The Optimal Degree of Discretion in Fiscal Policy

Guillaume Sublet

Université de Montréal and CIREQ †

August 3, 2023

The design of a fiscal rule involves a trade-off between committing governments to a fiscally responsible budget and giving governments the discretion to respond to shocks. What is the optimal degree of discretion for deficit-biased governments that are facing shocks to their fiscal needs? The tail of the distribution of shocks determines the optimal degree of discretion. If the tail is thin, an optimal rule features a cap on deficits enforced by off-equilibrium penalties. If the tail is thick, an optimal rule grants more discretion than a cap could achieve at the cost of on-equilibrium penalties. To measure the need for discretion, I use a tractable model of government spending to infer past fiscal needs from government finance data. Novel evidence of a Pareto tail of the distribution of shocks to fiscal needs for members of the European Union indicates a large need for discretion.

Keywords: Fiscal Rule, Delegation, Mechanism Design, Money Burning, Price vs. Quantity

JEL classification: D02, D82, E02, E6, F36, F45, F55, H1, H62, P16

† Email: guillaume.sublet@umontreal.ca.

This paper subsumes a previous paper circulated under the title "Fiscal Rules with Discretion for an Economic Union." I thank Manuel Amador and Pierre Yared for comments and suggestions that greatly improved the paper. I thank Davide Alonzo, Marco Bassetto, Anmol Bhandari, Roel Beetsma, Julien Bengui, Job Boerma, Rui Castro, Louphou Coulibaly, Alessandro Dovis, Mike Golosov, Barış Kaymak, Tim Kehoe, Rishabh Kirpalani, Ellen McGrattan, Thomas Phelan, Todd Schoellman, Smriti Upadhyay, Pierre-Olivier Weill and the audiences of academic seminars and conferences for helpful suggestions and comments. I thank Joan Gieseke for expert copyediting. I am responsible for any shortcomings of this paper.

Introduction

The European Union is currently debating the design of the Excessive Deficit Procedure of its fiscal rule (the Stability and Growth Pact). The debate centers on the fundamental trade-off in the design of a fiscal rule: how to optimally balance the need to impose discipline on a deficit-biased government with the need for discretion to respond to shocks to the country's fiscal needs. Should society impose discipline on its government with a fiscal rule featuring a graduated schedule of penalties on excessive deficits, a cap on deficits, or a combination of the two?

A main insight from the literature is that, under some conditions on the distribution of shocks to fiscal needs, a maximally enforced deficit limit is optimal (Halac and Yared (2022)). An assumption in much of the literature is that the cost of imposing discipline on a government is borne symmetrically by the government and the society. The sources of discipline on the government's fiscal policy, however, range from formal enforcement mechanisms such as the Excessive Deficit Procedure of the Stability and Growth Pact to informal mechanisms such as higher spreads on borrowing or negative coverage by the press. The point of departure of this paper is that the cost of imposing discipline on the government may be borne asymmetrically by the government and the society.

This paper makes two main contributions. First, I study the optimal design of a fiscal rule enforced by penalties whose burden is borne asymmetrically by the government and the society, as in the "leaky bucket" model of government finances (Amador and Bagwell (2013, 2020)). The analysis covers the full range of degrees of asymmetry in the cost of discipline. At one end of the spectrum, the asymmetry is maximal in the sense that the society is immune to the penalty on the government, and the optimal fiscal rule is a graduated schedule of penalties reminiscent of a Pigouvian tax. The other end of the spectrum nests the symmetric case, and the optimal fiscal rule is a cap on the deficit, in line with the main insight from the literature. The focus of this paper is on characterizing optimal fiscal rules for degrees of asymmetry between the two polar cases. While the main insight from the literature prevails under more stringent conditions, under a complementary set of conditions, a fiscal rule featuring a graduated schedule of penalties is optimal.

The second main contribution is a first step toward quantifying the trade-off between discipline and discretion in the design of a fiscal rule. I propose a tool to measure the need for discretion based on government finance data. As the theory shows, the thickness of the right tail of the distribution of shocks to fiscal needs provides a quantifiable measure of the need for discretion. I find evidence of a Pareto tail of the distribution of shocks to fiscal needs of members of the European Union, which indicates a large need for discretion.

I obtain these contributions by building on a line of research modeling the trade-off between discipline and discretion in the design of a fiscal rule (Amador, Werning, and Angeletos (2006), Halac and Yared (2014, 2022a, 2023)). Consider a small open economy whose government decides how much to spend in response to shocks to the country's fiscal needs. Suppose that while the government and the society agree on the need for discretion to respond to the shocks, they disagree on the relative value between present and future public spending. In particular, the government is shortsighted in the sense that it overvalues current spending. The shortsightedness captures, in a reduced form, the incentive to overborrow on the part of the government due to various reasons: political turnover (Alesina and Passalacqua (2016) offer a survey), heterogeneity in discounting among members of the society (Jackson and Yariv (2014, 2015)), or, for members of an economic union, the common pool problem caused by a common monetary authority that lacks the ability to commit (Beetsma and Uhlig (1999), Chari and Kehoe (2007), and Aguiar, Amador, Farhi, and Gopinath (2015)). The combination of shocks and a present-biased objective calls for both discretion and discipline. Lastly, a trade-off between discretion and discipline arises because the realization of the shock to fiscal needs is private information to the government, which implies that the prescription of the fiscal rule cannot be contingent upon the shock. I follow Halac and Yared (2022) in modeling a fiscal rule as a penalty schedule, which is a function of the government's choice.

The departure from the standard framework is that the cost of meting out a penalty on the government is borne asymmetrically by the government and the society. The problem of designing a fiscal rule maps to a mechanism design problem without transfers but with asymmetric "money burning," where money burning refers to a penalty that is meted out. The principal—here, the society—designs a penalty schedule to delegate the choice of fiscal policy to the agent—here, the government. While the literature has used either the first-order approach or global Lagrangian methods to solve similar problems, I show that the two methods complement one another. To start, I use insights from the standard first-order approach as a guide to guess a partial solution from locally optimal incentives.

The challenge, then, is to go from locally optimal incentives to global optimality. A commonly used approach in the literature consists of showing that, for any alternative to the optimal mechanism, there exists a valid perturbation that increases welfare. For mechanisms with money burning, however, the constraint on transfers makes this task particularly challenging (see the discussion in Athey, Atkeson, and Kehoe (2005), Ambrus and Egorov (2017), and Halac and Yared (2022)). To overcome this challenge, I use the powerful yet less commonly used global Lagrangian methods to obtain a guess for a full solution and determine the conditions under which the guess is indeed a solution. Because the Lagrangian accounts for the constraint on transfers, it suffices to study valid perturbations of the Lagrangian (see Appendix B for the global Lagrangian methods based on Luenberger (1969), Amador, Werning, and Angeletos (2006), and Amador and Bagwell (2013, 2020)). The upshot is a comprehensive characterization of optimal fiscal rules depending on easy-to-check and intuitive optimality conditions.

Using the first-order approach, I obtain a closed-form expression for a marginal penalty that balances the benefit of discipline with the incentive cost of limited discretion. The expression for the marginal penalty comes from the first-order condition of a triply relaxed problem, in which the monotonicity condition, the non-negativity constraint on penalties, and the non-negativity constraint on government spending are set aside. The closed-form expression obtains by rewriting the first-order condition as a wedge in the Euler equation of the government.

The wedge is the product of two terms with clear economic intuition. One term balances the need for discipline—measured by the degree of present bias—with the need for discretion—measured by the thickness of the tail of the distribution of shocks. A second term, which depends solely on the degree of asymmetry in the cost of meting out penalties, determines society's aversion to meting out penalties. As the degree of asymmetry goes to zero, society becomes infinitely averse to meting out penalties. To avoid meting out penalties while still imposing discipline, society finds it (locally) optimal to resort to extreme incentives—that is, either no penalty or prohibitively large penalties. Interestingly, this limiting case echoes the "bang-bang" result of Halac and Yared (2022). At the polar opposite, the degree of asymmetry is such that the society is immune to penalties levied on the government, in which case the wedge is reminiscent of a Pigouvian tax correcting for the present bias of the government.

The main theoretical result shows that, for a broad class of environments with asymmetric money burning, a hybrid fiscal rule featuring no penalties below a threshold and a graduated schedule of penalties conforming with the wedge above the threshold up to a cap on the deficit is optimal. Two insights emerge from the characterization of the optimal fiscal rule. The first insight relates to the optimal stringency of the cap. The second insight relates to the optimal leniency of on-equilibrium penalties on low levels of spending.

The optimal stringency of the cap depends on the balance between the need for discipline and the need for discretion. As suggested by the first-order approach, the thickness of the tail of the distribution of shocks to fiscal needs governs the need for discretion. If the tail is sufficiently thin, large fiscal needs are relatively unlikely, and the optimal fiscal rule imposes a cap on the deficit. In contrast, if the tail is thick, large fiscal needs are relatively likely, and the optimal rule does not impose a cap. Moreover, for a sufficiently thick tail, the absence of a fiscal rule is optimal because the need for discretion outweighs the need for discipline.

The second insight is specific to fiscal rules featuring on-equilibrium penalties. If the first-order approach suggests relying heavily on on-equilibrium penalties (i.e., the wedge schedule is positive and implements a monotonically increasing allocation), then it is optimal to exempt low levels of spending from on-equilibrium penalties. To gain intuition for the exemption, think of the penalty schedule as a level shifter plus the integral of the marginal penalty schedule. Because penalties cannot be negative, the level can only be shifted upward, which is undesirable. The exemption consists of setting the marginal penalty schedule at zero—instead of that conforming with the wedge—below a threshold as a partial substitute for shifting the level downward. Above the exemption threshold, however, the marginal penalty schedule conforms with the wedge schedule.

Intuitively, the benefit of the exemption is that it lowers the level of the penalty schedule above the exemption threshold. It does so while preserving the same marginal penalty schedule, and hence the same discipline, above the threshold. Below the threshold, however, there is a loss of discipline. Also, the exemption causes a kink in the penalty schedule, which is in contrast with the notch (i.e., a jump in level) in the penalty schedule of the Stability and Growth Pact.

The theory indicates that the thickness of the right tail of the distribution of fiscal needs is a quantifiable measure of the need for discretion and a key input to the design of an optimal fiscal rule. The task of measuring the need for discretion from government finance data has two challenges, however. First, fiscal needs are not directly observable. Second, the distribution of fiscal needs is an infinite dimensional object. A parametric assumption on the distribution would, to a large extent, presume the thickness of the tail instead of deducing it from the data (e.g., irrespectively of its mean and variance, a normal distribution has a thin tail).

In a first step, I use a tractable positive model to infer past fiscal needs from government finance data. The positive model is a standard model of (government) spending augmented with taste shocks that may be persistent. The taste shocks capture fiscal needs as a catchall for shocks to the need for a deficit, irrespectively of whether it originates in a shock to the need for public spending, a shock to the cost of servicing the debt, or a shock to government revenue, as in the normative theory. The identification of fiscal needs obtains from inverting the government policy function to recover the fiscal needs that rationalize the government finance data. A difficulty, however, is that the policy function depends on the distribution of fiscal needs, which is the object of interest to be inferred in the second step. Assuming log utility keeps the two-step approach tractable because I obtain a closed-form expression for the policy function that depends only on the first moment of the distribution of fiscal needs. As a result, the first step alone delivers the exact identification of past fiscal needs from government finance data.

In a second step, I infer the thickness of the tail of the distribution of measured fiscal needs. Because rare events featuring large observations—say, during recessions—are a common feature of government finance data, I use tools from heavy-tail analysis. Following Gabaix and Ibragimov (2011), a log-log plot of rank and size of measured spending needs can be used to detect evidence of a Pareto tail and to measure the thickness of the tail. I adopt a common practice in heavy-tail analysis by substantiating the analysis with a Hill plot (Resnick (2007), Chapter 4). I apply the two-step methodology to the case of the European Union and find novel evidence of a thick (Pareto) tail of the distribution of shocks to fiscal needs. The concluding section contains suggestions to reform the Stability and Growth Pact.

Related literature. More broadly, the theory relates to the literature on delegation, money burning, and the design of rules to discipline a policy-making authority to act in the best interest of the society (Holmström (1977), Melumad and Shibano (1991), Alonso and Matoushek (2008), Kováč and Mylovanov (2009), Athey, Atkeson, and Kehoe (2005), Amador, Werning, and Angeletos (2006), Ambrus and Egorov (2013, 2017), Amador and Bagwell (2013, 2020, 2022), Clayton and Schaab (2022), and Halac and Yared (2014, 2020a, 2020b, 2022a, 2022b, 2023)). The contribution of this paper is threefold. First, it contains a near-complete characterization of the optimal delegation contract with asymmetric money burning in a model with a multiplicative bias (Appendix A.1 considers the quadratic model with additive bias). Second, the analysis in this paper blends insights from the first-order approach, reminiscent of Diamond (1998)'s ABC formula, to obtain easy-to-check and intuitive optimality conditions from the global Lagrangian methods. The optimality conditions elicit the central role of the thickness of the tail of the distribution of shocks for the optimal balance between the needs for discipline and discretion.

Third, this paper takes a first step toward a quantitative application of the rich theory on the trade-off between discipline and discretion in the design of fiscal rules on deficits. The quantitative literature on fiscal rules has, so far, focused on inefficient debt levels—for example, as a result of the possibility of default (see, e.g., Azzimonti, Battaglini, and Coate (2016), Hatchondo, Martinez, and Roch (2020), Alfaro and Kanczuk (2019), and Aguiar, Amador, and Fourakis (2020); an exception is Bassetto and Sargent (2006)). Felli, Piguillem, and Shi (2021) show that the risk of default makes it optimal to introduce a default rule.

Because a Pigouvian tax is a price regulation and a cap is a quantity regulation, this paper also contributes to the literature on the choice of regulatory instrument. In an influential contribution, Weitzman (1974) finds that the relative curvature of the benefit and cost functions of economic activity determines whether fixing the price or fixing the quantity is preferable.[2] Perhaps surprisingly, the distribution of shocks does not affect the ranking of these two simple instruments. Allowing for a richer set of instruments, as I do in this paper, reveals the role of the distribution of shocks: a thick tail of the distribution of shocks calls for discretion, which favors on-equilibrium penalties.

1 Model

This section introduces a standard model used to analyze the design of a fiscal rule and relaxes the assumption of symmetry in the cost of burning money (Amador, Werning, and Angeletos (2006), Halac and Yared (2014, 2022a)).

[2]For instance, the consensus in environmental economics is that the marginal benefit curve of abatement is relatively flat, whereas the marginal cost curve is steep, which according to Weitzman (1974) favors price regulation over quantity regulation (see McKibbin Wilcoxen (2002)).

1.1 Economic environment

Consider a small open economy whose government decides how much to spend and borrow in response to shocks to the fiscal needs of its society. The government is present biased in the sense that it discounts the future at a higher rate than its society. The combination of shocks to fiscal needs and a present-biased objective creates a need for both discretion and discipline.

Formally, the preferences of the government over the allocation of public spending over time are represented by the objective function

$$\theta U(g) + \beta(W(x) - P(g)), \tag{1}$$

where $\theta U(\cdot)$ denotes the utility from government spending $g \geq 0$, and the continuation value $W(\cdot)$ is a function of future assets $x \in \mathbb{R}$.[3] Both U and W are twice continuously differentiable and strictly increasing. The index U is strictly concave and satisfies Inada conditions, and the index W is concave. Shocks to the country's *fiscal needs*, denoted θ, are private information to the government. The shocks follow a distribution F whose support is an interval Θ with lower bound $\underline{\theta} > 0$ and supremum $\bar{\theta} > \underline{\theta}$.[4] Although the supremum can be infinite, the first moment is assumed to be finite. The distribution F is twice continuously differentiable with density f. The *tail* of the distribution refers to $1 - F$.

The degree of *present bias* of the government is $1 - \beta \in (0, 1]$, and, to simplify the notation, the discount factor of the society is subsumed into the continuation value. A *fiscal rule* is a penalty schedule $P(\cdot)$ that is non-negative, $P(g) \geq 0$ for $g \geq 0$.

The government's budget constraint is $g + x = T$, where $T > 0$ denotes government revenues. For simplicity, the gross interest rate is exogenous and normalized to one. Substituting the budget constraint into the objective of the government (1) gives the government's problem

$$g(\theta) \in \arg\max_{g \geq 0} \; \theta U(g) + \beta(W(T - g) - P(g)). \tag{2}$$

In contrast, the society's objective is not present biased, and, from an ex ante perspective,

$$\int_{\Theta} [\theta U(g(\theta)) + W(T - g(\theta)) - \rho P(g(\theta))] \, dF(\theta), \tag{3}$$

[3] Amador, Werning, and Angeletos (2006) show that the results apply to a multi-period environment with iid shocks. Halac and Yared (2014) study an infinite horizon environment with persistent shocks.

[4] Shocks to government revenues map to shocks to fiscal needs with a CARA utility index (see Section 5.4 of Amador, Werning, and Angeletos (2006)).

where $1-\rho$ denotes the *asymmetry* in the cost of penalties for the economic union. This paper covers the full spectrum of the degrees of asymmetry in the cost of the penalties, $\rho \in [0,1]$. At one end of the spectrum (i.e., $\rho = 0$), the society is immune to the penalties on the government. At the other end of the spectrum (i.e., $\rho = 1$), penalties symmetrically affect the government and the society.[5] For $\rho \in (0,1)$, an on-equilibrium penalty mutually affects the society and the government, albeit asymmetrically.

Formally, the design of a fiscal rule consists of determining a penalty schedule $P(\cdot) \geq 0$ to maximize society's welfare (3) subject to the implementability constraint (2).

Economic union. Note that the analysis in this paper applies, at the cost of additional notation, to the design of a fiscal rule for an economic union that is not fiscally integrated. Appendix A.2 contains the additional notation needed for an economic union and the mapping to the environment of this section. In short, because the lack of fiscal integration limits the use of transfers between countries, the key trade-off remains between discretion and discipline—separately from insurance.[6]

Some useful definitions. An *allocation* is a contingent plan for government spending $g(\cdot)$ with $g(\theta) \geq 0$ for $\theta \in \Theta$. A fiscal rule $P(\cdot)$ *implements* an allocation $g(\cdot)$ if for $\theta \in \Theta$, $g(\theta)$ solves the government's problem (2). Consider a rule $P(\cdot)$ and the allocation it implements $g(\cdot)$. A strictly positive penalty on g is *on-equilibrium* if there exists $\theta \in \Theta$ such that, if θ realizes, the penalty is meted out (i.e., $g(\theta) = g$). Otherwise, the penalty is *off-equilibrium*. Let g_d denote the *discretionary allocation*, which is the allocation that solves the government problem (2) in the absence of a fiscal rule. Define the *wedge* Δ in the Euler equation of the government as follows: $(1-\Delta(g,\theta))\theta U'(g) = \beta W'(T-g)$. The wedge allows us to read the discipline imposed by a fiscal rule off the allocation it implements. The wedge evaluated at the discretionary allocation is $\Delta(g_d(\theta),\theta) = 0$. A rule that constrains the government to spend less (i.e., $g(\theta) \leq g_d(\theta)$) induces a positive wedge.

[5] The economic environment with $\rho = 1$ nests the model in Amador, Werning, and Angeletos (2006) and the model in Halac and Yared (2022a) without an upper bound on punishments (i.e., with unlimited enforcement).

[6] Harstad (2007) provides a rationale for limits on transfers to curb the incentive of members to strategically delay reaching an agreement while bargaining over the design of a common rule.

1.2 Designing a mechanism with asymmetric "money burning"

The optimal design of a fiscal rule maps to a mechanism design problem without transfers but with money burning. Using the revelation principle, the composition of the penalty schedule and the allocation gives the *money-burning* schedule $t(\theta) = P(g(\theta))$ for $\theta \in \Theta$. Note that on-equilibrium penalties impose discipline at the cost of "burning money," whereas off-equilibrium penalties do not "burn money."

Incentive compatibility constraints guarantee the implementability of the allocation as in (2). An allocation $g(\cdot)$ is *incentive compatible* given a money-burning schedule $t(\cdot)$ if

$$\theta U(g(\theta)) + \beta\left(W(T - g(\theta)) - t(\theta)\right) \geq \theta U(g(\hat{\theta})) + \beta\big(W(T - g(\hat{\theta})) - t(\hat{\theta})\big), \quad \text{for } \theta, \hat{\theta} \in \Theta. \quad \text{(IC)}$$

A fiscal rule is *optimal* if the allocation it implements $g(\cdot)$ and the associated schedule $t(\cdot)$ solve

$$\max_{g(\cdot),\ t(\cdot)} \left\{ \int_{\Theta} \big[\theta U(g(\theta)) + W(T - g(\theta)) - \rho t(\theta)\big] dF(\theta) \;\Big|\; \text{(IC) and } t(\theta) \geq 0 \text{ for } \theta \in \Theta \right\}. \quad (4)$$

The intercept of the schedule, if left implicit, is $t(\underline{\theta}) = 0$ and $P(g) = 0$ for $g \leq g(\underline{\theta})$.

The non-negativity constraint on money burning sets program (4) apart from the design of a mechanism with transfers because it rules out cross-subsidization across types.[7] The solution method exploits powerful Lagrangian techniques to allow for the non-negativity constraint on money burning (Section B of the Appendix contains a description of the solution method).

2 On the balance between discipline and discretion

In this section, I use the first-order approach to analyze the economics of on-equilibrium penalties. I first study a relaxed problem and then explore the conditions under which the relaxed constraints may be binding. The first-order approach provides an analytical formula for the (locally) optimal balance between the competing needs for discipline and discretion.

The following standard result exploits the incentive compatibility constraints to characterize, for a given intercept $t(\underline{\theta})$, the money-burning schedule associated with a non-decreasing allocation (Myerson (1981)). The proof is in Appendix C.1.

[7] Atkeson and Lucas (1992) study the case with transfers but without present bias, and Galperti (2015) studies the case with transfers and present bias.

Lemma 1 (Incentive compatible allocations)**.** *An allocation $g(\cdot)$ is incentive compatible given a money-burning schedule $t(\cdot)$ if and only if $g(\cdot)$ is non-decreasing and*

$$t(\theta) = t(\underline{\theta}) + \frac{\theta}{\beta}U(g(\theta)) + W(T - g(\theta)) - \frac{\underline{\theta}}{\beta}U(g(\underline{\theta})) - W(T - g(\underline{\theta})) - \frac{1}{\beta}\int_{\underline{\theta}}^{\theta} U(g(\tilde{\theta}))\, d\tilde{\theta}. \quad (5)$$

Lemma 1 is useful in substituting the money-burning schedule with a function of the allocation in the objective of program (4). The resulting objective functional reads as follows:

$$\int_{\Theta} \left[(1-\beta)W(T-g(\theta)) + (1-\tfrac{\rho}{\beta})\left(\theta U(g(\theta)) + \beta W(T-g(\theta))\right) + \tfrac{\rho}{\beta}\tfrac{1-F(\theta)}{f(\theta)}U(g(\theta))\right] dF(\theta) \quad (6)$$
$$- \tfrac{\rho}{\beta}\left(t(\underline{\theta}) - \underline{\theta}\, U(g(\underline{\theta})) - \beta W(T - g(\underline{\theta}))\right).$$

The first-order approach in this context consists of maximizing the objective (6) point-wise for $\theta > \underline{\theta}$ while ignoring three constraints (the usual monotonicity condition, the non-negativity constraint on money burning, and the non-negativity constraint on spending). Rearranging the first-order condition to express it as a wedge suggests the following definition.

Definition (foa-wedge)**.** *Let $\rho \in [0, 1)$. Define the foa-wedge as follows:*

$$\Delta_n(\theta) = \frac{1}{1-\rho}\left((1-\beta) - \rho\frac{1-F(\theta)}{\theta f(\theta)}\right). \quad (7)$$

The foa-wedge is the product of a scaling factor and a term capturing the key trade-off for on-equilibrium penalties. The scaling factor captures society's aversion to on-equilibrium penalties. It is the inverse of the degree of asymmetry in the cost of on-equilibrium penalties and ranges between 1 and infinity. It takes its smallest value for the largest degree of asymmetry. Indeed, for $\rho = 0$, society does not bear any cost associated with on-equilibrium penalties. In contrast, for symmetric penalties (i.e., $\rho = 1$), society is infinitely averse to on-equilibrium penalties, and only the sign of the term in parentheses matters.[8]

The term in parentheses in (7) balances the benefit and cost of imposing discipline with on-equilibrium penalties. The benefit of discipline is simply captured by $1 - \beta$. The term $\rho\frac{1-F(\theta)}{f(\theta)}$ captures the *incentive cost* of a marginal (on-equilibrium) penalty. To be compatible with incentives, on-equilibrium penalties are cumulative in the sense that a marginal penalty on

[8]This intuition echoes the bang-bang result of Halac and Yared (2022a): an optimal rule enforced by symmetric penalties necessarily relies on extreme—bang-bang—penalties. The foa-wedge (7) generalizes the function $Q(\theta) = -(1-\beta)\theta f(\theta) + 1 - F(\theta)$ studied in Halac and Yared (2022a) to study environments with asymmetric penalties.

spending $g_n(\theta)$ is borne with weight ρ and probability $1 - F(\theta)$ that the fiscal needs are above θ. For $\rho > 0$ and a Pareto distribution with tail parameter γ, the foa-wedge is constant because $\frac{1-F(\theta)}{\theta f(\theta)} = \frac{1}{\gamma}$. For $\rho = 0$, the foa-wedge is also constant and equal to the degree of present bias, which is reminiscent of a Pigouvian tax aligning the government's incentives with the society's objective.

