

Inattention and the Taxation Bias*

Jérémy Boccanfuso[†] Antoine Ferey[‡]

August 9, 2019

Abstract

This paper studies how information frictions in tax perceptions interact with the design of actual tax policy. It develops a positive theory of tax policy in which taxpayers endogenously choose their attention to taxes and the government sets tax policy taking taxpayers' attention choices into account. In equilibrium, we show that the government implements inefficiently high tax rates because it faces a commitment problem. This *taxation bias* directly depends on agents' attention and sheds a new light on the implications of tax misperceptions. Most notably, we find that the underestimation of tax rates is not welfare-improving if taxpayers are not sufficiently attentive to tax policy. Further, we highlight that the positive correlation between income and attention reduces the progressivity of actual tax schedules. Overall, our findings suggest that by inducing information frictions, tax complexity may lead to undesirable and regressive tax increases.

Keywords: Optimal taxation; inattention. JEL classification code: H21; D03.

*We are grateful to Pierre Boyer, Philippe Choné, Allan Drazen, Xavier Gabaix, Laurence Jacquet, Louis Kaplow, Etienne Lehmann, Benjamin Lockwood, Jean-Baptiste Michau, Xavier Ragot, Anasuya Raj, Daniel Reck, Alessandro Riboni, Emmanuel Saez, Stefanie Stantcheva, Dmitry Taubinsky, Clémence Tricaud, Nicolas Werquin, Mirko Wiederholt and audiences at AFSE 2018 Conference, UC Berkeley, CREST, CRED Taxation Group, LAGV 2018 Conference, NTA 2018 Conference and PET 2017 Conference for very helpful discussions and feedback. Any remaining errors are our own.

[†]PSE – EHESS, 48 Boulevard Jourdan, 75014 Paris, France. E-mail: jeremy.boccanfuso@psemail.eu

[‡]École Polytechnique – CREST, Route de Saclay, 91128 Palaiseau, France. E-mail: antoine.ferey@polytechnique.edu

1 Introduction

Tax systems are notoriously complex and hard to understand. Accordingly, a growing body of evidence documents substantial information frictions in agents' tax perceptions (Chetty, 2015; Bernheim and Taubinsky, 2018). In particular, taxpayers tend to misunderstand some characteristics of income tax schedules (Saez, 2010; Aghion et al., 2017), to rely on linearizing heuristics such as ironing¹ (Liebman and Zeckhauser, 2004; Rees-Jones and Taubinsky, 2016) and to partially ignore non-salient taxes and transfers (Chetty et al., 2009; Miller et al., 2015; Taubinsky and Rees-Jones, 2017). These findings indicate that taxpayers are not fully attentive to taxes and suggest the existence of systematic behavioral biases and misperceptions of tax rates.

In light of this evidence, a burgeoning normative literature analyzes how information frictions in tax perceptions affect the design of optimal tax policy. Goldin (2015) shows that a government may implement non-salient taxes to reduce the deadweight loss of taxation. Gerritsen (2016) highlights that tax misperceptions introduce a new corrective motive for taxation and derives adjusted optimal nonlinear tax formulas. Farhi and Gabaix (2018) provide a general treatment of this optimal tax literature with behavioral agents while stressing that "a difficulty confronting all behavioral policy approaches is a form of Lucas critique: how do the underlying biases change with policy?" (p.15).

So far, the dominant approach to address this difficulty has been to account for the endogeneity of taxpayers' attention to tax policies (Goldin, 2015; Reck, 2016; Farhi and Gabaix, 2018). Indeed, attention is key to determine taxpayers' responses to tax reforms and recent empirical evidence supports the idea that it results from a cost-benefit arbitrage. In an online shopping experiment, Taubinsky and Rees-Jones (2017) show that tripling the tax rate nearly doubles agents' attention to taxes. Using online search data about taxes, Hoopes et al. (2015) find that rational inattention motives and shocks to tax salience drive taxpayers' information search. Attention endogeneity is thus an important mechanism through which tax perceptions may adjust to changes in tax policy but it may not be the only one. As demonstrated in this paper, other indirect adjustment mechanisms conceal significant implications for tax policy.

This paper investigates how information frictions in tax perceptions affect, and in fact interact with, the design of actual tax policy. It develops a positive theory of tax policy in which taxpayers endogenously choose their attention to taxes and the government sets the tax policy to maximize welfare taking attention choices into account. We show

¹That is, agents linearize income tax schedules using their average tax rate rather than their marginal tax rate.

that information frictions generate a distortion in actual tax policy: the government implements inefficiently high tax rates. To explore the magnitude of this distortion we provide numerical simulations and a simple sufficient statistics formula that we bring to the data.

Our theoretical model contributes to the existing literature in several dimensions. First, we rely on a general model of tax perceptions featuring both inattention and systematic behavioral biases. We consider a population of heterogeneous and rationally inattentive agents who cannot freely observe the tax schedule and thus choose their earnings and consumption given their tax perceptions. We model agents' tax perceptions as resulting from a Bayesian learning model with a choice of costly information acquisition (Mackowiak et al., 2018; Gabaix, 2019). That is, taxpayers are endowed with a prior (or belief) about the tax policy and can acquire costly information in the form of a signal. The precision of this signal is endogenous and reflects attention to taxes: the more attentive a taxpayer is, the more precise her signal and the more accurate the posterior (or perceived tax). Importantly, we allow the prior to be systematically biased to capture potential perception biases thereby making a bridge between behavioral models with ad-hoc misperceptions and standard rational inattention frameworks. Hence, this model can for instance capture the use of biased rule-of-thumbs as default while allowing taxpayers to acquire additional information to improve their tax perceptions (Morrison and Taubinsky, 2019).

Second, we depart from the normative optimal tax literature by adopting a positive approach and analyzing the conduct of actual tax policy. We develop a positive theory of tax policy that we formalize as a simultaneous game in which agents choose their attention to taxes and the government sets the tax policy to maximize welfare taking attention choices into account. In equilibrium, (i) neither taxpayers nor the government has an incentive to deviate and (ii) taxpayers' actions and perceptions are mutually consistent with the government's choice of tax policy. The main contribution of the paper is to demonstrate that – irrespective of potential perception biases – inattention to taxes leads to inefficiently high equilibrium tax rates from a normative perspective: this is the taxation bias.

Third and central to this result, our framework exposes an important dichotomy between direct and indirect adjustments in tax perceptions following changes in tax policy. When they are partially attentive, taxpayers anchor their perception on their prior (or belief). Consequently, there are two margins through which perceptions may adjust: a *direct* margin capturing the attention agents devote to observing taxes and thus changes in tax policies, and an *indirect* margin capturing variations in the anchor. The empirical evidence supports the existence of this dichotomy. Sausgruber and Tyran (2005) and

Fochmann and Weimann (2013) show that, with time and experience, taxpayers tend to internalize the impact of non-salient taxes they initially ignored.² Moreover, if taxpayers act upon their perceptions this dichotomy should also be reflected in earnings choices. In effect, Chetty et al. (2011) document a systematic difference between micro – capturing direct adjustments – and macro – capturing total adjustments – estimates of the elasticity of taxable income. Chetty (2012) rationalizes this difference by the existence of adjustment frictions at the micro level and while he remains agnostic about their origin, information frictions are arguably a key element.

This dichotomy between direct and indirect adjustments induces equilibrium distortions in tax policy. Indeed, inattentive agents only perceive a fraction of tax changes. Their direct behavioral response following a tax reform is thus attenuated and the government targets a higher tax rate than it would otherwise have. Doing so, the government fails to internalize the indirect adjustment of perceptions (arising as an equilibrium mechanism) and implements inefficiently high tax rates. In other words, taxpayers' inattention to taxes creates the illusion that tax reforms induce lower efficiency costs than they actually do and ultimately prompts the government to misbehave from a normative perspective.

Fundamentally, this taxation bias reflects a commitment problem. By implicitly restricting the set of tax policies to precommitted policy rules, the aforementioned normative literature characterizes the tax policy under commitment. This commitment tax policy is by definition the optimal tax policy in the presence of information frictions. However, a side effect of information frictions is that a discretionary policymaker cannot credibly implement this policy. Indeed, starting from the commitment equilibrium, a discretionary government is systematically willing to implement a tax increase in order to leverage taxpayers' inattention and is thereby generating a taxation bias. We formally define the taxation bias as the difference in the equilibrium tax rates under discretion and commitment and establish the existence of a positive taxation bias under a mild general requirement.

We illustrate our theoretical results and explore their quantitative implications in a number of ways. In an attempt to gauge the empirical magnitude of tax policy distortions

²Indirect adjustments may also help explain why, when directly asked about their marginal tax rate, surveys of taxpayers' perceptions do not provide clear cut evidence of systematic underestimations – a finding that is *prima facie* hard to reconcile with e.g. ironing behaviors and low salience. Fochmann et al. (2010) review early surveys and show that one third of the studies concludes to the overestimation of marginal tax rates, one third to their underestimation and the rest to misperceptions without systematic over or under-estimation. Recently, Gideon (2017) also finds that reported marginal tax rates are accurate at the mean.

in the US, we derive a simple sufficient statistics formula for the taxation bias assuming that the income tax is linear. Beyond the elasticity of earnings with respect to changes in the perceived marginal net-of-tax rate, it shows that a key sufficient statistics is the income-weighted average attention in the population. This statistics captures the positive correlation between income and attention documented in Taubinsky and Rees-Jones (2017) and that emerges naturally in our model. Relying on the existing empirical literature to calibrate our sufficient statistics, we estimate that the taxation bias is equal to 3.66 percentage points in the actual US economy. This means that the linearized US income tax rate (29.46 percentage points) is more than 12% higher than what would be optimal holding the government’s objective constant.

Parametrizing our tax perceptions model with Gaussian distributions, we provide further theoretical results and numerical simulations. They indicate that even small information rigidities induce significant deviations in tax policy. Moreover, they allow to disentangle the implications of systematic perception biases (reflected in agents’ priors) and that of inattention to taxes. The commitment tax policy is mostly driven by equilibrium perception biases (reflected in agents’ posteriors). That is, the deviation in the commitment tax policy (Farhi and Gabaix, 2018) from a benchmark without information frictions (Saez, 2001) increases with the bias in the prior and decreases with attention. While the discretionary tax policy is similarly impacted by equilibrium perception biases, it also directly depends on taxpayers’ attention. The taxation bias, which measures the difference in the equilibrium tax rates under discretion and commitment, is thus primarily shaped by attention to taxes and relatively less by potential perception biases.

This conclusion holds important welfare implications. Situations in which information rigidities were previously thought to be welfare improving may actually turn out to be welfare decreasing once we account for the policy distortions induced by inattention. To illustrate this point, we carry out a welfare analysis in an economy where taxpayers’ priors systematically underestimate tax rates by 5 percentage points (e.g. salience bias). Unsurprisingly, we find that information frictions induce a welfare gain if the government was to implement the commitment tax policy. However, this policy cannot be credibly implemented by the government and actual tax policy features an additional welfare loss due to the taxation bias. As inattention grows, the welfare loss from policy distortions increases faster than the welfare gain from tax underestimation. Therefore, even if agents underestimate tax rates, information frictions can be ultimately detrimental to welfare when agents are not sufficiently attentive to tax policy.

Extending our analysis to nonlinear tax schedules, we show that the government is then able to leverage the positive correlation between taxpayers’ attention and income. That is, the taxation bias becomes income-specific and globally decreasing with income:

it is relatively large at low income levels and virtually nonexistent at top income levels. The taxation bias thus attenuates the U-shape pattern of marginal tax rates (Saez, 2001) and reduces the progressivity of actual income tax schedules. Hence, inattention to taxes leads to regressive tax increases which may have important redistributive implications.³

A broad lesson emerges from our analysis: inattention creates incentives to use discretionary policies. It is well-known that discretionary policies might lead to inefficient outcomes (Kydland and Prescott, 1977), thus justifying our use of the term *taxation bias* in analogy to the *inflation bias* (Barro and Gordon, 1983). A large body of evidence documents the existence of information frictions affecting e.g. consumers, firms or professional forecasters (Coibion and Gorodnichenko, 2015). Hence, our analysis suggests that policy distortions may exist in a wide variety of settings in which the portable framework developed in this paper could be fruitfully applied.

Policymaking is, at least to some extent, discretionary. In the realm of taxation, discretion is usually discussed in the context of capital levies in which there is indisputable historical evidence of discretionary policies (e.g. Japan post WWII, Italy in 1992, Cyprus in 2013). While less salient for income taxes, discretionary behaviors are likely reflected in the excessive complexity of existing tax schedules.⁴ As a result, it should not come as a surprise that individuals strongly oppose tax complexity, even after acknowledging the potential advantages of differential tax treatments (Blesse et al., 2019). Indeed, our findings suggest that by inducing information frictions, tax complexity may lead to undesirable and regressive tax increases.

The rest of the paper is organized as follows. To build up the intuition, we first characterize the taxation bias in a stylized representative agent model with exogenous attention (Section 2). Section 3 microfounds the behavior of heterogeneous and rationally inattentive taxpayers. In Section 4 we formalize our positive theory of tax policy and establish the existence of a taxation bias. In Section 5 we illustrate our theoretical results with numerical simulations and we derive a simple sufficient statistics formula for the taxation bias that we take to the data. Section 6 turns to the welfare implications of information frictions in tax perceptions and Section 7 provides an extension to nonlinear tax schedules. The last section concludes.

³This has also potentially important implications for the inverse optimum approach which aims at inferring the government's redistributive weights from the shape of actual tax schedules.

⁴The French constitutional court has several times repealed specific items of tax bills for their "excessive complexity" arguing that they would not be understood by taxpayers (Conseil Constitutionnel, 2005, 2012)

2 Taxation bias in a stylized model

This section derives the main result of the paper in a simple representative agent model with exogenous attention. Proofs and derivations are relegated to Appendix A.1.

Consider a canonical labor income taxation model where the government sets a linear tax rate τ to maximize tax revenue. Let $Y(1 - \tau)$ be the aggregate earnings function. The tax revenue function $\tau Y(1 - \tau)$ has an inverted U-shape and is nil when τ is equal to 0 or 100%. As is well-known (e.g. Piketty and Saez (2013)), the revenue maximizing tax rate follows an inverse elasticity rule and is equal to

$$\tau^r = \frac{1}{1 + e} \quad (1)$$

where e is the elasticity of aggregate earnings with respect to the net-of-tax rate.

Assume now that because of information frictions taxpayers are unable to perfectly observe the tax rate. They must nonetheless form an estimate of the latter to decide how much to work. Call this estimate the perceived tax rate $\tilde{\tau}$ and suppose it is determined by a convex combination of a common prior $\hat{\tau}$ and the actual tax rate τ

$$\tilde{\tau} = \xi\tau + (1 - \xi)\hat{\tau} \quad (2)$$

where the weight $\xi \in [0, 1]$ can be interpreted as a measure of taxpayers' attention to the actual tax rate τ .⁵ Indeed, when $\xi = 1$, taxpayers perfectly observe changes in the tax rate whereas, when $\xi = 0$, they are completely inattentive to tax changes and fully anchor their perception on their prior. Since individual earnings choices depend on their perceived tax rate, aggregate earnings now write $Y(1 - \tilde{\tau})$. The tax revenue function becomes $\tau Y(1 - \tilde{\tau})$ which remains concave with respect to the actual tax rate τ given the tax perception model in equation (2).

How do information frictions interact with the design of actual tax policy? Consider a situation in which taxpayers expect the government to implement the optimal tax rate in the absence of information frictions, that is $\hat{\tau} = \tau^r$. Now, suppose the government sets $\tau = \tau^r$ and consider the effect of a policy deviation that consists in a small increase in the tax rate $d\tau$. This mechanically increases tax revenues by $M = Y(1 - \tilde{\tau})|_{\tilde{\tau}=\tau^r} d\tau$ while it generates a behavioral response $dY = -\frac{\partial Y(1 - \tilde{\tau})}{\partial 1 - \tilde{\tau}}|_{\tilde{\tau}=\tau^r} \xi d\tau$ as inattentive taxpayers only observe a fraction ξ of the increase in the tax rate $d\tau$. By definition of τ^r , the mechanical effect M outweighs the fiscal externality $FE = \tau^r dY$ induced by the behavioral response when agents are not fully attentive ($\xi < 1$). As a result, the government systematically deviates from tax policy τ^r and ends up choosing a higher tax rate.

⁵In a recent Handbook chapter, Gabaix (2019) argues this is a unifying framework to model various behavioral biases and attention theories.

Conceptually, an important consequence of inattention is to anchor taxpayers' perceptions on their prior. Because of this anchoring, the government has an incentive to implement policy deviations that taxpayers are going to partially ignore. This is a form of discretionary policy which arises as a side-effect of information frictions in tax perceptions. The government thus chooses its tax policy taking agents' priors and attention into account. Specifically, tax policy $\tau(\xi, \hat{\tau})$ is decreasing in taxpayers' prior $\hat{\tau}$ and attention parameter ξ

$$\tau(\xi, \hat{\tau}) = \begin{cases} \frac{1-(1-\xi)\hat{\tau}}{\xi(1+e)} & \text{if } \hat{\tau} \geq 1 - \frac{\xi}{1-\xi}e \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

where the elasticity of aggregate earnings is defined with respect to the perceived net-of-tax rate $e \equiv \frac{1-\hat{\tau}}{Y} \frac{\partial Y}{\partial(1-\hat{\tau})} \geq 0$. The solution is interior whenever attention ξ or the prior $\hat{\tau}$ are high enough – otherwise the government finds it optimal to impose a 100% tax – and coincides with the inverse elasticity rule when agents are perfectly attentive ($\xi = 1$) and thus fully informed.

Figure 1 plots tax policy $\tau(\xi, \hat{\tau})$ as a function of agents' prior for different attention levels. It shows that small information frictions generate notable deviations in the government behavior. If agents' prior is that the government implements the inverse elasticity rule – $\tau^r = 75\%$ assuming $e = 0.33$ – the government chooses a tax rate of 82% (resp. 77%) when the attention parameter ξ is equal to 0.75 (resp. 0.90).⁶ This corresponds to point A (resp. B).


An equilibrium is as a situation in which (i) neither taxpayers nor the government has an incentive to deviate and (ii) taxpayers' actions and perceptions are mutually consistent with the government's choice of tax policy. We here focus on rational equilibria in which agents correctly anticipate the equilibrium tax policy $\hat{\tau} = \tau^{\text{eq}}$ and defer the introduction of biased equilibria to Section 4. Plugging this equilibrium condition in the government's choice of tax policy (3) we obtain that the equilibrium tax policy is

$$\tau^{\text{eq}} = \frac{1}{1 + \xi e} \quad (4)$$

Graphically, the rational equilibrium is represented by the point where the 45-degree line ($\hat{\tau} = \tau^{\text{eq}}$) intersects the government policy function. Again, small information rigidities lead to large deviations in equilibrium. In a rational equilibrium, the tax rate is 80%

⁶As an element of comparison, Gabaix (2019) states (p. 5) that "on average, the attention parameter estimated in the literature is 0.44, roughly halfway between no attention and full attention" while adding that "attention is higher when the incentives to pay attention are stronger" which should likely be the case when it comes to taxing 75% of one's income.

Figure 1: Optimal tax policy and equilibrium outcomes



NOTE: Optimal policy as a function of the prior $\hat{\tau}$ for different values of the attention parameter ξ . The elasticity of aggregate earnings with respect to the perceived net-of-tax rate is set to 0.33.

