the magnitude of the infinitesimal reform, $d\tau\epsilon$, since that term is common to each effect and cancels when the sum is set to zero.)

The mechanical effect, identical to that in Saez (2001), is straightforward:

$$dM = \int_{z^*}^{\infty} (1 - \bar{g}(s)) dH(s). \quad (33)$$

The local behavioral response effect on the intensive margin is composed of a fiscal externality $\mathbb{E} \left[\frac{dz(i)}{dT'(z(i))} \Big| z(i) = z^* \right] h(z^*) T'(z^*)$ and a welfare effect,

$$\mathbb{E} \left[\alpha(i) G'(U_i) \sum_{j=1}^J \frac{\partial U_i}{\partial \ell_j(i)} \frac{\partial \ell_j(i)}{dT'(z(i))} \Big| z(i) = z^* \right] h(z^*), \quad (34)$$

where U_i is normative utility as in (1), evaluated at the optimal choices of consumption and labor supply.

Note that

$$\frac{\partial U_i}{\partial \ell_j(i)} = -\delta^{-\tau(j)} v'_j(\ell_j(i)) + u'(c(i)) w_j(i) (1 - T'(z(i))). \quad (35)$$

From individual optimization, for all j such that $\tau(j) = 0$, (37) is equal to zero. For all other j ,

$$-\delta^{-\tau(j)} v'_j(\ell_j(i)) + \beta(i) u'(c(i)) w_j(i) (1 - T'(z(i))) = 0, \quad (36)$$

implying

$$\frac{\partial U_i}{\partial \ell_j(i)} = (1 - \beta(i))u'(c(i))(1 - T'(z(i))) \quad \text{for } j \text{ s.t. } \tau(j) > 0. \quad (37)$$

Therefore the local behavioral welfare effect in (38) can be rewritten

$$-\mathbb{E} [g(i)(1 - \beta(i))\phi(i)\varepsilon(i)|z(i) = z^*] h(z^*)z^* \quad (38)$$

Combining this welfare effect with the fiscal externality from the local behavioral response yields the total local intensive margin effect:

$$dL = -\bar{\varepsilon}(z^*)z^*h(z^*) \left[\frac{T'(z^*)}{1 - T'(z^*)} + \bar{g}(z^*)\bar{\phi}(z^*)(1 - \bar{\beta}(z^*)) \left(1 + \Sigma_{1-\beta,\varepsilon,g,\phi}^{(z^*)} \right) \right]. \quad (39)$$

There are also inframarginal behavioral responses due to the increased level of taxes for individuals with earnings above z^* , through income effects. This inframarginal effect dI can be decomposed into a fiscal externality component, $\mathbb{E} \left[\frac{dz(i)}{dT(z(i))} \Big| z(i) = z \right] h(z)T'(z)$ for each $z > z^*$, while the welfare effect is

$$\int_{z^*}^{\infty} \mathbb{E} \left[\alpha(i)G'(U_i) \sum_{j=1}^J \frac{\partial U_i}{\partial \ell_j(i)} \frac{\partial \ell_j(i)}{\partial T(z(i))} \Big| z(i) = s \right] dH(s) = \int_{z^*}^{\infty} \mathbb{E} [g(i)\phi(i)\eta(i)(1 - \beta(i))|z^*(i) = s] dH(s). \quad (40)$$

Here we have used the fact that $\sum_{\{j|\tau(j)>0\}} w_j \frac{d\ell_j(i)}{d(1-T')} \Big/ \frac{dz(i)}{d(1-T')} = \sum_{\{j|\tau(j)>0\}} w_j \frac{d\ell_j(i)}{dT} \Big/ \frac{dz(i)}{dT}$: the share of labor response to a tax perturbation which comes through labor choices prior to compensation is the same whether that perturbation concerns marginal tax rates or levels. The proof of this equivalence follows from the first-order condition for ℓ_j^{pb} , which can be written

$$\frac{v'(\ell_j^{pb})}{w_j \delta^{\tau(j)} B(j)} = u'(z - T(z))(1 - T'(z)), \quad (41)$$

where $B(j) = 1$ if $\tau(j) = 0$ and $B(j) = \beta$ if $\tau(j) > 0$. Consider a perturbation to the tax code dT which results in a vector of labor supply adjustments $d\ell_j^{pb}$. Employing (41), these changes satisfy

$$\frac{v''(\ell_j^{pb})}{(w_j \delta^{\tau(j)} B(j))^2} \frac{d\ell_j^{pb}}{dT} = K, \quad (42)$$

where K is the total derivative of the right side of (41) with respect to dT . Rearranging (42) gives

$$w_j \frac{d\ell_j^{pb}}{dT} = K \frac{w_j (w_j \delta^{\tau(j)} B(j))^2}{v''(\ell_j^{pb})} \quad (43)$$

Summing these equations over the j such that $\tau(j) > 0$, divided by the sum across all j , one finds that the term K cancels, so the share ϕ does not depend on the particular marginal source of the tax perturbation.