Proposition 1 (Optimal fiscal rule: Pigouvian benchmark)**.** *Suppose that* $\rho = 0$. *An optimal fiscal rule implements the* first-best *allocation* g_{fb} *defined by the constant wedge* $\Delta(g_{fb}(\theta), \theta) = 1 - \beta$.

Unlike the other propositions in this paper, this proposition follows from first principles. Without concern for the societal cost of meting out a penalty on the government, it suffices to verify that the first-best allocation is compatible with incentives and that the associated penalty schedule satisfies the non-negativity constraint on penalties. The first-best allocation, which satisfies society's Euler equation $\theta U'(g_{fb}(\theta)) = W'(T - g_{fb}(\theta))$ for $\theta \in \Theta$, is increasing and hence compatible with incentives. The marginal penalty is strictly positive for $g \geq g_{fb}(\underline{\theta})$, and for $g < g_{fb}(\underline{\theta})$, $P(g) = 0$.

Three constraints are relaxed in the first-order approach: the non-negativity constraint on government spending, the non-negativity constraint on penalties, and the monotonicity constraint. First, I use restrictions that the two non-negativity constraints place on the level of the foa-wedge to define three different degrees of present bias (relative to the incentive cost) depending on whether the foa-wedge is bigger than 1, smaller than 0, or in between 0 and 1.

The foa-wedge is bigger than 1, which is incompatible with non-negative public spending, if the *degree of present bias is high* in the sense that Assumption H holds.

Assumption H. $\rho\frac{1-F(\theta)}{\theta f(\theta)} \leq \rho - \beta$ *for* $\theta \in \Theta$.

If Assumption H is not satisfied at θ, denote by $g_n(\theta)$ the spending associated with the foa-wedge (i.e., $\Delta(g_n(\theta), \theta) = \Delta_n(\theta)$).

The *degree of present bias is low* at θ if the foa-wedge is smaller than 0, which is the case under the following assumption.

Assumption L. $1 - \beta \leq \rho\frac{1-F(\theta)}{\theta f(\theta)}$.

If Assumption L is not satisfied at θ, then the foa-wedge is positive and $g_n(\theta) < g_d(\theta)$. If the degree of present bias is neither high nor low at θ, then it is *intermediate* at θ.

Assumption I. $\rho - \beta < \rho\frac{1-F(\theta)}{\theta f(\theta)} < 1 - \beta$.

Assumption I implies that the foa-wedge is in between 0 and 1 at θ.

The monotonicity constraint places restrictions on the slope of the foa-wedge, which I use to define the (virtual) need for discretion. The following measure of the slope of an allocation g is indicative of the discretion granted at $g(\theta)$: $\frac{d}{d\theta}\theta(1 - \Delta(g(\theta), \theta))\frac{1}{\beta}$. For instance, implementing the first-best allocation grants constant discretion 1, which is lower than the constant discretion $1/\beta$ associated with the discretionary allocation. Following Myerson (1981), the term *virtual* qualifies a concept augmented by the incentive cost. I define the (virtual) *need for discretion* as the discretion granted by the foa-wedge as follows: $\frac{d}{d\theta}\theta(1 - \Delta_n(\theta))\frac{1}{\beta}$. The foa-wedge reveals that the slope of the inverse hazard rate governs the (virtual) need for discretion (need for discretion for short).

Lemma 2 (Monotonicity and the need for discretion)**.** *Suppose that $g_n(\theta)$ is well-defined for $\theta \in (\theta_*, \theta^*)$. Then, g_n is non-decreasing at $\theta \in (\theta_*, \theta^*)$ if and only if the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is not smaller than $\rho - \beta$.*

Appendix C.2 contains the proof. Lemma 2 shows that for on-equilibrium penalties to be part of a fiscal rule, the tail of the distribution of fiscal needs must be sufficiently thick, as measured by the slope of the inverse hazard rate. The threshold for the thickness of the tail depends on the degree of asymmetry and the degree of present bias. For $\rho = 0$, the threshold is trivially satisfied independently of the degree of present bias. For $\rho = \beta$, it suffices to check whether the inverse hazard rate is non-decreasing, which is the case if and only if $1 - F$ is log-convex. For instance, for $\rho = \beta$, the foa-wedge is a valid building block of an optimal fiscal rule for Pareto-distributed fiscal needs, not for normally distributed fiscal needs. Intuitively, the thick tail associated with Pareto-distributed fiscal needs calls for a blend of discretion and discipline that only on-equilibrium penalties can achieve. The (virtual) *need for discretion is high* at θ if the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is not smaller than $\rho - \beta$.

3 Optimal fiscal rules

This section contains the main theoretical results. The analysis is in three parts covering the cases of a low, an intermediate, and a high degree of present bias of the government. The insights from the previous section help conjecture the optimal fiscal rules. The optimality conditions of the global Lagrangian method then determine the features of the environment under which the conjecture is indeed a solution.

First, I propose a general fiscal rule featuring no penalty below a threshold and a graduated schedule of penalties from the threshold up to a point at which the marginal penalty jumps sufficiently to implement a cap. The penalty schedule is best described by the marginal penalty schedule because it is more telling of the discipline imposed (and the intercept set to $P(0) = 0$),

$$P'(g) = \begin{cases} 0 & \text{for } g < g_n, \\ U'(g)\Delta(g)/\beta & \text{for } g \in [g_n, g_p), \\ \infty & \text{for } g \geq g_p, \end{cases} \tag{8}$$

where $\Delta(\cdot) > 0$, g_n, and g_p remain to be determined. The infinite marginal penalty reflects the assumption of unlimited enforcement (see Halac and Yared (2022) for the implications of limited enforcement). If $\bar{\theta} < \infty$, then the marginal penalty needed to enforce a cap at g would be finite.

The fiscal rule (8) encompasses various cases observed in practice. First, the absence of a fiscal rule amounts to $g_n = g_p = \infty$. Second, for $g_n = g_p < \infty$, (8) amounts to a cap on spending. Third, for $g_n < g_p = \infty$, (8) can account for a graduated schedule of on-equilibrium penalties resembling a Pigouvian tax. Last, for $g_n < g_p < \infty$, (8) accounts for hybrid fiscal rules featuring a graduated schedule of on-equilibrium penalties up to a cap.

In the spirit of the revelation principle, the remainder of the analysis characterizes an optimal allocation from which an optimal fiscal rule can be inferred. For instance, an allocation with no bunching below the cap indicates that the marginal penalty schedule is continuous at g_n. In contrast, an allocation with bunching below g_n indicates a jump in the marginal penalty schedule (resulting in a kink in the penalty schedule). Starting with a continuous marginal penalty schedule up to a cap, define the *discretion, on-equilibrium, and off-equilibrium penalties* allocation as follows for $\theta \in \Theta$:

$$g(\theta) = \begin{cases} g_d(\theta) & \text{for } \theta < \theta_n, \\ g_n(\theta) & \text{for } \theta \in [\theta_n, \theta_p], \\ g_n(\theta_p) & \text{for } \theta > \theta_p, \end{cases} \tag{9}$$

where $\theta_n < \theta_p$. The fiscal needs θ_n and θ_p parametrize the threshold g_n and g_p, where the subscript n refers to on-equilibrium penalties above g_n and p refers to the prohibitive nature of off-equilibrium penalties above g_p. The allocation (9) is a partial guess for a solution because $g_n(\theta)$ denotes the level of public spending dictated by the foa-wedge. To complete the guess, it only remains to set the thresholds θ_n and θ_p.

To set the cap, I use a first-order condition of the Lagrangian method, given an allocation g,[9]

$$\theta_p = \inf \left\{ \tilde{\theta}_p \in \Theta \mid \text{Inequality (10) holds for } \hat{\theta} \geq \tilde{\theta}_p \right\},$$

$$\int_{\hat{\theta}}^{\bar{\theta}} \theta \left(\Delta(g(\tilde{\theta}_p), \theta) - (1 - \beta) \right) \mathrm{d}F(\theta) \leq \rho \, \hat{\theta} \Delta(g(\tilde{\theta}_p), \hat{\theta})(1 - F(\hat{\theta})). \tag{10}$$

The continuum of inequalities in the definition of θ_p are notoriously challenging to interpret, and the task of identifying θ_p may seem daunting. Below, I show that the first-order conditions used to set θ_p encapsulate two distinct optimality requirements, each with a compelling economic intuition.[10] The upshot of breaking down the definition into these two requirements is a two-step procedure that makes identifying θ_p a simple task.

Set θ_n at the lowest fiscal need such that g_n is well-defined and associated with a non-negative marginal penalty:

$$\theta_n = \inf \left\{ \tilde{\theta}_n \in \Theta \;\middle|\; \text{Assumption I holds for } \theta \in (\tilde{\theta}_n, \theta_p) \right\}.$$

If Assumption I holds for $\theta \in \Theta$, that is, $0 < g_n(\theta) \leq g_d(\theta)$ for $\theta \in \Theta$, then $\theta_n = \underline{\theta}$. If Assumption I holds for a non-empty subset in the interior of $[\underline{\theta}, \theta_p]$, then $\theta_n \in (\underline{\theta}, \theta_p)$. The continuity of both F and f implies that the threshold satisfies $g_n(\theta_n) = g_d(\theta_n)$ if the upper bound in Assumption I binds, which is the relevant case.[11] Lastly, if Assumption I does not hold for any θ, the set in

[9]To focus on the economic intuition, I postpone describing the solution method to Appendix B.

[10]To relate this contribution to the literature, note that for $\rho = 1$, the definition of θ_p corresponds to the condition in Proposition 2 in Amador, Werning, and Angeletos (2006).

[11]If the lower bound in Assumption I binds instead, then $\lim_{\theta \to \theta_n} g_n(\theta) = 0$.

the definition of θ_n is empty, and then $\theta_n = \bar{\theta}$. For instance, if shocks are Pareto distributed, $1 - F(\theta) = \theta^{-\gamma}$, and the elasticity of the tail $\gamma \in (1, \frac{\rho}{1-\beta})$, then the need for discretion is such that the candidate solution does not feature penalties since $\theta_n = \bar{\theta}$.

If the allocation (9) is optimal and $\theta_n < \theta_p$, then I say that the optimal fiscal rule features a graduated schedule of on-equilibrium penalties and a cap. The *cap binds* if $\theta_p < \bar{\theta}$.

3.1 Low degree of present bias

This section contains the first main result of this paper and a computed example.

Since Amador, Werning, and Angeletos (2006), a lower bound on the elasticity of the density of the distribution of shocks has been understood to imply that granting discretion below a threshold is optimal for symmetric penalties. For asymmetric penalties (i.e. $\rho < 1$), the lower bound is more stringent.

Assumption sL. $\frac{\theta f'(\theta)}{f(\theta)} \geq -\frac{1+\rho-\beta}{1-\beta}$.

Below, I show that, in the context of the next proposition, Assumption sL implies a low degree of present bias, as in Assumption L.

3.1.1 High need for discretion

If the degree of present bias is low relative to the incentive cost of discipline for fiscal needs below a threshold and the need for discretion is high for an intermediate range of fiscal needs, then a hybrid fiscal rule is optimal.

Proposition 2 (Optimal fiscal rule: low degree of present bias and high but decreasing need for discretion)**.** *Suppose $\underline{\theta} < \theta_n$ and Assumption sL holds for $\theta \leq \theta_n$. If the derivative of $\rho\frac{1-F}{f}$ is not smaller than $\rho - \beta$ for $\theta \in [\theta_n, \theta_p]$, then a fiscal rule that implements the* discretion, on-equilibrium, and off-equilibrium penalties *allocation is optimal.*

The proof is in Appendix C.3. The optimal fiscal rule has three parts. A computed example depicts the three parts and serves as a visual aid to the following discussion.

Discretion. A first part grants discretion below $g_d(\theta_n)$ because, as the next lemma shows, Assumption sL implies a low degree of present bias relative to the incentive cost of discipline.

Lemma 3 (Implications of Assumption sL)**.** *If $\theta_n \in (\underline{\theta}, \bar{\theta})$ and Assumption sL holds for $\theta \leq \theta_n$, then $1 - \beta \leq \rho \frac{1-F(\theta)}{\theta f(\theta)}$ for $\theta \leq \theta_n$.*

The proof is in Appendix C.4. Lemma 3 shows that the optimal fiscal rule grants discretion below $g_d(\theta_n)$ because the incentive cost of discipline is too high. Discipline is so costly below θ_n that the foa-wedge is negative.

On-equilibrium penalties. A second part of the optimal fiscal rule imposes on-equilibrium penalties on spending between $g_d(\theta_n)$ and $g_n(\theta_p)$ if the (virtual) need for discretion is sufficiently high to justify burning money by imposing a graduated schedule of on-equilibrium penalties. If g_n is decreasing over a subinterval of $[\theta_n, \theta_p]$, the solution would be to "iron" g_n, as in Myerson (1981). Ironing the allocation requires a jump in the marginal penalty instead of a continuous schedule of marginal penalties. The marginal penalty schedule jumps at $g_n(\theta_p)$ precisely for this reason, which relates to the cap.

Cap. A third part of an optimal fiscal rule imposes off-equilibrium penalties above $g_n(\theta_p)$. To gain intuition and simplify the task of determining θ_p, I decompose the definition of θ_p into two distinct optimality requirements. Note that inequality (10) is a function of two thresholds: $\tilde{\theta}_p$ identifies the threshold at which public spending bunches, and $\hat{\theta}$ determines the range over which the bunching is evaluated. The first requirement equates the marginal benefit to the marginal cost of the cap. It consists of identifying the set of fiscal needs at which the inequality is satisfied with equality for $\hat{\theta} = \tilde{\theta}_p$. Indeed, if the infimum is in the interior of Θ, a continuity argument implies that the inequality holds with equality at the infimum. This first requirement usually simplifies the identification of θ_p considerably because it suffices to find the root of an equation. The second requirement checks that off-equilibrium penalties dominate on-equilibrium penalties above $g(\tilde{\theta}_p)$. It asks to check that the inequality (10) holds for $\hat{\theta} \geq \tilde{\theta}_p$.

The first requirement, that is, inequality (10) holding with equality for $\hat{\theta} = \tilde{\theta}_p > \underline{\theta}$, determines the stringency of the cap as a root of the following equation:

$$\int_{\tilde{\theta}_p}^{\bar{\theta}} \theta \left(\Delta(g(\tilde{\theta}_p), \theta) - (1 - \beta) \right) \mathrm{d}F(\theta) = \rho\, \tilde{\theta}_p \Delta(g(\tilde{\theta}_p), \tilde{\theta}_p)(1 - F(\tilde{\theta}_p)).$$

It sets the cap to equate the average distortion at the top on the left-hand side to the marginal reduction in the burden of on-equilibrium penalties on the right-hand side. On the left-hand side, the difference between the wedge and the degree of present bias measures the distortion relative to the first best, and the integral takes a weighted average with the fiscal needs as the weights. On the right-hand side, the burden of on-equilibrium penalties depends on whether the fiscal rule features on-equilibrium penalties. For a rule that does not feature on-equilibrium penalties, the right-hand side is null because the wedge is null. Then the cap is such that there is no distortion at the top on average. If, instead, the cap comes on top of on-equilibrium penalties, then the right-hand side is positive. The cap equates the average distortion at the top on the left-hand side to the marginal benefit of not imposing the marginal penalty at $g(\tilde{\theta}_p)$ with probability $1-F(\tilde{\theta}_p)$.

The second requirement, that is, inequality (10) holding for $\hat{\theta} \geq \tilde{\theta}_p$, determines the structure of the fiscal rule above $g(\tilde{\theta}_p)$. It ensures that off-equilibrium penalties above $g(\tilde{\theta}_p)$ dominate on-equilibrium penalties at the margin. At any point $\hat{\theta} \geq \tilde{\theta}_p$, the mechanism designer has the choice to prolong the bunching of public spending or offer an alternative while preserving the bunching of public spending between $\tilde{\theta}_p$ and $\hat{\theta}$ (see Figure 1). On the left-hand side of inequality (10), the marginal cost of prolonging the bunching of public spending is the loss of discretion net of the marginal benefit of discipline at $\hat{\theta}$. The wedge captures the loss of discretion, and the degree of present bias captures the marginal benefit discipline. A kink in the penalty schedule preserves the bunching of public spending at $g(\tilde{\theta}_p)$ between θ_p and $\hat{\theta}$, but not beyond $\hat{\theta}$. The kink is caused by a jump in the marginal penalty schedule from zero to $\hat{\theta}\Delta(g_d(\tilde{\theta}_p), \hat{\theta})U'(g_d(\tilde{\theta}_p))$, and the government would choose to incur the marginal penalty with probability $1 - F(\hat{\theta})$. In sum, for $\hat{\theta} \geq \tilde{\theta}_p$, inequality (10) verifies that the marginal cost net of the marginal benefit of enforcing the cap at $\hat{\theta}$ is lower than the marginal cost of resorting to on-equilibrium penalties.

The second requirement is easy to check directly from three fundamentals of the economy, namely, the degree of present bias, the degree of asymmetry in the cost of on-equilibrium penalties, and the conditional tail expectation of the distribution of shocks. The next lemma shows that the second requirement is equivalent to an upper bound on the thickness of the tail of the distribution of shocks.

Lemma 4 (Cap and $1 - F$)**.** *Suppose that inequality* (10) *holds with equality for some* $\tilde{\theta}_p < \bar{\theta}$ *and* $\hat{\theta} = \tilde{\theta}_p$*. Then, inequality* (10) *holds for* $\hat{\theta} \in [\tilde{\theta}_p, \bar{\theta})$ *if and only if*

$$\beta E[\theta|\theta \geq \hat{\theta}] - \rho\hat{\theta} \leq \beta E[\theta|\theta \geq \tilde{\theta}_p] - \rho\tilde{\theta}_p.$$

Appendix C.5 contains the proof. Lemma 4 shows that, according to the first-order condition

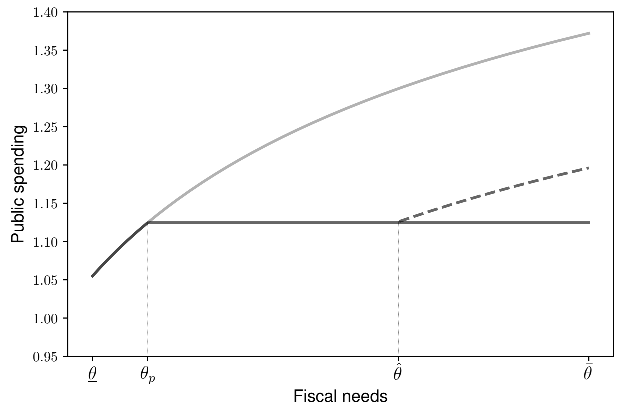

Figure 1: On the continuum of inequalities in the definition of the thresholds θ_p and θ_x.

Notes: The grey line depicts the discretionary allocation. The black line depicts an allocation g implemented by a fiscal rule that starts with a kink at $g(\theta_x)$ and a graduated schedule of on-equilibrium penalties up to $g(\theta_p)$ and above which penalties are prohibitively large to implement a cap. The dashed line depicts an alternative to the cap. To implement the alternative, the fiscal rule resorts to on-equilibrium penalties after a kink at $g(\hat{\theta})$. To implement the alternative, the fiscal rule features a graduated schedule of on-equilibrium penalties between $g(\underline{\theta})$ and $g(\theta_p)$ within which there is a kink at $g(\theta_x)$.

defining θ_p, a cap is desirable if the distribution of shocks is thin enough, and no cap is optimal otherwise. As in Weitzman (1974), the difference between the slope of the marginal benefit of discipline, governed by β, and the slope of the marginal cost of discipline, governed by ρ, matters for the choice of instrument (i.e., a binding cap or no cap). Lemma 4 highlights the importance of the distribution of shocks for the choice of instruments. For a sufficiently thin tail, the need for discretion is sufficiently low that a cap dominates on-equilibrium penalties above a threshold. For instance, for $\rho = \beta < 1$ and a log-concave tail, the optimal cap is binding and set to satisfy the equality condition of the first requirement (by Lemma 4, the log-concavity of $1 - F$ implies the second requirement). In contrast, for $\rho = \beta$ and a strictly log-convex tail, the second requirement is not satisfied for an interior threshold, and θ_p is either $\underline{\theta}$ or $\bar{\theta}$.

How does the extent of discretion granted by the cap, determined by θ_p, depend on the degree of present bias of the government and the degree of asymmetry in the cost of penalties? Lemma 4 already gives a partial answer. If the tail of the distribution of shocks is sufficiently thick, then $\theta_p \in \{\underline{\theta}, \bar{\theta}\}$, which does not change with a marginal change in β or ρ. Suppose instead that θ_p is in the interior of Θ.

Corollary 1 (The cap, β, and ρ)**.** *The threshold for the cap θ_p is weakly increasing in β and weakly decreasing in ρ.*

The proof is in Appendix C.6. Intuitively, β governs the benefit of discipline, and ρ governs the cost of granting discretion with on-equilibrium penalties. A subtlety of the corollary, however, arises for a cap on top of a graduated schedule of on-equilibrium penalties because β and ρ affect the foa-wedge. Recall that the cap is set such that there is no distortion, on average, over the bunching induced by the cap. While a lower β exacerbates the average distortion, it also increases the severity of the on-equilibrium penalty. Corollary 1 shows that β has a stronger effect on the average distortion above the cap than it has on the marginal on-equilibrium penalty below the cap. Similarly, more symmetric penalties (i.e., higher ρ) induce less discretion (i.e., lower θ_p) as a result of the balance between two forces. A higher ρ increases the benefit of resorting to a cap instead of on-equilibrium penalties. A higher ρ, however, may also lower the severity of on-equilibrium penalties, which lowers the benefit of resorting to a cap. The corollary shows that the former effect dominates, and hence, the optimal cap grants less discretion with more symmetric penalties.

Example (Low degree of present bias and a high but decreasing need for discretion)**.** *Suppose* $U(g) = \ln(g)$ *and* $W(x) = \ln(\omega + x)$*. The government revenues* T *are normalized to* 1*. The parameter* $\omega = 1$ *sets the average deficit as a percentage of fiscal revenues to* 12% *in the absence of a fiscal rule.*[12] *For simplicity, suppose* $\rho = \beta$ *so that the derivative of* $\rho\frac{1-F}{f}$ *is not smaller than* $\rho - \beta$ *if and only if* $1 - F$ *is log-convex. Suppose* $\beta = 0.8$*. Define the distribution of shocks* F_a *by its hazard rate, which is a convex combination of the hazard rates of the exponential distribution and the Pareto distribution,*[13] $h_a(\theta) = a\lambda + (1-a)\frac{\gamma}{\theta}$*, and* $a \in (0, 1)$*. The inverse hazard rate of the distribution* F_a *implies a positive but decreasing virtual need for discretion. In turn, this implies a marginal penalty schedule from the foa-wedge that is increasing in severity, which induces a widening gap between the discretionary allocation and the allocation associated with the foa-wedge in the right panel of Figure 2. To obtain a distribution that is log-convex below a threshold and log-cancave above the threshold, it suffices to truncate a log-convex distribution. For ease of comparison, the parameters of the truncated version of* F_a *are kept the same as the*

[12]The choice is based on data for the euro area. The average deficit for the three years prior to the implementation of the Stability and Growth Pact was 4.9% of GDP (Source: OECD (2021), General government deficit). Fiscal revenues averaged 40.8% of GDP (Source: Eurostat (gov_10a_taxag)). Combining these two moments gives a target average deficit as a percentage of fiscal revenues of 4.9%/40.8% = 12%.

[13]The hazard rate $h_a(\cdot)$ uniquely characterizes the distribution $F_a(\theta) = 1 - \exp\left(-\int_{\underline{\theta}}^{\theta} h_a(x)dx\right)$ for $\theta \in \Theta$.

ones for its non-truncated counterpart.[14] *Figure 2 depicts the optimal allocations in black to illustrate Proposition 2 and Lemma 4.*

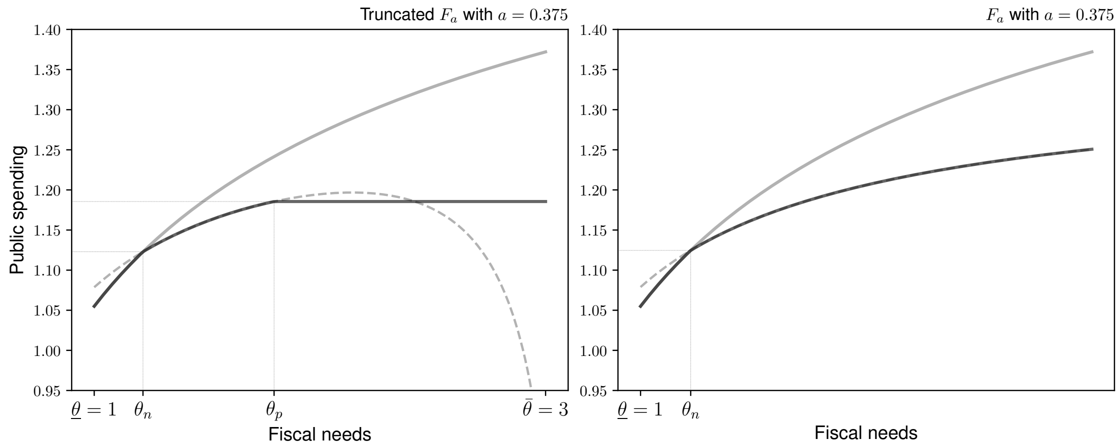

Figure 2: Low degree of present bias and high but decreasing need for discretion.