(resp. 77%) when attention is such that 75% (resp. 90%) of a tax reform is directly internalized by taxpayers. This corresponds to point C (resp. D).

The government is unable to reach the top of the Laffer curve: the equilibrium tax rate is inefficiently high. We refer to this phenomenon as the *taxation bias* in analogy to the *inflation bias* (Barro and Gordon, 1983). While the government internalizes the direct impact of its choice on agents' perceptions (proportional to attention ξ), it does not internalize the equilibrium impact associated with the adjustment of the prior (proportional to inattention $1 - \xi$). In Barro and Gordon's (1983) words, "the equality of policy expectations and realizations is a characteristic of equilibrium – not a prior constraint" (p. 591). Hence, inattention resurges the possibility that discretionary policies lead to inefficient outcomes (Kydland and Prescott, 1977).

To formalize this result, we characterize the optimal policy under commitment τ^* . This is the optimal policy of a government who can credibly commit to implement a tax level and thereby has to take into account the equilibrium effect of its choice of tax policy on perceptions. By definition, this is the optimal policy in the presence of information frictions and it coincides in this stylized rational equilibrium framework with the inverse elasticity rule $\tau^* = \tau^r$ (point E). However and as shown above, it cannot be an equilibrium

policy under discretion in the presence of inattention.

Defining the taxation bias as the difference between the tax rates under discretion and commitment we have

$$\tau^{\text{eq}} - \tau^* = \frac{(1 - \xi)e}{(1 + \xi e)(1 + e)} \geq 0 \quad (5)$$

Therefore, the taxation bias is strictly positive when agents are not fully attentive to taxes ($\xi < 1$). Moreover, the (absolute) size of the taxation bias increases with agents inattention $1 - \xi$ and with the elasticity e as they intuitively both make policy deviations relatively more attractive.

To highlight the mechanisms that lead to a taxation bias, we have analyzed in this section a stylized representative agent model in which agents' behavior and attention are exogenously given and in which the government implements a linear tax policy to maximize tax revenue. In the remainder of the paper we broaden the scope of the analysis by studying the problem of a welfare-maximizing government facing a heterogeneous population of agents whose individual behavior is fully micro-founded and whose attention is endogenous. We also extend the equilibrium concept to allow for perception biases in equilibrium and examine how the taxation bias affects (the progressivity of) non-linear tax schedules. These extensions provide valuable insights on the magnitude and implications of the taxation bias in a policy relevant environment.

3 Agents' behavior, perceptions and attention

This section describes the behavior of taxpayers in the economy. Because of information frictions, taxpayers may not freely observe the tax rate implemented by the government. They rely on a Bayesian learning model with costly information acquisition to form their perceptions about the tax schedule in order to decide how much to earn and consume.

3.1 Primitives and assumptions

We consider a population of agents with heterogeneous productivities w which are private information and distributed from a well-defined probability distribution function $f_w(w)$. We assume taxpayers have a utility function $U(c, y; w)$ where c is consumption and y earnings and where we impose $U(\cdot)$ to be continuously differentiable, increasing in consumption ($U_c > 0$), decreasing in effort ($U_y < 0$ and $U_w > 0$) and such that the Spence-Mirrlees condition holds (MRS_{yc} decreases with skill w). For simplicity, we consider a separable and quasi-linear utility $U(c, y; w) = c - v(y; w)$ in the body of the paper

and show in the Online Appendix how we can extend the analysis to more general utility functions.⁷

Agents choose their consumption c and earnings y subject to an income tax $T(y)$. Because of information frictions, we assume that taxpayers are unable to freely observe $T(y)$ and instead rely on individual-specific perceived income tax schedules denoted $\tilde{T}(y)$.

Assumption 1 (linear representation). *Individuals use a linear representation of the tax schedule $\tilde{T}(y) = \tilde{\tau}y - \tilde{R}$*

To make their consumption and earnings choices, individuals rely on their perceptions of the tax liability at each earnings level. Assumption 1 imposes that taxpayers use a linear representation of the tax schedule. Hence, agents only need to form estimates of the marginal tax rate $\tilde{\tau}$ and the intercept \tilde{R} thereby reducing the dimensionality of the perceptions formation problem to two parameters.⁸

In most of the paper, we consider that the actual tax schedule is also linear and denote by (τ_0, R_0) its slope and intercept. Consequently, we define (τ, R) as the associated random variables from the point of view of the agents. In Section 7 we extend the analysis to non-linear tax schedules.

3.2 Individual problem

Individuals jointly choose an allocation (c, y) and how much information to collect about the tax schedule. This one-step problem is equivalent to a two-step problem that we characterize. The first step identifies the optimal allocation choice given a perceived tax schedule $\tilde{T}(y)$ while the second step determines the optimal information acquisition taking into account how perceptions affect allocations.

Allocation choice Agents choose consumption c and earnings y to maximize their utility subject to their perceived budget constraint which depends on their perceptions of the tax schedule. This problem writes

$$\begin{aligned} \max_{c, y} \quad & \int_{\tau} U(c, y; w) \tilde{q}(\tau) d\tau \\ \text{s.t.} \quad & c \leq R + (1 - \tau)y \end{aligned} \tag{6}$$

⁷Separability between earnings and consumption preferences combined with quasi-linearity guarantees the absence of income effects in labor supply decisions and considerably simplifies the analysis.

⁸Beyond the fact that a linear approximation is usually a good approximation of existing tax schedules (Piketty and Saez, 2013), recent empirical evidence suggests that in practice taxpayers tend to use linear representations of tax schedules (Rees-Jones and Taubinsky, 2016).

where $\tilde{q}(\tau)$ is the perceived probability distribution of the marginal tax rate τ . With a separable and quasi-linear utility function, the first-order condition determining earnings writes

$$\frac{\partial v(y; w)}{\partial y} = 1 - \tilde{\tau} \quad (7)$$

with $\tilde{\tau} \equiv E_{\tilde{q}(\tau)}[\tau]$ the average perceived marginal tax rate. Consequently, the average perceived marginal tax rate $\tilde{\tau}$ is a sufficient statistics for labor supply and uniquely pins down optimal earnings $y^*(\tilde{\tau}; w)$. Hence, a direct implication of quasi-linear separable preferences is that tax liability, and in particular the perceived value of the demogrant \tilde{R} , is irrelevant for labor supply and only matters to determine agents' consumption levels.

Assumption 2 (slack budget). *Consumption adjusts such that agents exhausts their true budget i.e. $c^*(\tilde{\tau}; w) = R_0 + (1 - \tau_0)y^*(\tilde{\tau}; w)$*

We assume consumption adjusts to ensure that the true budget constraint holds ex post.⁹ The only parameter of interest for agents' allocation choice is thus the perceived marginal tax rate $\tilde{\tau}$.

Given this allocation choice, an agent's indirect utility is

$$V(\tilde{\tau}, \tau_0, R_0; w) = R_0 + (1 - \tau_0)y^*(\tilde{\tau}; w) - v(y^*(\tilde{\tau}; w); w) \quad (8)$$


Figure 2 summarizes the allocation choice in a y - c diagram. Perceptions of the tax schedule determine earnings (tangency condition with perceived budget line) while consumption adjusts to the true budget constraint (intersection with true budget line).

A natural observation from Figure 2 is that misperceptions induce utility misoptimization costs: the utility level \bar{U} associated with the choice under accurate perceptions (black dot) is higher than the utility level U^* associated with the choice under misperceptions (grey dot). However, when agents underestimate tax rates ($\tilde{\tau} < \tau_0$) misperceptions also induce efficiency gains: earnings y^* chosen under misperceptions are larger than earnings \bar{y} at the optimal allocation. As a result, tax underestimation may increase social welfare if efficiency gains dominate utility misoptimization costs.

Perceptions formation Tax perceptions here follow from a Bayesian learning model with a choice of information (in Gabaix's (2019) terminology). They result from the combination of an exogenous and free prior (also referred to as a belief or an anchor) and an endogenous and costly information acquisition process. We choose this model for its

⁹This assumption is used throughout the behavioral tax literature as emphasized in Reck (2016) who discusses different budget adjustment rules in misperception models. See also Farhi and Gabaix (2018) for a related discussion.

Figure 2: Allocation choice in a y - c diagram



NOTE: The figure displays an example of allocation choice when the agent underestimates the marginal tax rate $\tilde{\tau} < \tau_0$. The grey (resp. black) dot represents the allocation choice of an individual who misperceives (resp. correctly perceives) the tax schedule. The plain lines are the indifference curves and the dashed lines the budget constraints.

wide use in economics, its well-understood micro-foundations and the fact that – as we show – it generates predictions that are consistent with the empirical evidence.

Let $\hat{q}(\tau)$ be the prior probability distribution about the tax rate and $\hat{\tau} \equiv E_{\hat{q}}[\tau]$ the expected tax rate derived from the prior. This probability distribution accounts for sources of structural and subjective uncertainty which may be related to policy primitives (e.g. hidden tastes for redistribution), economic fundamentals (e.g. shocks to the government expenditure requirements), institutions (e.g. inability to implement a chosen policy), heuristic decision rules (e.g. ironing), etc.

In the following, we voluntarily remain agnostic about the origin of the prior and the sources of uncertainty it may capture for two reasons. First, the assumed ex ante uncertainty essentially represents a motive for taxpayers to learn in our setup and the main results of the paper will hold for a wide variety of well-defined smooth priors. Second, while the empirical literature clearly indicates that taxpayers tend to misperceive tax rates, there is yet no consensus on the exact rationale – or rationales – behind these misperceptions. Hence, we consider diverse situations ranging from priors that are correct on average to priors that are systematically biased due to cognitive or perception biases.

Information about the actual tax rate τ_0 takes the form of an unbiased Gaussian signal with precision $1/\sigma^2$. For a realization s of the signal, the posterior belief follows from

Bayes law

$$\tilde{q}(\tau|s; \sigma) \propto \phi(s; \tau, \sigma^2) \hat{q}(\tau) \quad (9)$$

where $\phi(s; \tau, \sigma^2)$ is the Gaussian pdf with mean τ and variance σ^2 . Building on the rational inattention literature (Sims, 2003), the information content transmitted through the signal is measured from the entropy reduction between the prior and the posterior

$$\mathcal{I}(\sigma) \equiv H(\hat{q}(\tau)) - E_{p(s)} \left[H(\tilde{q}(\tau|s; \sigma)) \right] \quad (10)$$

where $H(q(\tau)) \equiv - \int q(\tau) \log_2(q(\tau)) d\tau$ is the differential entropy (in bits) of the probability distribution $q(\tau)$ and $E_{p(s)}[\cdot]$ the expectation taken over the marginal distribution of signals $p(s) \equiv \int \phi(s; \tau, \sigma) \hat{q}(\tau) d\tau$. Intuitively, $\mathcal{I}(\sigma)$ is a measure of the expected amount of information transmitted through the signal. To account for the energy and time devoted to acquiring and processing information, taxpayers suffer a utility cost κ per unit (bit) of processed information.¹⁰ The attention strategy of a taxpayer with productivity w thus results from an arbitrage between improved private decisions thanks to more accurate information and the cost to acquire this information. More specifically, she chooses the signal's precision – or equivalently its standard error $\sigma^*(\hat{q}(\tau), \kappa, w)$ – to maximize her expected indirect utility

$$\max_{\sigma} \iint V(\tilde{\tau}(s, \sigma), \tau, R; w) \phi(s; \tau, \sigma) \hat{q}(\tau) ds d\tau - \kappa \mathcal{I}(\sigma) \quad (11)$$

where $\tilde{\tau}(s, \sigma) \equiv E_{\tilde{q}(\tau|s; \sigma)}[\tau]$ is the expected perceived marginal tax rate once the signal is observed and henceforth referred to as the perceived tax rate. Note that the decision to acquire information is here only based on the information contained in the prior distribution $\hat{q}(\tau)$ which ensures the internal consistency of this learning model.¹¹

In the following, we denote by $f_{\tilde{\tau}}(\tau|\tau_0, w)$ the posterior distribution of $\tilde{\tau}(s, \sigma^*)$ for a taxpayer with productivity w and signal s drawn from the Gaussian distribution with mean τ_0 and variance σ^* . This function summarizes the distribution of agent w perceptions in the economy. Moreover, for a given perceived tax rate $\tilde{\tau}$, agent w indirect utility net of information costs writes

$$\mathcal{V}(\tilde{\tau}, \tau_0, R_0; w, \kappa) = V(\tilde{\tau}, \tau_0, R_0; w) - \kappa \mathcal{I}(\sigma^*) \quad (12)$$

3.3 Tractable Gaussian model

The general Bayesian learning model presented above is generally intractable. We here focus on the Gaussian formulation in order to derive some predictions and implications

¹⁰Our results naturally extend to more general information cost functions.

¹¹In other words, agents "don't know what they don't know".

of the model. As highlighted in the inattention literature (Maćkowiak and Wiederholt, 2015; Mackowiak et al., 2018), a closed form solution to problem (11) can be obtained under the following assumption.

Assumption 3 (tractable Gaussian learning). *Let the prior $\hat{q}(\tau)$ be the Gaussian distribution with mean $\hat{\tau}$ and variance $\hat{\sigma}^2$ and assume that agents use a quadratic approximation of their indirect utility to choose their attention strategies.*

Under the assumption that the prior is Gaussian, the posterior will also be Gaussian and the information measure $\mathcal{I}(\sigma^*)$ takes a simple form.¹² Relying on a second-order approximation of indirect utility, the solution to this problem can then be derived.

Lemma 1. *In a tractable Gaussian learning model, the expected perceived marginal tax rate $\tilde{\tau}$ is given by*

$$\tilde{\tau}(s, \sigma^*) = \xi(\sigma^*)s + (1 - \xi(\sigma^*))\hat{\tau} \quad (13)$$

where $\xi(\sigma^*) \equiv \frac{\hat{\sigma}^2}{\hat{\sigma}^2 + \sigma^{*2}} \in [0, 1]$ is a measure of attention strategies.

Proof. See Appendix A.2 □

The perceived tax rate $\tilde{\tau}$ is given by a convex combination of the prior $\hat{\tau}$ and the realization of the signal s where the weight ξ is a measure of attention. Indeed, the lower the attention parameter ξ is, the more taxpayers rely on their prior $\hat{\tau}$ and the less attention they devote to acquiring information about the actual tax rate through the signal s . In other words, agents tend to choose to ignore their signal if they do not invest in information acquisition and the signal is hence relatively uninformative in comparison to the prior.

Lemma 2. *In a tractable Gaussian learning model, the optimal attention strategy ξ is given by*

$$\xi = \max \left(0, 1 + \frac{\kappa}{\hat{\sigma}^2 \int \frac{\partial^2 y^*}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau} \phi(\tau; \hat{\tau}, \hat{\sigma}) d\tau} \right). \quad (14)$$

Proof. See Appendix A.2 □

¹²One may instead consider that the prior is a truncated Gaussian with support on $[0, 1]$ in order to ensure that perceived tax rates $\tilde{\tau}$ always remain between zero and one. Doing so, the problem remains tractable but formulas become a lot less transparent. In practice, our simulations suggest that when the prior is sufficiently informative (low variance) and the tax rate not too extreme, the posterior support belongs to $[0, 1]$ a.s. with a Gaussian prior.

Each taxpayer's attention strategy ξ is characterized by equation (14). Attention decreases with the information cost κ and increases with the uncertainty in the prior $\hat{\sigma}^2$. It also depends on the responsiveness of agents labor supply decisions to changes in perceived tax rates through $\frac{\partial^2 y^*}{\partial \tilde{\tau}^2}$. Indeed, an agent's responsiveness to changes in perceptions determines the value of information acquisition. As a result, attention increases with earnings ability w and with expected prior tax rates. Intuitively, agents who are more productive have a greater latitude in their earnings choice and will thus be more attentive – and hence responsive – to taxes. In a similar fashion, responsiveness to tax changes and thus attention increases when expected tax rates increase because it shifts the labor supply function to regions with a larger curvature.

Predictions and implications Agents choose their attention strategies through the maximization of their expected indirect utility which is based on their prior. Attention is hence unaffected by an unanticipated change in the realized tax rate $d\tau_0$. As a result, an unanticipated change in the tax rate $d\tau_0$ induces a change in the distribution of the posterior $f_{\tilde{\tau}}(\tilde{\tau}|w)$ only through a change in the signal s which is now drawn from a novel distribution $\phi(s; \tau_0 + d\tau_0, \sigma^*)$. Therefore, taxpayers' perceived tax rate slowly adjusts to news and perceptions are anchored on the prior. Anchoring is a widely documented bias in the behavioral literature (Gabaix, 2019) which has two major implications in this context.

First, if agents' prior is biased ($\hat{\tau} \neq \tau_0$) and agents are not fully attentive ($\xi < 1$), the posterior and hence agents perceptions of the tax schedule will also be biased (almost surely). Indeed, taking expectations over signal realizations we have that

$$\tau_0 - E_{\phi(s)}[\tilde{\tau}(s, \sigma^*)] = \tau_0 - [\xi\tau_0 + (1 - \xi)\hat{\tau}] = (1 - \xi)(\tau_0 - \hat{\tau}). \quad (15)$$

Second, taxpayers labor supply responses to unanticipated changes in the tax rate are attenuated by anchoring. Indeed, taking the prior as given, responses to tax changes only transit through the variation of the signal which is weighted by attention ξ . Formally, this means $\frac{d\tilde{\tau}}{d\tau_0} = \xi$ such that

$$\frac{dy^*}{d\tau_0} = \xi \frac{dy^*}{d\tilde{\tau}} \quad (16)$$

Intuitively, $\frac{dy^*}{d\tilde{\tau}}$ captures agents' preferences while ξ is a dampening factor that captures the fraction of the tax change that agents perceive. As a result the elasticity of labor supply with respect to unanticipated changes in the tax rate decreases in the presence of inattention.

The predictions derived from a Bayesian learning model with a choice of information thus seem consistent with the bulk of the empirical evidence on tax perceptions and

behavioral responses to taxes. Most importantly, it can account for the presence of systematic perception biases and implies that elasticities will be lower when inattention is at play. In addition the model also generates dispersion in perceptions – through the noisiness of the signal – and features an increase in overall attention upon tax increases which hold potentially important welfare implications (Taubinsky and Rees-Jones, 2017).

4 Discretion, commitment and the taxation bias

This section introduces the problem of the government and formalizes our positive theory of tax policy. It characterizes tax policy under discretion and commitment and provides a formal definition of the equilibrium. A general result on the existence of a taxation bias concludes.

4.1 Government problem and discretion

We consider a welfarist government that maximizes a general social welfare function summing an increasing and weakly concave transformation $G(\cdot)$ of taxpayers' indirect utilities net of information costs. It chooses a target tax schedule (τ_g, R_g) , where τ_g is the marginal tax rate and R_g the demogrant, taking the distribution of skills $f_w(w)$ in the population as given.

Following Matějka and Tabellini (2017), we introduce implementation shocks ϑ as an underlying source of uncertainty in the model. The target tax rate is implemented up to a realization of this implementation shock such that the actual tax rate is $\tau_0 = \tau_g + \vartheta$ where ϑ is a white noise drawn from an exogenous distribution $f_\Theta(\vartheta)$ known to both taxpayers and the government. We assume the actual demogrant R_0 adjusts to the realization of the implementation shock ϑ to ensure that the government budget constraint is always binding. Conceptually, these implementation shocks are introduced to ensure that Bayesian taxpayers have an incentive to learn in equilibrium. They allow to formally close the model but have an otherwise negligible impact on the optimal tax policy. Hence, we sometimes use small shocks approximations in which case we explicitly disregard the small effects they may induce.

