

Reconciling Fiscal Ceilings with Macro Stabilization*

Regis Barnichon^(a) and *Geert Mesters*^(b)

^(a) Federal Reserve Bank of San Francisco and CEPR

^(b) Universitat Pompeu Fabra and BSE

May 4, 2022

Abstract

Fiscal rules are widely used to constrain policy decisions and promote fiscal discipline, but the design of flexible yet effective rules has proved a formidable task. In practice, fiscal rules take the form of fiscal ceilings —hard thresholds on fiscal variables— which have the benefit of simplicity but are rigid and frequently violated. In this paper, we show that there exists a class of fiscal rules —fiscal-macro targeting rules— that can simultaneously flexibilize fiscal ceilings —leave more room for macro stabilization— and increase overall fiscal discipline. Fiscal-macro targeting rules nest fiscal ceilings as a special case and offer the same benefits: they are simple, transparent, easy to monitor and can be set without reference to a specific model. We illustrate the workings of fiscal-macro targeting in the context of the EU Stability and Growth Pact.

*JEL classification:*E32, E62.

Keywords: fiscal rules, stability and growth pact, impulse responses, forecasting.

*We thank Roel Beetsma, Davide Debortoli, Luca Fornaro, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, Emi Nakamura, Adrien Schmidt, Adam Shapiro and seminar participants for helpful comments, and Lily Seitelman for superb research assistance. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Mesters acknowledge support from the Spanish Ministry of Economy and Competitiveness through the Ramon y Cajal fellowship (RYC2019-028287-I), the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), and the Netherlands Organization for Scientific Research (NWO) through the VENI research grant (016.Veni.195.036).

1 Introduction

If fiscal policy makers have a bias towards overusing debt financing, how can society restrain such tendency?¹ Today, fiscal rules are widely used to constrain fiscal policy discretion and promote fiscal discipline. More than 90 countries have implemented fiscal rules, either at the national or supranational level.

In practice, fiscal rules take the form of “fiscal ceilings”, i.e., hard thresholds on publicly observable fiscal variables, such as the size of the public deficit or the growth rate of public expenditures. Fiscal ceilings are attractive for a number of reasons: they are simple, transparent, easy to monitor and they can be set without any reference to a specific economic model. Unfortunately, in practice fiscal ceilings have had limited success at restraining debt financing as rule violations occur frequently (e.g., Eyraud et al., 2018).

The challenges facing fiscal ceilings can be traced back to the interaction of two issues: (i) the excessive rigidity of fiscal ceilings, and (ii) the limited enforceability of fiscal rules. First, hard thresholds on fiscal variables are inherently rigid, making them make them very costly at times, and prompting (ex-post) little buy-in from policy makers (e.g., Lledó et al., 2017). Second, fiscal rules suffer from a limited enforcement problem. Since fiscal policy is ultimately at the discretion of elected officials, imposing constraints on fiscal policy is difficult in a democratic society. As a result, sanctions for rule violation are limited in scope. The combination of limited sanctions and excessive rigidity can lead to frequent ceiling violations: in recessions, policy makers may prefer to break the fiscal ceiling and incur the sanction than to incur the macro costs of respecting the rule.

The problem of excessive rigidity of fiscal ceilings is well-known, and the common route to flexibilize fiscal ceilings has been to introduce addendum to fiscal ceilings such as “escape clauses” or “cyclical adjustments”.² Unfortunately, without a formal framework, such cyclical adjustments are necessarily ad-hoc and subject to interpretation and political interference. Ultimately, this can lead to lower credibility, lower compliance and lower fiscal discipline, as has been the case with the unsuccessful reforms of the EU Stability and Growth Pact (e.g., Eyraud et al., 2018; Larch and Santacroce, 2020).

In this work, we show that it is possible to simultaneously flexibilize fiscal ceilings — leave more room for macro stabilization— *and* increase overall fiscal discipline. Specifically, there exists a class of fiscal rules —fiscal-macro targeting rules— that nests fiscal ceilings as a special case but can be more flexible, have higher compliance, all the while remaining

¹For discussion on the many sources of deficit bias —time inconsistency, political cycles, bureaucratic behavior, among others—, see e.g. Drazen (2004).

²For example, the European Union (EU) Stability and Growth Pact (SGP), initially based on two fiscal ceilings (a 3% deficit ceiling and a 60% debt ceiling), has undergone a number of reforms including the introduction of escape clauses to flexibilize the SGP.

simple, transparent and not model-dependent.³

The key idea is to incorporate the policy makers’ own objectives —macro stabilization— into the problem. Specifically, instead of providing policy makers with a set of fiscal ceilings, the idea is to give policy makers both fiscal ceiling objectives *and* macro stabilization objectives. A fiscal-macro targeting rule then consists in providing policy makers with a loss function to be minimized, where the loss function combines both the fiscal and the macro objectives.⁴

By taking into account the macro stabilization objective, the fiscal-macro rule is inherently more flexible than a rule based on fiscal ceilings, and this naturally leads to better macro stabilization outcomes. Importantly however, this higher flexibility can *also* lead to greater fiscal discipline, that is to lower deviations from the fiscal ceilings on average. Intuitively, by incorporating the policy makers’ own objective, a fiscal-macro targeting rule increases policy makers’ incentive to comply with the rule, i.e., it reduces the instances of non-compliance compared to a fiscal ceiling rule. A well-designed fiscal-macro rule can then Pareto improve upon a fiscal ceilings rule by trading the rare but costly instances of non-compliance in exchange for frequent but benign deviations from the fiscal ceiling. Stated differently, instead of trying to force compliance via sanctioning —something not possible in practice—, a fiscal-macro targeting rule enhances compliance by cooperation.