Notes: The grey line depicts the discretionary allocation. The dashed line depicts the allocation g_n associated with the marginal penalty schedule induced by the foa-wedge. The black line depicts the allocation implemented by the optimal fiscal rule. The distribution of fiscal needs is displayed at the top right of each panel. In the right panel, $\theta_p = \bar{\theta} = \infty$.

The order of θ_n and θ_p depends on the distribution of shocks. For $\rho = \beta$ and a strictly log-convex tail, the need for discretion is such that $\theta_n \leq \theta_p$, as in the example above. However, $\theta_p \leq \theta_n$ for $\rho = \beta$ and a strictly log-concave tail (indicating a low need for discretion).

3.1.2 Low need for discretion

Lemma 4 showed that for a sufficiently thick tail, the optimal fiscal rule does not feature a cap. In this section, I show that for a sufficiently thin tail, a main insight from the literature extends to asymmetric penalties. For $\rho < 1$, under more stringent conditions than those found in the literature (i.e., for $\rho = 1$), the optimal fiscal rule imposes a cap only on public spending.

Definition. *The* discretion and off-equilibrium penalties *allocation, denoted $g_d^p(\cdot)$, is defined as follows for $\theta \in \Theta$:*

$$g(\theta) = \begin{cases} g_d(\theta_p) & \textit{for } \theta > \theta_p \\ g_d(\theta) & \textit{for } \theta \leq \theta_p. \end{cases}$$

[14] Truncating the Pareto and F_a distributions at $\bar{\theta} = 3$ does not significantly alter the mean.

Proposition 3 (Optimal fiscal rule: low degree of present bias and low need for discretion)**.** *Suppose $\underline{\theta} < \theta_p$ and Assumption sL holds for $\theta \leq \theta_p$. A fiscal rule that implements the* discretion and off-equilibrium penalties *allocation is optimal.*

The proof is in Appendix C.7. For symmetric penalties $\rho = 1$, the proposition nests Proposition 3 in Amador, Werning, and Angeletos (2006). The proposition holds for degrees of asymmetry in the cost of penalties that are not too strong. In particular, the following lemma shows that the assumption of Proposition 3 cannot hold for the extreme case $\rho = 0$, which is consistent with the previous finding that, for $\rho = 0$, the optimal fiscal rule implements the first-best allocation with on-equilibrium penalties.

Lemma 5. *If $\theta_p \in (\underline{\theta}, \bar{\theta})$ and Assumption sL holds for $\theta \leq \theta_p$, then $1 - \beta \leq \rho \frac{1-F(\theta)}{\theta f(\theta)}$ for $\theta \leq \theta_p$.*

Appendix C.4 contains the proof. Lemma 5 is the analog of Lemma 3 for fiscal rules that do not feature on-equilibrium penalties: again, under the conditions in Proposition 3, the optimal fiscal rule grants discretion below the cap because the degree of present bias is low relative to the incentive cost. The optimal cap may not bind, however, in which case it may be optimal to not impose any discipline.

Corollary 2 (Optimal fiscal rule: low degree of present bias)**.** *Suppose that Assumption L holds at $\underline{\theta}$ and Assumption sL holds for $\theta \in \Theta$. Then Assumption L holds for $\theta \in \Theta$ and the absence of a fiscal rule is optimal.*

The condition is trivially satisfied if $1 - \beta = 0$. It is also satisfied for a sufficiently thick tail of the distribution of shocks relative to the degree of present bias—for example, a Pareto distribution with tail parameter γ and $1 - \beta \leq \rho \frac{1}{\gamma}$.[15] For $1 - \beta > 0$ and a distribution on a compact support, however, Assumption L cannot be satisfied at $\bar{\theta}$.[16]

[15]This formalizes and extends an observation in footnote 6 in Amador, Werning, and Angeletos (2006) to asymmetric money burning and to a broader class of distributions relative to Pareto.

[16]This is reminiscent of the "no distortion at the top" in optimal mechanisms with transfers such as redistributive taxation. The analog for the design of corrective mechanisms without transfers on a compact type space Θ is that a cap is binding at the top such that that there is no distortion on average over the bunching induced by the cap (see Amador, Werning, and Angeletos (2008), Proposition 2, and Ambrus and Egorov (2013), Proposition 3, for symmetric money burning and this paper for asymmetric money burning).

Example (Low degree of present bias and a low need for discretion). *The economic environment is identical to the one in the previous example, with the exception that the distribution of shocks is exponential with parameter* $\lambda = 3$. *The parameters of the exponential distribution and the distribution* F_a *are set so that the two distributions have the same mean. The constant inverse hazard rate* $1/\lambda$ *of the exponential distribution implies that the virtual need for discretion is null for* $\rho = \beta$ *because* $0 = \frac{d}{d\theta}(\rho\frac{1}{\lambda} - (\rho - \beta)\theta)$, *which implies a constant allocation* g_n *(see Figure 3). Proposition 3 applies because* $1 - F$ *is log-concave, Assumption sL holds for* $\theta \leq \frac{1}{\lambda}\frac{1}{1-\beta}$, *and* $\theta_p < 1.6\bar{6} = \frac{1}{\lambda}\frac{1}{1-\beta}$.

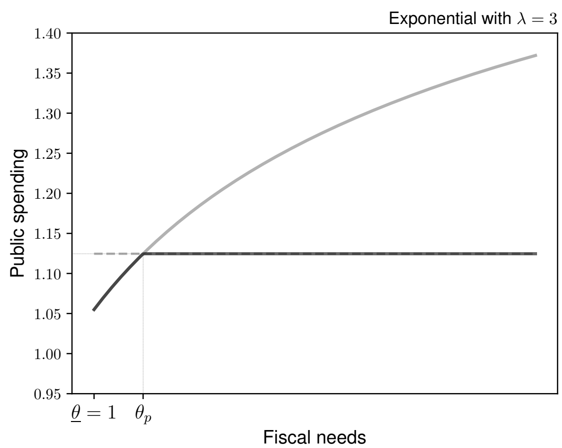

Figure 3: Low degree of present bias and a low need for discretion.

Notes: The grey line depicts the discretionary allocation. The dashed line, which is horizontal and partly covered by the black line, depicts the allocation g_n associated with the marginal penalty schedule induced by the foa-wedge. The black line depicts the allocation implemented by the optimal fiscal rule.

3.2 Intermediate degree of present bias

This section contains the second main theoretical result: an optimal fiscal rule may grant an exemption from on-equilibrium penalties below a threshold to lower the level of penalties above the threshold.

The previous subsection showed that for a low degree of present bias, to comply with the non-negativity constraint on penalties, it is optimal to set the marginal penalty to zero instead of that conforming with the foa-wedge below a threshold (which I refer to as truncating the foa-wedge schedule). For an intermediate degree of present bias, although the foa-wedge is positive, it may still be optimal to truncate part of the foa-wedge schedule.

To gain insights into how the non-negativity constraint on penalties matters for the design of the optimal fiscal rule, decompose the penalty schedule into two building blocks: the intercept and the marginal penalty schedule. While the marginal penalty schedule determines the extent of discipline, the intercept determines the overall burden of penalties. The non-negativity constraint on penalties forces the intercept to be non-negative; hence, the overall burden of penalties can only be shifted up, which is not desirable. It only leaves us with altering the marginal penalty schedule to lower the level of on-equilibrium penalties. Because lowering the marginal penalty on spending g lowers the level of penalties on all spending above g, altering the bottom of the marginal penalty schedule is the most effective substitute for a negative intercept.

Unlike a negative intercept, however, the truncation entails a loss of discipline because it grants an *exemption* from the candidate marginal penalty below a threshold. The exemption causes a jump in the marginal penalty schedule, and the resulting kink induces the bunching of government spending at the exemption threshold.

For this section, assume $\rho < 1$ so that if Assumption I holds at θ, then $g_n(\theta)$ is well-defined. For the next definition, suppose that Assumption I holds below θ_p.

The *exemption, on-equilibrium, and off-equilibrium penalties* allocation is defined for $\theta \in \Theta$ as follows:

$$g(\theta) = \begin{cases} g_n(\theta_p) & \text{for } \theta > \theta_p \\ g_n(\theta) & \text{for } \theta_x \leq \theta \leq \theta_p \\ g_n(\theta_x) & \text{for } \theta \leq \theta_x, \end{cases}$$

where the θ_p is defined in (10). For θ_x, the first-order condition of the Lagrangian method sets the threshold as the highest fiscal need below which a continuum of inequalities are satisfied:

$$\theta_x = \sup \left\{ \tilde{\theta}_x \in \Theta \mid \text{Inequality (11) holds for } \hat{\theta} \leq \tilde{\theta}_x \right\},$$

and

$$\int_{\underline{\theta}}^{\hat{\theta}} \theta \left((1-\beta) - \Delta(g_n(\tilde{\theta}_x), \theta) \right) dF(\theta) \leq \rho \, \hat{\theta} \Delta(g_n(\tilde{\theta}_x), \hat{\theta})(1 - F(\hat{\theta})). \tag{11}$$

I discuss the economics of exempting public spending below $g_n(\theta_x)$ from penalties in the context of the next proposition.

3.2.1 High need for discretion

The next proposition is the second main theoretical result of this paper. For an intermediate degree of present bias such that $g_n(\theta) < g_d(\theta)$ for $\theta \leq \theta_p$, if an optimal fiscal rule features on-equilibrium penalties, then it features a kink in the penalty schedule caused by an exemption from penalties below a threshold.

Proposition 4 (Optimal fiscal rule: intermediate degree of present bias and high need for discretion)**.** *Suppose that Assumption I holds for $\theta \leq \theta_p$. If the derivative of $\rho\frac{1-F}{f}$ is not smaller than $\rho - \beta$ for $\theta \in [\underline{\theta}, \theta_p]$, then a fiscal rule that implements the* exemption, on-equilibrium, and off-equilibrium penalties *allocation is optimal.*

Appendix C.8 contains the proof. Proposition 4 contains one novel insight besides the insight regarding the balance between the needs for discretion and discipline from Propositions 2 and 3. The novel insight—the optimality of granting an exemption below a threshold—matters if the degree of present bias is intermediate even for low realizations of fiscal needs (i.e., $\theta_n = \underline{\theta}$).

Exemption from penalties below a threshold. To gain intuition and simplify the task of determining θ_x, I decompose the definition of θ_x into two distinct optimality requirements. Inequality (11) is a function of two thresholds: $\tilde{\theta}_x$ determines the level at which public spending bunches, and $\hat{\theta}$ determines the range of the bunching. First, by continuity, inequality (11) holds with equality at $\hat{\theta} = \tilde{\theta}_x = \theta_x$ if $\theta_x \in (\underline{\theta}, \bar{\theta})$. Second, inequality (11) holds for $\hat{\theta} \leq \tilde{\theta}_x = \theta_x$.

The first requirement determines the leniency of the exemption. It sets θ_x to equate the marginal cost to the marginal benefit of the exemption. On the left-hand side of (11), the marginal cost of the exemption is the average distortion due to the loss of discipline for fiscal needs below the exemption threshold. On the right-hand side, the marginal benefit is the marginal reduction in the level of penalties. For $\hat{\theta} = \tilde{\theta}_x = \theta_x$, the exemption from the marginal penalty $U'(g_n(\theta_x))\theta_x\Delta_n(\theta_x)$ is beneficial with probability $1 - F(\theta_x)$.

The second requirement determines the structure of the fiscal rule below the threshold. It checks that for any $g \leq g_n(\theta_x)$, the exemption dominates any alternative. Instead of a zero marginal penalty below $g_n(\theta_x)$, the mechanism designer can impose a marginal penalty that a government with fiscal needs $\hat{\theta} < \theta_x$ would incur, while preserving the bunching of public spending between $\hat{\theta}$ and θ_x (see Figure 4 for an illustration). The marginal cost of such a switch to on-equilibrium penalties is the marginal penalty $U'(g_n(\theta_x))\hat{\theta}\Delta(g_n(\theta_x),\hat{\theta})$. The marginal benefit is the discipline from the marginal penalty net of the loss of discretion for governments with fiscal needs below $\hat{\theta}$. Inequality (11) implies that the marginal benefit of switching from the exemption to on-equilibrium penalties lies below the marginal cost at any point below the exemption threshold.

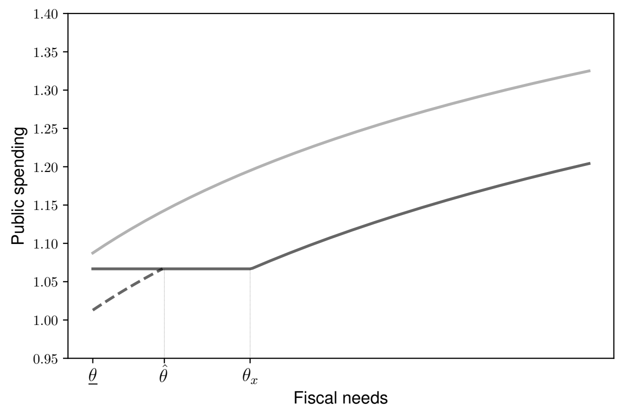

Figure 4: Alternative to an exemption from penalties on spending below g(θ_x).

Notes: The grey line depicts the discretionary allocation. The black line depicts an allocation implemented by a fiscal rule with no penalties below a kink at $g(\theta_x)$. The dashed line depicts an alternative to the exemption from penalties below $g(\theta_x)$. A graduated schedule of on-equilibrium penalties starting at $g(\underline{\theta})$ and a kink at $g(\theta_x)$ implements the alternative.

The second requirement is easy to check based on fundamentals of the economy. It is equivalent to a sufficiently thick right tail of the distribution of shocks, which, because it is at the left end of the range of fiscal needs, may loosely be thought of as a thin left tail.[17] As the next lemma shows, an interior exemption is optimal for a sufficiently "thin left tail," and no exemption is optimal otherwise.

Lemma 6 (The exemption and $1-F$)**.** *Suppose that inequality* (11) *holds with equality for some* $\tilde{\theta}_x > \underline{\theta}$ *and* $\hat{\theta} = \tilde{\theta}_x$*. Then, inequality* (11) *holds for* $\hat{\theta} \in [\underline{\theta}, \tilde{\theta}_x]$ *if and only if*

$$\int_{\hat{\theta}}^{\tilde{\theta}_x} \left[\left(\rho \tfrac{1-F(\theta)}{\theta f(\theta)} - (\rho - \beta) \right) \theta \right] dF(\theta) \leq \int_{\hat{\theta}}^{\tilde{\theta}_x} \left[\left(\rho \tfrac{1-F(\tilde{\theta}_x)}{\tilde{\theta}_x f(\tilde{\theta}_x)} - (\rho - \beta) \right) \tilde{\theta}_x \right] dF(\theta). \tag{12}$$

The proof is in Appendix C.9. The intuition for this result relates the thinness of the left tail of the distribution of shocks to the marginal benefits and costs of an exemption. The inverse

[17] Although a thick right tail $1-F$ is related to a thin left tail F, the two are not equivalent. F log-concave implies $1-F$ log-convex for a decreasing density. For an increasing density, $1-F$ log-convex implies F log-concave. Indeed, for a twice differentiable F, $1-F$ log-convex is equivalent to $\frac{d}{d\theta}\frac{F}{f} \leq \frac{d}{d\theta}\frac{1}{f}$, and F log-concave is equivalent to $\frac{d}{d\theta}\frac{F}{f} \leq 0$.

hazard rate governs the incentive cost that a marginal penalty on public spending g imposes on all public spending above g. The benefit of an exemption is to economize on the incentive cost of the marginal penalties between $\hat{\theta}$ and θ_x for $\hat{\theta} \leq \theta_x$. The thicker is the right tail of the distribution of shocks between $\hat{\theta}$ and θ_x for $\hat{\theta} \leq \theta_x$, the larger is the benefit. For instance, for $\rho = \beta$, a log-convex $1 - F$ over $[\underline{\theta}, \theta_x]$ is sufficient for condition (12) to hold. Similarly, for Pareto-distributed shocks with tail parameter γ, condition (12) simplifies to a lower bound on the thickness of the right tail $\rho\frac{1}{\gamma} \geq \rho - \beta$.

Although the inequality determining the exemption threshold resembles the inequalities determining the thresholds for the cap, they differ in economic content. The resemblance stems from the shared origin of these characterizations. Both come from the first-order conditions of the Lagrangian methods. The economics of the exemption, however, pertain to the non-negativity constraint on penalties, whereas the economics of the cap pertain to the trade-off between the need for discretion and the need for discipline.

The degree of leniency implied by the optimal exemption depends on both the government's present bias and the asymmetry in the cost of on-equilibrium penalties.

Corollary 3 (The exemption and β)**.** *The exemption threshold θ_x is decreasing in β.*

The proof is in Appendix C.10. The intuition is that the marginal benefit of lowering the level of the penalty above the exemption is larger for a higher degree of present bias. However, the marginal cost of forgoing discipline due to the exemption is also larger the higher is the degree of present bias. The corollary shows that the former effect dominates the latter.

In contrast, the effect of ρ on the optimal exemption is ambiguous. Although a higher ρ implies a higher benefit of an exemption from a given penalty schedule above the exemption, ρ also affects the foa-wedge.

Example (Intermediate degree of present bias and a high need for discretion)**.** *The economy is identical to the one in the previous examples, with two exceptions. First, $\beta = 0.7$ to bring the degree of present bias from low to intermediate (and ρ also changes to keep $\rho = \beta$). Second, the distribution of shocks is Pareto with parameter $\gamma = 4$ to have a thick tail with a constant elasticity. The parameter for the tail of the Pareto distribution is such that the mean is the same as that of the distribution F_a. The linear inverse hazard rate of the Pareto distribution implies a constant virtual need for discretion $\beta - \rho(1 - 1/\gamma)$, which induces a constant gap between the* *discretionary allocation and the allocation associated with the foa-wedge. Proposition 4 applies because for the Pareto distribution,* $\rho\frac{1-F(\theta)}{\theta f(\theta)} = \frac{\beta}{\gamma} = \frac{0.7}{4} < 0.3 = 1-\beta$*, and Assumption I holds for* $\theta \leq \theta_p$*. Figure 5 depicts the economics of an optimal exemption.*

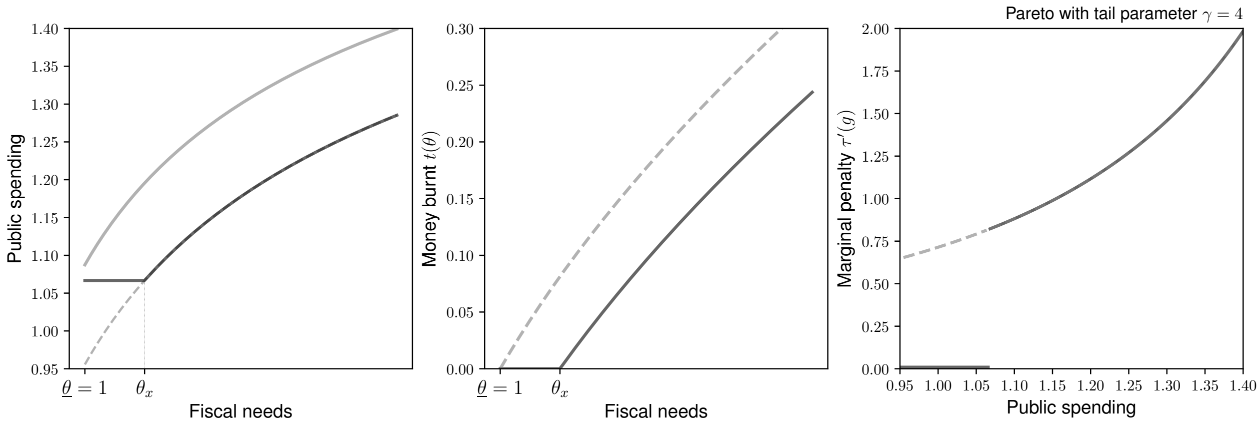

Figure 5: Intermediate degree of present bias and a high need for discretion.

Notes: The grey line in the left panel depicts the discretionary allocation. The dashed lines and the black lines depict the allocation (left panel), the money burnt (middle panel), and the marginal penalties (right panel) with and without the exemption from penalties for public spending below $g_n(\theta_x)$. While the exemption forgoes discipline below $g(\theta_x)$ (left panel) by truncating the marginal penalty schedule (right panel), it lowers the level of penalty (middle panel) while keeping the same level of discipline above $g(\theta_x)$ (left and right panels).

3.2.2 Low need for discretion

The result in this section complements Proposition 4. If the need for discretion is low instead of high, then the need for discipline may outweigh the need for discretion. In this case, a cap that fulfills the average fiscal need is optimal. Define the *tight cap* allocation to be $g_c(\theta) = g_c$ for $\theta \in \Theta$, where g_c fulfills the expected fiscal need, $W'(T-g_c) = \mathbb{E}[\theta]U'(g_c)$.

Proposition 5 (Optimal fiscal rule: intermediate degree of present bias and low need for discretion)**.** *Suppose that Assumption I holds for* $\underline{\theta}$*. If the derivative of* $\rho\frac{1-F}{f}$ *is smaller than* $\rho - \beta$ *for* $\theta \in \Theta$*, then a fiscal rule that implements the* tight cap *allocation is optimal.*

Appendix C.11 contains the proof.

Example (Intermediate degree of present bias and low need for discretion)**.** *The economy is identical to the one in the previous example, except that the distribution is exponential with* *parameter $\lambda = 3$. Unlike the distributions in the previous example, the exponential distribution is log-concave. Proposition 5 applies since the exponential distribution has a log-concave tail, $\beta = \rho$, and Assumption I holds for $\theta = \underline{\theta}$. The tight cap allocation is depicted in the left panel of Figure 6.*

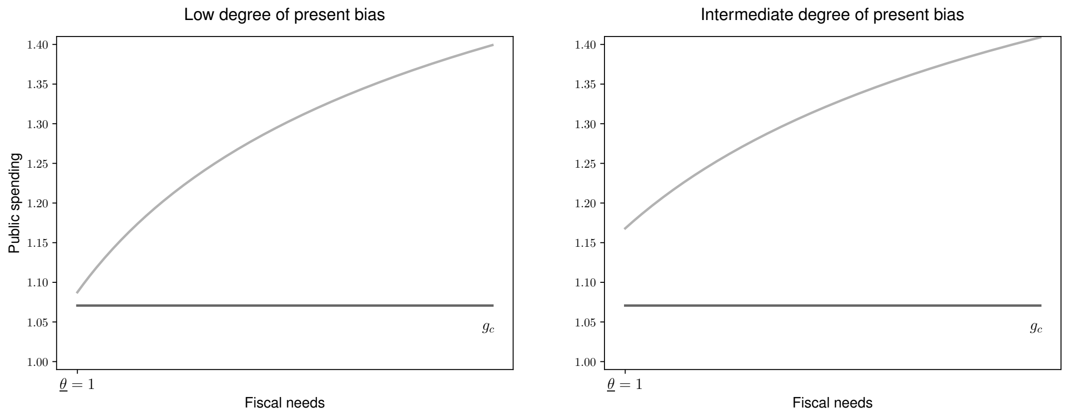

Figure 6: Tight caps.

Notes: The grey lines depict the discretionary allocation, which is higher in the right panel because the degree of present bias is higher. The black lines depict the optimal allocation at the tight cap, which fulfills the expected fiscal needs of the society.

3.3 High degree of present bias

Proposition 6 (Optimal fiscal rule: high degree of present bias)**.** *Suppose that Assumption H holds. A fiscal rule that implements the* tight cap *allocation is optimal.*

The proof of the following proposition is in Appendix C.12. The tight cap is so stringent that it binds for all possible shock realizations. Indeed, Assumption H implies that the tight cap constrains the government even for its lowest possible fiscal need (i.e., $g_c \leq g_d(\underline{\theta})$) (see Appendix C.14 for the formal statement and its proof).

Example (High degree of present bias)**.** *The economy is identical to the one in the previous example, except that the degree of present bias is high because $\beta = 0.4$. The condition of Proposition 6 holds because $\rho\frac{1}{\lambda} = 0.2\bar{3} \leq 0.3 = \rho - \beta$ for $\rho = 0.7$ and $\lambda = 3$. The tight cap allocation is depicted in the right panel of Figure 6.*

4 Measurement

The need for discretion, governed by the elasticity of the right tail of the distribution of shocks, is a rich object that is not directly observable. This section introduces a two-step method to measure the need for discretion. A first step uses a positive model to infer the unobserved past fiscal needs from government finance data. A second step uses tools from heavy-tail analysis to infer the thickness of the right tail of the distribution of fiscal needs measured in the first step.

4.1 A tractable positive model of government deficit

In this section, I enrich the normative model of Section 1 to obtain a positive model from which unobserved fiscal needs are exactly identified from government finance data.

I enrich the model by fleshing out the dynamics, and, because the yearly frequency of government finance data tends to be shorter than the cyclicality of government budgets, I allow for shocks to the fiscal needs to be persistent.[18] The government budget constraint governs the dynamics of the debt, $g = T - Rb + b'$. The problem of the government is

$$\begin{aligned} V(w,\theta) = \max_{g,b'\in\mathbb{R}_+\times(-\infty,\bar{b}]} \quad & \theta U(g) \ + \ \beta\delta E\big[V(w',\theta')|\theta\big] \\ \text{s.t.} \quad w \ &= \ g - b', \\ w' \ &= \ T - Rb', \\ \theta' \ &= \ \theta_e + \varphi\theta + \epsilon', \text{ and } \epsilon' \sim F, \end{aligned} \tag{13}$$

where, for analytical tractability, the borrowing capacity $\bar{b}$ is the natural borrowing limit, and $-1 < \varphi < 1$.