The government problem writes

$$\max_{\tau_g, R_g} E_\vartheta \left[\iint G\left(\mathcal{V}(\tilde{\tau}, \tau_0, R; \kappa, w)\right) f_{\tilde{\tau}}(\tau|\tau_0; w) f_w(w) d\tau dw \right] \quad (17)$$

$$\text{s.t.} \quad \iint \tau_0 y^*(\tilde{\tau}; w) f_{\tilde{\tau}}(\tau|\tau_0; w) f_w(w) d\tau dw \geq R_0 + E \quad (18)$$

where E is an exogenous expenditure requirement, the expectation is taken over the implementation shock ϑ and $f_{\tilde{\tau}}(\tau|\tau_0; w)$ is the posterior distribution of perceived rates for

a taxpayer with productivity w given the actual tax rate $\tau_0 = \tau_g + \vartheta$.

Discretionary policy The government's optimal tax policy solves problem (17). When doing so, it takes the prior distribution $\hat{q}(\tau)$ as given. This is a form of Nash conjecture used to compute the best response of the government. While the problem is fundamentally simultaneous, it can be equivalently described by the following sequence of events which we here layout for the sake of clarity:

0. Agents are endowed with a common prior $\hat{q}(\tau)$ and the distribution of skills is $f_w(w)$.
1. The government sets the target tax policy (τ_g, R_g) to maximize (17).
2. The actual tax rate $\tau_0 = \tau_g + \vartheta$ is implemented up to an implementation shock drawn from a known distribution $f_\vartheta(\vartheta)$ and the actual demogrant R_0 adjusts to the resource constraint.
3. Taxpayers choose their attention strategies using their common prior $\hat{q}(\tau)$, observe a Gaussian signal s about τ_0 which precision depends on their attention and decide how much to consume and earn.
4. The government levies taxes and redistributes through the demogrant.

The government understands that taxpayers will gather information and adjust their decisions in reaction to its choice of tax policy, it therefore plays "first" in the above-described sequence of events. However, it (i) treats the prior distribution $\hat{q}(\tau)$ and the skill distribution $f_w(w)$ as predetermined state variables and (ii) cannot directly influence agents' attention strategies since they are based on agents' predetermined prior. As a result, the government does not have a particular strategic advantage from playing "first" – thus reflecting the simultaneous nature of the problem. Importantly, the government is as rational and informed as in the standard Mirrlees (1971) model and the novelty relates to information frictions on the agents' side.

The tax policy of the government follows from Proposition 1, where first-order conditions have to hold in expectation of the realization of the implementation shock.

Proposition 1. *The discretionary tax policy (τ_g, R_g) is characterized by*

$$\begin{aligned}
 (\tau_g) : \quad E_{\vartheta} \left[\int \left\{ \underbrace{\int \left[-\frac{G'(\mathcal{V})}{p} y^* + y^* \right] f_{\tilde{\tau}}(\tau|\tau_0; w) d\tau}_{\text{mechanical \& welfare effects}} \right. \right. \\
 \left. \left. + \underbrace{\int \left[\frac{G(\mathcal{V})}{p} + \tau_0 y^* \right] \frac{df_{\tilde{\tau}}(\tau|\tau_0; w)}{d\tau_g} \Big|_{\hat{q}(\cdot)} d\tau}_{\text{direct behavioral responses}} \right\} f_w(w) dw \right] = 0 \quad (19)
 \end{aligned}$$

$$(R_g) : \quad E_{\vartheta} \left[\iint \left[\frac{G'(\mathcal{V})}{p} - 1 \right] f_{\tilde{\tau}}(\tau|\tau_0; w) f_w(w) d\tau dw \right] = 0 \quad (20)$$

together with the resource constraint (18) and where p represents the social marginal cost of public funds.

Proof. See appendix A.3 □

The first order condition (19) captures the (expected) effects of a marginal increase in the target tax rate $d\tau_g$. The first line measures the impact of the reform on allocations when the distribution of perceptions remains fixed. It corresponds to the standard mechanical and welfare effects: a marginal increase in the tax rate mechanically increases tax revenue by $y^* d\tau_g$ additional dollars but reduces taxpayers' consumption and thus welfare by $\frac{G'(\mathcal{V})}{p} y^* d\tau_g$ dollars (Piketty and Saez, 2013).

The second line in condition (19) relates to the impact of the reform on the distribution of perceptions and thus captures behavioral responses to the reform. Indeed, behavioral responses transit through variations in the posterior distribution $f_{\tilde{\tau}}(\tau|\tau_0, w)$ of perceived tax rates $\tilde{\tau}$ which reflect changes in the actual tax rate τ_0 . A marginal increase in the tax rate increases, on average, the perceived tax rate by $d\tilde{\tau}$ and thus reduces tax revenue by $\tau_0 y^*(\tilde{\tau}) d\tilde{\tau}$. This is a reformulation of the standard behavioral effect. Moreover, because agents misoptimize, the envelope theorem no longer applies and a marginal deviation from taxpayers' perceived rate induces a welfare cost equal to $\frac{G(\mathcal{V}(\tilde{\tau}))}{p} d\tilde{\tau}$. This new welfare effect introduces a corrective motive for taxation in the presence of misperceptions common to optimal tax models with behavioral agents (Gerritsen, 2016; Farhi and Gabaix, 2018).

Condition (20) states that in the absence of income effects, social marginal welfare weights $g \equiv \frac{G'(\mathcal{V})}{p}$ average to 1 at the optimum: the government is indifferent between having an additional dollar or redistributing an additional dollar (Saez, 2001).

4.2 Equilibrium definition

An equilibrium is a set of target tax policy, denoted $(\tau_g^{\text{eq}}, R_g^{\text{eq}})$, and a set of attention, consumption and earnings decisions such that neither the government nor taxpayers have an incentive to deviate. Moreover, in equilibrium agents' prior $\hat{q}(\tau)$ must be mutually

consistent with the government's target tax rate and with the uncertainty induced by the implementation shock.

As discussed in the introduction, there is a large body of evidence suggesting the existence of systematic perception biases. Therefore, we allow for a potential perception bias b in agents' common prior but remain agnostic on the origin of this potential bias. We henceforth call rational (resp. biased) an equilibrium in which agents correctly (resp. incorrectly) anticipate the target tax policy such that $b = 0$ (resp. $b \neq 0$).

Given the structure of the problem, the only free variables are the government's target tax rate τ_g , agents' attention strategies and the equilibrium distribution of the common prior $\hat{q}(\tau)$. Hence, for the sake of simplicity our formal definition of the discretionary equilibrium only involves these variables. Once they are set, all remaining variables may be mechanically deduced.

Definition 1 (equilibrium). *Given the distribution of the implementation shock $f_\Theta(\vartheta)$, the equilibrium is a set of target tax rate τ_g^{eq} chosen by the government and attention strategies chosen by the agents such that*

- (a) *The target tax rate $\tau_g^{\text{eq}} \in [0, 1]$ solves the government's problem (17) given the common prior distribution $\hat{q}(\tau)$.*
- (b) *Attention strategies solve agents' problem (11) given the prior distribution $\hat{q}(\tau)$.*
- (c) *The common prior distribution $\hat{q}(\tau)$ is the pdf of $\tau_g^{\text{eq}} + b + \vartheta$.*

Condition (a) and (b) guarantee that the government and the agents will not have an incentive to deviate while condition (c) ensures that agents' prior and actual tax policy are mutually consistent up to an arbitrary bias b . Indeed, condition (c) implies that the average prior is $\hat{\tau} = \tau_g^{\text{eq}} + b$ in equilibrium. Consequently, taxpayers correctly anticipate the government policy in the rational equilibrium ($b = 0$) and their attention strategies then reflect their willingness to observe the implementation shock ϑ – which is indeed the only information conveyed through the signals. Hence, implementation shocks are here essentially introduced to ensure that Bayesian taxpayers have an incentive to learn in equilibrium but do not otherwise play an economically meaningful role.

4.3 Commitment and the taxation bias

The discretionary equilibrium is socially suboptimal. To formalize this point, we characterize the welfare-maximizing feasible tax policy. It corresponds to the optimal policy that would be chosen by the government if it could credibly commit to a tax policy. We thus refer to it as the commitment tax policy.

Commitment policy The commitment tax policy is the policy that would be chosen by a benevolent social planner who has the same information as the government but internalizes all equilibrium effects of tax policy. By implicitly restricting the set of tax policies to precommitted policy rules, the normative literature (e.g. Farhi and Gabaix (2018)) characterizes this commitment tax policy which corresponds to the optimal tax policy in the presence of information frictions.

Formally, the commitment tax policy solves the government's problem (17) subject to the additional feasibility condition that agents' prior and actual tax policy realizations have to be mutually consistent in equilibrium (condition (c) in Definition 1). It is characterized by the following first order conditions.

Proposition 2. *The commitment tax policy (τ_g^*, R_g^*) is characterized by*

$$\begin{aligned}
 (\tau_g^*) : E_{\vartheta} \left[\int \left\{ \underbrace{\int \left[-\frac{G'(\mathcal{V})}{p} y^* + y^* \right] f_{\tilde{\tau}}(\tau|\tau_0; w) d\tau}_{\text{mechanical \& welfare effects}} \right. \right. \\
 \left. \left. + \underbrace{\int \left[\frac{G(\mathcal{V})}{p} + \tau_0 y^* \right] \frac{df_{\tilde{\tau}}(\tau|\tau_0; w)}{d\tau_g} d\tau}_{\text{direct \& equilibrium behavioral responses}} \right\} f_w(w) dw \right] = 0 \quad (21)
 \end{aligned}$$

$$(R_g^*) : E_{\vartheta} \left[\iint \left[\frac{G'(\mathcal{V})}{p} - 1 \right] f_{\tilde{\tau}}(\tau|\tau_0; w) f_w(w) d\tau dw \right] = 0 \quad (22)$$

together with the resource constraint (18) and where p represents the social marginal cost of public funds. This is the policy implemented in a commitment equilibrium.¹³

Proof. See Appendix A.3 □

As before, conditions have to hold in expectation because of the implementation shock ϑ . The main difference between Propositions 1 and 2 is that the derivative $\frac{df_{\tilde{\tau}}(\tilde{\tau}|\tau_0; w)}{d\tau_g}$ in equation (21) now reflects changes in the signal received (direct adjustment) as well as changes in the prior (equilibrium adjustment). Hence, equilibrium adjustments are here internalized in the choice of tax policy.

Taxation bias The discrepancy between the discretionary and commitment equilibria represents a taxation bias. It is a measure of the deviation from the welfare-maximizing feasible tax policy τ_g^* .

¹³The commitment tax policy is not an equilibrium policy in the sense of Definition 1 because tax policy solves a different problem under commitment. Hence, notions of equilibrium under commitment implicitly refer to the equilibrium of a game in which tax policy would solve the commitment problem. The equilibrium tax policy is then simply equal to the commitment tax policy since all equilibrium adjustments are internalized in the choice of tax policy through the feasibility constraint.

Definition 2 (taxation bias). *The taxation bias is the difference between the equilibrium tax rates under discretion τ_g^{eq} and commitment τ_g^* .*

The taxation bias arises as a consequence of the government’s inability to internalize equilibrium adjustments in its choice of tax policy which induces a commitment problem. Proposition 3 relates the existence of a positive taxation bias to the associated aggregate equilibrium behavioral responses.

Proposition 3. *When both equilibria exist and are unique, there is a positive taxation bias if and only if*

$$E_{\vartheta} \left[\iint \left(\frac{G(\mathcal{V})}{p} + (\tau_g^* + \vartheta)y^* \right) \left(\frac{df_{\bar{\tau}}(\tau|\tau_g^* + \vartheta; w)}{d\tau_g} - \frac{df_{\bar{\tau}}(\tau|\tau_g^* + \vartheta; w)}{d\tau_g} \Big|_{\hat{q}(\cdot)} \right) f(w) d\tau dw \right] \leq 0 \quad (23)$$

Proof. τ_g^* solves equation (21). Then, condition (23) implies that the left hand-side of equation (19) is ≥ 0 when evaluated at τ_g^* . Hence, it directly follows from the existence and uniqueness of the discretionary equilibrium that $\tau_g^{\text{eq}} \geq \tau_g^*$. \square

Equation (23) represents the expected change in welfare due to a marginal increase in the prior average. The term $G(\mathcal{V})/p$ stands for the welfare impact of the failure of the envelope condition. It is therefore of second order and can be overlooked when perception biases b are small. Therefore, the above condition holds whenever the expected change in aggregate tax revenue following a marginal increase in the prior average is negative (and of first order). In other words, when perception biases are small there is a positive taxation bias as long as agents tend to work less when they anticipate higher taxes – a mild condition. This shows that information frictions lead to upward distortions in actual tax policy: a discretionary government implements inefficiently high tax rates in equilibrium.

5 Gaussian illustration and sufficient statistics

This section presents an application to a setting with Gaussian implementation shocks. This allows us to derive simpler characterizations of the discretionary and commitment tax policies and to illustrate our findings with numerical simulations. We further provide a sufficient statistics formula for the taxation bias that we use to empirically assess its magnitude in the actual US economy.

5.1 Gaussian discretionary equilibrium

Let the implementation shocks be normally distributed, that is $f_{\Theta}(\vartheta)$ is the pdf of the Gaussian distribution $\mathcal{N}(0, \sigma_{\vartheta}^2)$. The common prior distribution $\hat{q}(\tau)$ is then also Gaussian

in equilibrium to ensure that priors are consistent with actual tax policy realizations (condition (c) in Definition 1). Because the Gaussian family is self-conjugate with respect to a Gaussian likelihood, agent w posterior distribution $f_{\tilde{\tau}}(\tau|\tau_0, w)$ is Gaussian as well with (type-specific) mean $\mu = \xi\tau_0 + (1 - \xi)(\tau_g + b)$ in equilibrium. Introducing these equilibrium conditions into Proposition 1, we characterize the discretionary equilibrium tax policy.

Proposition 1'. *Up to a first order approximation of the integrands in Proposition 1, the Gaussian discretionary equilibrium tax policy $(\tau_g^{\text{eq}}, R_g^{\text{eq}})$ solves*

$$E_{\vartheta} \left[\int \left\{ \underbrace{(1-g)y^*}_{\text{mech. \& wel. effects}} + \underbrace{\left(g(1-\xi)(b-\vartheta) + \tau_0 \right) \frac{dy^*}{d\tilde{\tau}} \xi}_{\text{direct behavioral responses}} \right\} \Big|_{\tilde{\tau}=\mu} dF_w(w) \right] = 0 \quad (24)$$

together with $E_{\vartheta} \left[\int g|_{\tilde{\tau}=\mu} dF(w) \right] = 1$, the government resource constraint (18) and where we have introduced social marginal welfare weights $g \equiv \frac{G'(\mathcal{V})}{p}$.

Proof. See Appendix A.4. □


Equation (24) provides a simple expression of taxpayers' direct behavioral responses to tax changes and their impact for the discretionary equilibrium tax policy. Indeed, taxpayers adjust their earnings choice according to changes in their perceived tax rate. Behavioral responses $\frac{dy^*}{d\tilde{\tau}}$ are thus attenuated by the type-specific attention parameter ξ measuring the fraction of the change in taxes that agents observe. Moreover, the new welfare effect associated with the failure of the envelope theorem is directly proportional to the average size of the error in the posteriors $\mu - \tau_0 = (1 - \xi)(b - \vartheta)$ multiplied by social welfare weights g . It again captures the corrective motive for taxation in the presence of perception biases.

Figure 3 plots the tax rates (left panel) and income weighted average attention levels (right panel) in the Gaussian discretionary equilibria for different values of the information cost parameter κ . In these simulations, the distribution of skills is calibrated from the 2016 March CPS data and extended with a Pareto tail for incomes above \$200,000. We assume that the government has a log objective and agents have iso-elastic work disutility given by $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ where we set $\epsilon = 1/e$ with the structural elasticity parameter $e = 0.33$ (Chetty, 2012). A detailed presentation of the simulation procedure and the calibration strategy is available in the Online appendix.¹⁴

These simulations highlight the importance of information rigidities for tax policy. Under discretion, the equilibrium tax rate (left panel) increases substantially when taxpayers

¹⁴Our simulations indicate that the loss in accuracy due to the approximation in Proposition 1' is very small (with our calibration). Comparing this tax rate to the one obtained directly from Proposition 1, the largest error is smaller than 1% in relative terms.

Figure 3: Gaussian discretionary equilibrium



NOTE: The left panel reports the equilibrium target tax rates for different values of the information cost κ expressed in \$ /bit /year. The right panel reports the average attention parameter ξ weighted by incomes. Low (resp. high) uncertainty corresponds to Gaussian implementation shocks with a standard deviation equal to 0.05 (resp. 0.1). b is the equilibrium perception bias in agents' prior. The government has a log social welfare function and its policy follows from Proposition 1. Taxpayers have an iso-elastic disutility to work $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ with $\epsilon = 1/0.33$. The distribution of skills $f_w(w)$ is calibrated using 2016 CPS data and a Pareto tail for high incomes.

are inattentive (right panel). As the information cost parameter κ increases, attention decreases and the equilibrium tax rate increases. For example, when the average attention level (weighted by incomes) is equal to 0.8, the tax rate at the rational equilibrium is 47.5% in comparison to a 44% tax rate without information frictions. Introducing a systematic downward bias of 5 percentage points in agents' priors further increases the equilibrium tax rate, for instance by 1 to 2 percentage points when $\kappa = \$60/\text{bit}/\text{year}$. The influence of the systematic perception bias b strengthens with taxpayers' inattention because it is the equilibrium perception bias that matters for tax policy.

Finally, taxpayers are *ceteris paribus* less attentive when there is little prior uncer-

tainty about the tax rate or, equivalently in equilibrium, when the variance of implementation shocks σ_ϑ^2 is small. In this case, the government has higher incentives to increase taxes and the discretionary equilibrium tax rate is higher. It should however be noted that the main effect of the parameter σ_ϑ^2 is to rescale the mapping between attention levels ξ and the information cost parameter κ . Indeed equilibrium attention strategies depend on the ratio $\kappa/\sigma_\vartheta^2$ as can be seen from equation (14). Therefore, once we consider pairs of $(\kappa, \sigma_\vartheta^2)$ that induce the same (income-weighted) average attention, tax rates in the low and high uncertainty equilibria are similar.¹⁵

5.2 Commitment and the taxation bias

With Gaussian implementation shocks, the characterization of the commitment tax policy can be simplified to

Proposition 2'. *Up to a first order approximation of the integrands in Proposition 2, the Gaussian commitment (equilibrium) tax policy (τ_g^*, R_g^*) solves*

$$E_\vartheta \left[\int \left\{ \underbrace{(1-g)y^*}_{\text{mech. \& wel. effects}} + \underbrace{\left(g(1-\xi)(b-\vartheta) + \tau_0 \right) \frac{dy^*}{d\tilde{\tau}} \left(1 - \frac{d\xi}{d\tau_g}(b-\vartheta) \right)}_{\text{direct \& equilibrium behavioral responses}} \right\} \Big|_{\tilde{\tau}=\mu} dF_w(w) \right] = 0 \quad (25)$$

together with $E_\vartheta \left[\int g|_{\tilde{\tau}=\mu} dF(w) \right] = 1$, the government resource constraint (18) and where we have introduced social marginal welfare weights $g \equiv \frac{G'(v)}{p}$.