Combining these effects yields

$$dI = - \int_{z^*}^{\infty} \bar{\eta}(z) \left[\frac{T'(z)}{1-T'(z)} + \bar{g}(z)\bar{\phi}(z)(1-\bar{\beta}(z)) \left(1 + \Sigma_{1-\beta, \eta, g, \phi}^{(z)} \right) \right] dH(z). \quad (44)$$

Using these terms, the first-order condition for the optimal tax policy requires $dM + dL + dI = 0$. Rearranging yields the expression in Proposition 4. \square

Proof of Proposition 5

The proof of Part 1 of the proposition follows immediately from the fact that the choice of earnings dimensions with $\tau(j) > 0$ does not depend on ψ . When $\phi(i) = 1$ for all i , effort is the only determinant of earnings, and the Part 1 of the proposition is implied.

Proof of Part 2. Consider the optimal tax when $\psi = 0$, which is characterized by Proposition 4, and suppose ψ is raised slightly, by $d\psi$. Any individuals with $\phi(i) = 1$ are insensitive to the change and can be ignored. Individuals with $\phi(i) < 1$ and $\beta(i)$ choosing some $\ell_{j'}$, where $\tau(j') = 0$, perceive the tax to be reduced by $d\psi(1-\beta(i))T'(z(i))$, and therefore raise $\ell_{j'}$ by $d\psi(1-\beta(i))\frac{d\ell_{j'}(i)}{d(1-T')}$, which (beginning from $\psi = 0$) has no first-order effect on welfare due to the envelope theorem, yet has a strictly positive fiscal externality, implying that the total first-order effect on social welfare of slightly raising ψ is strictly positive, proving the proposition. \square

Proof of Part 3. Extending the logic in Part 2, consider raising ψ by $d\psi$, but beginning from some $\psi > 0$. If the effect on social welfare of raising ψ remains positive at $\psi = 1$, the proposition is proved. There is still no effect on advance labor effort, so the reform generates a response in earnings through contemporaneous effort alone equal to $dz(i) = d\psi(1-\beta(i))\frac{T'}{1-T'}(1-\phi(i))z(i)\varepsilon(i)$. This behavioral response generates a fiscal externality equal to $dz(i)T'(z(i))$. However when $\psi > 0$, the envelope theorem no longer holds, so there is also a first-order effect on welfare, equal to $-dz(i)g(i)T'(z^*(i))\psi(1-\beta(i))$. Combining the two effects, the total effect from i 's hours response is $dz(i)T'(z(i))(1-\psi g(i)(1-\beta(i)))$, which is nonnegative for $\psi = 1$ if and only if $g(i)(1-\beta(i)) \leq 1$. If that inequality holds for all i , then fully delayed taxes ($\psi = 1$) are optimal, proving the proposition. \square

Appendix B Optimal tax condition with a participation margin

A participation margin can be added to the model in Section 2 by adding heterogeneous fixed costs of work χ_i to the individuals' utility function. Since such costs are conventionally understood as the costs of having a job at the time work is performed (such as transportation and child care) these costs are most naturally modeled as occurring contemporaneously with consumption, so that the normative utility function in Equation (1) is replaced by

$$U(c, \ell) = - \sum_{j=1}^J \delta^{-\tau(j)} v'_j(\ell_j) + u(c) - \chi \cdot 1\{z > 0\}. \quad (45)$$

Maintaining Assumption 1 (accurate income forecasts), the separability of fixed costs from the other terms in the utility function implies that present biased utility satisfies the first-order conditions in (2) if the resulting $u(w \cdot \ell^{pb} - T(w \cdot \ell^{pb})) > \chi$, and otherwise $\ell_j^{pb} = 0$ for all j .

Extensive margin labor supply responses can then be denoted in the conventional way under fixed cost models, with the participation elasticity defined as $\rho(z) = -\frac{dh(z)}{dT(z)} \frac{(z-T(z))+T(0)}{h(z)}$. The derivation of the optimal tax condition in this extension proceeds as in the proof of Proposition 4, except that there is an additional fiscal externality in the inframarginal effect dI , equal to

$$-\int_{s=z}^{\infty} \rho(s) \left(\frac{\bar{T}(s)}{1-\bar{T}(s)} \right) h(s) ds,$$

where $\bar{T}(z) = \frac{T(z)-T(0)}{z}$ denotes the participation tax rate. Since fixed costs are incurred contemporaneously with compensation, there is no present bias distortion on the participation margin, and thus the fiscal externality is the only addition to the optimality condition from 4.