The taste shock θ is a catchall for all sources of fluctuations in fiscal needs, irrespectively of whether it originates in a fluctuation in spending needs, in the cost of servicing the debt, or in revenue. Hence, government revenue is constant in the model precisely to attribute any fluctuation in government revenue in the data to a fluctuation in measured fiscal needs θ.[19] This

[18] Halac and Yared (2014) show that in a world with persistent shocks, although the sequentially optimal fiscal rule is static, the ex ante optimal fiscal rule is dynamic. Because the normative analysis above abstracts from issues of persistence in the shocks, it restricts attention to sequentially optimal fiscal rules.

[19] This interpretation finds support in Section 5.4 in Amador, Werning, and Angeletos (2006). With a CARA utility index $U(g) = e^{-\alpha g}$, the mapping of additive shocks to the government revenue $\tilde{T}$ into taste shocks has a closed-form expression: $\theta = e^{-\alpha\tilde{T}}$. With a log utility index, however, the mapping is not in closed form.

approach has the advantage of yielding a tractable model in the spirit of recent developments in the structural identification of uninsurable shocks in the labor market and at home (e.g., Heathcote, Storesletten, and Violante (2014) and Boerma and Karabarbounis (2022)).

Given that shocks may be serially correlated, I assume $U(g) = \ln(g)$ to preserve analytical tractability (more on the role of the elasticity of intertemporal substitution below). Also, although the formulation above assumes a constant gross interest rate R for simplicity, it is without loss of generality with log utility (see Barro (1999)).

Lastly, to simplify the exposition, the government discounts the future geometrically instead of quasi-hyperbolically. This is without loss for the positive model because, with log preferences, a hyperbolic discounter behaves like a geometric discounter (see Barro (1999)). The discount factor of the society is δ, whereas the present-biased government has discount factor $\beta\delta$.

At this stage, the task of using a dynamic model to infer the distribution of shocks may appear intractable because of the simultaneity in determining the value function and the distribution of spending needs. This is precisely the motivation to keep the model tractable. The model admits an analytical solution because, as shown below, the value function depends only on the first moment of the distribution of taste shocks, which can conveniently be normalized because taste shocks are in utils, $\theta_e = (1 - \beta\delta)(1 - \varphi\beta\delta)$.

The solution method is to guess and verify the solution. The guess is

$$V(w, \theta) = a(\theta) \ln(w + \bar{b}) + \nu(\theta), \tag{14}$$

where $a(\theta)$ and $\nu(\theta)$ capture the dependence of the continuation value on the realized fiscal need and on the first moment of the distribution of fiscal needs. In turn, the guess for the policy functions is linear in effective wealth,

$$g(w, \theta) = (1 - s(\theta))\ (w + \bar{b}), \tag{15}$$

and $b'(w, \theta) = -s(\theta)\ (w + \bar{b}) + \bar{b}$, where $s(\theta)$ denotes the savings rate. Note that the savings rate $s(\theta)$ refers to the ratio of the stock of unused borrowing capacity $\bar{b} - b'(w, \theta)$ over the stock of effective wealth $w + \bar{b}$ (unlike other common definitions of the savings rate based on the ratio of the flow of savings over the flow of income).

Proposition 7. *The value function* (14) *satisfies the Bellman equation. The policy function*

(15) *solves the government problem with the value function* (14), *and the savings rate is*

$$s(\theta) = \frac{\beta\delta\left(1 + \frac{\varphi}{1-\varphi\beta\delta}\theta\right)}{\theta + \beta\delta\left(1 + \frac{\varphi}{1-\varphi\beta\delta}\theta\right)}. \tag{16}$$

The proof is in Appendix C.13. The formula for the savings rate is only a function of the discount factors, the persistence in the distribution of fiscal needs, and, importantly, of the realized fiscal needs. In response to a persistent shock, the government behaves as if it was more patient than it would be if shocks were iid. As a result, the savings rate is higher, and, more importantly, the θ-elasticity of the savings rate is lower with persistent shocks.

Corollary 4 (Elasticity of savings rate and persistence)**.** *The elasticity of the government's savings rate to fiscal needs is lower with persistent shocks than that with iid shocks,*

$$\left|\frac{\partial \ln s(\theta; \varphi > 0)}{\partial \ln(\theta)}\right| \leq \left|\frac{\partial \ln s(\theta; \varphi = 0)}{\partial \ln(\theta)}\right|.$$

Corollary 4 implies that abstracting from the persistence in the shocks would give a conservative measure of the variation in fiscal needs as a function of the variation in the savings rate.[20]

Exact identification of the need for public spending. The key to the identification of fiscal needs obtains from inverting (16) to express θ as a function of s,

$$\theta(s) = \beta\delta\left(\frac{1-s}{s}\right)\left(\frac{1-\varphi\beta\delta}{1-\varphi\beta\delta\left(\frac{1}{s}\right)}\right). \tag{17}$$

Equation (17) identifies the unobserved fiscal needs that rationalize the savings rates observed in government finance data.

The degree of persistence φ used for the measurement ought to be consistent with the autocorrelation in the measured fiscal needs. I let the data determine φ by finding a fixed point of the function that maps φ to the estimated autocorrelation in the measured fiscal needs with φ (iterating over φ worked well in the application below). The persistence introduces a distinction between the measured fiscal needs θ and the shock to the fiscal needs ϵ. I use the residuals $\hat{\epsilon}$ from fitting an AR(1) process on the measured fiscal needs as an estimate of the shocks.

[20] Appendix D contains a sensitivity analysis of the measurement with respect to the persistence of the shocks.

4.2 Measuring the thickness of $1 - F$

In theory, kernel density estimation is a natural approach to infer the distribution from a sample. It amounts to smoothing out the histogram by placing a kernel density on each data point and summing them up to obtain the estimated density. In practice, however, because government finance datasets are finite samples containing rarely occurring large observations—say, during recessions—the kernel density estimator performs poorly in this context. A complementary approach draws from heavy-tail analysis to infer whether the tail of the distribution behaves asymptotically like a power function. That is, $1 - F$ is *heavy* if there exists $\gamma > 0$, the exponent of variation, such that $\lim_{t \to \infty} \frac{1-F(\theta t)}{1-F(t)} = \theta^{-\gamma}$. The inverse of the exponent of variation governs the tail thickness at the top (i.e., for $\theta \to \infty$). Interestingly, estimation methods from heavy-tail analysis are also informative about the behavior of the tail below the top (i.e., for θ above some threshold). The Pareto distribution is a notable example of a heavy tail because the tail is a power function $1 - F(\theta) = \theta^{-\gamma}$.

I use two methods to investigate the behavior of $1 - F$. The appeal of the first method—the tail empirical distribution—is that it is graphical and intuitive. The appeal of the second method—the Hill estimator—is that it is well-suited for serially correlated observations. A plot of the tail empirical distribution depicts the log of the rank on the log of the size of the observations. Intuitively, the motivation comes from taking the log on both sides of $1 - F(\theta) = \theta^{-\gamma}$, which gives a linear relationship whose slope is the tail exponent γ. The tail empirical distribution plots

$$\{\ln(\theta_{j:N}), \ln(1 - \hat{F}(\theta_{j:N})), j = 1, \dots N\}, \tag{18}$$

where $\theta_{j:N}$ refers to the j^{th}-order statistics from $\theta_{1:N} \leq \theta_{2:N} \leq \dots \theta_{N:N}$, and $\hat{F}(\theta_{j:N}) = \frac{j}{N}$ denotes the rank of the observation.

If the plot of the log of the rank on the log of the size of the measured θ depicts a linear relationship, the tail exponent can be estimated by OLS (see Gabaix and Ibragimov (2011)). Using the closed-form expression (17) makes the identification of the thickness of the tail from government finance data on savings rate $(s_t)_{t=1}^{N}$ transparent,

$$\ln(1 - \hat{F}(\theta_{j:N})) = \text{constant} - \gamma \ln\left(\frac{1 - s_{j:N}}{s_{j:N}}\right) + \gamma \ln\left(1 - \hat{\varphi}\beta\delta\frac{1}{s_{j:N}}\right), \tag{19}$$

which shows that for independent data, variations in the savings rate alone identify γ. The discount factors matter for the evaluation of the estimation of γ only for the term correcting for the serial dependence in fiscal needs.

The role of the elasticity of intertemporal substitution (EIS) is hidden in (19) because with a log utility index, the EIS equals 1. Appendix D shows that the estimation identifies $\frac{\gamma}{EIS}$. Hence, the measured tail thickness is inversely related to the EIS. Intuitively, with a smaller EIS, a larger variation in fiscal needs rationalizes the observed variation in savings rates.

The second method consists of estimating the tail parameter using the Hill estimator for serially correlated data. Intuitively, for iid Pareto-distributed data, the Hill estimator corresponds to the maximum likelihood estimator of the tail parameter. The Hill estimator remains a consistent estimator of the exponent of variation γ even if the the data are not Pareto distributed and exhibit serial correlation (see Resnick and Stărică (1995)). Although the Hill estimator is consistent irrespectively of whether the measured θ or the residual $\hat{\epsilon}$ from fitting an autoregressive process is used, for finite samples, Resnick and Stărică (1995) recommend using the residuals. The Hill estimator of $1/\gamma$ based on $N-k$ upper-order statistics is

$$H_k = \frac{1}{N-k}\sum_{t=k+1}^{N} \ln\left(\frac{\hat{\epsilon}_{t:N}}{\hat{\epsilon}_{k:N}}\right). \tag{20}$$

Computing the Hill estimator for different upper-order statistics is informative about the behavior of the tail above different thresholds.

5 Application: the case of the European Union

In this section, I use the two-step methodology to measure the need for discretion of members of the European Union (EU) over the past 28 years. I find evidence of a Pareto tail of the distribution of shocks.

5.1 Government finance data

The dataset is from EuroStat. It contains data on government expenditure, revenue, debt, and interest payable on government debt for the 27 current EU members between 1995 and 2022. EuroStat harmonizes the data across EU members and computes interest payable on an accrual basis (unlike the interest payments reported by the US government (see Hall and Sargent (2011)). Because the model abstracts from default, I do not include data covering periods of high spreads for the countries concerned (Italy, Greece, Portugal, and Spain between 2007 and 2014, and Bulgaria and Romania before 2010). The resulting sample has $N = 631$ observations.

Abstracting from the Excessive Deficit Procedure (EDP) of the European fiscal rule in the positive model, despite the EDP being in place from 1995 to 2019, makes sense for two reasons. First, note that it accords with the historical evidence that EU members largely disregarded the EDP. Figure 7 shows that as much as 34% of the deficits/GDP of EU members between 1995 and 2019 exceeded the 3% threshold above which the EDP would, in theory, impose penalties.[21] Second, if the EDP in place between 1995 and 2019 did incentivize EU members to have lower government deficits, then abstracting from the EDP gives a conservative measure of the spending needs.

Figure 7: EU deficits, 1995-2019

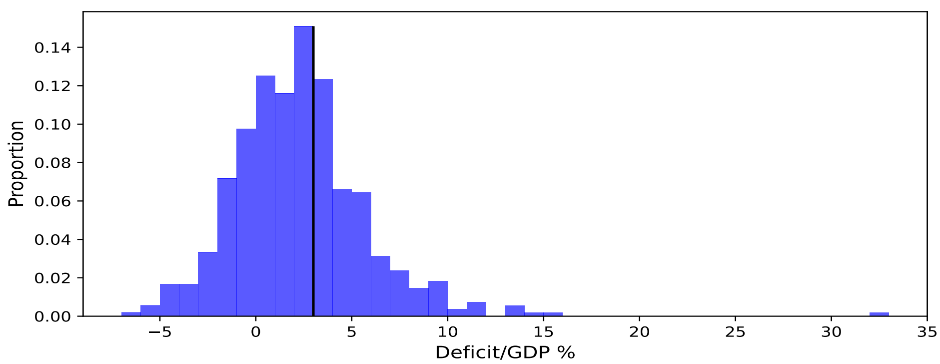

Application to an economic union. The measurement outlined in Section 4 applies without change to an economic union of homogeneous members. The member countries of the EU, however, are not homogeneous. Following the literature on rare events in macroeconomics, the key assumption in using the panel of measured shocks for heterogeneous members to infer the thickness of the tail of the distribution of shocks ϵ is that the distributions have the same exponent of variation (see, for instance, Barro and Jin (2011)). The member countries can, however, be heterogeneous in their other characteristics including their government revenues, the degree of

[21] Also, although the lack of enforcement of the EDP is commonly interpreted as evidence of the non-enforceability of supra-national fiscal rules, note that it does not mean that a better-designed fiscal rule is not enforceable (see Halac and Yared (2023) for a theory of self-enforcing fiscal rules and Dovis and Kirpalani (2020, 2021) for reputational concerns in the design of a fiscal rule with limited commitment).

present bias of their government, their cost of borrowing, the mean of their fiscal needs, and the persistence of shocks to their fiscal needs.

5.2 Step 1: measurement of past fiscal needs

The key input to measure fiscal needs is the empirical counterpart to the savings rate. Using (15), the savings rate is the complement to 1 to the spending rate, and the spending rate is the ratio of public spending as a fraction of the net present value of the government's revenue net of servicing the debt,

$$s_{it} = 1 - \frac{g_{it}}{\bar{b}_i + T_{it} - R_{it} d_{it-1}}, \tag{21}$$

where g_{it}, T_{it}, and d_{it} denote government i spending, revenue, and debt in year t. The natural borrowing limit $\bar{b}_i$ is set at $\frac{T_i}{R-1}$, where T_i and R denote averages over the sample period for the revenue of government i and the interest rate payable on the debt.[22] For a given country i, the savings rate varies over time because of variations in government spending, in the cost of servicing the debt, and in government revenues.

I obtain a panel of measured fiscal needs θ_{it} by using s_{it} from (21) in (17) with $\delta = 0.96$, and the present bias β_i is calibrated, for each country, to match the expected government spending in the model to the average government spending in the data. The average β is 0.86. The Greek government has the lowest β at 0.73 relative to the government of Luxembourg, which has the highest β at 1. Further details on the calibration and a sensitivity analysis with respect to the present bias are in Appendix D. The shocks $\hat{\epsilon}_{it}$ are the residuals from fitting country-specific AR(1) processes to the measured fiscal needs, with the location minimally shifted by adding $\min_{it} \hat{\epsilon}_{it}$ so that the shocks are positive (it is without loss because it amounts to a normalization of θ_{ei}). The shift in location does not affect the tail index. The Hill estimator, however, is not location invariant (Appendix D addresses this shortcoming of the Hill estimator with a sensitivity analysis).

5.3 Step 2: measurement of the tail thickness

In this section, I present evidence of a thick (Pareto) tail of the distribution of shocks on the need for government spending of EU members.

[22] Appendix D includes a sensitivity analysis to alternative measures of the borrowing limit.

Figure 8: Tail empirical distribution of fiscal needs of EU members, 1995-2022

(a) All measured fiscal needs

(b) Largest two-thirds of measured fiscal needs

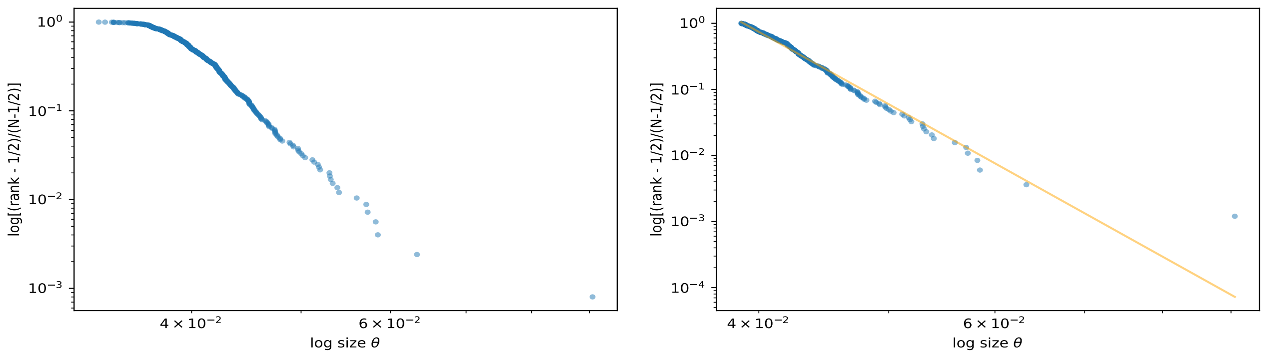

The tail empirical distribution. Figure 8a depicts the tail empirical distribution (18). The correction of the rank by -1/2 reduces a finite sample bias in estimating the tail exponent (see Gabaix and Ibragimov (2011)).

Figure 8b depicts evidence that the tail of the distribution of fiscal needs for EU members follows a Pareto distribution above a threshold. Figure 8b displays the top two-thirds of the measured fiscal needs, which amounts to truncating Figure 8a by retaining the 420 largest observations. The alignment of the log rank on log size indicates that a Pareto distribution fits the upper tail of the distribution of spending needs remarkably well.[23]

The slope of the line fitted to the top two-thirds of the observations is an estimate of the thickness of the Pareto tail parameter. The ordinary least squares estimate from a regression of the log of rank corrected by a half on the log of size gives $\hat{\gamma} = 11.3$, with standard error $0.78 = \sqrt{\frac{2}{(2/3)N}}\hat{\gamma}$.

The Hill plot. Figure 9 depicts the Hill plot $\{(k, H_k^{-1}), 1 \leq k \leq N\}$, where H_k are the Hill estimates of $1/\gamma$ computed with the Hill estimator (20) and the shocks $\hat{\epsilon}$.

The Hill plot substantiates the finding from the tail empirical distribution. There is evidence of a heavy tail above a threshold, which corresponds to the top two-thirds of the observations. At the left end of Figure 9, the Hill estimator based on all the residuals indicates no evidence of a Pareto tail, which is consistent with the curvature on the left end of the tail empirical distribution. At the opposite end, the Hill estimator becomes arbitrarily imprecise because it is based on a

[23] Figure 10 in Appendix D depicts each country's contribution to the empirical tail distribution in Figure 8.

Figure 9: Hill plot for the tail of the distribution of shocks to fiscal needs of EU members

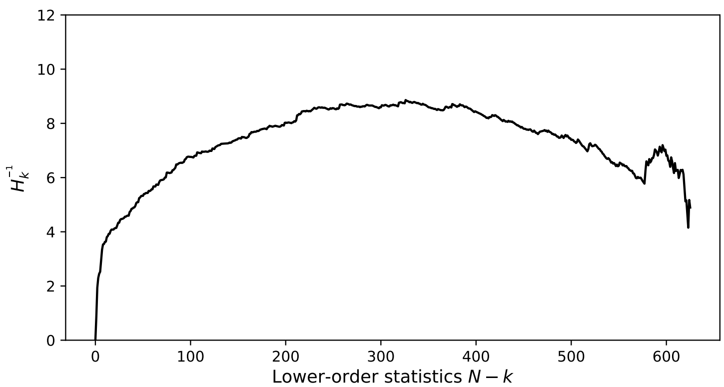

decreasing number of observations. In between, the Hill plot stabilizes with estimates of γ in a range between 5 and 7. The lower estimate from the Hill plot relative to the estimate from the tail empirical distribution accords with the implications of Corollary 4; that is, not accounting for persistence in the shock process biases the measure of the tail thickness downward.

6 Conclusion

To conclude, I use the findings from this paper to evaluate the Excessive Deficit Procedure (EDP) of the Stability and Growth Pact and to propose avenues for reforms. For some context, the following quote reflects the current penalty schedule of the EDP:

> A non-interest-bearing deposit of 0.2% of GDP may be requested from a euro area country that is placed in EDP. [...] In case of non-compliance with the initial recommendation for corrective action, this non-interest-bearing deposit will be converted into a fine. *—European Commission, "EU Economic Governance 'Six Pack' State of Play," Memo/11/647, September 28, 2011.*

First, there is a threshold below which discretion prevails and above which the country is placed in EDP. This paper lends support to this feature of the Stability and Growth Pact.

Second, the EDP features a jump in the level of penalties—a notch point—from 0 to a non-interest-bearing deposit of 0.2% of GDP. Initially, the penalty is the forgone interest on the deposit. A notch can be part of the optimal penalty schedule if the distribution of shocks has a sufficiently thin tail. In this case, the notch needs to be prohibitive to implement a cap. With France and Germany violating the EDP, history has taught us that the notch did not implement a cap. Besides, doubt remains about the enforceability of the penalties because the violations were left unpunished (see Dovis and Kirpalani (2020, 2021) and Halac and Yared (2022a, 2023)).

This paper suggests that the lack of enforcement may partly be due to a poor design of the penalty schedule in the EDP. Under the conditions outlined in Section 3.2, and given the novel evidence of a Pareto tail of the distribution of shocks to the fiscal needs of EU members, a fiscal rule featuring a graduated schedule of mild penalties is a promising avenue for reform. Although the EDP features some gradualism with the possibility to convert the deposit into a fine (i.e., the penalty would then be the deposit instead of the forgone interest on this deposit), the theory does not lend support to gradualism with two notches. Expressing the penalty as a percentage of the deficit above a threshold, instead of a percentage of GDP, would turn the notch into a kink or a smooth schedule.

Lastly, the global Lagrangian method is widely applicable to solve mechanism design problems with limited transfers in other contexts. Werning (2007) studies Pareto efficient income taxation, where the requirement that a reform be Pareto improving limits transfers. The method also applies to designing the cost of verifying the state to determine whether escape clauses apply (see Halac and Yared (2020b) for the design of rules with costly state verification). Another promising application is the study of the optimal illiquidity of retirement savings accounts for households who undersave for their retirement (see Laibson et al. (1998) and Beshears et al. (2020)).

References

Aguiar, M., Amador, M., Farhi, E. and Gopinath, G. (2015). Coordination and Crisis in Monetary Unions, Quarterly Journal of Economics, 130(4): 1727-1780

Aguiar, M., Amador, M., and Fourakis, S. (2020). On the Welfare Losses from External Sovereign Borrowing. IMF Economic Review, 68(1): 163-194.

Alesina, A., and Passalacqua, A. (2016). The Political Economy of Government Debt. Handbook of Macroeconomics, 2: 2599-2651. North Holland.

Alfaro, L. and Kanczuk, F. (2019). Fiscal Rules and Sovereign Default. Working paper.

Alonso, R., and Matouschek, N. (2008). Optimal Delegation. Review of Economic Studies, 75: 259-293

Amador, M., and Bagwell, K. (2022). Regulating a Monopolist With Uncertain Costs Without Transfers. Theoretical Economics, 17(4):1719-1760

Amador, M., and Bagwell, K. (2020). Money Burning in the Theory of Delegation. Games and Economic Behavior, 112: 382-412

Amador, M., Bagwell, K., and Frankel, A. (2018). A Note on Interval Delegation. Economic Theory Bulletin, 6(2): 239-249

Amador, M. and Bagwell, K. (2013). The Theory of Optimal Delegation With an Application to Tariff Caps. Econometrica, 81: 1541-1599

Amador, M., Werning, I. and Angeletos, G.-M. (2006). Commitment vs. Flexibility. Econometrica, 74: 365-396

Ambrus, A., and Egorov, G. (2017). Delegation and Nonmonetary Incentives, 171: 101-135

Ambrus, A., and Egorov, G. (2013). Comment on "Commitment vs. Flexibility", Econometrica, 81(5): 2113-2124

Athey, S., Atkeson, A., and Kehoe, P. (2005). The Optimal Degree of Discretion in Monetary Policy, Econometrica, 73(5): 1431-1475

Atkeson, A. and Lucas, R. E. (1992). On Efficient Distribution with Private Information. Review of Economic Studies, 59(3): 427-453

Azzimonti, M., Battaglini, M., and Coate, S. (2016). The Costs and Benefits of Balanced Budget Rules: Lessons from a Political Economy Model of Fiscal Policy, Journal of Public Economics, 136: 45-61

Bassetto, M. and Sargent, T. J. (2006). Politics and Efficiency of Separating Capital and Ordinary Government Budgets. Quarterly Journal of Economics, 121(4):1167-1210

Barro, R. J. (1999). Ramsey Meets Laibson in the Neoclassical Growth Model. Quarterly Journal of Economics, 114(4): 1125-1152

Barro, R. J. and Jin, T. (2011). On the Size Distribution of Macroeconomic Disasters. Econometrica, 79(5): 1567-1589

Beetsma, R. and Uhlig, H. (1999). An Analysis of the Stability and Growth Pact, Economic Journal, 109: 546-571

Beshears, J., Choi, J. J., Clayton, C., Harris, C., Laibson, D., and Madrian, B. C. (2020). Optimal Illiquidity. Working paper.

Boerma, J., and Karabarbounis, L. (2021). Inferring Inequality with Home Production. Econometrica 89(5): 2517-2556

Chari, V. V., and Kehoe, P. (2007). On the need for fiscal constraints in a monetary union. Journal of Monetary Economics, 54: 2399-2408

Clayton, C., and Schaab, A. (2022). A Theory of Dynamic Inflation Targets. Working paper.

Cooper, R., and Kempf, H. (2004). Overturning Mundell: Fiscal Policy in a Monetary Union. Review of Economic Studies, 71: 371-396

Diamond, P. (1998). Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates. American Economic Review, 88: 83-95

Dovis, A., and Kirpalani, R. (2020). Fiscal Rules, Bailouts, and Reputation in Federal Governments. American Economic Review 110(3): 860-888

Dovis, A., and Kirpalani, R. (2021). Rules without Commitment: Reputation and Incentives. Review of Economic Studies, 88(6): 2833-2856

Felli, C., Piguillem, F., and Shi, L. (2021). Fiscal Rules and Discretion with Risk of Default. Working paper.