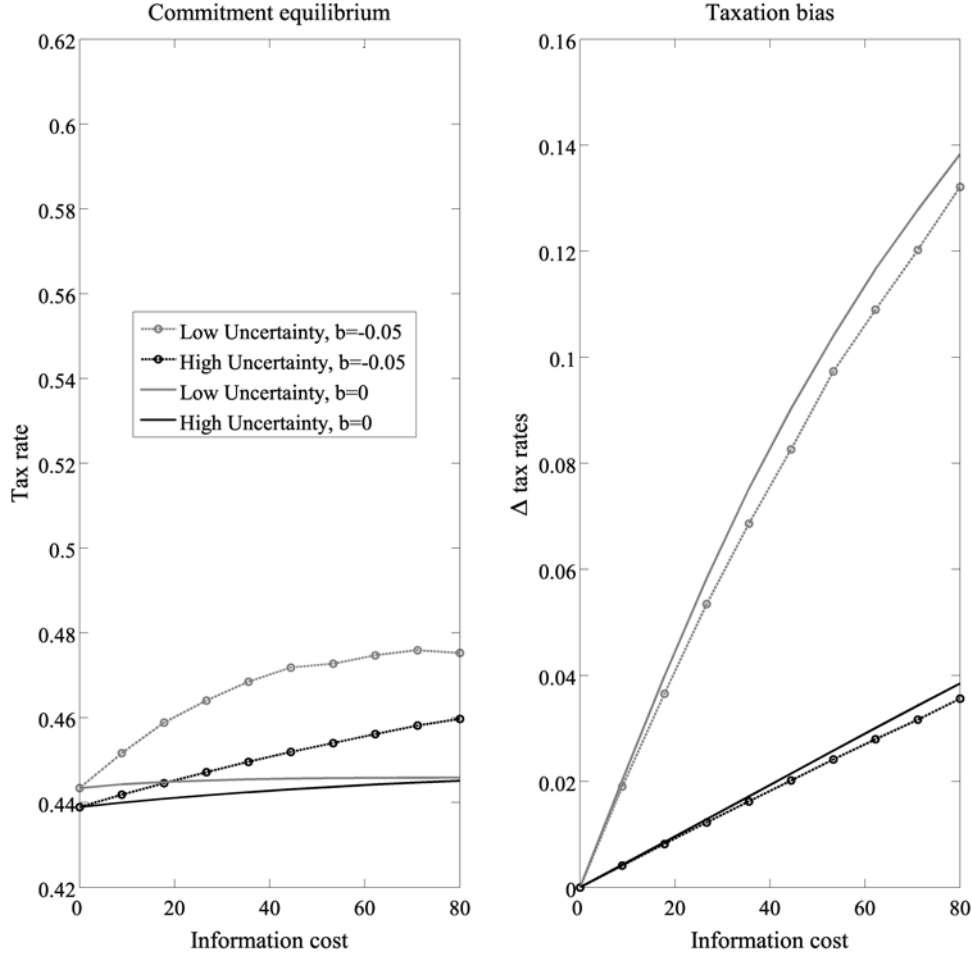
Proof. See Appendix A.4. □

The direct and equilibrium responses to a change in taxes is now captured through the term $\frac{dy^*}{d\tilde{\tau}} \left(1 - \frac{d\xi}{d\tau_g}(b-\vartheta) \right)$. In the latter, the factor 1 stands for the fact that the government accounts for the equilibrium adjustment of priors when setting its policy. The supplemental term $\frac{d\xi}{d\tau_g}(b-\vartheta)$ captures the effect of changes in attention ξ following a marginal increase in the tax rate on equilibrium tax perceptions. This term vanishes (in expectation) when agents correctly anticipate the tax rate ($b = 0$) and prompts the government to decrease the tax rate when agents underestimate tax rates ($b < 0$). Indeed, increasing the tax rate increases attention and thereby reduces tax underestimation in equilibrium which is detrimental to efficiency.

In the left panel of Figure 4 we report the commitment equilibrium tax rates for different values of the information cost parameter κ . In the rational equilibrium ($b = 0$), the tax rate is only marginally higher in the presence of information frictions. Indeed, the policymaker finds it optimal to marginally increase taxes to prompt taxpayers to be

¹⁵For example, we find that when the (income-weighted) average attention level is 80%, the difference between the tax rates in the low and high uncertainty equilibria is only equal to 0.2 percentage points.

Figure 4: Taxation bias



NOTE: Difference between positive and normative tax rates for different values of the information cost κ expressed in annual \$ / bit. Low (resp. high) uncertainty corresponds to Gaussian implementation shocks with a standard deviation equal to 0.05 (resp. 0.1). b is the equilibrium perception bias in agents' prior. The government has a log social welfare function and its policy follows from Proposition 2' for the commitment equilibrium. Taxpayers' behavior relies on Assumptions 3.1-3.5 and an iso-elastic disutility to work $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ with $\epsilon = 1/0.33$ (Chetty, 2012). The distribution of skills $f_w(w)$ is calibrated using 2016 CPS data.

more attentive to implementation shocks. In downward biased equilibria ($b < 0$), the government further increases the tax rate to exploit the efficiency gains from agents tax underestimation. Indeed, because taxpayers remain inattentive in equilibrium and the prior is downward biased, perceived tax rates are lower than the actual rate. Ultimately, this underestimation of tax rates reduces the efficiency costs of taxation allowing for tax increases. The leverage to increase tax rates is however limited here because the policymaker realizes that it also prompts increases in agents' prior and attention. This results in an increase of perceived tax rates which ultimately increases the efficiency costs of taxation. Consequently, commitment equilibrium tax rates are much smaller than

discretionary equilibrium tax rates depicted in Figure 3.

As a consequence, and as depicted in the right panel of Figure 4, the taxation bias increases as the information cost parameter κ grows. Even small information frictions generate a significant taxation bias. Our simulations indicate that in a rational equilibrium, there is a taxation bias of 4 (resp. 3.5) percentage points in the presence of low (resp. high) uncertainty when the income-weighted average attention parameter is 0.8. Moreover, the taxation bias is above 10 percentage points when the income-weighted average attention falls below 0.55.

The taxation bias can thus lead to significant upward distortions in actual tax rates when the income-weighted average attention turns out to be low. We now show theoretically that this is indeed a key sufficient statistic to empirically assess the magnitude of the taxation bias.

5.3 Sufficient statistics formulas and taxation bias in the US

We derive sufficient statistics formulas for the equilibrium tax policy under discretion and commitment that echo textbook optimal tax formulas and that we combine to obtain a sufficient statistics formula for the taxation bias. To obtain simple sufficient statistics formulas we further assume that preferences are iso-elastic such that the structural labor supply elasticity e – i.e. computed with respect to the perceived marginal net-of-tax rate – is constant and that implementation shocks and perception biases are small.

Corollary 1. *A sufficient statistics formula for the Gaussian discretionary equilibrium tax rate characterized in Proposition 1' is*

$$\tau_g^{\text{eq}} \simeq \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^*} + \overline{y^*\xi}} e - b \frac{\overline{g(1-\xi)y^*\xi}}{\overline{(1-g)y^*} + \overline{y^*\xi}} e \quad (26)$$

where all endogenous right hand side quantities are evaluated at τ_g^{eq} and we have introduced the mean operator $\bar{x} \equiv \int x(w)f(w)dw$.

Proof. See Appendix A.5. □

The first term in equation (26) corresponds to the textbook optimal linear tax formula up to the presence of the income weighted average attention $\overline{y^*\xi}$. The second term corresponds to the corrective motive of taxation in the presence of perception biases.

Corollary 2. *A sufficient statistics formula for the Gaussian commitment equilibrium tax rate characterized in Proposition 2' is*

$$\tau_g^* \simeq \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^*} + \overline{y^*}} e - b \frac{\overline{g(1-\xi)y^*}}{\overline{(1-g)y^*} + \overline{y^*}} e \quad (27)$$

where all endogenous right hand side quantities are evaluated at τ_g^* and we have introduced the mean operator $\bar{x} \equiv \int x(w)f(w)dw$.

Proof. See Appendix A.5. □

The first term now exactly coincides with the textbook optimal linear tax formula while the second term again corresponds to the corrective motive of taxation in the presence of perception biases. It should however be noted that corrective terms in equations (26) and (27) are not identical. Their sign are however identical and both are proportional to the perception bias b .

Focusing on near-rational equilibria – that is equilibria with small perception biases $b \simeq 0$ such that corrective terms are second-order – we obtain a simple sufficient statistics formula for the taxation bias.

Proposition 4. *A sufficient statistics formula for the taxation bias in Gaussian near-rational equilibria is*

$$\tau_g^{\text{eq}} - \tau_g^* \simeq \frac{\overline{(1 - \xi)y^*}}{\overline{(1 - g)y^*}} e t^2 \quad (28)$$

where all endogenous right hand side quantities are evaluated at the actual tax rate t .

Proof. See Appendix A.5. □

This simple formula for the taxation bias is reminiscent of the one provided Section 2 and is a generalization to a situation in which the government has tastes for redistribution and agents' attention is endogenous and hence type-specific.¹⁶ The income-weighted average attention $\overline{\xi y^*}$ – or equivalently inattention $\overline{(1 - \xi)y^*}$ – thus becomes a key sufficient statistic for the taxation bias. *Ceteris paribus*, the taxation bias increases with the structural elasticity of labor supply e , with the square of the actual tax rate t and decreases with the government redistributive tastes.¹⁷

In an attempt to gauge the empirical magnitude of the taxation bias in the actual US economy, we bring this sufficient statistics formula to the data. The meta-analysis of Gabaix (2019) combines existing measures of attention to sales taxes to trace out the evolution of average attention with the stakes. We find that income taxes in the US are well approximated by a linear tax schedule with a tax rate of $t = 29.46\%$ which

¹⁶Equation (28) can also be expressed in terms of covariances as

$$\tau_g^{\text{eq}} - \tau_g^* \simeq \frac{\text{cov}(\xi, y^*) - \overline{(1 - \xi)y^*}}{\text{cov}(g, y^*)} e t^2$$

¹⁷Intuitively, the taxation bias *increases* with the government's redistributive tastes as it relatively increases the incentives to implement unanticipated tax increases. However, this first-order effect here transits through an increase in the actual tax rate t and we get the inverse relationship controlling for t .

would correspond to an average attention parameter of about 0.70. Focusing on the US personal income tax, Rees-Jones and Taubinsky (2016) estimate that agents’ attention parameter to their marginal tax rate is equal to 0.81. Accordingly, we consider an average attention of 0.75 to taxes as our baseline. We are then able to compute the associated income-weighted average attention using our model of endogenous attention and the actual distribution of income.

Turning to other sufficient statistics, we take the structural elasticity parameter $e = 0.33$ estimated by Chetty (2012) and use an inverse optimum approach to deduce the US government’s redistributive tastes from the actual tax policy.¹⁸

Table 1: Estimated taxation bias in the actual US economy

Taxation bias (percentage points)		Average attention parameter		
		0.65	0.75	0.85
Benthamite	0.27	5.21	3.66	2.17
redistribution	0.50	3.11	2.18	1.29
parameter	1.00	1.85	1.30	0.77

NOTE: Our estimation of the taxation bias (in percentage points) follows from the characterization in Proposition 4. A larger Benthamite parameter corresponds to a more redistributive objective. The value in bold corresponds to our baseline estimate for the 2016 US economy.

In the actual US economy, we estimate that the taxation bias is roughly equal to 3.66 percentage points in our baseline calibration. This means that the US income tax rate is 12% higher than what would be optimal holding the government’s redistributive objective constant. Table 1 provides a sensitivity analysis varying average attention and the government’s redistributive objective. For the latter we use a Benthamite social welfare function for which we vary the value of the parameter that shapes the desire for redistribution. The value of 0.27 closely approximates the welfare weights we estimate using an inverse optimum approach and a value of 1 corresponds to a logarithmic social welfare function – which captures rather extreme redistributive tastes. For realistic redistributive tastes and attention parameters, the magnitude of the taxation bias in the actual US economy ranges from 1.29 to 5.21 percentage points and our baseline estimate of 3.66 lies in the middle of this range.

6 Welfare implications

This section analyzes the welfare implications of information frictions. It first decomposes the variation in aggregate social welfare between potential welfare gains that may

¹⁸That is, we deduce $(1 - g)y^*$ from equation (26) assuming $b \simeq 0$.

be attained with information rigidities (commitment) and the welfare losses associated to actual policy distortions (discretion). It then quantifies the relative importance of the different channels through which information rigidities ultimately affect welfare and redistribution at the individual level.

6.1 Information rigidities and aggregate welfare

Let $SW^{\text{eq}}(b, \kappa)$ be the social welfare from equation (17) evaluated at the discretionary equilibrium. The total welfare impact of information rigidities writes $\Delta SW^{\text{eq}}(b, \kappa) \equiv SW^{\text{eq}}(b, \kappa) - SW^{\text{eq}}(0, 0)$.¹⁹ It may be decomposed between the potential welfare gains from misperceptions and the welfare costs induced by the taxation bias as follows

$$\Delta SW^{\text{eq}}(b, \kappa) = \underbrace{SW^{\text{eq}}(b, \kappa) - SW^*(b, \kappa)}_{\text{Taxation bias } (\leq 0)} + \underbrace{SW^*(b, \kappa) - SW^*(0, 0)}_{\text{Potential gains}} \quad (29)$$

where $SW^*(b, \kappa)$ is the social welfare attained under commitment.

The welfare impact of the taxation bias is negative since the commitment tax policy is by definition the welfare-maximizing feasible policy. As a result, information rigidities are welfare improving if and only if this negative welfare impact is dominated by the welfare gains induced by misperceptions, that is $|SW^{\text{eq}}(b, \kappa) - SW^*(b, \kappa)| \leq SW^*(b, \kappa) - SW^{\text{eq}}(0, 0)$.

This condition requires a downward bias in priors $b < 0$ such that agents underestimate tax rates. This can be easily seen when looking at the sufficient statistics formula for the commitment tax rate in equation (27). Indeed, when $b = 0$ the commitment tax rate is equal to the optimal tax rate without information up to a first order approximation. Therefore, there cannot be first order gains from information rigidities. However, equation (28) indicates that the taxation bias is nonetheless positive. Consequently, the total welfare impact is negative. Panel B of Figure 5 illustrates this result for Gaussian implementation shocks.²⁰

¹⁹Note that $SW^{\text{eq}}(0, b) = SW^{\text{eq}}(0, 0)$ since as soon as the information cost κ is nil, agents have perfect information and whether priors are biased is irrelevant.

²⁰Strangely enough, the potential gain is first decreasing and then increasing when b is small or nil. While the magnitude of the potential gain is small and thus negligible in comparison to the impact of the taxation bias, it deserves to be briefly explained. Consider the two extreme cases where $\kappa = 0$ and $\kappa \mapsto \infty$. Hence, ξ is respectively equal to one or zero for each taxpayer. Everything else being equal, aggregate earnings are larger when $\kappa \mapsto \infty$ as agents behave as if there were no implementation shocks when deciding how much to earn (individual earnings are a concave function of the perceived rate), while they fully adjust to these shocks when $\kappa = 0$. Consequently, the potential gain converges to a positive value as κ tends to infinity. However, when κ is small but strictly positive, some taxpayers noisily observe the implementation shocks so that the variance of their earnings choices increases. Ultimately,

While a negative perception bias is necessary for information rigidities to be welfare improving, it is not a sufficient condition. Information rigidities should also not be too large to ensure taxpayers are sufficiently attentive to tax policy. Indeed, as the information cost parameter κ grows, the welfare losses induced by the taxation bias increase more rapidly than the welfare gains from the negative perception bias. Indeed, the former is convex while the latter is concave.

Panel A of Figure 5 illustrates this mechanism with Gaussian implementation shocks. When the downward equilibrium perception bias is equal to 5 percentage points, the welfare gains induced by information rigidities dominate the welfare cost induced by the taxation bias as long as the information cost parameter κ is lower than \$25/bit/year. Above this threshold, inattention and the associated deviation from the commitment policy become too important such that information rigidities are welfare decreasing. Therefore, downwards biases in perceived marginal tax rates will be typically associated with a *decrease* – rather than an *increase* – in aggregate social welfare when agents are not sufficiently attentive to tax policy.

6.2 Redistributive impacts

We now turn to an analysis of the welfare implications of information rigidities at the individual level. Let $\Delta\mathcal{V} \equiv \mathcal{V}_\kappa(\tau_g^{\text{eq}}(\kappa), w) - \mathcal{V}_0(\tau_g^*(0), w)$ be the variation in the expected utility of a taxpayer with skill w between the discretionary equilibrium when the information cost is κ and a counterfactual with perfect information. Using a quasi-linear and separable utility function allows us to decompose the variation in expected utility induced by information rigidities in the following way (see Appendix A.7 for precise definitions)

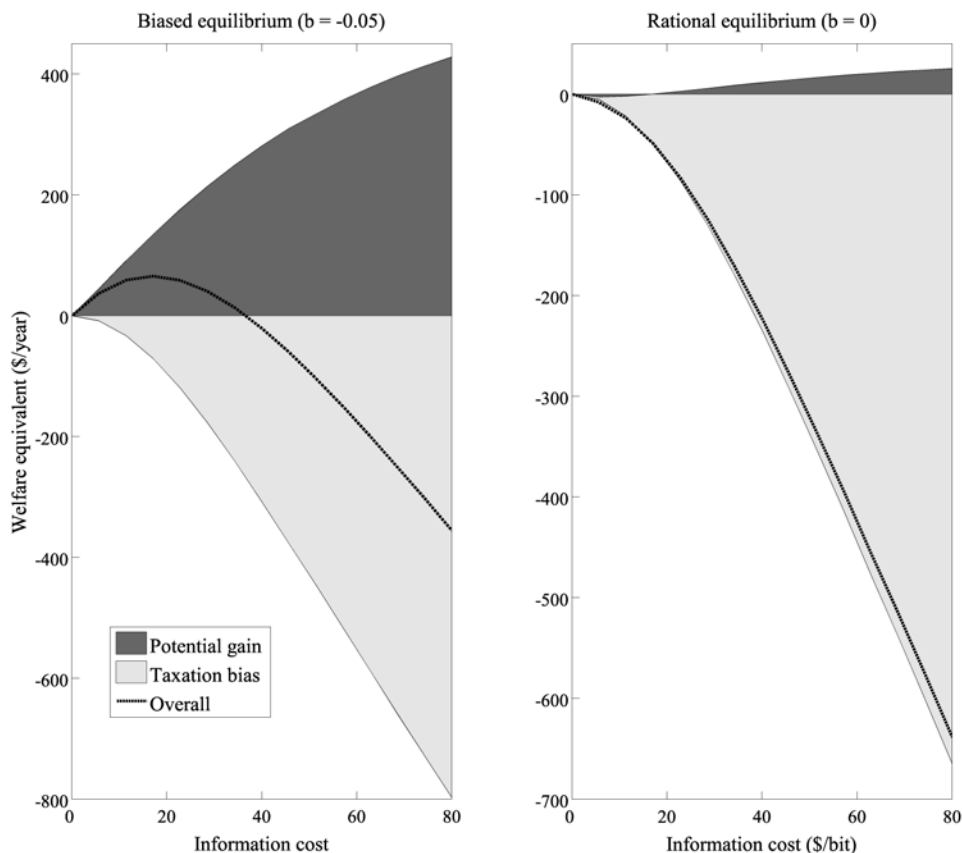
$$\Delta\mathcal{V} = \Delta^{\mathcal{V}}R + \Delta^{\mathcal{V}}\tau_g + \Delta^{\mathcal{V}}\text{info cost} + \Delta^{\mathcal{V}}b + \Delta^{\mathcal{V}}\text{uncertainty} \quad (30)$$

that is, the welfare impact of information rigidities at the individual level arises from a variation in the demogrant R , the tax rate τ_g , the cost of information acquisition $\kappa\mathcal{I}(\sigma^*)$, the misoptimization costs induced by potential perception bias b , and the change in overall uncertainty. The increase in the demogrant has a positive effect on expected utility while all other terms are negative.