Fiscal-macro targeting rules preserve the benefits of fiscal ceilings: transparency, simplicity, ease of monitoring and non-reliance on a specific model. In particular, compliance can be assessed without needing to agree on a specific model. Indeed, assessing compliance amounts to verifying whether the policy maker is minimizing her assigned loss function, which can be done by means of an objective statistical test as shown in Barnichon and Mesters (2022). Specifically, to assess compliance one only need to verify the orthogonality of two sufficient statistics: (i) the impulse responses of the policy objectives to fiscal policy shocks and (ii) the conditional forecasts for the policy objectives. The two statistics are already familiar to policy makers and are in fact routinely computed as part of the decision process. To assess compliance, one would only require an independent agency in charge of constructing the forecasts (a requirement already in place or called for e.g., Bénassy-Quéré et al. (2018)) and conducting the test.

In practice, many countries subject to fiscal ceilings are de facto already following fiscal-

³Although little emphasized in the academic literature, the ability to not rely on a specific underlying economic model is highly valued by policy makers, who face a very complex reality with many remaining unknowns and are wary of model mis-specification (e.g., Blanchard, Leandro and Zettelmeyer, 2020).

⁴The loss function can be any function that is strictly increasing in the deviations from the macro and fiscal variables and targets/ceilings. Assigning policy makers with a loss function to be minimized is sometimes called a “targeting rule”, going back to the work of Rogoff (1985); Walsh (2017); Svensson (1997) in the context of monetary policy. In fact, there is a parallel between our proposal to replace fiscal ceilings with a fiscal-macro targeting rule, and the way central banks switched from monetary ceilings (e.g., a 4.5% ceiling on money growth) to forecast targeting.

macro rules, though the underlying loss function is never explicitly defined. Indeed, while countries often deviate from hard fiscal ceilings such as the SGP 3% deficit ceiling, policy makers do try to stay somewhat “close” to these ceilings (e.g., Eyraud and Wu, 2015), though the notion of “being close” is never clearly defined. A fiscal-macro rule allows to formalize this notion of “closeness”: by making the underlying loss function explicit, it allows to define what constitutes an appropriate deviation from the fiscal ceiling, and it provides a means to rule ex-ante on the appropriate balance between fiscal discipline objectives and macro stabilization objectives.

To illustrate how fiscal-macro rules could be used to enforce fiscal discipline in practice, we consider a fictitious fiscal-macro rule for EU countries, where the loss function includes two objectives: stabilizing GDP growth at potential and keeping the budget deficit below 3%, and we evaluate whether fiscal policy decisions over 1998-2020 satisfy that fictitious fiscal-macro rule. For this purpose, we constructed a new database containing the economic forecasts provided by each Union member to the EU commission, as based on the records of the Stability and Growth Pact (SGP). Using these forecasts and impulse response estimates from Guajardo, Leigh and Pescatori (2014) that capture the effects from fiscal austerity packages, we show how one would test rule compliance for France and Germany. In addition, we use our fiscal-macro targeting rule to quantify the fiscal discipline of a given country, that is to quantify how much weight a given country (say France) is putting on the 3% deficit ceiling versus another country (say Germany).

We find that, compared to Germany, France’s fiscal policy puts more weight on a GDP stabilization objective and much less weight on the 3% deficit ceiling. In other words, after controlling for the economic outlook, France makes less effort than Germany to respect the 3% deficit rule. Looking across all EU countries, we proceed similarly to describe the different EU members in terms of the weight they place on the 3% deficit ceiling. We find substantial variations across EU countries: holding the economic outlook fixed, southern EU countries but also France and Belgium put the least weight on the fiscal objectives, i.e., make the least effort to satisfy the SGP constraints.

Relation to literature

A number of recent works have discussed the need for an overhaul of the EU’s Stability and Growth Pact and fiscal ceilings in general (e.g. Claeys, Darvas and Leandro, 2016; Bénassy-Quéré et al., 2018; Heinemann, 2018; Constâncio, 2020; Blanchard, Leandro and Zettelmeyer, 2020). The debate has so far mostly focused on either the most appropriate fiscal instrument to put a ceiling on, say the budget deficit vs. the growth rate of government spending (e.g., Bénassy-Quéré et al., 2018; Giavazzi et al., 2021; Martin, Pisani-Ferry and Ragot, 2021), or on the appropriate level of fiscal ceiling, e.g., when r falls below g (e.g.,

Blanchard, 2019; Furman and Summers, 2020). The present paper focuses on a related, but complementary, issue: how can we address the limitations of fiscal ceilings: limited room for macro stabilization and low compliance. And while the trade-off between fiscal discipline and macro stabilization objectives is often recognized as a crucial center piece of any desirable fiscal rule (e.g., Giavazzi et al., 2021), a contribution of this work is to provide a simple and transparent way to explicitly contract on a desired the macro stabilization—fiscal discipline trade-off.

Our paper also relates to a large academic literature that has derived elaborate state-contingent rules from specific macroeconomic models, i.e., rules stating how the fiscal instruments should be set for all time and state contingencies. Such rules do not suffer from the rigidity of fiscal ceilings and can promise superior outcome *if* the underlying model is well specified. Unfortunately, a worry among policy makers is that any assumed model structure may always be too stylized relative to the complexity and unknowns of the economy (e.g. Blanchard, Leandro and Zettelmeyer, 2020).⁵ In practice, such model-based state-contingent rules are seldom used.