Galperti, S. (2015). Commitment, Flexibility, and Optimal Screening of Time Inconsistency. Econometrica, 83: 1425-1465

Gabaix, X., and Ibragimov R. (2011). Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents. Journal of Business Economics and Statistics. 29(1): 24-39

Hall, G. J., and Sargent, T. J. (2011). Interest Rate Risk and Other Determinants of Post-WWII US Government Debt/GDP Dynamics. American Economic Journal: Macroeconomics, 3: 192-214

Halac, M., and Yared, P. (2014). Fiscal Rules and Discretion under Persistent Shocks. Econometrica, 82: 1557-1614

Halac, M., and Yared, P. (2018). Fiscal Rules and Discretion in a World Economy. American Economic Review, 108: 2305-2334

Halac, M., and Yared, P. (2020a). Inflation Targeting under Political Pressure. in Independence, Credibility, and Communication of Central Banking, edited by Ernesto Pastén and Ricardo Reis, Santiago, Chile: Central Bank of Chile, 7-27

Halac, M., and Yared, P. (2020b). Commitment vs. Flexibility with Costly Verification. Journal of Political Economy, 128: 4523-4573

Halac, M., and Yared, P. (2022a). Fiscal Rules and Discretion under Limited Enforcement. Econometrica, 90: 2093-2127

Halac, M., and Yared, P. (2022b). Instrument vs. Target-Based Rules. Review of Economic Studies, 89: 312-345

Halac, M., and Yared, P. (2023). A Theory of Fiscal Responsibility and Irresponsibility. *working paper*

Harstad, B. (2007). Harmonization and Side Payments in Political Cooperation. American Economic Review, 97(3): 871-889

Hatchondo, J.-C., Martinez, L., and Roch, F. (2020). Fiscal Rules and the Sovereign Default Premium. American Economic Journal: Macroeconomics, 14(4): 244-273

Heathcote, J., Storesletten, K., and Violante, G. (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. American Economic Review 104(7): 2075-2126

Holmström, B. (1977). On Incentives and Control in Organizations. Ph.D. thesis, Stanford Graduate School of Business

Kováč, E. and Mylovanov, T. (2009). Stochastic mechanisms in settings without monetary transfers: The regular case. Journal of Economic Theory, 144: 1373-1395

Laibson, D., I. Repetto, A., Tobacman, J., Hall, R., Gale, W. and Akerlof, G., A. (1998). Self-Control and Saving for Retirement. Brookings Papers on Economic Activity 91-196

Luenberger, D. G. (1969). *Optimization by Vector Space Methods.* Series in Decision and Control. New York: Wiley

McKibbin W. J., and Wilcoxen, P. J. (2002). The Role of Economics in Climate Change Policy. Journal of Economic Perspectives, 16(2): 107-129

Melumad, N. D., and Shibano, T. (1991). Communication in Settings with No Transfers. The Rand Journal of Economics, 22: 173-198

Myerson, R. B. (1981). Optimal Auction Design. Mathematics of Operations Research, 6: 58-73

Resnick, S. (2007). Heavy-Tail Phenomena: Probabilistic and Statistical Modeling. Springer-Verlag, New York

Resnick, S., and Stărică, C. (1995). Consistency of Hill's Estimator for Dependent Data. Journal of Applied Probability, 32(1): 139-167

Weitzman, M. (1974). Prices vs. Quantities. Review of Economic Studies, 41(4): 477-491

Werning, I. (2007). Pareto Efficient Income Taxation. working paper.

Yared, P. (2019). Rising Government Debt: Causes and Solutions for a Decades-Old Trend. Journal of Economic Perspectives, 33: 115-140

Online Appendix

A Model

A.1 Quadratic preferences with additive bias

Although the economic environment in this paper does not nest the quadratic preferences model with additive bias, the analysis of this paper applies with a minor change to the definition of the wedge.

The social welfare function for the quadratic preferences model with additive bias is $\theta g - \frac{g^2}{2}$, and the governments' objective is $\theta g + (1-\beta)g - \frac{g^2}{2}$. The discretionary allocation differs from the first-best allocation due to an additive bias, $g_d(\theta) = \theta + (1-\beta)$.

While the natural definition of the wedge in the environment with a multiplicative bias is a proportional wedge in the Euler equation, for an environment with an additive bias, an additive wedge is analytically more tractable. Define the *additive wedge* Δ^a in the Euler equation of the government as follows: $\Delta^a(g,\theta) = \theta + (1-\beta) - g$. The additive wedge at the discretionary allocation is null and positive for an allocation that spends less than the discretionary allocation.

The design of a rule with an additive bias is

$$\max_{g(\cdot),\ t(\cdot)} \left\{ \int_\Theta \left[\theta g(\theta) - \tfrac{1}{2} g(\theta)^2 - \rho t(\theta)\right] dF(\theta) \;\middle|\; \text{(IC) and } t(\theta) \geq 0 \text{ for } \theta \in \Theta \right\}, \tag{22}$$

where the incentive compatibility constraints are

$$\theta g(\theta) + (1-\beta)g(\theta) - \tfrac{1}{2}g(\theta)^2 - t(\theta) \geq \theta g(\hat\theta) + (1-\beta)g(\hat\theta) - \tfrac{1}{2}g(\hat\theta)^2 - t(\hat\theta), \quad \text{for } \theta, \hat\theta \in \Theta. \tag{IC}$$

Lemma 7 (Incentive compatible allocations)**.** *An allocation $g(\cdot)$ is incentive compatible given a money-burning schedule $t(\cdot)$ if and only if $g(\cdot)$ is non-decreasing and*

$$t(\theta) = t(\underline{\theta}) + \theta g(\theta) + (1-\beta)g(\theta) - \tfrac{1}{2}g(\theta)^2 - \underline{\theta} g(\underline{\theta}) - (1-\beta)g(\underline{\theta}) + \tfrac{1}{2}g(\underline{\theta})^2 - \int_{\underline{\theta}}^{\theta} g(\tilde\theta)\, d\tilde\theta. \tag{23}$$

Substituting the money burning schedule (23) in the objective in (22) gives

$$\int_\Theta \left[-(1-\beta)g(\theta) + (1-\rho)\left(\theta g(\theta) + (1-\beta)g(\theta) - \tfrac{1}{2}g(\theta)^2\right) + \rho \tfrac{1-F(\theta)}{f(\theta)} g(\theta)\right] dF(\theta) \tag{24}$$
$$- \rho\left(t(\underline{\theta}) - \underline{\theta} g(\underline{\theta}) - (1-\beta)g(\underline{\theta}) + \tfrac{1}{2}g(\underline{\theta})^2\right).$$

Maximizing (24) pointwise for $\theta > \underline{\theta}$ and expressing the first-order condition as an additive wedge gives the following definition:

$$\Delta_n^a(\theta) = \frac{1}{1-\rho}\left((1-\beta) - \rho\frac{1-F(\theta)}{f(\theta)}\right). \qquad \text{(additive foa-wedge)}$$

The analysis in the main text applies with this minor modification to the foa-wedge to characterize optimal fiscal rules for the model with quadratic preferences and an additive bias.

A.2 Economic Union

This section introduces the additional notation needed to design a fiscal rule for an economic union.[24] The design a fiscal rule for an economic union that is not fiscally integrated separates into the design of country-specific penalty schedules as studied in Sections 1 to 3.

The economic union has $\mathcal{N}$ countries indexed by $i = 1, \ldots, \mathcal{N}$. Each country has its own government choosing its spending according to (2), where θ denotes the idiosyncratic shock to the country's fiscal needs.

A fiscal rule is a tuple of penalty schedules $(P_i(\cdot))_{i=1}^{\mathcal{N}}$. In the context of an economic union, the non-negativity constraint on penalties captures the lack of fiscal integration because a negative penalty could model a transfer across members. The lack of fiscal integration also implies that the budget constraints of the different members are independent due to the absence of transfers, $g + \frac{x}{R_i} = T_i$. The country-specific interest rates are assumed to be exogenous which implicitly assumes that each country is small and the economic union is small as well. Halac and Yared (2018) show that for a large economic union, however, optimal coordinated and uncoordinated fiscal rules differ due to the disciplining and redistributive effects of the endogenous response of the interest rate to the fiscal rule.

The three key determinants of an optimal fiscal rule, namely the distribution of shocks F_i, the degree of present bias $1-\beta_i$, and the degree of asymmetry in the cost of meting out a penalty $1-\rho_i$, can be country specific. The welfare of the economic union is the aggregation of the

[24] A long line of research studies fiscal rules in economic unions with limited commitment (see Beetsma and Uhlig (1999), Cooper and Kempf (2004), Chari and Kehoe (2007), Chari and Kehoe (2008), Aguiar, Amador, Farhi, and Gopinath (2015), Dovis and Kirpalani (2020), and Dovis and Kirpalani (2021)). Yared (2019) surveys the literature. As mentioned in the introduction, some of the literature provides the micro-foundations for modeling the deficit bias on the part of members of an economic union with a present bias and focuses on the optimal stringency of a cap on deficit. This paper focuses on the optimal structure of a fiscal rule.

welfare of each country:

$$\sum_{i=1}^{\mathcal{N}} \int_{\Theta_i} \left[\theta U(g_i(\theta)) + W_i(T_i - g_i(\theta)) - \rho_i P_i(g_i(\theta))\right] dF_i(\theta),$$

where the country-specific interest rates are subsumed into the continuation values $W_i(T_i - g(\theta)) = W(R_i(T_i - g(\theta)))$. Setting equal weights on the welfare of each country is without loss of generality precisely because, for strictly positive weights, the problem for an economic union separates into $\mathcal{N}$ country-specific problems independently of the weights.

B Solution method for the design of rules

The non-negativity constraint on the penalty schedule sets program (4) apart from mechanism design problems in which transfers are possible. Unlike the incentive compatibility constraints, the non-negativity constraint on the penalty schedule cannot be easily summarized in the objective function to be maximized point-wise without resorting to Lagrangian methods. This section outlines how I use the first-order conditions of the Lagrangian method to identify a candidate solution and to find conditions for global optimality.

The first step uses a Lagrange multiplier function to combine the non-negativity constraint on the penalty schedule and the objective in a Lagrangian. Let $\Lambda : \Theta \mapsto [0,1]$ be a non-decreasing function such that $\lim_{\theta \to \bar{\theta}} \Lambda(\theta) = 1$ and $1 - \Lambda$ is integrable. A valid Lagrange multiplier function is non-decreasing, which is the infinite dimensional analog of a non-negative Lagrange multiplier for the Kuhn-Tucker theorem with finitely many inequality constraints. Define the Lagrangian, with Lagrange multiplier function Λ, as a functional on $\Phi \equiv \{(u, \underline{t}) \mid u : \Theta \mapsto \mathbb{R}_+ \text{ is non-decreasing, and } \underline{t} \in \mathbb{R}_+\}$ as follows:

$$\mathcal{L}(u, \underline{t}|\Lambda) \equiv \int_{\Theta} \left[\theta u(\theta) + W(T - U^{-1}(u(\theta))) - \rho t(\theta, u, \underline{t})\right] dF(\theta) + \int_{\Theta} t(\theta, u, \underline{t}) d\Lambda(\theta), \tag{25}$$

where $t(\theta, u, \underline{t})$ is the schedule associated with allocation $U^{-1}(u(\cdot))$ in Lemma 1 and $t(\underline{\theta}) = \underline{t}$.

The Gateaux derivative in the direction $(h, h_t) \in \Phi$ is defined as follows:[25]

$$\partial\mathcal{L}(u, \underline{t}, h, h_t|\Lambda) \equiv \lim_{\alpha \downarrow 0} \frac{1}{\alpha}\left[\mathcal{L}(u + \alpha h, \underline{t} + \alpha h_t|\Lambda) - \mathcal{L}(u, \underline{t}|\Lambda)\right]. \tag{26}$$

The next lemma gives optimality conditions in terms of the Gateaux derivative evaluated at the solution. The optimality conditions are that the Gateaux derivative evaluated at the solution is null in the direction of the solution and non-positive in any non-decreasing direction.

[25]The existence of the Gateaux differential follows from Lemma A.1 p. 390 of Amador, Werning, and Angeletos (2006).

Lemma 8 (Lemma of optimality)**.** *If there exists a non-decreasing* $u^* \equiv U(g^*)$ *and* $\underline{t}^*$ *in the convex cone* Φ *and a non-decreasing function* $\Lambda^* : \Theta \mapsto [0,1]$ *such that* $\lim_{\theta \to \bar{\theta}} \Lambda^*(\theta) = 1$ *and* $1 - \Lambda^*$ *is integrable, and if*

$$\partial \mathcal{L}(u^*, \underline{t}^*, u^*, \underline{t}^* | \Lambda^*) = 0, \quad \textit{and} \quad \partial \mathcal{L}(u^*, \underline{t}^*, h, h_t | \Lambda^*) \leq 0 \quad \textit{for all } (h, h_t) \in \Phi,$$

then $g^* \equiv U^{-1}(u^*)$ *and the associated money-burning schedule* t^* *characterized by* (5) *with* $t^*(\underline{\theta}) = \underline{t}^*$ *solve the mechanism design problem* (4).

The proof is an application of the global theory of constrained optimization (Chapter 8 in Luenberger (1969), Lemma 1 p.227, and Theorem 1 p.220), as used in Lemma A.2 in Amador, Werning, and Angeletos (2006) and Theorem 1 in Amador and Bagwell (2013). A part of the proof shows that the degree of concavity of the Lagrangian depends on ρ. The Lagrangian is strictly concave for $\rho \in [0, 1)$ and a non-decreasing Lagrange multiplier function (see the proof of Lemma 8). The Lagrangian is linear for $\rho = 1$.

Proof of Lemma 8. Lemma A.2 in Amador, Werning, and Angeletos (2006) implies that if the Lagrangian with Lagrange multipliers Λ^* is concave, then the equality and inequality conditions in terms of Gateaux derivatives imply that the Lagrangian is maximized at $u^*, \underline{t}^*$:

$$\mathcal{L}(u^*, \underline{t}^* | \Lambda^*) \geq \mathcal{L}(u, \underline{t} | \Lambda^*) \quad \text{for all } (u, t) \in \Phi.$$

To show the concavity of the Lagrangian with Lagrange multipliers Λ^*, it is convenient to spell out the Lagrangian (25) and factorize the non-linear terms as follows:

$$\begin{aligned}
\mathcal{L}(u, \underline{t} | \Lambda^*) \equiv & \int_\Theta \left[u(\theta) \ \left((1 - F(\theta)) - \theta \tfrac{\rho - \beta}{\rho} f(\theta) \right) \right] d\theta \\
& - \int_\Theta \left[u(\theta)(1 - \Lambda^*(\theta)) \right] d\theta + (\underline{\theta} \ u(\underline{\theta}) - \underline{t}) \Lambda^*(\underline{\theta}) + \int_\Theta \left[\theta \ u(\theta) \right] d\Lambda^*(\theta) \\
& + \int_\Theta \left[\beta W(T - U^{-1}(u(\theta))) \right] d \left(\tfrac{1-\rho}{\rho} F(\theta) + \Lambda^*(\theta) \right) \\
& + \beta W(T - U^{-1}(u(\underline{\theta}))) \Lambda^*(\underline{\theta}).
\end{aligned}$$

The integrands for the integrals in the first two lines are linear in u. For the terms in the remaining two lines, to show that the function $u \mapsto W(T - U^{-1}(u)))$ is concave, note that the utility index U is strictly increasing and concave so its inverse U^{-1} is strictly increasing and convex ($U^{-1\prime}(U(x)) = 1/U'(x)$ and $U^{-1\prime\prime}(U(x)) = -U^{-1\prime}(U(x))U''(x)/U'(x)^2$). Since W is increasing and concave and $-U^{-1}$ is concave, the composition $u \mapsto W(T - U^{-1}(u))$ is concave. A sufficient condition for the Lagrangian to be concave is that the function $\frac{1-\rho}{\rho} F(\theta) + \Lambda^*(\theta)$ be non-decreasing, which is the case since $0 \leq \rho \leq 1$ and F and Λ^* are both non-decreasing.

It remains to show that the maximizer of a concave Lagrangian at a valid Lagrange multiplier is the solution to the constrained optimization problem of interest. This is precisely what the global theory of constrained optimization does for us.

The following notation maps the environment studied in this paper to Theorem 1 in Amador and Bagwell (2013) p.1575: $X = \{u, \underline{t} \mid u : \Theta \mapsto \mathbb{R},\ \underline{t} \in \mathbb{R}\}$, $Z = \{z \mid z : \Theta \mapsto \mathbb{R}\}$ with norm $||z||= \sup_{\theta\in\Theta}|z(\theta)|$, $\Omega = \{(u, \underline{t}) \in X|\ u$ is non-decreasing, $\underline{t} \geq 0\}$, and $P = \{z \in Z \mid z(\theta) \geq 0$ for $\theta \in \Theta\}$. The objective is a functional $f : \Omega \mapsto \mathbb{R}$ defined as follows:

$$f(u, \underline{t}) = -\int_\Theta \left[u(\theta)\rho\tfrac{1-F(\theta)}{f(\theta)} + \beta(1-\beta)W(T - U^{-1}(u(\theta)))\right] dF(\theta)$$
$$- (\rho - \beta)\int_\Theta \left[\theta u(\theta) + \beta W(T - U^{-1}(u(\theta)))\right] dF(\theta).$$

The constraints on limited transfers are defined as follows: $G : \Omega \mapsto Z$,

$$G(u, \underline{t}) = -\left(\underline{t} + \theta\ u(\theta) + \beta W(T - U^{-1}(u(\theta))) - \underline{\theta}\ u(\underline{\theta}) - \beta W(T - U^{-1}(u(\underline{\theta}))) - \int_{\underline{\theta}}^{\theta} u(\tilde{\theta})\ d\tilde{\theta}\right),$$

and their contributions to the Lagrangian are given by $T : Z \mapsto \mathbb{R}$,

$$T(z) = \int_\Theta z(\theta)d\Lambda^*(\theta),$$

which satisfies $T(z) \geq 0$ for all $z \in P$ since Λ^* is non-decreasing. Since $\mathcal{L}(u|\Lambda^*) = -f(u) - T(G(u))$, Theorem 1 from Amador and Bagwell (2013) implies that $(u^*, \underline{t}^*)$ solves

$$\min_{(u,\underline{t})\in\Omega} \{f(u, \underline{t})| - G(u, \underline{t}) \in P\}.$$

Inverting the above mapping from the environment of this paper to Theorem 1 in Amador and Bagwell (2013) and using $t(\cdot)$ defined in (5) as a function of $g(\cdot)$, the allocation $g^* = U^{-1}(u^*)$ and the initial level $\underline{t}^*$ solve the optimization problem:

$$\max_{g\in\Omega,\ \underline{t}\geq 0} \int_\Theta \left[\theta U(g(\theta)) + W(T - g(\theta)) - \rho t(\theta)\right] dF(\theta)$$

s.t. for all $\theta \in \Theta$:

$$\beta t(\theta) = \beta\underline{t} + \theta U(g(\theta)) + \beta W(T - g(\theta)) - \underline{\theta} U(g(\underline{\theta})) - \beta W(T - g(\underline{\theta})) - \int_{\underline{\theta}}^{\theta} U(g(\tilde{\theta}))d\tilde{\theta}$$

g is non-decreasing

$t(\theta) \geq 0.$

The characterization of incentive compatible allocations in Lemma 1 implies that (g^*, t^*) in which

$$\beta t^*(\theta) \equiv \beta\underline{t}^* + \theta\ U(g^*(\theta)) + \beta W(T - g^*(\theta)) - \underline{\theta}\ U(g^*(\underline{\theta})) - \beta W(T - g^*(\underline{\theta})) - \int_{\underline{\theta}}^{\theta} U(g^*(\tilde{\theta}))\ d\tilde{\theta}$$

solve the mechanism design problem (4). □

The solution method appears to ask the designer to guess the solution and verify that it satisfies the optimality conditions in Lemma 8. Guessing the solution amounts to the arrangement of three building blocks. The first building block obtains from ignoring the monotonicity and the non-negativity constraint on the penalty schedule to determine the spending g_n for $\rho < 1$. The second building block is the discretionary allocation. It is a natural candidate because of the non-negativity constraint on the penalty schedule. Third, the allocation may be constant over subintervals of Θ.

I use the optimality conditions of Lemma 8 to determine the arrangement of the three building blocks. The first optimality condition sets the Gateaux derivative of the Lagrangian to zero. It is the first requirement in the definition of the thresholds θ_p and θ_x. It also determines the Lagrange multiplier function. In turn, Assumption sL is precisely the condition needed for the Lagrange multiplier function to be non-decreasing. The second optimality condition verifies that the Gateaux derivative of the Lagrangian in any non-decreasing direction is negative. It is the second requirement in the definition of the thresholds θ_p and θ_x.

For reference, the Lagrangian (25), after rescaling the objective by β and the Lagrange multiplier by $\rho\beta$, reads

$$\begin{aligned}\mathcal{L}(u,\underline{t}|\Lambda) = &\int_\Theta \left[\beta(1-\beta)W(T-U^{-1}(u(\theta))) + \rho\tfrac{1-F(\theta)}{f(\theta)}u(\theta)\right] dF(\theta)\\ &+ (\beta-\rho)\int_\Theta \left[\theta u(\theta) + \beta W(T-U^{-1}(u(\theta)))\right] dF(\theta)\\ &+ \rho\left(\underline{\theta}u(\underline{\theta}) + \beta W(T-U^{-1}(u(\underline{\theta}))) - \underline{t}\right)\Lambda(\underline{\theta})\\ &+ \rho\int_\Theta \left[\theta u(\theta) + \beta W(T-U^{-1}(u(\theta)))\right] \mathrm{d}\Lambda(\theta) - \rho\int_\Theta \left[u(\theta)(1-\Lambda(\theta))\right] d\theta.\end{aligned}$$

The Gateaux derivative in the direction (h, h_t), defined in (26), reads as follows:[26]

$$\begin{aligned}\partial\mathcal{L}(u,\underline{t},h,h_t|\Lambda) = &\int_\Theta \left[\left(-(1-\beta)\theta + \rho\tfrac{1-F(\theta)}{f(\theta)} + (1-\rho)\Delta(U^{-1}(u(\theta)),\theta)\theta\right) h(\theta)\right] dF(\theta)\\ &+ \rho\left(\underline{\theta}\Delta(U^{-1}(u(\underline{\theta})),\underline{\theta})h(\underline{\theta}) - h_t\right)\Lambda(\underline{\theta})\\ &+ \rho\int_\Theta \left[\theta\Delta(U^{-1}(u(\theta)),\theta)h(\theta)\right] \mathrm{d}\Lambda(\theta) - \rho\int_\Theta \left[h(\theta)(1-\Lambda(\theta))\right] d\theta.\end{aligned} \quad (27)$$

C Proofs

C.1 Lemma 1 on incentive compatible money-burning schedules

Proof of Lemma 1. The proof follows the argument in Myerson (1981). Suppose that $g(\cdot)$ is incentive compatible given a money-burning schedule $t(\cdot)$. Define $V(\theta) = \theta U(g(\theta)) + \beta W(T -$

[26] The existence of the Gateaux differential follows from Lemma A.1 p. 390 of Amador, Werning, and Angeletos (2006).

$g(\theta)) - \beta t(\theta)$ and $u(\theta) = U(g(\theta))$. Consider $\theta > \hat{\theta}$, incentive compatibility implies,

$$V(\theta) \geq V(\hat{\theta}) + (\theta - \hat{\theta})u(\hat{\theta}), \quad \text{and} \quad V(\hat{\theta}) \geq V(\theta) + (\hat{\theta} - \theta)u(\theta).$$

The inequalities combined imply that $u(\cdot)$ is non-decreasing,

$$u(\theta) \geq \frac{V(\theta) - V(\hat{\theta})}{\theta - \hat{\theta}} \geq u(\hat{\theta}).$$

Since U is strictly increasing and $u(\cdot)$ is non-decreasing, g is also non-decreasing and $V(\cdot)$ is continuous and differentiable almost everywhere. Taking the limit, $V'(\theta) = u(\theta)$. Integrating from $\underline{\theta}$ to θ gives $V(\theta) = V(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} u(\theta)$. Replacing V and u by their definitions gives (5).