Figure 6 plots the above expected utility decomposition with Gaussian implementation shocks ($\sigma_\vartheta = 0.1$), no perception bias ($b = 0$) and an information cost κ of \$50/bit per year.

it lowers aggregate earnings and generates a negative and decreasing potential gain for small values of κ . Simulations indicate that the above described variations in aggregate earnings dominate other second order effects (e.g. misoptimization costs).

Figure 5: Welfare decomposition




NOTE: Welfare decomposition from equation (29) for different values of the information cost κ . The standard deviation for the Gaussian implementation shocks is equal to 0.05. b is the equilibrium perception bias in agents' prior. The government has a log social welfare function and its policy follows from Proposition ???. Taxpayers have an iso-elastic disutility to work $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ with $\epsilon = 1/0.33$ (Chetty, 2012). The distribution of skills $f_w(w)$ is calibrated using 2016 CPS data.

Information rigidities induce policy distortions that create losers and winners. Indeed, the redistributive impact of information rigidities is driven by the increase in the tax rate and thus in redistribution through the demogrant. This naturally benefits low skill workers at the extent of high skill workers.

Somewhat surprisingly, information costs represent a relatively small deadweight loss for society in comparison to the large indirect impact of information frictions on tax policy and welfare. Moreover, it turns out that these information costs are higher for high skill workers because they have higher incentives to collect information and are thus more attentive. Extending our analysis to non-linear tax schedules, we show in the next section that this regressivity of attention has an impact on actual tax progressivity.

Figure 6: Variation in expected utility



NOTE: Expected utility decomposition from equation (30) by deciles of the productivity level w . The standard deviation for the Gaussian implementation shocks is equal to 0.1. There is no perception bias in agents' prior $b = 0$ and the information cost is $\kappa = 50$. The government has a log social welfare function and its policy follows from Proposition 1'. Taxpayers' behavior relies on Assumptions 3.1-3.5 and an iso-elastic disutility to work $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ with $\epsilon = 1/0.33$ (Chetty, 2012). The distribution of skills $f_w(w)$ is calibrated using 2016 CPS data.

7 Tax progressivity and the taxation bias

In this section we extend the analysis to non-linear tax schedules. We find that the taxation bias becomes heterogeneous across income levels and ultimately reduces tax progressivity.

7.1 Introducing nonlinear tax schedules

We allow the government to use a nonlinear tax schedule $T(y)$ but the setup introduced in Section 3 is otherwise unchanged. In particular, we maintain Assumption 1 that individuals use a linear representation of the tax schedule $\tilde{T}(y) = \tilde{\tau}y - \tilde{R}$ which now raises a new question: in the continuum of marginal tax rates $\{T'(y)\}_y$, what is the marginal tax rate $T'(y_w)$ agent w gathers information about?

Absent income effects, the perceived marginal tax rate $\tilde{\tau}_w$ remains a sufficient statistics for labor supply and uniquely pins down earnings $y^*(\tilde{\tau}_w; w)$.²¹ Using this mapping, we define an agent w ex ante – before information acquisition – optimal earnings level $\hat{y}_w =$

²¹While the perceived marginal tax rate $\tilde{\tau}$ was already type-specific in the previous sections – through type-specific attention choices –, we here introduce the subscript w to emphasize that it will in addition be type-specific through agents' type-specific priors (see Assumption 4).

$y^*(\hat{\tau}_w; w)$ and make the following assumption.

Assumption 4 (prior reliance). *Taxpayer w gathers information about the actual marginal tax rate $\tau_0(\hat{\tau}_w, w) \equiv T'_0(\hat{y}_w)$ at her ex-ante optimal earnings level \hat{y}_w .*

Essentially, Assumption 4 guarantees the internal consistency of the perception formation process by ensuring agents have no additional information ex ante than that contained in their prior. Moreover, it gives a novel allocative role for the prior as taxpayers now linearize the tax schedule around the income level that they deem optimal ex ante.²²

The presence of a nonlinear tax schedule does not fundamentally affect equilibrium concepts. Two refinements are nevertheless necessary. First, for the sake of simplicity we assume that implementation shocks ϑ uniformly affect marginal tax rates at all earnings levels y such that $T'_0(y) = T'_g(y) + \vartheta$. Second, the equilibrium condition (c) from Definition 1 – which characterizes the equilibrium adjustment of priors – now becomes

(c') *The type-specific prior distribution $\hat{q}_w(\tau)$ is the pdf of $T'_g(\hat{y}_w) + b + \vartheta$.*

That is, each taxpayer's prior is consistent with her marginal tax rate of interest up to an arbitrary perception bias b . Incidentally, the prior average $\hat{\tau}_w \equiv E_{\hat{q}_w(\cdot)}[\tau]$ is thus necessarily type-specific in equilibrium when the government implements a nonlinear tax schedule. While natural in our context, this poses a potential challenge for the resolution of this nonlinear tax model.

We rely on a perturbation approach in order to derive the optimal tax schedule. Following Jacquet and Lehmann (2017), one needs three assumptions to solve for the optimal non-linear tax schedule using a tax perturbation approach. (i) The tax function $T_g(\cdot)$ must be twice differentiable. (ii) The optimization program of each taxpayer must admit a unique global maximum. (iii) Agents' second-order conditions must hold strictly. While (i) and (ii) are generic requirements to ensure the global smoothness of the problem so that tax perturbations will not induce individuals to jump between different maxima, condition (iii) has less intuitive consequences.

²²To illustrate this new allocative role, consider the limit where the information cost κ goes to zero. Perceptions are then perfect $\tilde{\tau}_w = \tau_0(\hat{\tau}_w, w)$ and each agent chooses earnings $y_{w|\hat{\tau}_w} = y^*(\tau_0(\hat{\tau}_w, w); w)$. This is in contrast to the full information case in which earnings are the solution to a fixed-point problem characterized by $y_w = y^*(T'_0(y^*(\cdot)); w)$. In a rational equilibrium ($b = 0$), both income concepts coincide. They will however differ in biased equilibria $b \neq 0$.

In standard models, condition (iii) – combined with a single crossing assumption on individuals preferences – ensures the existence of an increasing mapping between earnings y and skills w . This is known as a monotonicity condition on allocations.²³ It is a requirement for the tax perturbation approach which disciplines the curvature of the tax function $T_g''(\cdot)$. Here, allowing for type-specific priors and a perception bias b poses a potential threat to the existence of an increasing mapping between earnings y and skills w . In the Online Appendix, we show that under our assumptions the monotonicity condition is also expected to hold when $T_g''(\cdot)$ is smooth enough. As a result, we solve for the optimal tax schedule assuming the monotonicity condition is verified and check ex post that it holds at the optimum.

7.2 ABCD tax formula

We can now solve for the optimal nonlinear tax schedule. The government chooses a target nonlinear tax schedule $T_g(\cdot)$ that consists in a continuum of marginal tax rates $\{T_g'(y)\}_y$ and a tax level indexed by the demogrant $T_g(0)$. It is implemented up to an implementation shock ϑ on marginal tax rates and the tax level adjusts such as to satisfy the government budget constraint ex post. The government problem writes

$$\max_{T_g'(\cdot), T_g(0)} E_\vartheta \left[\iint G(\mathcal{V}(\tilde{\tau}_w, T_0(\cdot); \kappa, w)) f_{\tilde{\tau}}(\tau | \tau_0(\hat{\tau}_w, w); w) f_w(w) d\tau dw \right] \quad (31)$$

$$\text{s.t.} \quad \iint T_0(y^*(\tilde{\tau}_w; w)) f_{\tilde{\tau}}(\tau | \tau_0(\hat{\tau}_w, w); w) f_w(w) d\tau dw \geq E \quad (32)$$

where E is an exogenous expenditure requirement, $f_{\tilde{\tau}}(\tau | \tau_0(\hat{\tau}, w); w)$ is the posterior distribution of agent w perceived tax rate and with the indirect utility function

$$\mathcal{V}(\tilde{\tau}_w, T_0(\cdot); \kappa, w) = y^*(\tilde{\tau}_w; w) - T_0(y^*(\tilde{\tau}_w; w)) - v(y^*(\tilde{\tau}_w; w); w) - \kappa \mathcal{I}(\sigma^*) \quad (33)$$

We solve this problem using a perturbation approach. Namely, we consider the effect of a reform that consists in a small increase $\Delta\tau^r$ in marginal tax rates in a small bandwidth of earnings $[y^r - \Delta y, y^r]$ and characterize its impact on the objective function of the government. Following the tax perturbation literature, this reform may be apprehended through three mechanisms: a mechanical effect, a welfare effect and a behavioral effect. However, analyzing the impact of a reform in this setting with information frictions in tax perceptions calls for a careful identification of the agents affected by the reform.

The standard mechanical and welfare effects capture the change in taxes and welfare for individuals w whose earnings are higher than y^r given their perceived tax rates $\tilde{\tau}_w$.

²³It follows from agents incentive constraints in a mechanism design approach.

Following from the aforementioned monotonicity condition, it corresponds to all agents with a productivity $w \geq w^r$ where $y^*(\tilde{\tau}_{w^r}; w^r) \equiv y^r$. In contrast, the behavioral effect comes from taxpayers who are learning the marginal tax rate affected by the reform. That is, all agents whose ex ante optimal earnings level \hat{y} belong to $[y^r - \Delta y, y^r]$. Again using the monotonicity condition, we can equivalently identify these agents as those with a productivity $w \in [\hat{w}^r - \Delta \hat{w}, \hat{w}^r]$ where $y^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r) \equiv y^r$. Note that the two cut-offs w^r and \hat{w}^r differ almost surely.²⁴

We then characterize the discretionary and commitment equilibrium tax schedules assuming Gaussian implementation shocks. As before, these conditions are easier to interpret after applying a small implementation shocks approximation which is what we report in Proposition 5, relegating general conditions to the Online Appendix.

Proposition 5 (ABCD formula). *Assuming small Gaussian implementation shocks, the equilibrium non-linear tax schedule is to a first-order approximation characterized by*

$$\begin{aligned} & \frac{T'_g(y^*(\mu_{\hat{w}^r}; \hat{w}^r)) + g(\hat{w}^r)|_{\tilde{\tau}=\mu_{\hat{w}^r}} \left(\mu_{\hat{w}^r} - T'_g(y^*(\mu_{\hat{w}^r}; \hat{w}^r)) \right)}{1 - \mu_{\hat{w}^r}} \\ &= \frac{1}{e^{\frac{d\mu_{\hat{w}^r}}{d\tau_g}|_{\tilde{\tau}=\mu_{\hat{w}^r}}} y^*(\mu_{\hat{w}^r}; \hat{w}^r)} \frac{1}{f_w(\hat{w}^r)} \frac{dy^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r)}{dw} \int_{w^r}^{\infty} \left(1 - g(w)|_{\tilde{\tau}=\mu_w} \right) f_w(w) dw \end{aligned} \quad (34)$$

together with the transversality condition $\int g(w)|_{\tilde{\tau}_w=\mu_w} dF(w) = 1$ and the government budget constraint (32), where all endogenous quantities are evaluated at their equilibrium values.

Moreover the ex-post average perceived marginal tax rate is $\mu_w \equiv \xi \tau_0(\hat{\tau}_w, w) + [1 - \xi] \hat{\tau}_w$ such that $\frac{d\mu_w}{d\tau_g} = \xi$ under discretion and $\frac{d\mu_w}{d\tau_g} = 1 + \frac{d\xi}{d\tau_g} [\tau_0(\hat{\tau}_w, w) - \hat{\tau}_w]$ under commitment.

Proof. See Online Appendix. □

Under commitment and absent perception biases ($b = 0$), the ABCD formula boils down to the ABC formula derived in Diamond (1998) and the standard interpretation prevails.²⁵ The presence of perception biases ($b \neq 0$) has several effects. First, it creates a disconnect between w^r and \hat{w}^r new to this non-linear setting. Second, it adds a welfare effect ($(g(\hat{w}^r)(1 - \xi)b$ in the numerator of the LHS) related to the failure of the envelope theorem when agents misoptimize. Third, it adds an efficiency term (fraction on the RHS with $\frac{d\mu_w}{d\tau_g} = 1 + \frac{d\xi}{d\tau_g} [\tau_0(\hat{\tau}_w, w) - \hat{\tau}_w]$ in the denominator) accounting for the variation

²⁴The two cut-offs coincide only when $\hat{\tau}_w = \tilde{\tau}_w$. That is, when $b = 0$ and $\vartheta = 0$. Since we focus on Gaussian implementation shocks here, it is never the case (a.s.).

²⁵The additional term on the RHS disappears since $\frac{d\mu_w}{d\tau_g} = 1$ while we have that $\hat{\tau}_w^{\text{eq}} \approx \tilde{\tau}_w^{\text{eq}}$ ensuring both $w^r \approx \hat{w}^r$ and $\tilde{\tau}_w^{\text{eq}} \approx T'_g(y^*(\mu_{\hat{w}^r}^{\text{eq}}; \hat{w}^r))$.

in agents equilibrium misperception when their attention ξ changes in response to tax reforms.

As before, the emergence of a taxation bias comes from the discrepancy between the estimated impact of a reform under discretion and under commitment. Under discretion, the government fails to internalize the equilibrium impact of the reform on perceptions and accordingly considers that an increase $\Delta\tau^r$ in marginal tax rates only increases perceived marginal tax rates by $\xi\Delta\tau^r \leq \Delta\tau^r$. Increasing marginal tax rates is thus perceived as less costly in terms of efficiency than it really is. As a result, marginal tax rates are in a discretion equilibrium higher than in a commitment equilibrium. In other words, marginal tax rates are higher than they should be from a normative perspective: this is the taxation bias.

What is new to this non-linear setting is that the taxation bias affecting the marginal tax rate $T'_g(y^r)$ at a given level of earnings y^r is driven by the attention level of agents of type \hat{w}^r . Surprisingly, agents \hat{w}^r may not even be located at earnings y^r in the presence of perception biases. More importantly, if attention levels vary across the earnings distribution, the taxation bias will have an impact of the progressivity of the tax schedule.

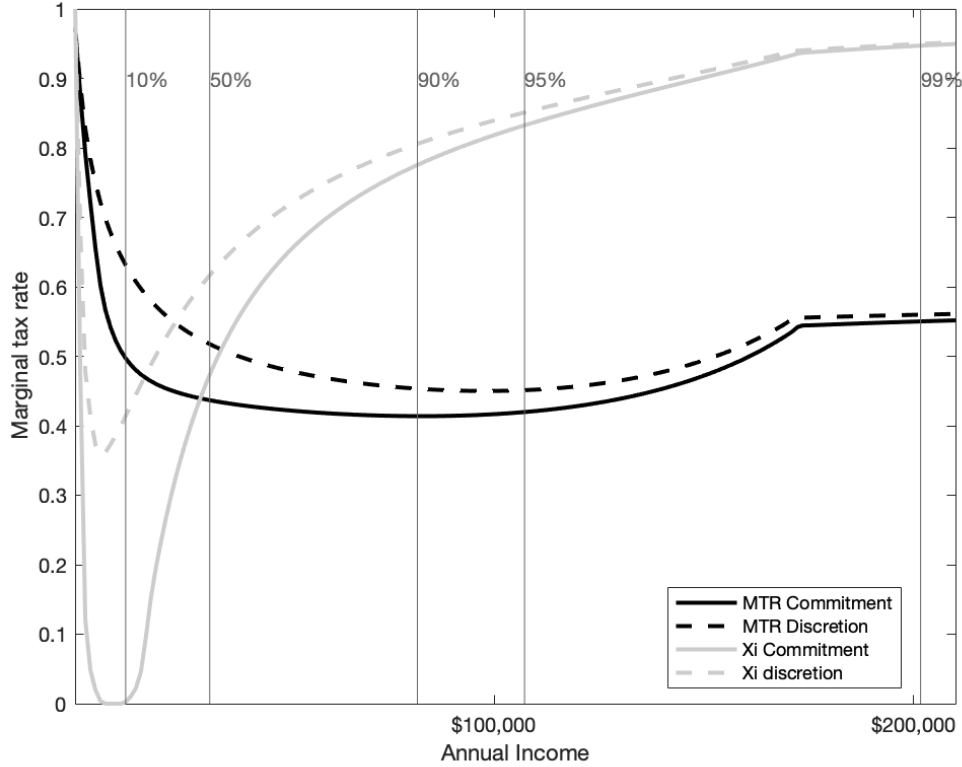
7.3 Numerical illustration

To illustrate this property we represent in Figure 7 the nonlinear tax schedule implemented under discretion (dashed black line) and commitment (full black line). Simulations are carried out absent systematic perception biases ($b = 0$) such that the optimal non-linear tax schedule under commitment corresponds to the textbook optimal non-linear tax schedule of Saez (2001). We thus naturally retrieve the known U-shape pattern of marginal tax rates.

Because of the taxation bias, marginal tax rates are higher under discretion than under commitment. Strikingly, this difference in marginal tax rates is not constant across earnings levels. For instance, agents located at the first decile (resp. the median) of the earnings distribution face a marginal tax rate of 50% (resp. 44%) under commitment and a marginal tax rate of 63% (resp. 52%) under discretion. In contrast, the marginal tax rate faced by individuals in the top decile (resp. top percentile) increases by at most 4 (resp. 1) percentage points. This reflects the impact of the taxation bias on the progressivity of the tax schedule coming from the variation in attention ξ across earnings.

Attention levels represented in Figure 7 (grey lines) are indeed generally increasing in earnings. In our model, more productive agents have intuitively more latitude to choose the earnings level they see fit and attach thus a higher value to being informed about the tax schedule. As a result, attention globally increases with productivity and thus –

Figure 7: Target non-linear tax schedules in equilibrium



NOTE: Equilibrium target non-linear tax schedules under discretion and commitment a value of the information cost $\kappa = 30\$/bit/yr$. The government has a log social welfare function and follows the optimal policy from Proposition 5. Taxpayers' have an iso-elastic disutility to work $v(y, w) = (y/w)^{1+\epsilon}/(1+\epsilon)$ with $\epsilon = 1/0.33$ (Chetty, 2012). The distribution of skills $f_w(w)$ is calibrated using 2016 CPS data and extended with a Pareto-tail of parameter $a = 2$ (Saez, 2001).

through the monotonicity condition – with earnings.²⁶ Note that this pattern is obtained assuming all individuals have the same cost of information κ . Therefore, assuming that more able workers are also more efficient at collecting information would only reinforce the striking result that, because it decreases with income, inattention to taxes induces regressive tax increases.