The remainder of this paper is organized as follows. In Section 2 we consider a simple environment that allows us to explain the main ideas that underlie the benefits of targeting rules in an intuitive manner. Section 3 then generalizes these ideas for a generic macroeconomic environment. The evaluation of compliance with a targeting rule using hypothesis testing is discussed in Section 5. The general practical implementation of fiscal-macro targeting is discussed in Section 6. The results from the empirical analysis of fiscal discipline in the EU is discussed in Section 7 and Section 8 concludes.

2 Illustrative example

In this section we illustrate how a fiscal-macro targeting rule can Pareto improve upon a hard fiscal ceiling, by leaving more room for macro stabilization and delivering higher fiscal discipline. A general treatment will follow in section 3.

Environment

There are two decisions makers: a policy maker that decides on fiscal policy and a higher-level legislator. The legislator should be understood broadly, it can be society as a whole,

⁵See e.g., Galí and Monacelli (2008); Halac and Yared (2014, 2019) for examples of model-based instrument rules. In the context of monetary policy, many policy makers have noted the practical limitations of following strict instrument rules, see Bernanke (2015) for a vivid discussion. The same limitations apply in the context of fiscal policy. As Blanchard, Leandro and Zettelmeyer (2020) put it, designing an instrument rule that captures ex-ante all relevant contingencies may simply not be possible.

the writers of a constitution, or a higher level organization like a monetary union.

The policy maker aims to stabilize output y around potential y^* using the fiscal instrument p . In this example, we can think of p as government spending. The policy maker's loss function \mathcal{L}^y and the economy are described by

$$\mathcal{L}^y = (y - y^*)^2, \quad y - y^* = \mathcal{R}p + \varepsilon, \quad \varepsilon = h(w),$$

where \mathcal{R} captures the effect of the fiscal instrument on the output gap and ε incorporates all non-policy factors w via the function $h(\cdot)$. The distribution function of ε is denoted by \mathcal{F}_ε . We can think about the model for the output gap as describing one equation from a general simultaneous equations model that also includes equations for p and w , which are unspecified, but also unrestricted, in our setting.

The legislator would like to restrain public spending and ensure that the fiscal instrument satisfies $p \leq \bar{p}$. In this paper, we do not take a stand on the reasons underlying this motive, only taking it at a starting point. Whenever the policy maker exceeds the \bar{p} limit, the legislator incurs a loss

$$\mathcal{L}^x = (p - \bar{p})_+^2,$$

where $(\cdot)_+$ takes the positive part of the function.⁶

The problem

In this example we study how a rule agreed between the legislator and the policy maker can best achieve the two objectives —keeping output close to potential *and* restraining public spending—. We study this question under two practical restrictions on the rule:

1. The maximum sanction cost for rule non-compliance is *finite*: $\bar{S} < \infty$
2. $h(w)$ cannot be contracted upon.

The first restriction stems from that fact that it is not possible to enforce arbitrarily large sanctions on fiscal policy makers. This implies that ensuring $p \leq \bar{p}$ at all time may not be possible as any rule can be breached —a limited enforcement problem.

The second restriction stems from the inherent complexity of the underlying economy, and the fact that there is no universally accepted model of the economy.⁷

⁶Specifically, we have $(p - \bar{p})_+^2 = \mathbf{1}(p \geq \bar{p})(p - \bar{p})^2$ where $\mathbf{1}(p \geq \bar{p})$ is equal to one whenever $p \geq \bar{p}$ and zero else.

⁷If $h(\cdot)$ was known and contractible, the problem would become trivial: one would only need to solve the model, derive the welfare function and then find the policy rule that delivers the highest welfare.

A fiscal ceiling rule

The approach commonly used in practice consists in setting up a rule that stipulates that the policy maker must satisfy $p \leq \bar{p}$ or face a non-compliance sanction S . Under such a “fiscal ceiling” rule \mathcal{C}_ℓ (with an ℓ for “limit”) the policy maker would solve

$$\mathcal{C}_\ell : \begin{cases} \min_p (y - y^*)^2 & \text{s.t. } p \leq \bar{p} & \text{if } (y - y^*)^2 \leq S \\ \min_p (y - y^*)^2 + S & & \text{else} \end{cases} . \quad (1)$$

Note how the fiscal ceiling \bar{p} is constant and does not depend on the output gap y , that is the ceiling \bar{p} is “rigid”. This rigidity has two unfortunate implications.

First, a fiscal ceiling can be very costly for the policy maker. With a strictly increasing function \mathcal{L}^y , that \mathcal{L}^y “macro” cost can become very large in large recessions (realizations of ε in the left-tail of \mathcal{F}_ε), as illustrated in Figure 1a, top panel.

Second, with limited enforcement a fiscal ceiling rule may end up delivering poor fiscal discipline (high $\mathbb{E}\mathcal{L}^x$) when the sanction cost S is not large enough, as illustrated in Figure 1b. Since rule compliance is more \mathcal{L}^y -costly in bad times, the policy maker will breach the rule as soon as \mathcal{L}^y exceeds S (top-middle panel), leading to large deviations from the fiscal constraint $p \leq \bar{p}$ (bottom-middle panel). While these deviations may be rare as they happen only in the tail of the \mathcal{F}_ε distribution, they are also the most costly in terms of fiscal discipline as they lead to larger losses in \mathcal{L}^x when \mathcal{L}^x is strictly increasing (bottom-middle panel). As a result, a fiscal ceiling rule can deliver poor fiscal discipline —a high $\mathbb{E}\mathcal{L}^x$ — if the upper-limit \bar{S} on sanctions is too low.