Suppose instead that $g(\cdot)$ is non-decreasing and, for a given $t(\underline{\theta})$, define $t(\cdot)$ according to (5). Using the definitions of V and u, rewrite (5) as follows: $V(\theta) = V(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} u(\theta)d\theta$. For $\theta \geq \hat{\theta}$,

$$V(\theta) - V(\hat{\theta}) = \int_{\hat{\theta}}^{\theta} u(\theta)d\theta \geq \int_{\hat{\theta}}^{\theta} u(\hat{\theta})d\theta = (\theta - \hat{\theta})u(\hat{\theta}).$$

The inequality holds because a non-decreasing $g(\cdot)$ implies that $u(\cdot)$ is also non-decreasing. Substituting the definitions of V and u gives the incentive compatibility constraints. □

C.2 Lemma 2 on the monotonicity of g_n

Proof of Lemma 2. Combining the definition of the (virtual) need for discretion and the definition of g_n gives:

$$\theta(1 - \Delta_n(\theta))\frac{1}{\beta} = \frac{W'(T - g_n(\theta))}{U'(g_n(\theta))}.$$

Substituting the foa-wedge with its formula (7) gives:

$$\frac{1}{1-\rho}\left(\rho\frac{1 - F(\theta)}{f(\theta)} - (\rho - \beta)\theta\right)\frac{1}{\beta} = \frac{W'(T - g_n(\theta))}{U'(g_n(\theta))}.$$

Since W is concave and the utility index U is strictly concave, the ratio $\frac{W'(T-g_n(\theta))}{U'(g_n)}$ is unambiguously increasing in the argument g_n. The left-hand side is non-decreasing in θ if and only if the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is not smaller than $\rho - \beta$. □

C.3 Proposition 2 on the optimality of hybrid rules

Proof of Proposition 2. The proof consists of applying Lemma 8. A valid allocation is non-decreasing. The discretion, on-equilibrium, and off-equilibrium penalties allocation, denoted by g_d^{np}, is continuous since $g_d(\theta_n) = g_n(\theta_n)$ for $\theta_n > \underline{\theta}$. The discretionary allocation g_d is strictly increasing. The government spending $g_n(\theta)$ is defined for $\theta \in (\theta_n, \theta_p)$ because Assumption I holds for $\theta \in (\theta_n, \theta_p)$, by definition of θ_n. Lemma 2 implies that g_n is non-decreasing over (θ_n, θ_p) because the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is not smaller than $\rho - \beta$. The utility index U is strictly increasing so the utility profile from the *discretion, on-equilibrium, and off-equilibrium penalties* allocation $u^*(\theta) = U(g_d^{np}(\theta))$ for $\theta \in \Theta$ and $\underline{t}^* = 0$ satisfies $(u^*, \underline{t}^*) \in \Phi$.

The Lagrange multiplier function is

$$\rho\Lambda^*(\theta) = \begin{cases} \rho & \text{for } \theta \geq \theta_n \\ \rho F(\theta) + (1-\beta)\theta f(\theta) & \text{for } \theta \in (\underline{\theta}, \theta_n) \\ 0 & \text{for } \theta = \underline{\theta}. \end{cases}$$

A valid Lagrange multiplier function is non-decreasing. The lower bound on the elasticity of the density in Assumption sL holding for $\theta \leq \theta_n$ is equivalent to the Lagrange multiplier Λ^* being non-decreasing on $(\underline{\theta}, \theta_n)$. The jumps at $\underline{\theta}$ and θ_n must be non-negative. The jump at $\underline{\theta}$ is non-negative since f is non-negative, $0 < \beta \leq 1$, and $\rho > 0$. The jump at θ_n is non-negative since either $\theta_n < \bar{\theta}$ in which case $g_n(\theta_n) = g_d(\theta_n)$ and the Lagrange multiplier is continuous at θ_n, or $\theta_n = \bar{\theta}$ in which case Assumption I in the definition of θ_n implies that $g_n(\theta_n) \leq g_d(\theta_n)$ and the jump is non-negative. Note also that $1 - \Lambda^*$ is integrable because $1 - F$ is integrable and the expectation exists.

The Gateaux derivative (27), with a Lagrange multiplier function equal to 1 for $\theta \geq \theta_n$, reads

$$\partial\mathcal{L}(u, \underline{t}, h, h_t|\Lambda^*) = \int_{\underline{\theta}}^{\theta_n} \rho \left[\theta\Delta(U^{-1}(u(\theta)), \theta)h(\theta)\right] d\Lambda^*(\theta) \tag{28}$$

$$+ \int_{\underline{\theta}}^{\theta_n} \left[\left(-(1-\beta)\theta + \rho\tfrac{1-F(\theta)}{f(\theta)} + (1-\rho)\theta\Delta(U^{-1}(u(\theta)), \theta) - \rho(1-\Lambda^*(\theta))\right) h(\theta)\right] dF(\theta) \tag{29}$$

$$+ \int_{\theta_n}^{\theta_p} \left[\left(-(1-\beta)\theta + \rho\tfrac{1-F(\theta)}{f(\theta)} + (1-\rho)\theta\Delta(U^{-1}(u(\theta)), \theta)\right) h(\theta)\right] \; dF(\theta) \tag{30}$$

$$+ \int_{\theta_p}^{\bar{\theta}} \left[\left(-(1-\beta)\theta + \rho\tfrac{1-F(\theta)}{f(\theta)} + (1-\rho)\theta\Delta(U^{-1}(u(\theta)), \theta)\right) h(\theta)\right] \; dF(\theta). \tag{31}$$

The last step shows that the conditions in terms of Gateaux derivatives in Lemma 8 are met. The term (28) evaluated at u^* is null irrespectively of the direction of the Gateaux derivative (h, h_t) because $g_d^{np}(\theta) = g_d(\theta)$ for $\theta \leq \theta_n$ implies $\Delta(U^{-1}(u^*(\theta)), \theta) = 0$ for $\theta \leq \theta_n$. The choice of Lagrange multiplier over $(\underline{\theta}, \theta_n)$ is precisely the condition needed for the term (29) to be null irrespectively of the direction (h, h_t) of the Gateaux derivative. The definition of g_n in (7) implies that the term (30) evaluated at u^* is null irrespectively of the direction (h, h_t) of the Gateaux derivative.

Using the definition of the wedge to get $\theta\Delta(U^{-1}(u(\theta)), \theta) = \theta - \frac{\beta W'(T - U^{-1}(u(\theta)))}{U'(U^{-1}(u(\theta)))}$, and the following characterization of u^* above θ_p: $\frac{\beta W'(T - U^{-1}(u(\theta)))}{U'(U^{-1}(u(\theta)))} = \frac{1}{1-\rho}\left(\rho\frac{1-F(\theta_p)}{f(\theta_p)} + (\beta - \rho)\theta_p\right)$, for $\theta \geq \theta_p$, line (31) reads

$$\int_{\theta_p}^{\bar{\theta}} \left[\left(\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta - \rho)(\theta - \theta_p)\right) h(\theta)\right] \; dF(\theta). \tag{32}$$

Integrating (32) by parts gives

$$\int_{\theta_p}^{\bar\theta} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)\ h(\theta_p) \tag{32.1}$$

$$+\int_{\theta_p}^{\bar\theta} \left[\int_{\hat\theta}^{\bar\theta} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)\right] dh(\hat\theta), \tag{32.2}$$

where I used that $\lim_{\hat\theta\to\bar\theta}\int_{\hat\theta}^{\bar\theta}\left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)h(\hat\theta)$ is zero since h is bounded if $\bar\theta<\infty$. If $\bar\theta=\infty$, the result follows from taking the limit of a sequence of environments with compact support as shown below.

The next claim shows that the definition of θ_p is precisely so that the inner integral in (32.2) is negative for $\hat\theta \geq \theta_p$ and null for $\hat\theta=\theta_p$ (which also implies that (32.1) is null).

Claim 1. *Inequality* (10) *is equivalent to*

$$\int_{\hat\theta}^{\bar\theta}\left[\rho\frac{1-F(\theta)}{f(\theta)} + (\beta-\rho)(\theta-\tilde\theta)\right] dF(\theta) \leq \rho\frac{1-F(\tilde\theta)}{f(\tilde\theta)}(1-F(\hat\theta)). \tag{33}$$

Proof of Claim 1. The definition of the wedge implies the following identity:

$$\hat\theta\Delta(g,\hat\theta) = \hat\theta - \frac{\beta W'(T-g)}{U'(g)} = \hat\theta - \tilde\theta + \tilde\theta\Delta(g,\tilde\theta).$$

Substituting this identity for the wedge (evaluated at $g_n(\tilde\theta)$) in inequality (10) gives

$$\int_{\hat\theta}^{\bar\theta}\left[\rho(\tilde\theta-\hat\theta)+\beta(\theta-\tilde\theta)\right] dF(\theta) \leq \left((1-\beta)\tilde\theta - (1-\rho)\tilde\theta\Delta_n(\tilde\theta)\right)(1-F(\hat\theta)).$$

Using the definition of Δ_n, i.e., equation (7), on the right-hand side gives

$$\int_{\hat\theta}^{\bar\theta}\left[\rho(\tilde\theta-\hat\theta)+\beta(\theta-\tilde\theta)\right] dF(\theta) \leq \rho\frac{1-F(\tilde\theta)}{f(\tilde\theta)}(1-F(\hat\theta)). \tag{34}$$

Add and subtract $\rho\theta$ to the integrand on the left-hand side and rearrange terms to get

$$\int_{\hat\theta}^{\bar\theta}\left[\rho(\theta-\hat\theta)+(\beta-\rho)(\theta-\tilde\theta)\right] dF(\theta) \leq \rho\frac{1-F(\tilde\theta)}{f(\tilde\theta)}(1-F(\hat\theta)).$$

Integrating the left-hand side by parts gives inequality (33). □

For $(h,h_t)\in\Phi$, h is non-decreasing and because the inner integral from (32.2) is non-positive, the integral (32.2) is non-positive for $(h,h_t)\in\Phi$. The Gateaux derivative of the Lagrangian at the candidate solution $(u^*,\underline{t}^*)$ is null in the direction $(h,h_t)=(u^*,\underline{t}^*)$ and negative in all directions $(h,h_t)\in\Phi$.

For $\bar\theta=\infty$, consider a sequence of environments, indexed by $m\in\mathbb{N}$, with $\Theta=[\underline\theta,\bar\theta_m]$, $\bar\theta_m<\infty$, and $\lim_{m\to\infty}\bar\theta_m=\infty$. Denote by F_m the truncation of F on $[\underline\theta,\bar\theta_m]$, defined as follows: $F_m(\theta)=\frac{F_\theta}{F(\theta_m)}$ for $\theta\in[\underline\theta,\bar\theta_m]$ and $F_m(\theta)=1$ for $\theta\geq\bar\theta_m$. Note that F_m converges weakly to F.

Also, since F is twice continuously differentiable, f_m is continuous and it converges point-wise to f. For each $m \in \mathbb{N}$, denote the solution of the environment with the truncated distribution F_m by $g_d^{np}(\cdot; m)$ and the threshold at which the cap binds by $\theta_p^{(m)}$.

First, note that $g_d^{np}(\cdot; m)$ converges point-wise to $g_d^{np}(\cdot)$ if $\lim_{m\to\infty} \theta_p^{(m)} = \theta_p$. By assumption $\theta_p > \underline{\theta}$. For $\theta_p < \bar{\theta}$, it is the lowest fiscal need that solves $\beta\ \mathbb{E}[\theta|\theta \geq \theta_p] = \theta_p$, and satisfies (see Lemma 4), $\beta E[\theta|\theta \geq \hat{\theta}] - \rho\hat{\theta} \leq \beta E[\theta|\theta \geq \theta_p] - \rho\theta_p$ for $\hat{\theta} \geq \theta_p$. The threshold $\theta_p^{(m)}$ is characterized analogously. For $\theta_p^{(m)} < \bar{\theta}$, it is the lowest fiscal need that solves $\beta\ \mathbb{E}_m[\theta|\theta \geq \theta_p^{(m)}] = \theta_p^{(m)}$, and, by Lemma 4, $\beta E_m[\theta|\theta \geq \hat{\theta}] - \rho\hat{\theta} \leq \beta E_m[\theta|\theta \geq \theta_p^{(m)}] - \rho\theta_p^{(m)}$ for $\hat{\theta} \geq \theta_p^{(m)}$. For any $\hat{\theta} > \underline{\theta}$, because F_m converges weakly to F, $\lim_{m\to\infty} \mathbb{E}_m\left[\frac{\theta}{\hat{\theta}}|\theta \geq \hat{\theta}\right] = \mathbb{E}\left[\frac{\theta}{\hat{\theta}}|\theta \geq \hat{\theta}\right]$. Because taking the limit preserves weak inequalities, $\lim_{m\to\infty} \theta_p^{(m)} \geq \theta_p$. Since F_m is a right-truncation of F, $\mathbb{E}_m\left[\frac{\theta}{\hat{\theta}}|\theta \geq \hat{\theta}\right] \leq \mathbb{E}\left[\frac{\theta}{\hat{\theta}}|\theta \geq \hat{\theta}\right]$. Hence $\theta_p^{(m)} \leq \theta_p$ for every m. It follows that $\lim_{m\to\infty} \theta_p^{(m)} \leq \theta_p$. Combining the two inequalities, $\lim_{m\to\infty} \theta_p^{(m)} = \theta_p$, and $g_d^{np}(\cdot, m)$ converges pointwise to $g_d^{np}(\cdot)$.

Second, note that since $(g_d^{np}(\cdot, m))_{m\in\mathbb{N}}$ is a uniformly bounded sequence and $(f_m)_{m\in\mathbb{N}}$ is bounded by an integrable density, the dominated convergence theorem implies that the sequence of social welfare (with the incentive compatible $t(\cdot; m)$ substituted in) resulting from the sequence of truncated economies converges to the social welfare of the non-truncated economy. Hence a fiscal rule with a null intercept that implements $g_d^{np}(\cdot)$ is optimal for the non-truncated economy. □

C.4 Lemma 3 and Lemma 5 on the implications of Assumption sL

The proofs of Lemma 3 and Lemma 5 use the following Claim.

Claim 2. *Suppose that Assumption sL holds for $\theta \leq \theta^*$ and there exists $\theta_* \leq \theta^*$ such that $\rho\frac{1-F(\theta_*)}{\theta_* f(\theta_*)} < 1 - \beta$. Then $\rho\frac{1-F(\theta)}{\theta f(\theta)} < 1 - \beta$ for $\theta \in [\theta_*, \theta^*]$.*

Proof of Claim 2 . For any $\theta \leq \theta^*$,

$$\begin{aligned}\tfrac{d}{d\theta}\left(\rho\tfrac{1-F(\theta)}{f(\theta)} - \theta(1-\beta)\right) &= -\rho - \rho\tfrac{1-F(\theta)}{\theta f(\theta)}\tfrac{\theta f'(\theta)}{f(\theta)} - (1-\beta) \\ &\leq \rho\tfrac{1-F(\theta)}{\theta f(\theta)}\tfrac{1-\beta+\rho}{1-\beta} - (1-\beta+\rho) \\ &= \tfrac{1-\beta+\rho}{1-\beta}\left(\rho\tfrac{1-F(\theta)}{\theta f(\theta)} - (1-\beta)\right),\end{aligned}$$

in which the inequality follows from Assumption sL. By assumption, $\rho\frac{1-F(\theta_*)}{\theta_* f(\theta_*)} - 1 - \beta < 0$ for $\theta_* \leq \theta^*$. Given that $\frac{1-\beta+\rho}{1-\beta} \geq 0$, combining the two inequalities implies $\frac{d}{d\theta}\left(\rho\frac{1-F(\theta)}{f(\theta)} - \theta(1-\beta)\right) < 0$ for $\theta \in [\theta_*, \theta^*]$. □

Proof of Lemma 3. Note that $\theta_n \in (\underline{\theta}, \bar{\theta})$ implies that $\rho\frac{1-F(\theta_n)}{\theta_n f(\theta_n)} = 1 - \beta$. The argument is by contradiction. Suppose that there exists $\theta_* \leq \theta_n$ such that $\rho\frac{1-F(\theta_*)}{\theta_* f(\theta_*)} < 1 - \beta$. Claim 2, implies that $\rho\frac{1-F(\theta)}{\theta f(\theta)} < 1 - \beta$ for $\theta \in [\theta_*, \theta_n]$, which contradicts $\rho\frac{1-F(\theta_n)}{\theta_n f(\theta_n)} = 1 - \beta$. □

Proof of Lemma 5. The proof proceeds in three steps. The first step shows that the definition of θ_p implies $\rho\frac{1-F(\theta_p)}{\theta_p f(\theta_p)} \geq 1-\beta$ if $\theta_p \in (\underline{\theta}, \bar{\theta})$. The definition of θ_p implies the following inequality on the conditional tail expectation (see Lemma 4): $\beta E[\theta|\theta \geq \hat{\theta}] - \rho\hat{\theta} \leq \beta E[\theta|\theta \geq \theta_p] - \rho\theta_p$ for $\hat{\theta} \geq \theta_p$. The derivative of the left-hand side with respect to $\hat{\theta}$ reads $-\beta\frac{\hat{\theta} f(\hat{\theta})}{1-F(\hat{\theta})} + \beta\mathbb{E}[\theta|\theta \geq \hat{\theta}]\frac{f(\hat{\theta})}{1-F(\hat{\theta})} - \rho$. The derivative must be negative at $\hat{\theta} = \theta_p$ because the inequality holds with equality at θ_p. Using that $\beta\mathbb{E}[\theta|\theta \geq \theta_p] = \theta_p$ for an interior θ_p gives $-\beta\frac{\theta_p f(\theta_p)}{1-F(\theta_p)} + \theta_p\frac{f(\theta_p)}{1-F(\theta_p)} - \rho \leq 0$, which completes the first step. The second step is Claim 2.

The last step shows that $\rho\frac{1-F(\theta)}{\theta f(\theta)} \geq 1-\beta$ for $\theta \leq \theta_p$ by contradiction. Suppose not, so there exists $\theta_* < \theta_p$ such that $\rho\frac{1-F(\theta_*)}{\theta_* f(\theta_*)} < 1-\beta$. Claim 2 implies that $\rho\frac{1-F(\theta)}{\theta f(\theta)} < 1-\beta$ for $\theta \in [\theta_*, \theta_p]$, which contradicts $\rho\frac{1-F(\theta_p)}{\theta_p f(\theta_p)} \geq 1-\beta$. □

C.5 Lemma 4 on the second requirement in the definition of θ_p

Proof of Lemma 4. A first step consists of rewriting inequality (10) as one of the first-order conditions of the Lagrangian method. The first step is Claim 1 in the proof of Proposition 2 in Appendix C.3. I repeat the claim here for convenience.

Claim 1. *Inequality* (10) *with* $g = g_n$ *is equivalent to*

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho\frac{1-F(\theta)}{f(\theta)} + (\beta-\rho)(\theta-\tilde{\theta})\right] dF(\theta) \leq \rho\frac{1-F(\tilde{\theta})}{f(\tilde{\theta})}(1-F(\hat{\theta})). \tag{33}$$

The proof of Claim 1 is in the proof of Proposition 2 in Appendix C.3. It also shows that inequality (33) and inequality (34) are equivalent.

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho(\tilde{\theta}-\hat{\theta}) + \beta(\theta-\tilde{\theta})\right]\ dF(\theta) \leq \rho\frac{1-F(\tilde{\theta})}{f(\tilde{\theta})}(1-F(\hat{\theta})). \tag{34}$$

The second step uses the equivalence between inequalities (10) and (34) and the assumption that inequality (10) holds with equality for some $\tilde{\theta} < \bar{\theta}$ and $\hat{\theta} = \tilde{\theta}$ to get:

$$\int_{\tilde{\theta}}^{\bar{\theta}} \left[\beta(\theta-\tilde{\theta})\right]\ dF(\theta) = \rho\frac{1-F(\tilde{\theta})}{f(\tilde{\theta})}(1-F(\tilde{\theta})). \tag{35}$$

Multiplying both sides of (34) by $(1-F(\tilde{\theta}))$ and using (35) for the right-hand side gives

$$(1-F(\tilde{\theta}))\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho(\tilde{\theta}-\hat{\theta}) + \beta(\theta-\tilde{\theta})\right] dF(\theta) \leq (1-F(\hat{\theta}))\int_{\tilde{\theta}}^{\bar{\theta}} \left[\beta(\theta-\tilde{\theta})\right] dF(\theta),$$

which simplifies to

$$(1-F(\tilde{\theta}))\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho(\tilde{\theta}-\hat{\theta}) + \beta\theta\right] dF(\theta) \leq (1-F(\hat{\theta}))\int_{\tilde{\theta}}^{\bar{\theta}} \left[\beta\theta\right] dF(\theta).$$

Rearranging terms give:

$$(1-F(\tilde{\theta}))\beta\int_{\hat{\theta}}^{\bar{\theta}} \theta dF(\theta) - (1-F(\tilde{\theta}))(1-F(\hat{\theta}))\rho\hat{\theta} \leq (1-F(\hat{\theta}))\beta\int_{\tilde{\theta}}^{\bar{\theta}} \theta dF(\theta) - (1-F(\hat{\theta}))(1-F(\tilde{\theta}))\rho\tilde{\theta}.$$

Since $\tilde{\theta} \leq \hat{\theta} < \bar{\theta}$, both sides can be divided by $(1-F(\tilde{\theta}))(1-F(\hat{\theta}))$ to give the result. □

C.6 Corollary 1 on the comparative statics of θ_p

Proof of Corollary 1. Consider two economies that are identical except for the degree of present bias of their governments, $\beta_l < \beta_h$, and let $\theta_{p,\beta}$ denote the respective threshold fiscal needs above which the cap binds. To show that $\theta_{p,\beta_l} \leq \theta_{p,\beta_h}$, it suffices to show that if inequality (10) holds for an economy with β_h, it also holds—all else equal—for an economy with β_l. Substituting the following identity $\theta\Delta(g,\theta) = \theta - \hat{\theta} + \hat{\theta}\Delta(g,\hat{\theta})$, which holds for any g, θ, and $\hat{\theta}$, in (10) gives

$$\int_{\hat{\theta}}^{\bar{\theta}} \Big(\theta - \hat{\theta} - \theta(1-\beta)\Big)\,\mathrm{d}F(\theta) \leq (\rho-1)(\hat{\theta} - \tilde{\theta}_p + \tilde{\theta}_p\Delta(g(\tilde{\theta}_p),\tilde{\theta}_p))(1-F(\hat{\theta}))$$

which reduces to

$$\beta E[\theta|\theta \geq \hat{\theta}] - \hat{\theta} \leq (\rho-1)\Big(\hat{\theta} - \tilde{\theta}_p + \tilde{\theta}_p\Delta(g(\tilde{\theta}_p),\tilde{\theta}_p)\Big). \tag{36}$$

For the discretionary allocation, the wedge $\Delta(g(\tilde{\theta}_p),\tilde{\theta}_p))$ is null irrespectively of β. Hence, if inequality (36) is satisfied for β_h, it is also satisfied for $\beta_l < \beta_h$.

For the allocation implemented by on-equilibrium penalties, the foa-wedge is increasing in the degree of present bias. Substituting the foa-wedge (7) in (36) gives

$$\beta(E[\theta|\theta \geq \hat{\theta}] - \tilde{\theta}_p) - \hat{\theta} \leq (\rho-1)(\hat{\theta} - \tilde{\theta}_p) - \tilde{\theta}_p + \rho\frac{1-F(\tilde{\theta}_p)}{f(\tilde{\theta}_p)}. \tag{37}$$

Because $\hat{\theta} \geq \tilde{\theta}_p$, if inequality (37) is satisfied for β_h, it is also satisfied for $\beta_l < \beta_h$. Because θ_{p,β_l} is the infimum of a superset of the set defining θ_{p,β_h}, we have $\theta_{p,\beta_l} \leq \theta_{p,\beta_h}$.

Consider two economies that are identical except for the asymmetry in the cost of penalties, $\rho_l < \rho_h$, and let $\theta_{p,\rho}$ denote the threshold fiscal needs above which the cap binds. I show that the set of $\tilde{\theta}_p$ satisfying inequality (10) for $\hat{\theta} \geq \tilde{\theta}_p$ in the economy with degree of asymmetry ρ_h, of which θ_p is the infimum, is a superset of the set for the economy with degree of asymmetry ρ_l.

Inequality (10), reduces to inequality (36). Because $\hat{\theta} \geq \tilde{\theta}_p$ and the wedge is null for the discretionary allocation, if (36) is satisfied for ρ_l, it is also satisfied for $\rho_h > \rho_l$. Hence, $\theta_{p,\rho_h} \leq \theta_{p,\rho_l}$ for a cap on the discretionary allocation.

For the allocation implemented by the foa-wedge, inequality (36) reduces to inequality (37). Likewise, because $\hat{\theta} \geq \tilde{\theta}_p$ and the inverse hazard rate is positive, if (37) is satisfied for ρ_l, it is also satisfied for $\rho_h > \rho_l$. Hence, $\theta_{p,\rho_h} \leq \theta_{p,\rho_l}$ for a cap on the allocation implemented by the foa-wedge. □

C.7 Proposition 3 on the optimality of a cap

Proof of Proposition 3. The proof consists of applying Lemma 8. Let g_d^p denote the discretion and off-equilibrium penalties allocation, $u^*(\theta) = U(g_d^p(\theta))$ for $\theta \in \Theta$, and $\underline{t}^* = 0$. Since g_d^p is non-decreasing, $(u^*, \underline{t}^*) \in \Phi$. The Lagrange multiplier function is

$$\rho\Lambda^*(\theta) = \begin{cases} \rho & \text{for } \theta \geq \theta_p \\ \rho F(\theta) + (1-\beta)\theta f(\theta) & \text{for } \theta \in (\underline{\theta}, \theta_p) \\ 0 & \text{for } \theta = \underline{\theta}. \end{cases}$$

A valid Lagrange multiplier function is non-decreasing. The lower bound on the elasticity of the density in Assumption sL holding for $\theta \leq \theta_p$ is equivalent to the Lagrange multiplier Λ^* being non-decreasing on $(\underline{\theta}, \theta_p)$. The jumps at $\underline{\theta}$ and θ_p must be non-decreasing. The jump at $\underline{\theta}$ is non-negative since f is non-negative, $0 < \beta \leq 1$, and $\rho > 0$. The jump at θ_p is non-negative, as shown in Lemma 5. Lemma 5 shows that Assumption sL holding for $\theta \leq \theta_p$ and $\beta E[\theta|\theta \geq \hat{\theta}] - \rho\hat{\theta} \leq \beta E[\theta|\theta \geq \theta_p] - \rho\theta_p$ holding for $\hat{\theta} \geq \theta_p$ implies $\rho\frac{1-F(\theta)}{\theta f(\theta)} \geq 1 - \beta$ for $\theta \leq \theta_p$, which is a non-negative jump of Λ^* at θ_p. Note also that $1 - \Lambda^*$ is integrable because $1 - F$ is integrable and the expectation exists.