Conclusion

The paper develops a positive theory of tax policy with information frictions and shows that it leads to a taxation bias. Our results suggest that the welfare gains from using

²⁶The pattern reverts at the very beginning of the earnings distribution because as the marginal tax rate approaches one, we approach the origin of the labor supply function where earnings become infinitely responsive to changes in the marginal tax rates. As a result, very low productive agents end up choosing very high attention levels.

precommitted tax rules could be large. Interestingly, the model identifies a key parameter to limit the government’s deviations from the welfare-maximizing tax policy: the information cost κ . Therefore, it would be interesting for future research to investigate the determinants behind this information cost. Indeed, if the latter is related to the complexity of the tax system, monitoring and restricting tax complexity may be a simple way to prevent the implementation of inefficient and regressive tax increases.

While some of our results may be model-specific, our analysis sheds a new light on the welfare consequences of information frictions in tax perceptions. It underlines that downward biases in tax perceptions are not necessarily welfare improving. They do lower the efficiency costs of taxation in existing tax systems, but existing tax systems without misperceptions are arguably not the right counterfactual to use for welfare analysis. Indeed, there may be other (equilibrium) effects at play – which take the form of a taxation bias in this paper. As we show, the welfare consequences of such effects may be dominant, overturning previous welfare implications. One should thus be very careful with welfare implications drawn from the measurement of misperceptions. We believe that this general lesson applies outside of the realm of taxation.

References

- Aghion, P., Akcigit, U., Lequien, M., and Stantcheva, S. (2017). Tax simplicity and heterogeneous learning. Technical report, National Bureau of Economic Research. 1
- Barro, R. J. and Gordon, D. B. (1983). A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91(4):589–610. 5, 8
- Bernheim, B. D. and Taubinsky, D. (2018). Behavioral public economics. Technical report, National Bureau of Economic Research. 1
- Blesse, S., Buhlmann, F., and Doerrenberg, P. (2019). Do people really want a simple tax system? Evidence on preferences towards income tax simplification. *Forthcoming*. 5
- Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018. 3, 22, 25, 28, 31, 32, 37, 49
- Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *American Economic Review*, 105(5):1–33. 1

- Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–75. 3
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77. 1
- Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78. 5
- Conseil Constitutionnel (2005). Décision n. 2005-530 DC du 29 décembre 2005. 5
- Conseil Constitutionnel (2012). Décision n. 2012-662 DC du 29 décembre 2012. 5
- Diamond, P. A. (1998). Optimal income taxation: An example with a U-shaped pattern of optimal marginal tax rates. *American Economic Review*, pages 83–95. 35
- Farhi, E. and Gabaix, X. (2018). Optimal taxation with behavioral agents. *Forthcoming*. 1, 4, 11, 18, 20
- Fochmann, M., Kiesewetter, D., Blaufus, K., Hundsdoerfer, J., and Weimann, J. (2010). Tax perception - an empirical survey. *Arqus Discussion Paper*. 3
- Fochmann, M. and Weimann, J. (2013). The effects of tax salience and tax experience on individual work efforts in a framed field experiment. *FinanzArchiv: Public Finance Analysis*, 69(4):511–542. 3
- Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics, Vol 2*, edited by D Bernheim, S DellaVigna and D Laibson, *Forthcoming*. Elsevier. 2, 6, 7, 11, 15, 27
- Gerritsen, A. (2016). Optimal taxation when people do not maximize well-being. *Journal of Public Economics*, 144:122–139. 1, 18
- Gideon, M. (2017). Do individuals perceive income tax rates correctly? *Public Finance Review*, 45(1):97–117. 3
- Goldin, J. (2015). Optimal tax salience. *Journal of Public Economics*, 131:115–123. 1
- Hoopes, J. L., Reck, D. H., and Slemrod, J. (2015). Taxpayer search for information: Implications for rational attention. *American Economic Journal: Economic Policy*, 7(3):177–208. 1

- Jacquet, L. and Lehmann, E. (2017). Optimal income taxation with composition effects. 33
- Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491. 5, 8
- Liebman, J. B. and Zeckhauser, R. J. (2004). Schmeduling. Technical report, Working Paper. 1
- Maćkowiak, B. and Wiederholt, M. (2015). Business cycle dynamics under rational inattention. *The Review of Economic Studies*, 82(4):1502–1532. 14
- Mackowiak, B. A., Matejka, F., Wiederholt, M., et al. (2018). Survey: Rational inattention, a disciplined behavioral model. Technical report, CEPR Discussion Papers. 2, 14, 42
- Matějka, F. and Tabellini, G. (2017). Electoral competition with rationally inattentive voters. *Forthcoming*. 16
- Miller, B. M., Mumford, K., et al. (2015). The salience of complex tax changes: Evidence from the child and dependent care credit expansion. *National Tax Journal*, 68(3):477–510. 1
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208. 17
- Morrison, W. and Taubinsky, D. (2019). Rules of thumb and attention elasticities: Evidence from under- and overreaction to taxes. *Mimeo*. 2
- Piketty, T. and Saez, E. (2013). Chapter 7 - optimal labor income taxation. In Alan J. Auerbach, Raj Chetty, M. F. and Saez, E., editors, *Handbook of Public Economics*, vol. 5, volume 5 of *Handbook of Public Economics*, pages 391 – 474. Elsevier. 6, 10, 18, 47
- Reck, D. (2016). Taxes and mistakes: What’s in a sufficient statistic? *Forthcoming*. 1, 11
- Rees-Jones, A. and Taubinsky, D. (2016). Heuristic perceptions of the income tax: Evidence and implications for debiasing. Technical report, National Bureau of Economic Research. 1, 10, 28
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, 68(1):205–229. 4, 5, 18, 36, 37, 49

- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2(3):180–212. 1
- Sausgruber, R. and Tyran, J.-R. (2005). Testing the Mill hypothesis of fiscal illusion. *Public choice*, 122(1-2):39–68. 2
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690. 13
- Taubinsky, D. and Rees-Jones, A. (2017). Attention variation and welfare: theory and evidence from a tax salience experiment. *The Review of Economic Studies*, 85(4):2462–2496. 1, 4, 16

A Appendix

A.1 Solution to the Toy Model with Imperfect Information

The government seeks to maximize tax revenue taking the prior as given. Its problem writes $\max_{\tau} \tau Y(1 - \tilde{\tau})$ such that $\tilde{\tau} = \xi\tau + (1 - \xi)\hat{\tau}$, $\{\tau, \hat{\tau}\} \in [0, 1]^2$ and $\xi \in (0, 1)$. The associated Lagrangian is $\mathcal{L}(\tau, \lambda) = \tau Y(1 - \xi\tau - (1 - \xi)\hat{\tau}) + \lambda(\tau - 1)$. Following from the first order Kuhn and Tucker conditions, $\tau = 1$ if and only if $\hat{\tau} \leq 1 - \frac{\xi}{1-\xi}e$ and $\tau = \frac{1-(1-\xi)\hat{\tau}}{\xi(1+e)}$ otherwise. These conditions are also sufficient since the problem is convex under the assumption that $\tau Y(1 - \tau)$ is concave.

At the rational equilibrium, the prior is correct $\hat{\tau} = \tau^*$. Guess that the rational equilibrium is interior. Hence, $\tau^* = \frac{1}{1+\xi e}$. Because $e > 0$, it implies that $\hat{\tau} > 1 - \frac{\xi}{1-\xi}e$ in equilibrium, thus confirming the guess. It is then straightforward to prove that $\tau^* Y(1 - \tau^*) < \tau^r Y(1 - \tau^r)$ where $\tau^r \equiv \frac{1}{1+e}$ as $\tau^r = \arg \max_{\tau \in [0,1]} \tau Y(1 - \tau)$. Moreover, the taxation bias $\tau^* - \tau^r = \frac{(1-\xi)e}{(1+\xi e)(1+e)}$ is strictly positive for all $\xi \in (0, 1)$.

A.2 Reformulation of the Tractable Gaussian Learning

The indirect utility of a taxpayer is given by equation (12). Performing a second order Taylor approximation of the latter around τ_0 gives

$$V_{\tau_0}^2(\tilde{\tau}, \tau_0, R_0; w) = V(\tau_0, \tau_0, R_0; w) + (\tilde{\tau} - \tau_0) \frac{\partial V}{\partial \tilde{\tau}} \Big|_{\tilde{\tau}=\tau_0} + \frac{(\tilde{\tau} - \tau_0)^2}{2} \frac{\partial^2 V}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau_0} \quad (\text{A.1})$$

where $\frac{\partial V}{\partial \tilde{\tau}} \Big|_{\tilde{\tau}=\tau_0} = 0$ and $\frac{\partial^2 V}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau_0} = \frac{\partial^2 y^*}{\partial \tilde{\tau}^2}$ from (7). Hence,

$$\iint V_{\tau}^2(\tilde{\tau}, \tau, R; w) \phi(s; \tau, \sigma) \phi(\tau; \hat{\tau}, \hat{\sigma}) ds d\tau = \int \left[V(\tau, \tau, R; w) + \frac{\tilde{\sigma}^2}{2} \frac{\partial^2 y^*}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau} \right] \phi(\tau; \hat{\tau}, \hat{\sigma}) d\tau$$

where $\tilde{\sigma}^2$ is the posterior variance and we are using the fact that with a Gaussian prior and a Gaussian signal, the posterior is also Gaussian. Accordingly, the expected information reduction writes

$$\mathcal{I}(\sigma) = \frac{1}{2} (\log(2\pi e \hat{\sigma}^2) - \log(2\pi e \tilde{\sigma}^2)) = \frac{1}{2} \log \frac{\hat{\sigma}^2}{\tilde{\sigma}^2} \quad (\text{A.2})$$

where $\frac{1}{2} \log(2\pi e \sigma^2)$ is the differential entropy (in bits) of a Gaussian distribution with variance σ^2 . Therefore, in a Gaussian model, problem (11) becomes

$$\max_{\tilde{\sigma} \geq \hat{\sigma}} \tilde{\sigma}^2 \int \frac{\partial^2 y^*}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau} \phi(\tau; \hat{\tau}, \hat{\sigma}) d\tau - \kappa \log \frac{\hat{\sigma}^2}{\tilde{\sigma}^2} \quad (\text{A.3})$$

This problem has been extensively studied in the literature. For instance, a step-by-step derivation of the solution is provided in Mackowiak et al. (2018). It shows that the

perceived tax rate is $\tilde{\tau} = \xi s + (1 - \xi)\hat{\tau}$ where $\xi \in [0, 1]$ is a measure of the attention level set optimally to

$$\xi = \max \left(0, 1 + \frac{\kappa}{\hat{\sigma}^2 \int \frac{\partial^2 y^*}{\partial \tilde{\tau}^2} \Big|_{\tilde{\tau}=\tau} \phi(\tau; \hat{\tau}, \hat{\sigma}) d\tau} \right) \quad (\text{A.4})$$

A.3 Proofs of Proposition 1 and 2

We here prove both propositions at the same time since the only difference between the two problems is in the nature of responses to tax changes that are taken into account. We thus solve the general problem where all agents' responses are taken into account (including equilibrium adjustments) to obtain Proposition 2 and from which Proposition 1 naturally follows.

The Lagrangian associated to problem (17) writes

$$\begin{aligned} \mathcal{L}(\tau_g, R, p) = & E_{\vartheta} \left[\iint \left[G(\mathcal{V}(\tilde{\tau}, \tau_g + \vartheta, R, \kappa; w)) \right. \right. \\ & \left. \left. + p \left((\tau_g + \vartheta) y^*(\tilde{\tau}; w) - R_0 - E \right) \right] f_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w) f_w(w) d\tilde{\tau} dw \right] \end{aligned} \quad (\text{A.5})$$

The first-order condition associated with the choice of the marginal tax rate τ_g is

$$\begin{aligned} \frac{1}{p} \frac{d\mathcal{L}}{d\tau_g} = & E_{\vartheta} \left[\int \left\{ \int \left[\frac{G'(\mathcal{V})}{p} \frac{d\mathcal{V}}{d\tau_g} + y^* \right] f_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w) d\tilde{\tau} \right. \right. \\ & \left. \left. + \int \left[\frac{G(\mathcal{V})}{p} + (\tau_g + \vartheta) y^* - R_0 - E \right] \frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g} d\tilde{\tau} \right\} f_w(w) dw \right] \end{aligned} \quad (\text{A.6})$$

where $\frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g}$ is the change in the posterior distribution of perceived tax rate for type w and captures agents' responses to tax changes.

By definition $\int f_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w) d\tilde{\tau} = 1$ thus $\int \frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g} d\tilde{\tau} = 0$. Moreover, the quasi-linearity of utility implies that $\frac{d\mathcal{V}}{d\tau_g} = -y^*(\tilde{\tau}; w)$. Therefore, the optimality condition $\frac{1}{p} \frac{d\mathcal{L}}{d\tau_g} = 0$ writes

$$\begin{aligned} E_{\vartheta} \left[\int \left\{ \int \left[-\frac{G'(\mathcal{V})}{p} y^* + y^* \right] f_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w) d\tilde{\tau} \right. \right. \\ \left. \left. + \int \left[\frac{G(\mathcal{V})}{p} + (\tau_g + \vartheta) y^* \right] \frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g} d\tilde{\tau} \right\} f_w(w) dw \right] = 0 \end{aligned} \quad (\text{A.7})$$

This is equation (21) from Proposition 1 and characterizes the commitment tax rate. Equation (19) from Proposition 2 which characterizes the tax rate chosen by a discretionary government is obtained by when agents' responses to a change in the tax rate is computed holding agents' prior \hat{q} constant. That is $\frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g} \Big|_{\hat{q}(\cdot)}$ replaces $\frac{df_{\tilde{\tau}}(\tilde{\tau} | \tau_g + \vartheta; w)}{d\tau_g}$ in equation (A.7).

The first-order condition associated with the choice of the demogrant R is

$$\frac{1}{p} \frac{d\mathcal{L}}{dR} = E_{\vartheta} \left[\iint \left[\frac{G'(\mathcal{V})}{p} \frac{d\mathcal{V}}{dR} - 1 \right] f_{\tilde{\tau}}(\tilde{\tau}|\tau_g + \vartheta; w) f_w(w) d\tilde{\tau} dw \right] \quad (\text{A.8})$$

By quasi-linearity we have $\frac{d\mathcal{V}}{dR} = 1$. The optimality condition $\frac{1}{p} \frac{d\mathcal{L}}{dR} = 0$ thus writes

$$E_{\vartheta} \left[\iint \left[\frac{G'(\mathcal{V})}{p} - 1 \right] f_{\tilde{\tau}}(\tilde{\tau}|\tau_g + \vartheta; w) f_w(w) d\tilde{\tau} dw \right] = 0 \quad (\text{A.9})$$

This is equation (20) from Proposition 1 and equation (22) from Proposition 2.

A.4 Optimal policies in tractable Gaussian case

Conditions (A.7) and (A.9) apply to any learning leading to a differentiable posterior distribution of perceptions $f_{\tilde{\tau}}(\tau|\tau_0; w)$ with positive support on $[0, 1]$, where $\tau_0 = \tau_g + \vartheta$. Further insights may be gained by using a tractable Gaussian learning (Assumption 3). Indeed, in this case $f_{\tilde{\tau}}(\tau|\tau_0; w)$ is a Gaussian pdf $\phi(\tau; \mu, \sigma^2)$ with mean $\mu = \xi\tau_0 + (1 - \xi)\hat{\tau}$ and variance $\sigma^2 = \sigma^{*2}$. We can thus express agents' responses to tax reforms in terms of changes in the true tax rate τ_0 , changes in the prior mean $\hat{\tau}$ and induced changes in attention ξ that correspond to changes in the precision of the signal σ^* . To do so, we use a first-order approximation of the objective at the mean μ and exploit the following Lemma.

Lemma 3. *Let $\psi(x)$ be a differentiable real-valued function, $\psi_a(x) = \psi(a) + (x - a)\psi'(a)$ its first-order Taylor approximation evaluated at a and $\phi(x; \mu, \sigma^2)$ the pdf of the Gaussian distribution with mean μ and variance σ^2 . Then,*

$$\int_{\mathbb{R}} \psi_{\mu}(x) \phi(x; \mu, \sigma^2) dx = \psi(\mu) \quad (\text{A.10})$$

$$\int_{\mathbb{R}} \psi_{\mu}(x) \frac{\partial \phi(x; \mu, \sigma^2)}{\partial \mu} dx = \psi'(\mu) \quad (\text{A.11})$$

$$\int_{\mathbb{R}} \psi_{\mu}(x) \frac{\partial \phi(x; \mu, \sigma^2)}{\partial \sigma} dx = 0 \quad (\text{A.12})$$

Proof. Equation (A.10) directly follows from $\int_{\mathbb{R}} (x - \mu) \phi(x; \mu, \sigma^2) dx = 0$ by definition of the mean. To prove equation (A.11), realize that $\int_{\mathbb{R}} \frac{\partial \phi(x; \mu, \sigma^2)}{\partial \mu} dx = 0$ and $\frac{\partial \phi(x; \mu, \sigma^2)}{\partial \mu} = \frac{x - \mu}{\sigma^2} \phi(x; \mu, \sigma^2)$ so that $\int_{\mathbb{R}} (\psi(\mu) + (x - \mu)\psi'(\mu)) \frac{\partial \phi(x; \mu, \sigma^2)}{\partial \mu} dx = \frac{\psi'(\mu)}{\sigma^2} \int_{\mathbb{R}} (x - \mu)^2 \phi(x; \mu, \sigma^2) dx = \psi'(\mu)$. Equation (A.12) follows from the fact that $\int_{\mathbb{R}} \frac{\partial \phi(x; \mu, \sigma^2)}{\partial \sigma} dx = 0$ such that the integral of a constant is nil and that $\frac{\partial \phi(x; \mu, \sigma^2)}{\partial \sigma}$ is symmetric such that the integral of x also nil by a symmetry argument. \square

Rewriting equation (A.7) as

$$E_{\vartheta} \left[\int \left\{ \int \left[-\frac{G'(\mathcal{V})}{p} y^* + y^* \right] \phi(\tau; \mu, \sigma^2) d\tau \right. \right. \quad (\text{A.13}) \\ \left. \left. + \int \left[\frac{G(\mathcal{V})}{p} + \tau_0 y^* \right] \left(\frac{d\phi(\tau; \mu, \sigma^2)}{d\mu} \frac{d\mu}{d\tau_g} + \frac{d\phi(\tau; \mu, \sigma^2)}{d\sigma} \frac{d\sigma}{d\tau_g} \right) d\tau \right\} f_w(w) dw \right] = 0$$

allows us to apply Lemma 3 to obtain with $\mu = \xi\tau_0 + (1 - \xi)\hat{\tau}$

$$E_{\vartheta} \left[\int \left\{ \left[-\frac{G'(\mathcal{V})}{p}y^* + y^* \right] \Big|_{\tilde{\tau}=\mu} + \left[\left(\frac{G'(\mathcal{V})}{p}(\tilde{\tau} - \tau_0) + \tau_0 \right) \frac{dy^*}{d\tilde{\tau}} \frac{d\mu}{d\tau_g} \right] \Big|_{\tilde{\tau}=\mu} \right\} f_w(w) dw \right] = 0 \quad (\text{A.14})$$

since taking $\psi(\tau) = \left[\frac{G(\mathcal{V})}{p} + \tau_0 y^* \right](\tau)$ implies $\psi'(\mu) = \left[\left(\frac{G'(\mathcal{V})}{p}(\tilde{\tau} - \tau_0) + \tau_0 \right) \frac{dy^*}{d\tilde{\tau}} \right](\mu)$ by the modified envelope condition. Recall that $\mu = \xi\tau_0 + (1 - \xi)\hat{\tau}$. Now, in equilibrium we have by definition that $\hat{\tau} = \tau_g + b$ meaning $\mu = \tau_g + \xi\vartheta + (1 - \xi)b$ and $\mu - \tau_0 = (1 - \xi)(b - \vartheta)$. Hence, in equilibrium,

$$E_{\vartheta} \left[\int \left\{ \left[-\frac{G'(\mathcal{V})}{p}y^* + y^* \right] \Big|_{\tilde{\tau}=\tau_g+\xi\vartheta+(1-\xi)b} + \left[\left(\frac{G'(\mathcal{V})}{p}(1 - \xi)(b - \vartheta) + \tau_g + \vartheta \right) \frac{dy^*}{d\tilde{\tau}} \frac{d\mu}{d\tau_g} \right] \Big|_{\tilde{\tau}=\tau_g+\xi\vartheta+(1-\xi)b} \right\} f_w(w) dw \right] = 0 \quad (\text{A.15})$$

Last, we characterize taxpayers' average response to tax reforms $\frac{d\mu}{d\tau_g}$ as computed under discretion and commitment. Under discretion, the policymaker takes agent's priors and thus attention strategies as given, hence $\frac{d\mu}{d\tau_g} = \xi \frac{d\tau_0}{d\tau_g} = \xi$ which yields equation (24). Under commitment, the policymaker internalizes the equilibrium condition that priors and thus attention strategies adjust to the tax policy such that $\frac{d\mu}{d\tau_g} = \xi \frac{d\tau_0}{d\tau_g} + (1 - \xi) \frac{d\hat{\tau}}{d\tau_g} + \frac{d\xi}{d\tau_g}(\tau_0 - \hat{\tau}) = 1 + \frac{d\xi}{d\tau_g}(\vartheta - b)$ in equilibrium. This yields equation (25).