Therefore the optimal tax condition, accounting for a participation margin effect, is as in Proposition 4, except with

$$\mathcal{B}(z) = \int_{s=z}^{\infty} \left[1 - \bar{g}(s) - \bar{\eta}(s) \left(\frac{T'(s)}{1-T'(s)} + \bar{g}(s)\bar{\phi}(s)(1-\bar{\beta}(s)) \left(1 + \Sigma_{1-\beta, \eta, g, \phi}^{(s)} \right) \right) - \rho(s) \left(\frac{\bar{T}(s)}{1-\bar{T}(s)} \right) \right] dH(s). \quad (46)$$

Appendix C Details of Table 1 and Figure 2

Table 1 cites several papers which estimate present bias (or substantial short-run discounting) in contexts which are informative for the calibration of β in Section 3. Figure 2 relates these estimates to incomes, when possible. This appendix discusses these sources, and describes the construction of Table 1 and Figure 2.

Augenblick, Niederle and Sprenger (2015) and Augenblick and Rabin (2015) both analyze laboratory experiments with college students at UC Berkeley, who are asked to make decisions about real effort tasks. Augenblick et al. (2015) presents a lab experiment in which student participants face a fixed amount of effort to be performed within a given period. Individuals without commitment devices exhibit an apparent discount rate of about 11% per week. If they individuals were time consistent (with no discounting) beyond one week, this would suggest a misoptimization wedge of 0.89—this is the estimate reported in Table 1. Augenblick and Rabin (2015) estimates β explicitly from effort-for-money choices at various time horizons. Both studies find evidence of commitment demand, which is correlated with individual-specific measures of present bias. These papers do not study the relationship between β and any measure of income. Figure 2 places both estimates at \$55,535, the annual income for graduates of Berkeley after 10 years,²⁶ according to the US Department of Education’s College Scorecard.²⁷

Kaur, Kremer and Mullainathan (2015) measures the labor supply responses of employees in an Indian data entry center who were exposed to a number of treatments during a year-long experiment. Two findings are of particular interest. First, workers generated more output on payday, with production rising smoothly over the weekly pay cycle as payday approached. This “payday effect” suggests a daily discount factor of about 5%. Their results are inconsistent with a strict $\beta \delta$ model of quasi-hyperbolic discounting, as effort rises smoothly as payday approaches, rather than jumping upward discretely. Because the time horizon in Kaur et al. (2015) is shorter than in Augenblick et al. (2015) or Augenblick and Rabin (2015), they need not be inconsistent, if

²⁶All incomes are converted into 2010 dollars using the CPI-U.

²⁷See <https://collegescorecard.ed.gov/school/?110635-University-of-California-Berkeley>.

individuals have a daily discount rate of 0.05 for the upcoming week, with no discounting thereafter. Interpreted as such, Kaur et al. (2015)’s results suggest a β equal to implied discount factor at a one week horizon, equal to discount factor of $(\frac{1}{1.05})^7 = 0.71$, which is the value reported in Table 1. Like Augenblick and Rabin (2015), Kaur et al. (2015) finds demand for commitment which is correlated with individual-specific present bias. Since comparisons between incomes among Indian data center workers and US EITC recipients are difficult (and since the authors do not report annual earnings) I exclude this estimate from Figure 2.

Meier and Sprenger (2015) presents a field experiment wherein EITC filers in Boston are given choices between intertemporal tradeoffs between monetary payments at different horizons. Subjects exhibit greater impatience at shorter horizons, suggestive of present bias. The estimated β for the full sample is 0.69; that is the value reported in 1. Monetary tradeoffs (as opposed to effort tradeoffs) may generate upward-biased estimates of β (understating the degree of present bias) if payments are not immediately converted into consumption (i.e., in the presence of saving or borrowing). If individuals are liquidity constrained, however, monetary payments may be consumed promptly, consistent with this paper’s substantial measured present bias in this study. Moreover, an advantage of Meier and Sprenger (2015), relative to the preceding experimental studies, is that it studies precisely the population of interest for understanding the implications of present biased behavior for low-income work subsidies: low income EITC recipients in the US. This is one of the few studies which reports the covariation of β with income—the paper finds a strong positive correlation, with β rising by about 0.05 for every \$10,000 of income.²⁸ To generate an approximate plot of β estimates across income for 2, I plot the overall average estimate of β (0.69) at the sample’s mean income of \$16,603. In addition, I plot incomes approximately one standard deviation above and below the mean income (where I use \$14,000 to approximate the standard deviation, see Table 1 of that paper), with corresponding values of β computed using their linear best fit estimate of 0.05 per \$10,000 of income.