The rest of the proof checks that the conditions in terms of Gateaux derivatives in Lemma 8 are satisfied. The Gateaux derivative (26) in the direction of h, with Lagrange multiplier function Λ^* reads

$$\begin{aligned} \partial\mathcal{L}(u, \underline{t}, h, h_t|\Lambda^*) = &\int_\Theta \Big[\big(- (1-\beta)\theta(1 - \Delta(U^{-1}(u(\theta)))) + \rho\tfrac{1-F(\theta)}{f(\theta)}\big)h(\theta)\Big]\, dF(\theta) \\ &+ (\beta - \rho)\int_\Theta \big[\theta\Delta(U^{-1}(u(\theta)), \theta)h(\theta)\big]\, dF(\theta) \\ &+ \int_{\underline{\theta}}^{\theta_p} \big[\theta\Delta(U^{-1}(u(\theta)), \theta)h(\theta)\big]\, \rho d\Lambda^*(\theta) - \rho\int_{\underline{\theta}}^{\theta_p} \left[(1 - \Lambda^*(\theta))\ h(\theta)\right] d\theta. \end{aligned}$$

Rewriting the Euler equations characterizing g_d gives

$$\theta\Delta(U^{-1}(u^*(\theta)), \theta) = \begin{cases} \theta - \theta_p & \text{for } \theta > \theta_p \\ 0 & \text{for } \theta \leq \theta_p. \end{cases}$$

After substitution of this expression, the Gateaux derivative evaluated at u^* simplifies to

$$\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) = \int_{\underline{\theta}}^{\theta_p} \Big[\Big(-(1-\beta)\theta + \rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-\Lambda^*(\theta)}{f(\theta)}\Big)\, f(\theta)h(\theta)\Big]\, d\theta \tag{38}$$

$$+ \int_{\theta_p}^{\bar{\theta}} \Big[\Big(-(1-\beta)\theta_p + \rho\tfrac{1-F(\theta)}{f(\theta)} + (\beta - \rho)\,(\theta - \theta_p)\Big)\, f(\theta)h(\theta)\Big]\, d\theta. \tag{39}$$

The Lagrange multiplier Λ^* over $(\underline{\theta}, \theta_p]$ is defined so that the integral (38) is null, for $(h, h_t) \in \Phi$. Suppose that $\bar{\theta} < \infty$ so that for any $(h, h_t) \in \Phi$, h is bounded (the case $\bar{\theta} = \infty$ is addressed below). For h bounded, the following term is null:

$$\lim_{\hat{\theta} \to \bar{\theta}} \int_{\hat{\theta}}^{\bar{\theta}} \Big[\Big(-(1-\beta)\theta_p + \rho\tfrac{1-F(\theta)}{f(\theta)} + (\beta - \rho)(\theta - \theta_p)\Big)\, f(\theta)\Big]\, d\theta\ h(\hat{\theta}) = 0,$$

hence, integrating (39) by parts gives

$$\int_{\theta_p}^{\bar{\theta}} \left[\left(-(1-\beta)\theta_p + \rho \tfrac{1-F(\theta)}{f(\theta)} + (\beta-\rho)(\theta-\theta_p) \right) f(\theta) \right] d\theta \; h(\theta_p) \tag{39.1}$$

$$+ \int_{\theta_p}^{\bar{\theta}} \left[\int_{\hat{\theta}}^{\bar{\theta}} \left[\left(-(1-\beta)\theta_p + \rho \tfrac{1-F(\theta)}{f(\theta)} + (\beta-\rho)(\theta-\theta_p) \right) f(\theta) \right] d\theta \right] dh(\hat{\theta}). \tag{39.2}$$

As the next claim shows, θ_p is defined so that the inner integral in (39.2) is negative for $\hat{\theta} \geq \theta_p$ and null for $\hat{\theta} = \theta_p$ (which also implies that (39.1) is null).

Claim 3. *Inequality* (10) *with* $g = g_d$ *is equivalent to*

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\left(-(1-\beta)\tilde{\theta} + \rho \tfrac{1-F(\theta)}{f(\theta)} + (\beta-\rho)(\theta-\tilde{\theta}) \right) f(\theta) \right] d\theta \leq 0. \tag{40}$$

Proof of Claim 3. After substitution of the wedge evaluated at the discretionary allocation, i.e., $\hat{\theta}\Delta(g_d(\tilde{\theta}), \hat{\theta}) = \hat{\theta} - \tilde{\theta}$, and rearranging, inequality (10) reads as follows:

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho(\tilde{\theta} - \hat{\theta}) + \beta(\theta - \tilde{\theta}) \right] dF(\theta) \leq (1-\beta)\tilde{\theta}(1 - F(\hat{\theta})). \tag{41}$$

Adding and subtracting $\rho(\theta - \tilde{\theta})$ to the integrand on the left-hand side gives

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho(\theta - \hat{\theta}) + (\beta-\rho)(\theta - \tilde{\theta}) \right] dF(\theta) \leq (1-\beta)\tilde{\theta}(1 - F(\hat{\theta})).$$

Integrating the left-hand side by parts gives inequality (40). □

Since $\theta_p > \underline{\theta}$, inequality (10) holds with equality for $\hat{\theta} = \tilde{\theta} = \theta_p$, hence (40) holds with equality for $\hat{\theta} = \tilde{\theta} = \theta$. By definition of $\theta_p > \underline{\theta}$, (39.1) is null and the inner integral of (39.2) is negative for $\hat{\theta} \geq \theta_p$.

For $\bar{\theta} = \infty$, the argument follows the one in the proof of Proposition 2 in Appendix C.3, with g_d^p instead of g_d^{np}.

Consider $(h, h_t) = (u^*, \underline{t}^*)$. The Gateaux derivative $\partial\mathcal{L}(u^*, \underline{t}^*, u^*, \underline{t}^* | \Lambda^*)$ is null since (39.1) and (39.2) are both null. Line (39.2) is null for $h = u^*$ because $dh(\theta) = du^*(\theta) = 0$ for $\theta \geq \theta_p$.

Consider any $(h, h_t) \in \Phi$. The Gateaux derivative $\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t | \Lambda^*)$ is negative since (39.1) is null and (39.2) is negative. Line (39.2) is negative because $dh \geq 0$ since $(h, h_t) \in \Phi$ and (40) is negative. □

C.8 Proposition 4 on the optimality of an exemption

Proof of Proposition 4. The proof consists of applying Lemma 8. Let g_x^{np} denote the exemption, on-equilibrium, and off-equilibrium penalties allocation, $u^*(\theta) = U(g_x^{np}(\theta))$ for $\theta \in \Theta$, and $\underline{t}^* = 0$. The lower bound on the derivative of the inverse hazard rate implies that $g_x^{np}(\theta)$ is non-decreasing, hence $(u^*, \underline{t}^*) \in \Phi$. The Lagrange multiplier function is $\Lambda^*(\theta) = 1$ for $\theta \in \Theta$ is valid since it is non-decreasing and $1 - \Lambda^*$ is integrable.

The definition of g_x^{np} implies the following wedge schedule:

$$(1-\rho)\theta(1-\Delta(U^{-1}(u^*(\theta)),\theta)) = \begin{cases} \rho\frac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)\theta_p & \text{for } \theta \geq \theta_p \\ \rho\frac{1-F(\theta)}{f(\theta)} + (\beta-\rho)\theta & \text{for } \theta_x \leq \theta \leq \theta_p \\ \rho\frac{1-F(\theta_x)}{f(\theta_x)} + (\beta-\rho)\theta_x & \text{for } \theta \leq \theta_x \ . \end{cases}$$

The Gateaux derivative (27), evaluated at the allocation u^* and the constant Lagrange multiplier function Λ^*, reduces to

$$\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) = \int_{\theta_p}^{\bar{\theta}} \left[\left(\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right) h(\theta)\right] dF(\theta) \tag{42}$$

$$+ \int_{\underline{\theta}}^{\theta_x} \left[\left(\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_x)}{f(\theta_x)} + (\beta-\rho)(\theta-\theta_x)\right) h(\theta)\right] dF(\theta) \tag{43}$$

$$+ \left(\left(\underline{\theta} - \tfrac{\rho}{1-\rho}\tfrac{1-F(\theta_x)}{f(\theta_x)} - \tfrac{\beta-\rho}{1-\rho}\theta_x\right) h(\underline{\theta}) - h_t\right)\rho. \tag{44}$$

Integrating (42) and (43) by parts and grouping terms,

$$\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) = -\rho h_t \tag{45}$$

$$+ \int_{\theta_p}^{\bar{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)\ h(\theta_p) \tag{46}$$

$$+ \int_{\theta_p}^{\bar{\theta}} \left[\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)\right] dh(\hat{\theta}) \tag{47}$$

$$+ \left(\int_{\underline{\theta}}^{\theta_x} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_x)}{f(\theta_x)} + (\beta-\rho)(\theta-\theta_x)\right] dF(\theta) + \left(\underline{\theta} - \tfrac{\rho}{1-\rho}\tfrac{1-F(\theta_x)}{f(\theta_x)} - \tfrac{\beta-\rho}{1-\rho}\theta_x\right)\rho\right) h(\theta_x) \tag{48}$$

$$- \int_{\underline{\theta}}^{\theta_x} \left[\int_{\underline{\theta}}^{\hat{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_x)}{f(\theta_x)} + (\beta-\rho)(\theta-\theta_x)\right] dF(\theta) + \left(\underline{\theta} - \tfrac{\rho}{1-\rho}\tfrac{1-F(\theta_x)}{f(\theta_x)} + \tfrac{\beta-\rho}{1-\rho}\theta_x\right)\rho\right] dh(\hat{\theta}), \tag{49}$$

using that $\lim_{\hat{\theta}\to\bar{\theta}} \int_{\hat{\theta}}^{\bar{\theta}} \left[\rho\frac{1-F(\theta)}{f(\theta)} - \rho\frac{1-F(\theta_p)}{f(\theta_p)} + (\beta-\rho)(\theta-\theta_p)\right] dF(\theta)h(\hat{\theta})$ is zero since h is bounded for $\bar{\theta} < \infty$. If $\bar{\theta} = \infty$, the result follows from taking the limit of a sequence of environments with compact support as in the proof of Proposition 3 in Appendix C.7.

The terms (46) and (47) with $\hat{\theta} = \theta_p$ are null and (47) is non-positive for $\hat{\theta} \geq \theta_p$ (see Claim 1 in the proof of Proposition 2 in Appendix C.3).

Claim 4. *For $\hat{\theta} \leq \theta_x$,*

$$\int_{\underline{\theta}}^{\hat{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\theta_x)}{f(\theta_x)} - (\rho-\beta)(\theta-\theta_x)\right]\ dF(\theta) \geq -\rho\underline{\theta} + \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\theta_x)}{f(\theta_x)} - \tfrac{\rho-\beta}{1-\rho}\theta_x\right), \tag{50}$$

with equality at $\hat{\theta} = \theta_x$.

Proof of Claim 4. Inequality (50) holds with equality at $\hat{\theta} = \theta_x$ precisely because of the definition of an interior θ_x in (11).

Since the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is not smaller than $\rho-\beta$ over $\theta \in [\underline{\theta}, \theta_p]$ and $\theta_x \leq \theta_p$, integrating from θ to θ_x gives $\rho\frac{1-F(\theta)}{f(\theta)} - \rho\frac{1-F(\theta_x)}{f(\theta_x)} \leq (\rho - \beta)(\theta - \theta_x)$. The left-hand side of inequality (50) is hence decreasing as a function of $\hat{\theta}$, and the inequality holds with equality at $\hat{\theta} = \theta_x$. It follows that inequalities (50) hold for $\hat{\theta} \leq \theta_x$ as claimed. □

The last step of the proof shows that the conditions in terms of Gateaux derivatives in Lemma 8 are satisfied. First, consider the different terms (45)-(49) of the Gateaux derivative in the direction of the solution $(h, h_t) = (u^*, \underline{t}^*)$. Since $\underline{t}^* = 0$, the term on line (45) is zero. Lines (46) and (48) are zero by Claim 1 and Claim 4. Lines (47) and (49) are zero because u^* is constant above θ_p and below θ_x. Hence, $\partial\mathcal{L}(u^*, \underline{t}^*, u^*, \underline{t}^*|\Lambda^*) = 0$ as desired.

Second, consider the Gateaux derivative (45)-(49) in any direction $(h, h_t) \in \Phi$. Since $h_t \geq 0$, the term on line (45) is negative. Claim 1 and Claim 4 imply that lines (47) and (49) are negative since h is non-decreasing and lines (46) and (48) are zero. Hence, $\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) \leq 0$ for $(h, h_t) \in \Phi$ as desired. □

C.9 Lemma 6 on the second requirement in the definition of θ_x

Proof of Lemma 6. A first step consists of rewriting inequality (11) as the first-order condition of the Lagrangian method. I record this step in the following claim.

Claim 5. *Inequality* (11) *is equivalent to*

$$\int_{\underline{\theta}}^{\hat{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - (\rho - \beta)(\theta - \tilde{\theta})\right] \, dF(\theta) \geq \rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})}F(\hat{\theta}) - \rho\underline{\theta} + \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - \tfrac{\rho-\beta}{1-\rho}\tilde{\theta}\right). \tag{51}$$

Proof of Claim 5. Substituting the identity $\hat{\theta}\Delta(g, \hat{\theta}) = \hat{\theta} - \tilde{\theta} + \tilde{\theta}\Delta(g, \tilde{\theta})$, and the foa-wedge (7) in inequality (11) gives

$$\rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})}F(\hat{\theta}) \leq \int_{\underline{\theta}}^{\hat{\theta}} [\beta(\theta - \tilde{\theta}) - \rho(\hat{\theta} - \tilde{\theta})] dF(\theta) + \rho(\hat{\theta} - \tilde{\theta}) + \rho\tilde{\theta}\Delta(g_n(\tilde{\theta}), \tilde{\theta}).$$

Again, using the foa-wedge (7), $\tilde{\theta}\Delta(g_n(\tilde{\theta}), \tilde{\theta}) = \tilde{\theta} - \frac{1}{1-\rho}\left(\rho\frac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - (\rho - \beta)\tilde{\theta}\right)$, which, upon substitution in the previous inequality gives

$$\rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})}F(\hat{\theta}) \leq \int_{\underline{\theta}}^{\hat{\theta}} [\beta(\theta - \tilde{\theta}) - \rho(\hat{\theta} - \tilde{\theta})] dF(\theta) + \rho(\hat{\theta} - \tilde{\theta}) + \rho\tilde{\theta} - \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - \tfrac{\rho-\beta}{1-\rho}\tilde{\theta}\right).$$

Upon subtracting $\rho\underline{\theta}$ on both sides and rearranging terms, the inequality reads

$$\rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})}F(\hat{\theta}) - \rho\underline{\theta} + \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - \tfrac{\rho-\beta}{1-\rho}\tilde{\theta}\right) \leq \int_{\underline{\theta}}^{\hat{\theta}} [\beta(\theta - \tilde{\theta}) - \rho(\hat{\theta} - \tilde{\theta})] dF(\theta) + \rho\hat{\theta} - \rho\underline{\theta}.$$

Adding and subtracting $\rho\theta$ to the integrand on the right-hand side gives

$$\rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})}F(\hat{\theta}) - \rho\underline{\theta} + \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - \tfrac{\rho-\beta}{1-\rho}\tilde{\theta}\right) \leq \int_{\underline{\theta}}^{\hat{\theta}} \left[\rho(\theta-\hat{\theta}) - (\rho-\beta)(\theta-\tilde{\theta})\right] dF(\theta) + \rho\hat{\theta} - \rho\underline{\theta}.$$

Integration by parts gives the following identity:

$$\rho\int_{\underline{\theta}}^{\hat{\theta}}[\theta-\hat{\theta}]dF(\theta) + \rho\hat{\theta} - \rho\underline{\theta} = \int_{\underline{\theta}}^{\hat{\theta}}\left[\rho\tfrac{1-F(\theta)}{f(\theta)}\right] dF(\theta),$$

which, after substitution in the right-hand side of the previous inequality, gives (51). □

The second step uses the equivalence in Claim 5 and the assumption that inequality (11) holds with equality for some $\tilde{\theta} < \bar{\theta}$ and $\hat{\theta} = \tilde{\theta}$ to get:

$$\int_{\underline{\theta}}^{\tilde{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - (\rho-\beta)(\theta-\tilde{\theta})\right] dF(\theta) = -\rho\underline{\theta} + \rho\left(\tfrac{\rho}{1-\rho}\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - \tfrac{\rho-\beta}{1-\rho}\tilde{\theta}\right). \tag{52}$$

Substituting (52) in the right-hand side of (51) gives

$$\int_{\underline{\theta}}^{\hat{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - (\rho-\beta)(\theta-\tilde{\theta})\right] dF(\theta) \geq \int_{\underline{\theta}}^{\tilde{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - (\rho-\beta)(\theta-\tilde{\theta})\right] dF(\theta).$$

Subtracting the left-hand side on both sides gives

$$0 \geq \int_{\hat{\theta}}^{\tilde{\theta}} \left[\rho\tfrac{1-F(\theta)}{f(\theta)} - \rho\tfrac{1-F(\tilde{\theta})}{f(\tilde{\theta})} - (\rho-\beta)(\theta-\tilde{\theta})\right] dF(\theta),$$

which gives (12) after rearranging terms. □

C.10 Corollary 3 on the comparative statics of θ_x

Proof of Corollary 3. Consider two economies that are identical except for the degree of present-bias of their governments, $\beta_l < \beta_h$, and let $\theta_{x,\beta}$ denote the respective threshold fiscal needs for the exemption.

Substituting the following identity $\theta\Delta(g,\theta) = \theta - \hat{\theta} + \hat{\theta}\Delta(g,\hat{\theta})$, which holds for any g, θ, and $\hat{\theta}$, in (11) gives

$$\int_{\underline{\theta}}^{\hat{\theta}} \left(\hat{\theta} - \beta\theta\right) dF(\theta) \leq \left(\hat{\theta} - \tilde{\theta}_x\right)\left(\rho(1-F(\hat{\theta})) + F(\hat{\theta})\right) + \tilde{\theta}_x\Delta(g_n(\tilde{\theta}_x),\tilde{\theta}_x)\left(\rho(1-F(\hat{\theta})) + F(\hat{\theta})\right), \tag{53}$$

and for $\hat{\theta} = \tilde{\theta}_x = \theta_x \in (\underline{\theta}, \bar{\theta})$, the inequality holds with equality:

$$\int_{\underline{\theta}}^{\tilde{\theta}_x} \left(\tilde{\theta}_x - \beta\theta\right) dF(\theta) = \tilde{\theta}_x\Delta(g_n(\tilde{\theta}_x),\tilde{\theta}_x)\left(\rho(1-F(\hat{\theta})) + F(\hat{\theta})\right).$$

Note that, as expected, there is an exemption only if the wedge is strictly positive. If the wedge is null, the inequality (11) cannot be satisfied for $\hat{\theta} = \tilde{\theta}_x$. If the wedge is strictly positive, then it is the foa-wedge,

$$\Delta_n(\theta) = \frac{1}{1-\rho}\left((1-\beta) - \rho\frac{1-F(\theta)}{\theta f(\theta)}\right).$$

Substituting the foa-wedge in (53), gives

$$\begin{aligned} &\int_{\underline{\theta}}^{\hat{\theta}} \left(\hat{\theta} - \beta\theta\right) dF(\theta) && (54) \\ &\quad \leq \left(\hat{\theta} - \tilde{\theta}_x\right)\left(\rho + (1-\rho)F(\hat{\theta})\right) + \left(\frac{\rho}{1-\rho} + F(\hat{\theta})\right)\left(\tilde{\theta}_x(1-\beta) - \rho\frac{1-F(\tilde{\theta}_x)}{f(\tilde{\theta}_x)}\right). \end{aligned}$$

Because $\hat{\theta} \leq \tilde{\theta}_x$, note that $\int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta) \leq \tilde{\theta}_x F(\hat{\theta})$ which, together with $\beta_l < \beta_h$, implies

$$\beta_h \int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta) - \left(\frac{\rho}{1-\rho} + F(\hat{\theta})\right)\tilde{\theta}_x\beta_h < \beta_l \int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta) - \left(\frac{\rho}{1-\rho} + F(\hat{\theta})\right)\tilde{\theta}_x\beta_l.$$

As a result, if the inequality (54) is satisfied with β_h, then it is also satisfied with $\beta_l < \beta_h$. Hence, the set used to define θ_{x,β_h} is a subset of the set used to define θ_{x,β_l}, and $\theta_{x,\beta_h} \leq \theta_{x,\beta_l}$. □

C.11 Proposition 5 on the optimality of a tight cap

Proof of Proposition 5. The proof consists of applying Lemma 8. Let $u^*(\theta) = U(g_c)$ for $\theta \in \Theta$ and $\underline{t}^* = 0$, hence $(u^*, \underline{t}^*) \in \Phi$. The Lagrange multiplier function $\Lambda^*(\theta) = 1$ for $\theta \in \Theta$ is valid since it is non-decreasing and $1 - \Lambda^*$ is integrable.

The wedge evaluated at g_c is $\theta\Delta(g_c, \theta)) = \theta - \beta \int_\Theta \tilde{\theta} dF(\tilde{\theta})$. The Gateaux derivative (27), evaluated at $(u^*, \underline{t}^*)$ and in the direction (h, h_t), given the constant Lagrange multiplier Λ^*, reads

$$\begin{aligned} \partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) = &\int_\Theta \left[\left(\rho\tfrac{1-F(\theta)}{f(\theta)} - \beta(1-\rho)\int_\Theta \tilde{\theta}\, dF(\tilde{\theta}) + (\beta - \rho)\theta\right) h(\theta)\right] dF(\theta) && (55) \\ &+ \left(\left(\underline{\theta} - \beta\int_\Theta \tilde{\theta}\, dF(\tilde{\theta})\right) h(\underline{\theta}) - h_t\right)\rho. && (56) \end{aligned}$$

Integrating the inverse hazard rate by parts gives the following: $\int_{\underline{\theta}}^{\hat{\theta}} \frac{1-F(\theta)}{f(\theta)}\, dF(\theta) = \hat{\theta}(1 - F(\hat{\theta})) - \underline{\theta} + \int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta)$. Since the expectation of θ exists, $\hat{\theta}(1 - F(\hat{\theta})) \to 0$ as $\hat{\theta} \to \bar{\theta}$. Substitute $\underline{\theta}$ in (56),

$$\left(\underline{\theta} - \beta\int_\Theta \tilde{\theta}\, dF(\tilde{\theta})\right) h(\underline{\theta}) - h_t = -\left(\int_\Theta \left[\tfrac{1-F(\theta)}{f(\theta)} - (1-\beta)\theta\right]\, dF(\theta)\right) h(\underline{\theta}) - h_t.$$

First, I show that the Gateaux derivative in the direction of the candidate solution is null when evaluated at the candidate solution. Consider the direction $h(\theta) = u^*(\theta)$ for $\theta \in \Theta$ and $h_t = \underline{t}^* = 0$. Since h is constant, $h(\theta) = h(\underline{\theta})$ for $\theta \in \Theta$ and $h(\underline{\theta})$ can be taken out of the expectation in (55). Since $-\beta(1-\rho)+\beta-\rho=-\rho(1-\beta)$, it follows that $\partial\mathcal{L}(u^*,\underline{t}^*,u^*,\underline{t}^*|\Lambda^*)=0$ as claimed.

It remains to show that the Gateaux derivative evaluated at $(u^*,\underline{t}^*)$ in any direction $(h,h_t)\in\Phi$ is non-positive. Integrating (55) by parts gives

$$\int_\Theta\left[\rho\tfrac{1-F(\theta)}{f(\theta)}-\beta(1-\rho)\int_\Theta\tilde{\theta}dF(\tilde{\theta})+(\beta-\rho)\theta\right]dF(\theta)h(\underline{\theta}) \tag{55.1}$$

$$+\int_\Theta\left[\int_{\hat{\theta}}^{\bar{\theta}}\left[\rho\tfrac{1-F(\theta)}{f(\theta)}-\beta(1-\rho)\int_\Theta\tilde{\theta}\;dF(\tilde{\theta})+(\beta-\rho)\theta\right]dF(\theta)\right]dh(\hat{\theta}). \tag{55.2}$$

Claim 6. *Suppose Assumption I holds for $\underline{\theta}$ and the derivative of $\rho\frac{1-F(\theta)}{f(\theta)}$ is smaller than $\rho-\beta$ for $\theta\in\Theta$, then*

$$\int_{\hat{\theta}}^{\bar{\theta}}\left[\rho\tfrac{1-F(\theta)}{f(\theta)}-\beta(1-\rho)\int_\Theta\tilde{\theta}\;dF(\tilde{\theta})+(\beta-\rho)\theta\right]dF(\theta)\leq 0 \tag{55.2.i}$$

for all $\hat{\theta}\in\Theta$ and $\int_\Theta\left[\frac{1-F(\theta)}{f(\theta)}-(1-\beta)\theta\right]\,dF(\theta)\leq 0$.