Transversality conditions follow from a direct application of Lemma 3 to equation A.9 with again $\mu = \tau_g + \xi\vartheta + (1 - \xi)b$:

$$E_{\vartheta} \left[\int \frac{G'(\mathcal{V})}{p} \Big|_{\tilde{\tau}=\mu} f(w) dw \right] = 1 \quad (\text{A.16})$$

A.5 Sufficient statistics formulas in tractable Gaussian case

Taking a small noise approximation, characterizations of equilibrium tax rates under discretion τ_g^{eq} and commitment τ_g^* in this tractable Gaussian model write

$$\begin{aligned} & \int \left[\left(1 - \frac{G'(\mathcal{V})}{p} \right) y^* + \left(\frac{G'(\mathcal{V})}{p}(1 - \xi)b + \tau_g^{\text{eq}} \right) \frac{dy^*}{d\tilde{\tau}} \xi \right] \Big|_{\tilde{\tau}=\tau_g^{\text{eq}}+(1-\xi)b} f_w(w) dw = 0 \\ & \int \left[\left(1 - \frac{G'(\mathcal{V})}{p} \right) y^* + \left(\frac{G'(\mathcal{V})}{p}(1 - \xi)b + \tau_g^* \right) \frac{dy^*}{d\tilde{\tau}} \left(1 - \frac{d\xi}{d\tau_g} b \right) \right] \Big|_{\tilde{\tau}=\tau_g^*+(1-\xi)b} f_w(w) dw = 0 \end{aligned}$$

Assuming preferences are iso-elastic, $U(c, y; w) = c - \frac{(y/w)^{1+\varepsilon}}{1+\varepsilon}$, the elasticity of earnings with respect to the perceived marginal net-of-tax rate e is constant

$$\forall \tilde{\tau}, w, e \equiv \frac{1 - \tilde{\tau}}{y^*} \frac{dy^*}{d(1 - \tilde{\tau})} = \frac{1}{\varepsilon} \iff \frac{dy^*}{d\tilde{\tau}} = -e \frac{y^*}{1 - \tilde{\tau}} \quad (\text{A.17})$$

Plugging in e we get

$$\begin{aligned} & \int \left[\left(1 - \frac{G'(\mathcal{V})}{p}\right) y^* - \left(\frac{G'(\mathcal{V})}{p}(1 - \xi)b + \tau_g^{\text{eq}}\right) e \frac{y^*}{1 - \tilde{\tau}} \xi \right] \Big|_{\tilde{\tau}=\tau_g^{\text{eq}}+(1-\xi)b} f_w(w) dw = 0 \\ & \int \left[\left(1 - \frac{G'(\mathcal{V})}{p}\right) y^* - \left(\frac{G'(\mathcal{V})}{p}(1 - \xi)b + \tau_g^*\right) e \frac{y^*}{1 - \tilde{\tau}} \left(1 - \frac{d\xi}{d\tau_g} b\right) \right] \Big|_{\tilde{\tau}=\tau_g^*+(1-\xi)b} f_w(w) dw = 0 \end{aligned}$$

To further simplify these formulas we now make a small perception bias approximation $b \ll 1$. This allows us to use the approximation $\frac{1}{1 - \tau_g - (1 - \xi)b} \approx \frac{1}{1 - \tau_g}$ and to assume $\frac{d\xi}{d\tau_g} b \ll 1$ to simplify some terms¹. Defining social marginal welfare weights $g(w) \equiv \frac{G'(\mathcal{V})}{p}$ and the mean operator $\bar{x} = \int x(w) f(w) dw$ we get

$$\begin{aligned} & \left\{ \overline{\left(1 - g\right) y^*} - \frac{\tau_g^{\text{eq}}}{1 - \tau_g^{\text{eq}}} \overline{y^* \xi} e - \frac{b}{1 - \tau_g^{\text{eq}}} \overline{g(1 - \xi) y^* \xi} e \right\} \Big|_{\tilde{\tau}=\tau_g^{\text{eq}}+(1-\xi)b} = 0 \\ & \left\{ \overline{\left(1 - g\right) y^*} - \frac{\tau_g^*}{1 - \tau_g^*} \overline{y^*} e - \frac{b}{1 - \tau_g^*} \overline{g(1 - \xi) y^*} e \right\} \Big|_{\tilde{\tau}=\tau_g^*+(1-\xi)b} = 0 \end{aligned}$$

which simplify to the compact sufficient statistics formulas

$$\tau_g^{\text{eq}} = \frac{\overline{(1 - g) y^*}}{\overline{(1 - g) y^*} + \overline{y^* \xi} e} - b \frac{\overline{g(1 - \xi) y^* \xi} e}{\overline{(1 - g) y^*} + \overline{y^* \xi} e} \quad (\text{A.18})$$

$$\tau_g^* = \frac{\overline{(1 - g) y^*}}{\overline{(1 - g) y^*} + \overline{y^*} e} - b \frac{\overline{g(1 - \xi) y^*} e}{\overline{(1 - g) y^*} + \overline{y^*} e} \quad (\text{A.19})$$

where all endogenous quantities on the right hand-side of the equations are evaluated at respectively $\tilde{\tau} = \tau_g^{\text{eq}} + (1 - \xi)b$ and $\tilde{\tau} = \tau_g^* + (1 - \xi)b$. In other words formulas are expressed in terms of sufficient statistics evaluated at the optimum.

A.6 Taxation bias in tractable Gaussian case

A difficulty in comparing τ_g^{eq} and τ_g^* is that some right-hand side quantities are endogenous to the tax rate and thus evaluated at different tax rates. To overcome this difficulty, we use a small taxation bias approximation $\tau_g^{\text{eq}} \approx \tau_g^* = t$ such that quantities can be evaluated to a first-order approximation at the same tax rate. Furthermore, we assume that corrective

¹In our simulations we do check that $\frac{d\xi}{d\tau_g}$ does not take large values (it takes values between 0.2 and 1 in equilibrium) as a way to confirm the validity of this approximation.

motives associated to the presence of a perception bias b are evaluated at tax rate t such that we can finally directly compare

$$\tau_g^{\text{eq}} = \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^* + y^*\xi} e} - b \frac{\overline{g(1-\xi)y^*\xi} e}{\overline{(1-g)y^* + y^*\xi} e} \quad (\text{A.20})$$

$$\tau_g^* = \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^* + \bar{y}^*} e} - b \frac{\overline{g(1-\xi)y^*} e}{\overline{(1-g)y^* + \bar{y}^*} e} \quad (\text{A.21})$$

The first terms on the right-hand side corresponds to the standard optimal tax formula (e.g. Piketty and Saez (2013)) whereas the second are corrective terms associated to the existence of a perception bias b . For small perception biases, these corrective terms are second-order and go in the same direction for both positive and normative tax rates. They are thus not driving the difference between the two and we disregard them to derive the following simple sufficient statistics formula for the taxation bias

$$\tau_g^{\text{eq}} - \tau_g^* = \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^* + \bar{y}^*\xi} e} - \frac{\overline{(1-g)y^*}}{\overline{(1-g)y^* + \bar{y}^*} e} \quad (\text{A.22})$$

$$= \frac{\overline{(1-g)y^*} \left(\overline{(1-g)y^* + \bar{y}^*} e \right) - \overline{(1-g)y^*} \left(\overline{(1-g)y^* + \bar{y}^*\xi} e \right)}{\left(\overline{(1-g)y^* + \bar{y}^*\xi} e \right) \left(\overline{(1-g)y^* + \bar{y}^*} e \right)} \quad (\text{A.23})$$

$$= \frac{e \tau_g^{\text{eq}} \tau_g^*}{\overline{(1-g)y^*}} (\bar{y}^* - \bar{y}^*\xi) \quad (\text{A.24})$$

$$\simeq \frac{\overline{(1-\xi)y^*}}{\overline{(1-g)y^*}} e t^2 \quad (\text{A.25})$$

A.7 Utility decomposition

Let $\mathcal{V}_\kappa(\tau_g^*(\kappa), w)$ be the expected utility of taxpayer w at the positive equilibrium when the information cost is κ and the optimal target tax rate of the government is $\tau_g^*(\kappa)$. Then, with a separable utility,

$$\begin{aligned} \mathcal{V}_0(\tau_g^*(0), w) &= E_{\tau_0|\tau_g^*(0)} \left[R_0 + (1-\tau_0)y^*(\tau_0; w) - v(y^*(\tau_0; w); w) \right] \\ \mathcal{V}_\kappa(\tau_g^*(\kappa), w) &= E_{\tau_0|\tau_g^*(\kappa)} \left[\int \left(R_0 + (1-\tau_0)y^*(\tau; w) - v(y^*(\tau; w); w) \right) f_{\hat{\tau}}(\tau|\tau_0, w) d\tau \right] - \kappa \mathcal{I}(\sigma^*(\hat{q}(\tau), \kappa, w)) \end{aligned}$$

Using straightforward algebra,

$$\mathcal{V}_\kappa(\tau_g^*(\kappa), w) - \mathcal{V}_0(\tau_g^*(0), w) = \Delta^\nu R + \Delta^\nu \tau + \Delta^\nu b + \Delta^\nu \text{uncertainty} + \Delta^\nu \text{info cost} \quad (\text{A.26})$$

where

$$\Delta^\nu R \equiv E_{\tau_0|\tau_g^*(\kappa)}[R_0] - E_{\tau_0|\tau_g^*(0)}[R_0]$$

is the change in the average demogrant,

$$\Delta^{\mathcal{V}}\tau \equiv E_{\tau_0|\tau_g^*(\kappa)}[(1 - \tau_0)y^*(\tau_0; w) - v(y^*(\tau_0; w); w)] - E_{\tau_0|\tau_g^*(0)}[(1 - \tau_0)y^*(\tau_0; w) - v(y^*(\tau_0; w); w)]$$

is the change in the expected utility due to the change in the tax target τ_g^* ,

$$\begin{aligned} \Delta^{\mathcal{V}}b \equiv & E_{\tau_0|\tau_g^*(\kappa)}[(1 - \tau_0)(y^*(\tau_0 + (1 - \xi)b; w) - y^*(\tau_0; w)) \\ & - (v(y^*(\tau_0 + (1 - \xi)b; w); w) - v(y^*(\tau_0; w); w))] \end{aligned}$$

is the change in the expected utility due to the bias b ,

$$\begin{aligned} \Delta^{\mathcal{V}}\text{uncertainty} \equiv & E_{\tau_0|\tau_g^*(\kappa)} \left[\int (1 - \tau_0)(y^*(\tau; w) - v(y^*(\tau; w); w)) \right. \\ & \left. \phi(\tau; \xi\tau_0 + (1 - \xi)(\tau_g^*(\kappa) + b), (\xi\sigma^*)^2) d\tau \right. \\ & \left. - ((1 - \tau_0)(y^*(\tau_0 + (1 - \xi)b; w) - v(y^*(\tau_0 + (1 - \xi)b; w); w))) \right] \end{aligned}$$

is the change in the expected utility due to noisy information and $\Delta^{\mathcal{V}}\text{info cost} = -\kappa\mathcal{I}(\sigma^*(\hat{q}(\tau), \kappa, w))$.

Online Appendix

(Not for publication)

This online appendix first provides detailed information on the numerical simulations. Second, it gives all proofs and derivations for the extension to nonlinear taxation. Third, we show how to incorporate income effects in the analysis.

B Numerical simulations

Simulations are implemented using Matlab and the algorithm may be summarized as follows. We first estimate a log-normal distribution of skills that we extend with a Pareto tail. This distribution of skills is then binned into a discrete approximation and taken as given for the rest of the exercise. Second, we find the optimal policy of the government using an iterative routine. Starting with a guess for the optimal policy, we compute the optimal attention strategies and allocations in equilibrium (i.e. when the priors are adjusted). We then compute a new optimal policy given taxpayers' choices and iterate until convergence to a fixed point solution.

This appendix provides details on these different steps. We first present the calibration strategy for the skill distribution. Second, we explain how to solve for the optimal attention strategies and allocations for a given tax schedule. Finally, we discuss how the government's problem is solved in the linear tax setting before turning to the nonlinear case.

B.1 Skill distribution

Simulations require an exogenous distribution of skills $f_w(\cdot)$. We fit the adjusted gross incomes from the 2016 March CPS data to a log-normal distribution. The parameters of the log-normal are chosen to exactly match the mean and median of the observed distribution. Following Saez (2001), we extend the log-normal distribution with a Pareto tail ($k = 2$) for annual incomes above \$200,000. We then discretize the income distribution using evenly distributed bins over the $[200; 200,000]$ interval and evenly distributed bins (in ln scale) over the $[200,000; 4,000,000]$ interval. This allows us to approximate integrals with Riemann sums.

To translate this income distribution into a skill distribution, we invert agents' first-order conditions for labor supply. We first use OECD data on 2016 labor taxes in the US and fit a linear tax schedule $\{\tau_{obs}, R_{obs}\}$. Then, we impose a quasi-linear utility specification $u(c, y; w) = c - (y/w)^{1+\epsilon}/(1+\epsilon)$ with $e = 1/\epsilon = 0.33$ (Chetty, 2012).

Assuming we are in a no bias equilibrium (i.e. rational expectation) such that agents' perceived tax rate coincide with the observed one τ_{obs} , this allows us to compute skills through $w = (y^\epsilon / (1 - \tau_{obs}))^{\frac{1}{1+\epsilon}}$. We also use the estimated linear tax system $\{\tau_{obs}, R_{obs}\}$ together with the actual distribution of earnings to deduce an exogenous expenditure requirement E for the government budget constraint.

B.2 Taxpayers' behavior

Taxpayers' choices are presented in Section 3. For the simulations, we consider Gaussian implementation shocks. Under this assumption, the equilibrium prior distribution is Gaussian as well. Consequently, one may easily compute the attention parameter (ξ), income (y) and consumption (c) for each taxpayer. Given an attention cost κ , a marginal tax rate τ – that potentially varies for each individual – and an uncertainty parameter σ_ϑ , the attention strategy in equilibrium follows from equation (14). Gaussian integrals are approximated using Gauss-Hermite quadratures. Using an agent's first-order condition (7) and budget constraint, we compute her income, consumption and utility for different signal realizations. These computations are made for each type of agent w . The demogrant R is computed from the government budget constraint.

B.3 Optimal linear tax

Unless stated otherwise, we assume throughout our numerical exercise that the social planner has a log objective $G(\cdot) = \log(\cdot)$.

In order to compute the optimal linear tax under discretion, we start with a guess $\tau_{g,0}$. Using this guess, we can deduce each taxpayer's attention strategy when the prior is adjusted to the guess $\hat{\tau}_0 = \tau_{g,0} + b$. We then consider this distribution of attention strategies as constant and use a Matlab optimization routine to find a new $\tau_{g,1}$ which maximizes social welfare for these attention strategies. We then update the prior $\hat{\tau}_1 = \tau_{g,1} + b$, recompute the attention strategies and re-optimize until convergence $|\hat{\tau}_i - \tau_{g,i+1}| \leq 1e^{-5}$. This method is intuitive and captures the essence of the discretionary policy: the government maximizes its objective taking attention strategies as fixed.

We also implement an alternative algorithm where instead of maximizing social welfare numerically we directly pick a new tax rate using the government FOCs in Proposition 1 under a small signals approximation. We find comparable equilibrium rates. Similarly, we compute the optimal policy under commitment using the FOCs in Proposition 2.

B.4 Optimal nonlinear tax

In order to compute the optimal nonlinear tax, we again use an iterative routine. We start with a guess – namely, a constant marginal rate – and iterate until convergence of the nonlinear tax schedule. We only present results for the unbiased equilibrium $b = 0$. We proceed in the same spirit as for the linear tax schedule:

1. Start with a guess for the nonlinear tax schedule
2. Compute the attention strategies $(\forall w)$ for a given adjusted prior $\hat{\tau}_w$
3. Compute allocations given attention strategies and tax schedule
4. Solve for the government FOCs at each w to deduce a new tax schedule
5. Repeat steps 1-4 until convergence.

To maintain the numerical stability of the algorithm we impose a slow adjustment of attention strategies ξ at each iteration. Indeed, marginal tax rates being sensitive to attention, one shall avoid large jumps in the attention parameter. The convergence criteria we use is the infinite norm for both marginal tax rates and attention strategies.

C Proofs for the extension to non-linear taxation

We here provide the proofs on the monotonicity condition and Proposition 5 (ABCD formula) of the main text.

C.1 Monotonicity

In this section, we demonstrate that the monotonicity condition is expected to hold for the quasi-linear and iso-elastic separable utility function that we consider in our simulations. For alternative specifications, we recommend to proceed using a guess-and-verify method. The latter is already implemented in our code and a warning is automatically displayed when the monotonicity does not hold ex post.