A fiscal-macro targeting rule

The key problem of the fiscal ceiling \bar{p} is that it does not take into account the macro cost —the \mathcal{L}^y -cost— of satisfying the fiscal constraint. By incorporating the policy maker’s own objective —macro stabilization— into the problem, we will see that it becomes possible to Pareto improve upon the fiscal ceiling rule, that is to lower both the \mathcal{L}^y loss *and* the \mathcal{L}^x loss.

Consider a rule \mathcal{C}_t that stipulates an *auxiliary* loss function \mathbf{L} that the policy maker should minimize (or face sanction S), where \mathbf{L} is a weighted average of the losses of the policy maker and the legislator:

$$\begin{aligned} \mathbf{L} &= \mathcal{L}^y + \lambda \mathcal{L}^x \\ &= (y - y^*)^2 + \lambda (p - \bar{p})_+^2 \end{aligned}$$

Figure 1: CEILINGS (\mathcal{C}_ℓ) VS. FLEXIBLE CEILINGS (\mathcal{C}_t)



Notes: Panel (a): With a large sanction for non-compliance ($S \rightarrow \infty$), the fiscal limit is always respected but at a high macro cost (high $\mathbb{E}\mathcal{L}_{\mathcal{C}_\ell}^y$) when \mathcal{L}^y is strictly increasing. Panel (b): Under limited enforcement ($S < \infty$), the fiscal ceiling is no longer respected when \mathcal{L}^y reaches S , leading to poor fiscal discipline (high $\mathbb{E}\mathcal{L}_{\mathcal{C}_\ell}^x$) because \mathcal{L}^x is strictly increasing. Panel (c): the flexibility allowed by fiscal-macro targeting can lower both \mathcal{L}^y (green area, top panel) and \mathcal{L}^x by trading the rare but large and costly deviations from \bar{p} under the fiscal ceiling rule (green area, bottom panel) with smaller (but more frequent) deviations (red area, bottom panel).

such that the policy maker's problem becomes

$$\mathcal{C}_t : \begin{cases} \min_p L & \text{if } (y - y^*)^2 \leq S \\ \min_p (y - y^*)^2 + S & \text{else} \end{cases} . \quad (2)$$

Note first that the fiscal-macro targeting rule nests the fiscal ceiling rule as a special case: when $\lambda \rightarrow \infty$.⁸ For finite values of λ , the rule incorporates the policy maker's own objective


⁸In fact, fiscal-macro targeting can be seen as a generalization of fiscal ceiling rules. Indeed, the constrained optimization problem implied by a fiscal ceiling can be represented as the minimization the Lagrangian $L = (y - y^*)^2 + \mu(p - \bar{p})$ where μ is the Lagrange multiplier μ . Comparing L with the auxiliary loss function L , our approach can be seen as substituting the Lagrange multiplier with $\lambda(p - \bar{p})$ where λ is a choice parameter controlling the constraint relaxation.

into the problem. Specifically, the parameter λ controls the relaxation of the fiscal constraint, i.e., the balance between the fiscal discipline objective and the macro stabilization objective. This is illustrated in Figure 2a, which plots the “macro stabilization–fiscal discipline frontier” offered by the \mathcal{C}_t rule under perfect rule enforcement (S infinite), i.e., it plots $(\mathbb{E}\mathcal{L}^y, \mathbb{E}\mathcal{L}^x)$, as we vary λ between 0 —an unconstrained policy— and ∞ —a fiscal ceiling rule—.

Pareto improvement

Our main result is then the following: under limited enforcement the constraint relaxation offered by a fiscal-macro targeting rule $\mathcal{C}_t(\lambda)$ can be Pareto improving over a fiscal ceiling rule \mathcal{C}_ℓ , reducing both the expected loss of the policy maker *and* the expected loss of the legislator. This is illustrated in Figure 2b. Starting from the fiscal ceiling rule ($\mathcal{C}_t(\infty)$) and relaxing the constraint (lowering λ) improves *both* the stabilization objective and the fiscal discipline objective: the frontier moves in a south-west direction — a Pareto improvement.

Figure 2: THE FISCAL DISCIPLINE – MACRO STABILIZATION FRONTIER



Notes: The two lines display the discipline–stabilization frontier allowed by the $\mathcal{C}_t(\lambda)$ rule under high sanction ($S \rightarrow \infty$, panel a) and finite sanction ($S < \infty$, panel b).