Laibson et al. (2017) uses the method of simulated moments to perform a calibration using data on income, wealth, and credit use. They report β computed separately for three partitions of education: those who did not finish high school, those who completed high school but not college, and those who completed college (with β values of 0.40, 0.51, and 0.74, respectively). The Bureau of Labor Statistics reports average weekly incomes within each of these education bins,²⁹ corresponding to annual incomes of \$23,939, \$32,868, and \$54,907, respectively. These income values are used to plot the points for Laibson et al. (2017) in Figure 2.

Paserman (2008) estimates a structural model of job search with quasi-hyperbolic preferences using the NLSY, extending the approach of DellaVigna and Paserman (2005) which finds strong evidence of present biased search behavior. The paper reports β estimated separately for three partitions of the wage distribution: the bottom quartile of the wage distribution, the middle half, and the top quartile, with estimates of 0.40, 0.65 (averaging the lognormal and normal specifications) and 0.89, respectively. I convert the mean re-employment weekly wage for each group into annual incomes of \$20,822, \$30,717, and \$53,409; these are the income values used to plot the points in Figure 2.

Fang and Silverman (2009) calibrates a quasi-hyperbolic model of model of welfare takeup and labor supply using data from the NLSY. The resulting estimate of β is 0.34, and the sample has an average income \$22,179.

Finally, DellaVigna and Paserman (2005) estimates a model of reference-dependent job search

²⁸One possible confound is that higher income EITC recipients may be less liquidity constrained, and may therefore exhibit less present bias over money payments. This possibility points to the value of further effort-based present bias experiments on populations with heterogeneous incomes.

²⁹See https://www.bls.gov/emp/ep_table_001.htm.

effort using Hungarian administrative data. The paper finds that admitting quasi-hyperbolic preferences (with $\beta < 1$) improves the fit significantly. The reported best-fit estimate of β is 0.58; this value is reported in Table 1.

Appendix D Calibration of income distribution and tax schedule

For the simulations in Section 3.2, I draw the income from the the 2010 Current Population Survey, restricted to households with positive total income, and I use kernel density estimation to calibrate the density across incomes.³⁰ I assume present bias is skill-specific, with the profile plotted in Figure 2. The first-order condition for effort choice can then be inverted to compute the implicit skill distribution. The first-order condition depends on the individual’s marginal tax rate, which is estimated from CPS and NBER’s TAXSIM. Specifically, I use TAXSIM’s estimated net federal marginal tax rate, including employer and employee portions of payroll taxes, based on wage income, number of dependents, marital status, and age. I average this value across individuals at each level of income, and I construct an approximate implicit marginal tax rate from the phaseout of benefits using CPS data by performing a kernel regression of the value of food stamps and welfare income on market income, then differentiating the resulting schedule. I use a bandwidth of \$2000 for the computation of marginal tax rates, and \$5000 for the density estimation, where a greater degree of smoothing is useful for generating smooth schedules of simulated optimal tax rates.

The inverse optimum exercise in Section 3.3 restricts to a smaller sub-population, as it focuses specifically on families who receive the EITC (or who could receive it if their income declined). For this exercise I use the most recent year for which IPUMS has complete data and TAXSIM can compute marginal tax rates (2015). (All figures are reported in 2010 dollars, adjusted using the CPI-U.) I restrict to households for whom the respondent is the head of household (and nonfamily householders), where the respondent is between the ages of 25 and 55. I further and I restrict to households with two children and with positive total family income. A continuous income distribution is constructed by discretizing the income space into \$2500 bins and using a fifth order polynomial regression on the number of households in each bin to generate a smooth density with a continuously differentiable derivative. The schedule of marginal tax rates is drawn from the National Bureau of Economic Research’s TAXSIM model. To compute the marginal tax rate at each point in the income distribution, I submit data on year, filing status, and the number and age of dependent children to TAXSIM, which provides an effective marginal tax rate on additional earnings, accounting for credits and deductions. I include the marginal tax rate from payroll taxes (both the employer and employee portions). I then compute approximate implicit marginal tax rates from the phaseout of benefits by performing a local polynomial regression of benefits (the sum of food stamps, housing subsidies, heat and energy subsidies, and other welfare income) on a measure of market income (total family income less any social security, unemployment, and welfare income). The local derivative is interpreted as the implicit marginal tax rate from phaseouts, which is added to the marginal tax rate from TAXSIM. I then average these marginal tax rates within each \$1000 bin, and use these averages to compute marginal social welfare weights.

³⁰All data comes from University of Minnesota’s IPUMS database (Ruggles et al., 2015).