With a quasi-linear and iso-elastic separable utility function the first-order condition defining $y^*(\tilde{\tau}_w; w)$ is

$$(FOC)_y : 1 - \tilde{\tau}_w - \frac{1}{w} \left(\frac{y^*}{w} \right)^\epsilon = 0 \quad (27)$$

Differentiating this equation with respect to w yields

$$\frac{\epsilon}{w^2} \left(\frac{y^*}{w} \right)^{\epsilon-1} \frac{dy^*(\tilde{\tau}_w; w)}{dw} = \frac{1 + \epsilon}{w^2} \left(\frac{y^*}{w} \right)^\epsilon - \frac{d\tilde{\tau}_w}{dw} \quad (28)$$

Now – in expectation of the realization of the implementation shock ϑ – we also have $\tilde{\tau}_w = T'_g(y^*(\hat{\tau}_w; w)) + (1 - \xi)b$ which allows us to get

$$\frac{d\tilde{\tau}_w}{dw} = T''_g(y^*(\hat{\tau}_w; w)) \frac{dy^*(\hat{\tau}_w; w)}{dw} + \frac{d}{dw} [(1 - \xi)b] \quad (29)$$

and we can show that

1. If agents correctly perceive marginal tax rates ($b = 0$), the equilibrium condition $\hat{\tau}_w = T'_g(y^*(\hat{\tau}_w; w)) + b$ becomes $\hat{\tau}_w = T'_g(y^*(\hat{\tau}_w; w)) = \tilde{\tau}$. We then have $\frac{dy^*(\hat{\tau}_w; w)}{dw} = \frac{dy^*(\tilde{\tau}_w; w)}{dw}$ such that plugging (29) with $b = 0$ into (28) the monotonicity condition boils down to

$$\frac{dy^*}{dw} = \frac{\frac{1+\epsilon}{w^2} \left(\frac{y^*}{w}\right)^\epsilon}{\frac{\epsilon}{w^2} \left(\frac{y^*}{w}\right)^{\epsilon-1} + T''_g(y^*)} \geq 0 \iff -T''_g(y^*) \leq \frac{\epsilon}{w^2} \left(\frac{y^*}{w}\right)^{\epsilon-1} \quad (30)$$

2. If agents exhibit a small perception bias ($b \approx 0$) such that we have $\frac{dy^*(\hat{\tau}_w; w)}{dw} \approx \frac{dy^*(\tilde{\tau}_w; w)}{dw}$ plugging (29) into (28) the monotonicity condition rewrites

$$\frac{dy^*}{dw} = \frac{\frac{1+\epsilon}{w^2} \left(\frac{y^*}{w}\right)^\epsilon - \frac{d}{dw} [(1 - \xi)b]}{\frac{\epsilon}{w^2} \left(\frac{y^*}{w}\right)^{\epsilon-1} + T''_g(y^*)} \geq 0 \iff \begin{cases} -T''_g(y^*) \leq \frac{\epsilon}{w^2} \left(\frac{y^*}{w}\right)^{\epsilon-1} \\ \frac{d}{dw} [(1 - \xi)b] \leq \frac{1+\epsilon}{w^2} \left(\frac{y^*}{w}\right)^\epsilon \end{cases} \quad (31)$$

where the equivalence comes from the fact that the other case in which we would have $-T''_g(y^*) \geq \frac{\epsilon}{w^2} \left(\frac{y^*}{w}\right)^{\epsilon-1}$ is infeasible.

Hence, the monotonicity condition will hold if the tax function $T_g(y)$ is sufficiently smooth such that its second derivative is bounded (in absolute value).

C.2 Proposition 5 (ABCD formula)

We proceed with a tax perturbation approach in order to characterize the nonlinear tax schedule chosen under discretion and under commitment. Consider a tax schedule $T_g(\cdot)$ and a reform that consists in a small increase $\Delta\tau^r$ in marginal tax rates in a small bandwidth of earnings $[y^r - \Delta y^r, y^r]$ and let us compute its impact on the government's objective (written in Lagrangian form)

$$\mathcal{L} = E_\vartheta \left[\iint \left\{ G\left(\mathcal{V}(\tilde{\tau}_w, T_0(\cdot); \kappa, w)\right) + p\left(T_0(y^*(\tilde{\tau}_w; w)) - E\right) \right\} f_{\tilde{\tau}_w}(\tau|\tau_0; w) f_w(w) d\tau dw \right] \quad (32)$$

where p is the multiplier associated to the government's budget constraint and is equal to the social marginal value of public funds at the optimum.

Impact of the reform For a given target tax schedule $T_g(\cdot)$, the reform has

- a mechanical effect dM and a welfare effect dW that translate the lump-sum increase of $\Delta\tau^r \Delta y^r$ in the tax liabilities of agents $w \in [w^r, \infty[$ defined by $y^*(\tilde{\tau}_{w^r}; w^r) \equiv y^r$ where $E_s[\tilde{\tau}_w | \hat{\tau}_w] = \xi T'_0(y^*(\hat{\tau}_w; w)) + (1 - \xi)\hat{\tau}_w$ with $T'_0 = T'_g + \vartheta$
- a labor supply or behavioral effect dB that translates an increase $\Delta\tau^r$ in marginal tax rates that impacts the perceived marginal tax rates $\tilde{\tau}_w$ of agents $w \in [\hat{w}^r - \Delta\hat{w}^r, \hat{w}^r]$ defined by $y^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r) \equiv y^r$ and $y^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r - \Delta\hat{w}^r) \equiv y^r - \Delta y^r$

such that the total impact on the government's objective is

$$\frac{d\mathcal{L}}{p} = \frac{dM}{p} + \frac{dW}{p} + \frac{dB}{p} \quad (33)$$

with

$$\begin{aligned} \frac{dM}{p} + \frac{dW}{p} &= \int_{w^r}^{\infty} E_{\vartheta} \left[\int \left\{ \Delta\tau^r \Delta y^r - \frac{G'(\mathcal{V}(w))}{p} \frac{\partial U}{\partial c} \Delta\tau^r \Delta y^r \right\} f_{\tilde{\tau}}(\tau | \tau_0; w) d\tau \right] f_w(w) dw \\ &= \int_{w^r}^{\infty} E_{\vartheta} \left[\int (1 - g(w)) \Delta\tau^r \Delta y^r f_{\tilde{\tau}}(\tau | \tau_0; w) d\tau \right] f_w(w) dw \end{aligned} \quad (34)$$

since we here have, holding $\tilde{\tau}_w$ constant,

$$d\mathcal{V} = \frac{d}{dc} \left\{ U \left(\underbrace{y^*(\tilde{\tau}_w; w) - T_0(y^*(\tilde{\tau}_w; w))}_c, y^*(\tilde{\tau}_w; w); w \right) - \kappa \mathcal{I}(\sigma^*) \right\} dc = -\frac{\partial U}{\partial c} dT_0 \quad (35)$$

and

$$\begin{aligned} \frac{dB}{p} &= \int_{\hat{w}^r - \Delta\hat{w}^r}^{\hat{w}^r} E_{\vartheta} \left[\int \left\{ \frac{G(\mathcal{V}(w))}{p} + T_0(y^*(\tilde{\tau}; w)) \right\} \frac{df_{\tilde{\tau}}(\tau | \tau_0; w)}{d\tau_g} \Delta\tau^r d\tau \right] f_w(w) dw \\ &\approx E_{\vartheta} \left[\int \left\{ \frac{G(\mathcal{V}(\hat{w}^r))}{p} + T_0(y^*(\tilde{\tau}; \hat{w}^r)) \right\} \frac{df_{\tilde{\tau}}(\tau | \tau_0; \hat{w}^r)}{d\tau_g} \Delta\tau^r d\tau \right] f_w(\hat{w}^r) \Delta\hat{w}^r \end{aligned} \quad (36)$$

since we here have, holding $\tilde{\tau}_w$ constant,

$$d\mathcal{V} = \frac{d}{d\tau_g} \left\{ U \left(y^*(\tilde{\tau}_w; w) - T_0(y^*(\tilde{\tau}_w; w)), y^*(\tilde{\tau}_w; w); w \right) - \kappa \mathcal{I}(\sigma^*) \right\} d\tau_g = 0 \quad (37)$$

Characterization of tax policy The optimality condition for the choice of tax policy

$\frac{d\mathcal{L}}{p} = 0$ thus writes

$$\begin{aligned} E_{\vartheta} \left[\int \left\{ \frac{G(\mathcal{V}(\hat{w}^r))}{p} + T_0(y^*(\tilde{\tau}_{\hat{w}^r}; \hat{w}^r)) \right\} \frac{df_{\tilde{\tau}_{\hat{w}^r}}(\tau | \tau_0; \hat{w}^r)}{d\tau_g} d\tau \right] \frac{f_w(\hat{w}^r)}{\frac{dy^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r)}{dw}} \\ + \int_{w^r}^{\infty} E_{\vartheta} \left[\int (1 - g(w)) f_{\tilde{\tau}_w}(\tau | \tau_0; w) d\tau \right] f_w(w) dw = 0 \end{aligned} \quad (38)$$

where we have simplified through by $\Delta\tau^r \Delta y^r$ noting that

$$y^*(\hat{\tau}; \hat{w}^r - \Delta\hat{w}^r) \equiv y^r - \Delta y^r \quad \implies \quad \Delta\hat{w}^r \frac{dy^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r)}{dw} \approx \Delta y^r$$

Assuming we are in the tractable Gaussian case, the ex post (after learning) distribution of the perceived marginal tax rate is Gaussian $f_{\tilde{\tau}_w}(\tau|\tau_0; w) \sim \mathcal{N}(\mu_w, \sigma^2)$ with mean $\mu_w = \xi\tau_0 + (1 - \xi)\hat{\tau}_w$ and variance parameter $\sigma = \sigma^*$. Applying Lemma 3 we can thus rewrite the optimality condition as

$$E_\vartheta \left[\left[\left\{ \frac{G'(\mathcal{V}(\hat{w}^r))}{p} (\tilde{\tau}_{\hat{w}^r} - T'_0(y^*(\tilde{\tau}_{\hat{w}^r}; \hat{w}^r))) \right\} + T'_0(y^*(\tilde{\tau}_{\hat{w}^r}; \hat{w}^r)) \right] \frac{dy^*}{d\tilde{\tau}} \frac{d\mu_{\hat{w}^r}}{d\tau_g} \Big|_{\tilde{\tau}=\mu_{\hat{w}^r}} \frac{f_w(\hat{w}^r)}{\frac{dy^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r)}{dw}} \right. \\ \left. + \int_{w^r}^{\infty} [1 - g(w)] \Big|_{\tilde{\tau}=\mu_w} f_w(w) dw \right] = 0 \quad (39)$$

where $\frac{d\mu_w}{d\tau_g} = \xi \frac{d\tau_0}{d\tau_g} = \xi$ under discretion since the government takes agents' priors as given whereas $\frac{d\mu_w}{d\tau_g} = \xi \frac{d\tau_0}{d\tau_g} + (1 - \xi) \frac{d\hat{\tau}_w}{d\tau_g} + \frac{d\xi}{d\tau_g} (\tau_0 - \hat{\tau}_w)$ under commitment since the government internalizes that priors adjust to the choice of tax policy and are thus an endogenous object.

In addition, the Lagrange multiplier is – absent income effects – determined by the same transversality condition as before

$$E_\vartheta \left[\int_0^{\infty} [1 - g(w)] \Big|_{\tilde{\tau}=\mu_w} f_w(w) dw \right] = 0 \quad (40)$$

which can be obtained in a perturbation approach by computing the impact of a uniform lump-sum increase in taxes.

ABCD formula To obtain our ABCD formula from equation (39), let us introduce $e = \frac{1 - \tilde{\tau}_w}{y^*(\tilde{\tau}_w; w)} \frac{dy^*(\tilde{\tau}_w; w)}{d(1 - \tilde{\tau}_w)}$ and assume that the shock ϑ is small to use $E_\vartheta[\psi(\vartheta)] \approx \psi(E_\vartheta[\vartheta])$ regardless of function ψ 's curvature such that $E_\vartheta[\psi(\tau_0)] \approx \psi(\tau_0)$. This yields

$$\frac{T'_g(y^*(\mu_{\hat{w}^r}; \hat{w}^r)) + g(\hat{w}^r) \Big|_{\tilde{\tau}=\mu_{\hat{w}^r}} (\mu_{\hat{w}^r} - T'_g(y^*(\mu_{\hat{w}^r}; \hat{w}^r)))}{1 - \mu_{\hat{w}^r}} \\ = \frac{1}{e \frac{d\mu_{\hat{w}^r}}{d\tau_g} \Big|_{\tilde{\tau}=\mu_{\hat{w}^r}}} \frac{1}{y^*(\mu_{\hat{w}^r}; \hat{w}^r)} \frac{\frac{dy^*(\hat{\tau}_{\hat{w}^r}; \hat{w}^r)}{dw}}{f_w(\hat{w}^r)} \int_{w^r}^{\infty} (1 - g(w)) \Big|_{\tilde{\tau}=\mu_w} f_w(w) dw \quad (41)$$

where $\mu_w = \xi T'_g(y^*(\hat{\tau}_w; w)) + (1 - \xi)\hat{\tau}_w$ and $\frac{d\mu_w}{d\tau_g} = \xi \frac{dT'_g(y^*(\hat{\tau}_w; w))}{d\tau_g} = \xi$ under discretion since the government takes agents' priors as given whereas $\frac{d\mu_w}{d\tau_g} = \xi + (1 - \xi) \frac{d\hat{\tau}_w}{d\tau_g} + \frac{d\xi}{d\tau_g} (T'_g(y^*(\hat{\tau}_w; w)) - \hat{\tau}_w)$ under commitment since the government internalizes that priors adjust to the policy rule and are thus an endogenous object.

Note that with a quasi-linear and iso-elastic separable utility function we have $y^*(\tilde{\tau}_w; w) = w^{1+\frac{1}{\epsilon}}(1 - \tilde{\tau}_w)^{\frac{1}{\epsilon}}$ and $e = \frac{1}{\epsilon}$ such that

$$\frac{dy^*(\tilde{\tau}_w; w)}{dw} = \left(1 + \frac{1}{\epsilon}\right) w^{\frac{1}{\epsilon}} (1 - \tilde{\tau}_w)^{\frac{1}{\epsilon}} = \frac{1 + e}{w} y^*(\tilde{\tau}_w; w) \quad (42)$$

Assuming small (or no) perception biases such that $\tilde{\tau}_w \approx \hat{\tau}_w$ and $\frac{dy^*(\tilde{\tau}_w; w)}{dw} \approx \frac{dy^*(\hat{\tau}_w; w)}{dw}$ yields

$$\frac{T'_g(y^*(\mu_{\hat{w}^r}^{\text{eq}}; \hat{w}^r)) + g(\hat{w}^r)(1 - \xi)b}{1 - \tilde{\tau}_{\hat{w}^r}^{\text{eq}}} = \frac{1}{\frac{d\mu_{\hat{w}^r}}{d\tau_g}} \frac{1 + e}{e} \frac{1}{\hat{w}^r f_w(\hat{w}^r)} \int_{w^r}^{\infty} (1 - g(w)) f_w(w) dw \quad (43)$$

D Income effects

In this final section of the Online Appendix, we illustrate how the (linear tax) model in the paper could be extended to account for income effects and accordingly characterize tax policy under discretion and commitment. We now have to account for the fact that the average posterior tax rate is no longer a sufficient statistics for taxpayers' earnings choices. This requires a mere reformulation of the initial problem without income effects: integration in the government's problem is now with respect to the signal distribution.

In order to introduce income effects, it will prove useful to slightly reformulate taxpayers' problem introduced in Section 3. To this end, consider that there is a continuum of individuals at each skill w of size $f(w)$ and let $Y(\tau_0) \equiv \iint y^*(\cdot) \phi(s; \tau_0, \sigma^*) ds dF(w)$ be the aggregate earnings. Then, because the government budget constraint is binding at the optimum, the demogrant writes $R(\tau_0) = \tau_0 Y(\tau_0) - E$ as the overall population remains of size one. Further, and given that a taxpayer's budget constraint binds ex post, consumption adjusts such that $c_0 = R(\tau_0) + (1 - \tau_0)y$. Therefore, an agent's utility is $u(R(\tau_0) + (1 - \tau_0)y, y)$ for a realization τ_0 and earnings choice y .

Given the above reformulation, the only uncertainty arises from the randomness in the realized tax rate. An individual therefore chooses the signal precision σ and income y to maximize her expected utility

$$\sup_{\sigma, y|s} \iint u(R(\tau) + (1 - \tau)y, y; w) \phi(s; \tau, \sigma) \hat{q}(\tau) ds d\tau - \kappa \mathcal{I}(\sigma) \quad (44)$$

where admissible earnings policies for this individual's choice may depend on the signal s . Now, guess that the optimal attention strategy σ^* depends only on w , $\hat{q}(\cdot)$, and κ . As a consequence, the optimal earnings choice $y^*(s, w; \sigma^*, \hat{q}(\cdot))$ now solves

$$\int [(1 - \tau)u_c(R(\tau) + (1 - \tau)y^*, y^*; w) + u_y(R(\tau) + (1 - \tau)y^*, y^*; w)] f(\tau|s; \sigma^*, \hat{q}(\cdot)) d\tau = 0 \quad (45)$$

where $f(\tau|s; \sigma^*, \hat{q}(\cdot)) = \frac{\phi(s; \tau, \sigma^*) \hat{q}(\tau)}{\int \phi(s; \tau, \sigma^*) \hat{q}(\tau) d\tau}$ from Bayes rule. Assume that a solution to equation (45) exists. In turn, it implies that

$$\sigma^*(w, \hat{q}(\cdot), \kappa) = \arg \sup_{\sigma} \iint u(R(\tau) + (1 - \tau)y^*, y^*; w) \phi(s; \tau, \sigma) \hat{q}(\tau) ds d\tau - \kappa \mathcal{I}(\sigma) \quad (46)$$

thus confirming the guess on σ^* (when it exists). We can now define agents' indirect utility function

$$\mathcal{V}(s, \tau_0; w, \kappa, \hat{q}(\cdot)) \equiv u(R(\tau_0) + (1 - \tau_0)y^*, y^*; w) - \kappa\mathcal{I}(\sigma^*) \quad (47)$$

Turning to the government problem, it requires a mere variation from (17)

$$\max_{\tau_g} E_{\vartheta} \left[\iint G(\mathcal{V}(s, \tau_0; w, \kappa, \hat{q}(\cdot))) \phi(s; \tau_0, \sigma^*) f_w(w) d\tau dw \right] \quad (48)$$

Note that the inner integration is now with respect to the signal distribution $\phi(s; \tau_0, \sigma^*)$ and no longer with respect to the posterior distribution of perceived rates. This is because the perceived tax rate $\tilde{\tau}$ is no longer a sufficient statistics for earnings choices.

The first order condition for the target tax rate under discretion writes

$$E_{\vartheta} \left[\int \left\{ \int \left[G'(\mathcal{V}) \frac{d\mathcal{V}}{d\tau_g} \phi(s; \tau_0, \sigma^*) ds + \int G(\mathcal{V}) \frac{d\phi(s; \tau_0, \sigma^*)}{d\tau_g} ds \right] f_w(w) dw \right\} \right] = 0 \quad (49)$$

and the first order condition for the target tax rate under commitment writes

$$E_{\vartheta} \left[\int \left\{ \int \left[G'(\mathcal{V}) \left(\frac{d\mathcal{V}}{d\tau_g} + \frac{d\mathcal{V}}{d\hat{q}(\cdot)} \frac{d\hat{q}(\cdot)}{d\tau_g} \right) \phi(s; \tau_0, \sigma^*) ds + \int G(\mathcal{V}) \frac{d\phi(s; \tau_0, \sigma^*)}{d\tau_g} ds \right] f_w(w) dw \right\} \right] = 0 \quad (50)$$

This characterizes tax policy under discretion and commitment in the presence of income effects. The key difference between the two equations is the fact that the commitment tax policy takes into account the adjustment in the prior $\frac{d\hat{q}(\cdot)}{d\tau_g}$ whereas the discretion tax policy does not. This leads to a taxation bias.