The intuition is as follows and is illustrated in Figure 1(c). A fiscal-macro targeting rule Pareto improves upon a fiscal ceiling, if it lowers both $\mathbb{E}\mathcal{L}^y$ and $\mathbb{E}\mathcal{L}^x$. While the policy maker is (by definition) better off under the more flexible fiscal-macro targeting rule — $\mathbb{E}\mathcal{L}^y$ is lower (the green area in the upper panel of Figure 1(c))—, the legislator sees two offsetting effects on its loss function $\mathbb{E}\mathcal{L}^x$. On the one hand, the flexibility of a fiscal-macro targeting rule has a cost in terms of lower fiscal discipline, since a policy maker facing a negative shock always deviates a little from the ceiling (the red area in the bottom-right panel of Figure

1(c)). On the other hand, under a fiscal-macro targeting rule non-compliance is less likely for large adverse shocks and $\mathbb{E}\mathcal{L}^x$ is lower (the green area in the bottom-right panel of Figure 1(c)). Indeed, faced with a large negative shock, a policy maker under fiscal-macro targeting is allowed to deviate from the fiscal constraint in order to stabilize the economy and avoid large macro \mathcal{L}^y costs. As a result, the $\mathcal{L}_{\mathcal{C}_t}^y(\varepsilon)$ curve is less steep than the $\mathcal{L}_{\mathcal{C}_t}^y(\varepsilon)$ curve, and it crosses the non-compliance threshold S later, i.e., for much larger adverse shocks. In other words, rule compliance is more likely under a fiscal-fiscal rule. When the green area dominates the red area, the second effect —higher rule compliance in the face of large shocks— dominates the first effect —frequent small deviations from the fiscal ceiling—, and the fiscal-macro targeting rule lowers $\mathbb{E}\mathcal{L}^x$ compared to the fiscal ceiling rule. In that case, the fiscal-macro targeting rule Pareto improves the fiscal ceiling rule: lowering both $\mathbb{E}\mathcal{L}^y$ and $\mathbb{E}\mathcal{L}^x$. In a nutshell, the fiscal-macro targeting rule can be seen as trading the rare but large and costly deviations from \bar{p} under the fiscal ceiling rule with smaller (but more frequent) deviations.

As we will prove formally in the general treatment, as long as rule violation occurs with positive probability under the fiscal ceiling rule, we can always find a set of λ s for which the fiscal-macro rule offers a Pareto improvement, see Theorem 1 below.

The remainder of this paper generalizes the fiscal-macro targeting rule for a generic dynamic macro environment and shows that the attractive properties of targeting carry over to this general setting.

3 General environment

In this section we generalize the setup of Section 2 to a generic economic environment. We allow for arbitrary loss functions and consider a generic macro modeling framework.

3.1 Loss functions

The policy maker is interested in stabilizing the economy by controlling M_y macroeconomic objectives, such as the output gap, unemployment gap, etc. Specifically we impose that at time t the policy maker aims to control the objectives $y_{i,t+h}$, for $i = 1, \dots, M_y$, over horizons h . We define $\mathbf{Y}_t = (y'_t, y'_{t+1}, \dots)'$ as the path of the policy objectives, where $y_t = (y_{1,t}, \dots, y_{M_y,t})'$. The loss that the policy maker incurs is measured by

$$\mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) , \tag{3}$$

where $\mathcal{L}^y(\cdot) : \mathbb{R}^\infty \rightarrow \mathbb{R}^+$ is a strictly increasing function and \mathbb{E}_t denotes the expectation with respect to the time t information set \mathcal{F}_t , i.e., $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot | \mathcal{F}_t)$. Importantly, the fiscal-macro

rules that we consider in this paper do not rely on specific choices for \mathbf{Y}_t or \mathcal{L}^y , and the properties that we derive will hold for any strictly increasing loss function that the policy maker considers.

To minimize the loss function the policy maker chooses an expected fiscal policy path, for instance current and future values of taxes, transfers and spending. In general the policy maker has J instruments and the expected fiscal policy path is denoted by $\mathbf{P}_t^e = \mathbb{E}_t(p'_t, p'_{t+1}, \dots)$, where p_t is the $J \times 1$ vector of policy instruments for time t .

The legislator wants to restrain the policy makers' actions and ensure that some fiscal variables $x_{i,t+h}$, for instance the budget deficit or government spending, satisfy constraints of the form

$$x_{i,t+h} \leq \bar{x}_{i,t+h}, \quad i = 1, \dots, M_x, \quad h = 0, 1, \dots,$$

where $\bar{x}_{i,t+h}$ is some threshold and there are M_x fiscal variables to control over H horizons.⁹ We stack the fiscal variables in $\mathbf{X}_t = (x'_t, x'_{t+1}, \dots)$ and $\bar{\mathbf{X}}_t = (\bar{x}'_t, \bar{x}'_{t+1}, \dots)'$, with $x_t = (x_{1,t}, \dots, x_{M_x,t})'$ and $\bar{x}_t = (\bar{x}_{1,t}, \dots, \bar{x}_{M_x,t})'$.

The loss incurred by the legislator when $\mathbf{X}_t > \bar{\mathbf{X}}_t$ is given by

$$\mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+), \tag{4}$$

where $\mathcal{L}^x(\cdot) : \mathbb{R}^\infty \rightarrow \mathbb{R}^+$ is strictly increasing for positive values and $(\mathbf{X}_t - \bar{\mathbf{X}}_t)_+$ has elements $(x_{i,t+h} - \bar{x}_i)_+ = \mathbf{1}(x_{i,t+h} > \bar{x}_i)(x_{i,t+h} - \bar{x}_i)$. A simple example for \mathcal{L}^x is $\mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) = \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_+^\nu$ where ν captures the degree of risk aversion of the legislator towards \mathbf{X}_t exceeding $\bar{\mathbf{X}}_t$, see e.g., Killian and Manganelli (2008). Taking $\nu = 2$ gives the quadratic loss function used in the illustrative example from section 2.