Proof of Claim 6. The proof consists of first showing that the inequality (55.2.i) holds for $\hat{\theta}=\underline{\theta}$. The condition on the slope of the inverse hazard rate implies $\rho\frac{1-F(\theta)}{f(\theta)}-(\rho-\beta)\theta\leq\rho\frac{1-F(\underline{\theta})}{f(\underline{\theta})}-(\rho-\beta)\underline{\theta}$ for $\theta\in\Theta$. Assumption I for $\underline{\theta}$ implies $\rho\frac{1-F(\underline{\theta})}{f(\underline{\theta})}\leq(1-\beta)\underline{\theta}$. Combining the two inequalities gives $\rho\frac{1-F(\theta)}{f(\theta)}-(\rho-\beta)\theta\leq(1-\rho)\underline{\theta}$. Taking expectations on both sides gives

$$\rho\int_\Theta\tfrac{1-F(\theta)}{f(\theta)}dF(\theta)-(\rho-\beta)\int_\Theta\theta dF(\theta)\leq(1-\rho)\underline{\theta}. \tag{57}$$

Integrating the inverse hazard rate by parts to substitute $\underline{\theta}$ in (57) gives

$$\rho\int_\Theta\tfrac{1-F(\theta)}{f(\theta)}dF(\theta)-(\rho-\beta)\int_\Theta\theta dF(\theta)\leq(1-\rho)\left(\int_\Theta\theta dF(\theta)-\int_\Theta\tfrac{1-F(\theta)}{f(\theta)}\;dF(\theta)\right),$$

which simplifies to $\int_\Theta\frac{1-F(\theta)}{f(\theta)}dF(\theta)\leq(1-\beta)\int_\Theta\theta dF(\theta)$. Note that for $\hat{\theta}=\underline{\theta}$, inequality (55.2.i) reduces to $\int_\Theta\left[\frac{1-F(\theta)}{f(\theta)}-(1-\beta)\theta\right]dF(\theta)$.

Second, I show that given that inequality (55.2.i) holds for $\hat{\theta}=\underline{\theta}$, then the condition on the derivative of the inverse hazard rate implies that inequality (55.2.i) holds for $\hat{\theta}\in\Theta$. Rewrite inequality (55.2.i) as follows: $E\left[\rho\frac{1-F(\theta)}{f(\theta)}-(\rho-\beta)\theta\;\middle|\;\theta\geq\hat{\theta}\right]\leq\beta(1-\rho)\int_\Theta\theta\,dF(\theta)$. The condition on the slope of the inverse hazard rate implies that $\rho\frac{1-F(\theta)}{f(\theta)}-(\rho-\beta)\theta$ is decreasing. It follows that the conditional expectation $E\left[\rho\frac{1-F(\theta)}{f(\theta)}-(\rho-\beta)\theta\;\middle|\;\theta\geq\hat{\theta}\right]$ is a decreasing function of $\hat{\theta}$. Since the inequality holds for $\hat{\theta}=\underline{\theta}$, it follows that it holds for $\hat{\theta}\in\Theta$. □

For any $(h,h_t)\in\Phi$, h is non-decreasing so inequality (55.2.i) implies that line (55.2) is non-positive. Claim 6 also shows that line (55.1) is non-positive. □

C.12 Proposition 6 on the optimality of a tight cap

Proof of Proposition 6. The proof is identical to the proof of Proposition 5 up to Claim 6. The rest of the proof uses implications from the definition of a high degree of present bias to imply the same conclusion as the one in Claim 6. Note that

$$\int_{\hat{\theta}}^{\bar{\theta}} \left[\rho\frac{1-F(\theta)}{f(\theta)} - \beta(1-\rho)\int_{\Theta}\theta\, dF(\theta) - (\rho-\beta)\theta\right] dF(\theta) \leq -\beta(1-\rho)\int_{\Theta}\theta\, dF(\theta) \leq 0,$$

where the first inequality uses Assumption H, $\rho\frac{1-F(\theta)}{f(\theta)} \leq (\rho-\beta)\theta$, for $\theta \in \Theta$, and the second inequality follows from $\rho \leq 1$. A high degree of present bias also implies $\frac{1-F(\theta)}{\theta f(\theta)} \leq 1-\beta$ for $\theta \in \Theta$ because $\frac{\rho-\beta}{\rho} \leq 1-\beta$ for $0 < \rho \leq 1$. Hence $\int_{\hat{\theta}}^{\bar{\theta}} \left[\frac{1-F(\theta)}{f(\theta)} - (1-\beta)\theta\right] dF(\theta) \leq 0$ for $\hat{\theta} \in \Theta$. For $(h, h_t) \in \Phi$, h is non-decreasing and $h_t \geq 0$. Hence $\partial\mathcal{L}(u^*, \underline{t}^*, h, h_t|\Lambda^*) \leq 0$ for $(h, h_t) \in \Phi$. $\square$

C.13 Proposition 7 on the measurement of fiscal needs

Proof of Proposition 7. To verify that the guess is a solution, it suffices to check that the value function is a fixed point of the Bellman equation and that the policy functions solve the optimization problem with the value function.

First, to check that the policy functions solve the optimization problem with the continuation value (14), take the first-order condition to get $\theta(w+b')^{-1} = \beta\delta RE[a(\theta')|\theta]\left(T - Rb' + \bar{b}\right)^{-1}$. Substituting the policy function (15) for spending on the left-hand side and the policy function for borrowing on the right-hand side gives $\theta\left((1-s(\theta))\ (w+\bar{b})\right)^{-1} = \beta\delta RE[a(\theta')|\theta](s(\theta)(w+\bar{b})R)^{-1}$. The terms involving the effective wealth $w+\bar{b}$ cancel out. Rearranging terms to isolate the savings rate gives $s(\theta) = \frac{\beta\delta E[a(\theta')|\theta]}{\theta+\beta\delta E[a(\theta')|\theta]}$.

Having verified that the policy functions (15) with the savings rate (16) solve the maximization problem in the Bellman equation given the value function (14), it only remains to verify that the value function solves the Bellman equation.

To check that the value function is a fixed point of the Bellman equation, substitute the value function (14) on both sides of the Bellman equation (13) and, after substitution of the policy functions, the terms involving the effective wealth $w + \bar{b}$ cancel out (because $a(\theta) = \theta + \beta\delta E\left[a(\theta')|\theta\right]$) to give $\nu(\theta) = \theta\ln(1-s(\theta)) + \beta\delta E\left[a(\theta')|\theta\right]\ln\left((s(\theta)R)\right) + \beta\delta E[\nu(\theta')|\theta]$. The term $\nu(\theta)$ solves the following recursion:

$$\nu(\theta) = \theta\ln\left(\frac{\theta}{\theta+\beta\delta E\left[a(\theta')|\theta\right]}\right) + \beta\delta E\left[a(\theta')|\theta\right]\ln\left(\frac{\beta\delta E\left[a(\theta')|\theta\right]}{\theta+\beta\delta E\left[a(\theta')|\theta\right]}\right) + \beta\delta\ln\left(R\right) + \beta\delta E[\nu(\theta')].$$

I guess that the conditional expectation is affine in θ (and verify the guess below), $E\left[a(\theta')|\theta\right] = a_0 + a_1\theta$. Then, $a(\theta) = \theta(1+\beta\delta a_1)+\beta\delta a_0$. To verify the guess that the conditional expectation of $a(\cdot)$ is affine, substituting in the expectation, $E\left[a(\theta')|\theta\right] = \theta_e(1+\beta\delta a_1)+\beta\delta a_0+\varphi(1+\beta\delta a_1)\theta$, which confirms that the conditional expectation is affine and the parameters are $a_1 = \varphi/(1-\varphi\beta\delta)$ and

$a_0 = 1$, where the last equality follows from the normalization. Substituting the linear conditional expectation simplifies the formula for the savings rate $s(\cdot)$ which gives (16).

This confirms that the value function (14) with $a(\theta)$ solves the Bellman equation. □

C.14 Assumption H implies that the cap is tight

This section contains the formal statement and the proof of an observation made in Section 3.3.

Claim 7. *Assumption H implies that the tight cap allocation lies below the discretionary allocation; that is, $g_c \leq g_d(\theta)$ for $\theta \in \Theta$.*

Proof of Claim 7. Integrate the inequality in Assumption H to get

$$\int_{\underline{\theta}}^{\bar{\theta}} (1 - F(\theta))d\theta \leq (1 - \beta) \int_{\underline{\theta}}^{\bar{\theta}} \theta f(\theta) d\theta.$$

Integrating $1 - F$ by parts gives

$$\int_{\underline{\theta}}^{\bar{\theta}} \theta f(\theta) d\theta - \underline{\theta} + \lim_{\theta \to \bar{\theta}} \theta(1 - F(\theta)) \leq (1 - \beta) \int_{\underline{\theta}}^{\bar{\theta}} \theta f(\theta) d\theta.$$

Since the expectation of θ is finite and the distribution of θ has a density, $\lim_{\theta \to \bar{\theta}} \theta(1 - F(\theta)) = 0$. The inequality reduces to $\beta\ \mathbb{E}[\theta] \leq \underline{\theta}$, which is equivalent to $g_c \leq g_d(\underline{\theta})$. The discretionary allocation is non-decreasing and hence $g_c \leq g_d(\underline{\theta}) \leq g_d(\theta)$. □

D Sensitivity analysis of the measurement

In this section, I study the sensitivity of the measurement to the elasticity of intertemporal substitution, the present bias of the governments, the persistence of shocks to fiscal needs, the borrowing limit, and the location parameter for the residuals.

Elasticity of Intertemporal Substitution. To elicit the role of the Elasticity of Intertemporal Substitution (EIS) in the measurement, I relax the assumption of a unit EIS and, to preserve analytical tractability, I assume that the interest rate is constant. Let the utility index be CRRA with EIS $1/\eta$.

Proposition 8. *The value function $V(w, \theta) = a(\theta) \frac{\left(w + \bar{b}\right)^{1-\eta}}{1-\eta}$ satisfies the Bellman equation* (13) *for $a(\theta) = (\theta^{1/\eta} + (\beta\delta R^{1-\eta})^{1/\eta})^{\eta}$. The policy function* (15) *solves the government problem with the value function and the savings rate is only a function of the intertemporal elasticity of substitution, the interest rate, the discount factor, and the realized fiscal needs as follows:*

$$s(\theta) = \frac{(\beta\delta)^{\frac{1}{\eta}} R^{\frac{1-\eta}{\eta}}}{\theta^{\frac{1}{\eta}} + (\beta\delta)^{\frac{1}{\eta}} R^{\frac{1-\eta}{\eta}}}. \tag{58}$$

Proof of Proposition 8. Note that $E[a(\theta)] = E[(\theta^{1/\eta} + (\beta\delta R^{1-\eta})^{1/\eta})^{\eta}] = 1$, where the first inequality uses the guess for $a(\theta)$ and the second equality follows from the normalization of the public spending needs.

First, to check that the policy functions solve the optimization problem with the continuation value V, take the first-order condition to get $\theta(w+b')^{-\eta} = \beta\delta R E[a(\theta)]\left(T - Rb' + \bar{b}\right)^{-\eta}$. Substituting the policy function (15), $\theta\left((1-s(\theta))\ (w+\bar{b})\right)^{-\eta} = \beta\delta R E[a(\theta)](s(\theta)(w+\bar{b})R)^{-\eta}$. Rearranging terms to isolate the savings rate gives (58) because $E[a(\theta)] = 1$.

To check that the value function is a fixed point of the Bellman equation, substitute the value function on both sides of the Bellman equation (13) to get, after substituting the policy function,

$$a(\theta)\frac{\left((w+\bar{b})\right)^{1-\eta}}{1-\eta} = \theta\frac{((1-s(\theta))\ (w+\bar{b}))^{1-\eta}}{1-\eta} + \beta\delta E\left[a(\theta')\right]\frac{\left((s(\theta)\ (w+\bar{b})R - \bar{b} + \bar{b})\right)^{1-\eta}}{1-\eta}.$$

Substituting the saving rate gives and $a(\theta) = \left(\theta^{\frac{1}{\eta}} + R^{\frac{1-\eta}{\eta}}(\beta\delta)^{\frac{1}{\eta}}E\left[a(\theta')\right]^{\frac{1}{\eta}}\right)^{\eta}$, as claimed. □

Rewriting (58) to express the fiscal needs as a function of the saving rate gives $\theta(s) = \beta\delta\left(\frac{1-s}{s}\right)^{\eta}R^{1-\eta}$. The empirical framework identifies the product of the tail parameter and the inverse of the elasticity of intertemporal substitution from variations in the savings rate across countries and time,

$$\ln(1 - F(\theta(s_{it}))) = -\gamma\eta\ln\left(\frac{1-s_{it}}{s_{it}}\right) - \gamma(1-\eta)\ln(R) - \gamma\ln(\beta\delta).$$

Hence, the measured tail parameter γ is inversely related to η. The intuition is the following. A larger elasticity of intertemporal substitution—a lower η—implies that a smaller variation in fiscal needs rationalizes the observed saving rates, and hence the tail is thinner (i.e., a larger γ).

The degrees of present bias of the governments. The calibration of the degree of present bias is such that, each country's average spending in the model matches the average spending in the data. Because the model does not separately identify δ, β, and the average fiscal needs, I set the discount factor of the society $\delta = 0.96$ and normalize the level of the average fiscal needs such that the country with the highest estimated β, namely Luxembourg, has a $\beta = 1$. This calibration attributes heterogeneity in average government spending across countries to heterogeneity in the degrees of present bias of the governments. The calibrated discount factors β are, 0.73 for Greece, 0.77 for Italy, 0.79 for Portugal, 0.80 for Hungary, 0.81 for Romania, 0.81 for Malta, 0.81 for Cyprus, 0.82 for Spain, 0.82 for Slovakia, 0.83 for Belgium, 0.83 for Ireland, 0.83 for France, 0.83 for Poland, 0.85 for Croatia, 0.85 for Austria, 0.86 for Slovenia, 0.87 for Germany, 0.87 for Czechia, 0.87 for Lithuania, 0.88 for the Netherlands, 0.88 for Latvia, 0.91

Figure 10: Tail empirical distribution by country

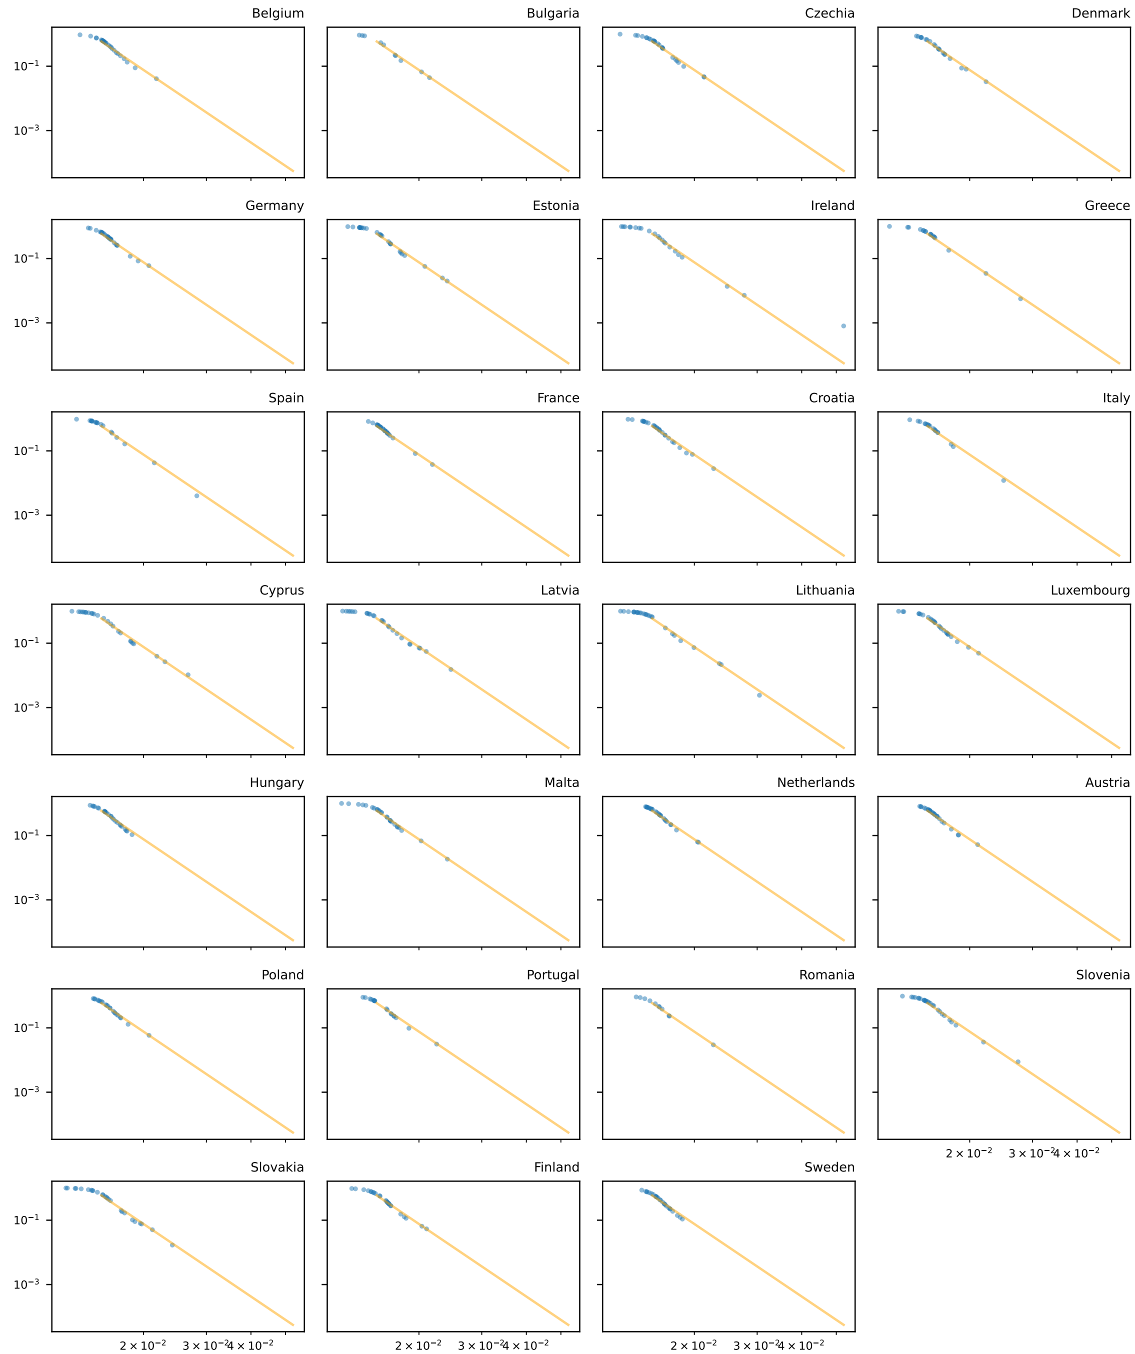

Notes: This figure decomposes the tail empirical distribution depicted in Figure 8a by country. The orange lines are the same as the orange line in Figure 8b.

for Bulgaria, 0.92 for Finland, 0.93 for Sweden, 0.97 for Denmark, 0.97 for Estonia, and 1 for Luxembourg.

To elicit the sensitivity of the measurement to heterogeneity in the degrees of present bias, consider an alternative calibration in which the governments have the same degree of present bias (set at the average of the estimates above). This alternative calibration attributes heterogeneity in average government spending across countries to heterogeneity in the average fiscal needs of the countries. Figures 11a and 11b show that the the tail of the distribution of shocks to fiscal needs also display evidence of a heavy tail with thickness $\hat{\gamma} = 7.59$ and a standard error of 0.56.

Figure 11: Sensitivity analysis of the tail empirical distribution

(a) All measured shocks to fiscal needs ($\beta_i = \beta_j$)

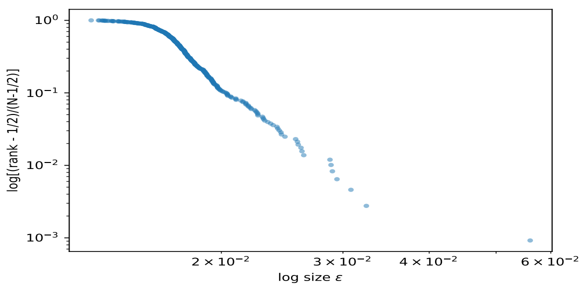

(b) Top two-thirds of shocks to fiscal needs ($\beta_i = \beta_j$)

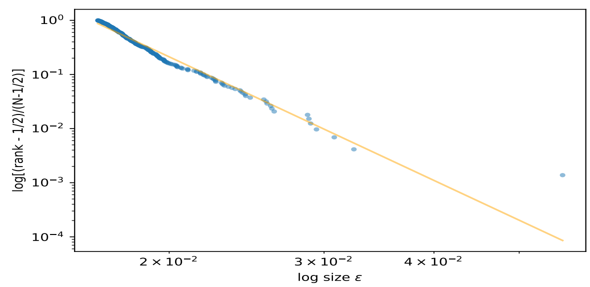

(c) All measured fiscal needs ($\varphi = 0$)

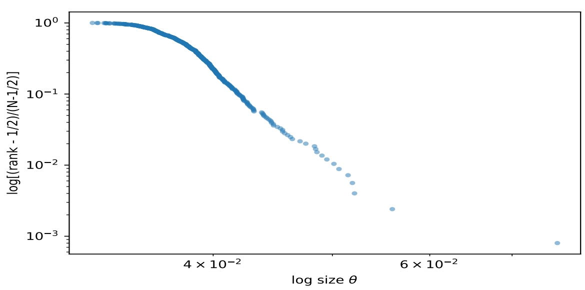

(d) Top two-thirds of fiscal needs ($\varphi = 0$)

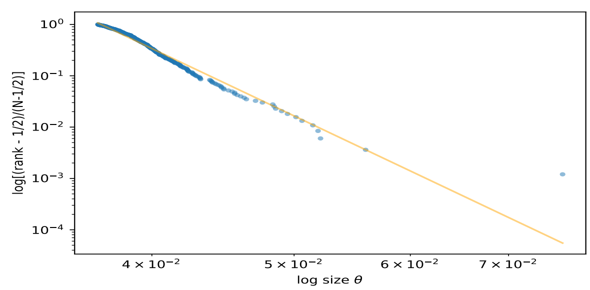

(e) All measured shocks to fiscal needs ($R_i \neq R_j$)

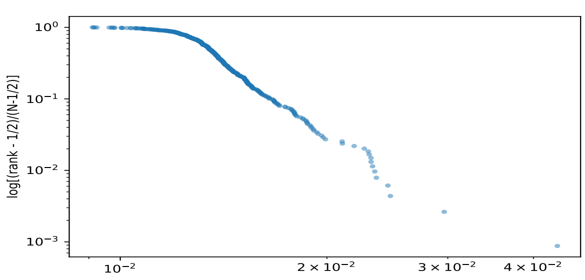

(f) Top two-thirds of shocks to fiscal needs ($R_i \neq R_j$)

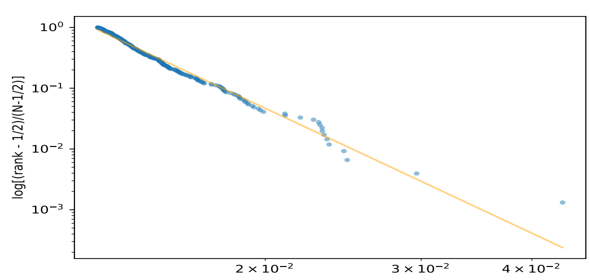

Persistence of shocks. The analysis in the main text jointly infers the fiscal needs and the associated persistence of the fiscal needs. Corollary 4 shows that persistence in the shocks to the fiscal needs amplifies the variation in the measured fiscal needs. To evaluate the sensitivity of the results to this amplification mechanism, Figures 11c and 11d depict the tail empirical distribution of fiscal needs measured with $\varphi = 0$. In line with Corollary 4, not modeling the persistence in the shock biases the measurement of the tail thickness downward, $\hat{\gamma} = 13.54$ and a standard error of 0.94.

Borrowing limit. The analysis in the main text computes each government's borrowing limit based on the average government revenues for each country and on the average interest rate for the EU. Because small differences in interest rates can lead to large differences in the discounted value of future cash flows, Figures 11e and 11f depict the tail empirical distribution based on a measure of the borrowing limit that uses each countries average interest rate instead of the EU average. The estimate is $\hat{\gamma} = 6.83$, with a standard error of 0.50.

Figure 12: Sensitivity of the Hill estimates to the location of the distribution

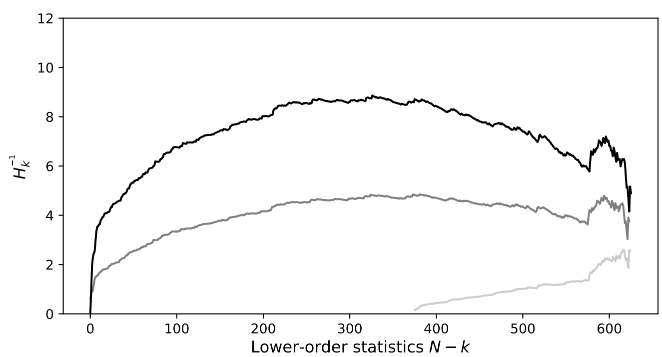

Notes: The black line depicts the same Hill plot as in Figure 9. The light grey line depicts the Hill plot for the residuals centered around 0. The dark grey line depicts the Hill plot for the residuals with location shifted by half of the shift of location for the black line.

Location shifter of the residuals. In the main text, because the Hill estimator uses the log of the shocks, the residuals are minimally shifted to have all estimates of the shocks be positive. Although this does not affect the thickness of the tail of a distribution in theory, it does affect the Hill estimator. Figure 12 shows that the analysis in the main text gives a conservative estimate of the tail thickness.