We note that the constraints are imposed on the fiscal variables \mathbf{X}_t and not on the fiscal policy paths \mathbf{P}_t directly. This allows for cases where the policy maker does not have complete control over the fiscal variables. For instance, if \mathbf{X}_t includes the budget deficit, there may be effects of risk premium shocks that affect the debt servicing cost. Alternatively, there can be mechanical cyclicalities in \mathbf{X}_t through the automatic stabilizers: in recessions the tax base shrinks and the deficit increases. A special case occurs when $\mathbf{X}_t = \mathbf{P}_t$, which was the setting considered in the simple example.

⁹The thresholds can be time and horizon specific, although in practice fiscal ceilings are constant across time and horizon, for instance a 3% deficit ceiling (e.g. Lledó et al., 2017).

3.2 Generic model

We consider a linear environment which can be justified for small fluctuations around a steady-state. A generic model for the expected paths of \mathbf{Y}_t , \mathbf{P}_t and \mathbf{X}_t is given by

$$\begin{cases} \mathcal{A}_{yy}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{yx}\mathbb{E}_t\mathbf{X}_t - \mathcal{A}_{yp}\mathbf{P}_t^e + \mathcal{A}_{yw}\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{y\xi}\mathbb{E}_t\boldsymbol{\Xi}_t \\ \mathcal{A}_{xx}\mathbb{E}_t\mathbf{X}_t - \mathcal{A}_{xy}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{xp}\mathbf{P}_t^e + \mathcal{A}_{xw}\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{x\xi}\mathbb{E}_t\boldsymbol{\Xi}_t \\ \mathcal{A}_{ww}\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{wy}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{wx}\mathbb{E}_t\mathbf{X}_t - \mathcal{A}_{wp}\mathbf{P}_t^e = \mathcal{B}_{w\xi}\mathbb{E}_t\boldsymbol{\Xi}_t \end{cases}, \quad (5)$$

where $\mathbf{W}_t = (w'_t, w'_{t+1}, \dots)'$ is a path for additional endogenous variables and $\boldsymbol{\Xi}_t = (\xi'_t, \xi'_{t+1}, \dots)'$ denotes the path of the structural shocks ξ_t and may possibly also include any initial conditions. The linear maps \mathcal{A}_\cdot and \mathcal{B}_\cdot are conformable. After taking expectations we can interpret $\mathbb{E}_t\boldsymbol{\Xi}_t$ as shocks (including news shocks) to the fundamentals of the economy that are released at time t (e.g. Chahrour and Jurado, 2018).

This model is general and it accommodates a large class models found in the literature, not only standard New-Keynesian (NK) models (e.g., Smets and Wouters, 2007), but also modern heterogeneous agents NK models (Auclert et al., 2021). Numerous specific examples can be found in Woodford (2003) and Walsh (2017).

3.3 The policy problem

The objective in this paper is to design a targeting rule that balances minimizing the policy makers loss \mathcal{L}^y and the legislators loss \mathcal{L}^x . To do so, we consider the environment, as characterized by equations (3)-(5), with two limitations on the rule design: (i) limited enforcement, and (ii) model uncertainty:¹⁰

Assumption 1 (Limited enforcement). *The maximum sanction cost (in units of $\mathbb{E}_t\mathcal{L}^y$) for rule non-compliance is finite: $\bar{S} < \infty$.*

Assumption 2 (Model uncertainty). *The maps \mathcal{A}_\cdot and \mathcal{B}_\cdot in (5) are unknown and cannot be used in the design of the rule.*

The first assumption allows for limited enforcement of the fiscal rule: if the cost of non-compliance is finite, the policy maker can choose to violate the constraints $\mathbf{X}_t \leq \bar{\mathbf{X}}_t$. That

¹⁰There is no asymmetry of information in our setup (and unlike the Principal-Agent literature (e.g. Bolton and Dewatripont, 2004, Part I)): the policy maker's preferences \mathcal{L}^y and policy choice \mathbf{P}_t^e are common knowledge, as is \mathcal{L}^x . Without private information, incentives issues disappear, such that *if* the underlying model could be explicitly written down, the principal could simply propose a rule that perfectly controls the agent. That is, the principal could specify a payment function that maps payments to the agent as a function of observed policy choices. In practice, the list of contingencies to take into account would be very long, complex, prone to disagreement, and even likely incomplete because of Knightian uncertainty, see Blanchard, Leandro and Zettelmeyer (2020). In the context of our generic model, this is captured by the fact that the maps \mathcal{A}_\cdot and \mathcal{B}_\cdot are unknown and thus cannot be contracted upon (Assumption 2).

assumption captures the fact that in practice it is hard to severely punish policy makers who choose not to respect a fiscal rule, i.e., the non-compliance sanction S cannot be arbitrarily large.¹¹

The second limitation captures the fact that in practice the specific model underlying the economy is highly complex and cannot be written down explicitly in all its details. As a result, it is difficult to agree on an exact model specification and set of model coefficients, and thus to agree on some “optimal” rule as implied by a specific model. Consistent with this assumption, the fiscal rules observed in practice do not depend on a specific underlying economic model (e.g., Lledó et al., 2017).

With these limitations in place, we will prove that there exists a non-empty set of fiscal-macro targeting rules that Pareto improve upon the fiscal ceiling rule. To formally rank rules in terms of performance and define the concept of Pareto improvement, we adopt the following notation and definitions. For a given rule \mathcal{C} , we let $\mathbb{E}_t \mathcal{L}_{\mathcal{C}}^i$, for $i = x, y$, denote the expected losses that result from the legislator and policy maker agreeing on rule \mathcal{C} . In general, we define the following criteria for ranking any two rules.

Definition 1 (Fiscal Discipline). *Given two rules \mathcal{C}_1 and \mathcal{C}_2 , fiscal discipline is higher under \mathcal{C}_2 if $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_2}^x < \mathbb{E}_t \mathcal{L}_{\mathcal{C}_1}^x$.*

Definition 2 (Pareto improvement). *Given two rules \mathcal{C}_1 and \mathcal{C}_2 , the \mathcal{C}_2 rule is a Pareto improvement over the \mathcal{C}_1 rule if $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_2}^y \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_1}^y$ and $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_2}^x \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_1}^x$.*

Definition 1 allows us to formally define fiscal discipline and compare rules in terms of their ability to induce policy makers to respect the fiscal constraints. Definition 2 allows to rank different rules in terms of their ability to jointly achieve the macro stabilization objectives and the fiscal discipline objective.

4 Fiscal-macro targeting rules

In this section we first discuss the fiscal ceilings rule as a benchmark rule that is commonly implemented in practice. Then we introduce our general class of fiscal-macro targeting rules and show how these rules can Pareto improve upon fiscal ceiling rules.

4.1 Fiscal ceilings

The common approach to ensure $\mathbf{X}_t \leq \bar{\mathbf{X}}_t$ is to directly impose the fiscal ceilings on the policy maker’s program. Formally, we define a fiscal ceiling rule (\mathcal{C}_ℓ) as follows:

¹¹Improvements on the sanction mechanisms are also of great interest but are outside of the scope of this paper.

Definition 3 (\mathcal{C}_ℓ rule). *A fiscal ceiling rule is defined by: (i) the requirement for the policy maker to satisfy $\mathbb{E}_t \mathbf{X}_t \leq \bar{\mathbf{X}}_t$, and (ii) a non-compliance sanction $S \leq \bar{S}$.*

Clearly, such a fiscal ceilings rule has the benefit of transparency and the vast majority of fiscal rules found in practice can be described by such a rule, see e.g., Lledó et al. (2017) for examples from over 90 countries. A prominent example is the EU SGP with a 3% deficit ceiling and a 60% debt-GDP ceiling.

As we already saw in Section 2, a rule based on fiscal ceilings has two related limitations: (i) rigidity which leads to high \mathcal{L}^y -cost (poor macro stabilization), and (ii) low compliance which leads to high \mathcal{L}^x -cost (poor fiscal discipline).¹²

Under the \mathcal{C}_ℓ rule, the policy maker solves the following problem

$$\begin{cases} \min_{\mathbf{Y}_t, \mathbf{X}_t, \mathbf{P}_t, \mathbf{W}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) & \text{s.t. (5) and } \mathbb{E}_t \mathbf{X}_t \leq \bar{\mathbf{X}}_t & \text{if } \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) \leq S \\ \min_{\mathbf{Y}_t, \mathbf{X}_t, \mathbf{P}_t, \mathbf{W}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) + S & & \text{else} \end{cases} \quad (6)$$

We denote an optimal solution for the expected paths by $\mathbb{E}_t(\mathbf{Y}_t^{\mathcal{C}_\ell}, \mathbf{X}_t^{\mathcal{C}_\ell}, \mathbf{P}_t^{\mathcal{C}_\ell}, \mathbf{W}_t^{\mathcal{C}_\ell})$. To ease on notations, we will refer to the expected policy path as $\mathbf{P}_t^{e\mathcal{C}_\ell}$.

4.2 Fiscal-macro targeting

As in Section 2, the fiscal-macro targeting rule stipulates an auxiliary loss function that the policy maker should minimize:¹³

$$\mathbb{E}_t \mathbf{L} = \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) + \lambda \mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+)$$

for some fixed constant $\lambda > 0$. Formally, we define a fiscal-macro targeting rule (\mathcal{C}_t) as follows:

Definition 4 ($\mathcal{C}_t(\lambda)$ rule). *A fiscal-macro targeting rule is defined by: (i) the requirement for the policy maker to minimize the loss function $\mathbb{E}_t \mathbf{L}$ for a given λ , and (ii) a non-compliance sanction S .*

A rule $\mathcal{C}_t(\lambda)$ nests as special cases (i) the fiscal ceilings rule when $\lambda \rightarrow \infty$ and (ii) the unconstrained solution when $\lambda = 0$. Clearly, there exists a range of rigidity in the fiscal

¹²Of course, if the underlying model was known, choosing a more elaborate and more appropriate (e.g., more flexible) fiscal rule would not be an issue; one would only have to solve the model and devise a rule that can approximate the planner's solution. However, this approach would violate Assumption 2 –model uncertainty–, and the goal of this paper to propose a fiscal rule that does not rely on a specific model.

¹³Again unlike in the principal agent literature, there is no private information. \mathcal{L}^y and \mathcal{L}^x are common knowledge and can be contracted upon.

constraint for λ in between these polar cases: the parameter λ controls the relaxation of the fiscal constraint.

Under a $\mathcal{C}_t(\lambda)$ rule, the policy maker solves the following problem

$$\begin{cases} \min_{\mathbf{Y}_t, \mathbf{X}_t, \mathbf{P}_t, \mathbf{W}_t} \mathbb{E}_t \mathbf{L} & \text{s.t. (5)} & \text{if } \mathbb{E}_t \mathcal{L}^y(Y_t) \leq S \\ \min_{\mathbf{Y}_t, \mathbf{X}_t, \mathbf{P}_t, \mathbf{W}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) + S & \text{s.t. (5)} & \text{else} \end{cases}, \quad (7)$$

The optimal solutions are denoted by $\mathbb{E}_t(\mathbf{Y}_t^{\mathcal{C}_t}, \mathbf{X}_t^{\mathcal{C}_t}, \mathbf{P}_t^{\mathcal{C}_t}, \mathbf{W}_t^{\mathcal{C}_t})$, and to ease on notations, we will refer to the corresponding expected policy path as $\mathbf{P}_t^{e\mathcal{C}_t}$.

4.3 Pareto improving relaxation

In this section we compare the expected losses — $\mathbb{E}_t \mathcal{L}^y$ and $\mathbb{E}_t \mathcal{L}^x$ — for the fiscal ceiling and fiscal-macro targeting rules. Specifically, given that the economy can be written as in (5), we are interested in establishing under which conditions a fiscal-macro rule yields a Pareto improvement over fiscal ceilings.

The following theorem establishes the main result.

Theorem 1. *Suppose that Ξ_t takes values in $\Gamma \subseteq \mathbb{R}^\infty$. Given assumption 1 and that the set $\{\Xi_t \in \Gamma : \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) > S\}$ is non-empty, we have that there exists a $\bar{\lambda}$ such that*

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^y \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y \quad \text{and} \quad \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^x \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x \quad \text{for all } \lambda \in [\bar{\lambda}, \infty) \quad (8)$$

The theorem states that there exists a range of values for λ that ensure that the \mathcal{C}_t rule Pareto dominates the \mathcal{C}_ℓ rule. The main assumption is that non-compliance can happen with positive probability for the \mathcal{C}_ℓ rule ($\{\Xi_t \in \Gamma : \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) > S\}$ is non-empty). The intuition is identical to the one described in Section 2, and we do not repeat it. We emphasize however the generality of the theorem: it holds for any loss function (as long as they are strictly increasing with the distance from the target) and any macro model that can be written as (5).

5 Assessing compliance with a fiscal-macro targeting rule

So far we established that a fiscal-macro targeting rule can improve upon fiscal ceilings under minimal assumptions —Theorem 1—, but showing that it can be implemented in practice under similarly modest assumptions is equally important from a practical perspective. This

section discusses how compliance with a targeting rule —whether the policy maker is minimizing some assigned loss function— can be evaluated without assuming a specific structure for the economy.

Policy proposal

Without loss of generality a proposed fiscal policy can be written as the sum of two terms, a component determined in response to the state of the economy captured by all time- t measurable variables —the instrument rule— and an exogenous component. Specifically, we posit that the expected policy path \mathbf{P}_t^e is determined by the generic system of equation

$$\mathcal{A}_{pp}\mathbf{P}_t^e - \mathcal{A}_{py}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{px}\mathbb{E}_t\mathbf{X}_t - \mathcal{A}_{pw}\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{p\xi}\mathbb{E}_t\boldsymbol{\Xi}_t + \boldsymbol{\epsilon}_t^e, \quad (9)$$

where $\boldsymbol{\epsilon}_t^e = \mathbb{E}_t\boldsymbol{\epsilon}_t$, $\boldsymbol{\epsilon}_t = (\epsilon'_t, \epsilon'_{t+1}, \dots)$, with $\epsilon_t = (\epsilon_{1,t}, \dots, \epsilon'_{M_p,t})'$. $\boldsymbol{\epsilon}_t^e$ are shocks to the expected policy paths \mathbf{P}_t^e , comprising shocks to the current values of the policy instruments as well as news shocks about the expected future values of the policy instruments. These shocks $\boldsymbol{\epsilon}_t^e$ are assumed to be uncorrelated with all other structural shocks, both current and future.

We collect all the elements of the instrument rule in $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$. A useful result from Barnichon and Mesters (2022) is that if a rule ϕ implies a unique equilibrium we have that

$$\mathbb{E}_t\mathbf{Y}_t = \mathcal{R}^y\boldsymbol{\epsilon}_t^e + \mathcal{C}_y\mathbb{E}_t\boldsymbol{\Xi}_t \quad \text{and} \quad \mathbf{X}_t = \mathcal{R}^x\boldsymbol{\epsilon}_t^e + \mathcal{C}_x\mathbb{E}_t\boldsymbol{\Xi}_t, \quad (10)$$

where the maps \mathcal{R}^y and \mathcal{R}^x capture the impulse responses of y_t and x_t to the fiscal policy shocks.¹⁴ The maps \mathcal{R} and \mathcal{C} depend on the underlying maps \mathcal{A} and \mathcal{B} , but we will not require exact knowledge of the mapping.

A gradient test

Consider a policy proposal \mathbf{P}_t^{e0} determined by (9), i.e., by a pair $(\phi^0, \boldsymbol{\epsilon}_t^0)$. Assessing compliance with a fiscal-macro targeting rule amounts to verifying that \mathbf{P}_t^{e0} minimizes the expected auxiliary loss $\mathbb{E}_t\mathbf{L}$. To do so, we will draw on Barnichon and Mesters (2022) who show that a necessary condition for \mathbf{P}_t^{e0} to be minimizing $\mathbb{E}_t\mathbf{L}$ is that the gradient of the loss function with respect to policy shocks is zero, i.e., that $\nabla_{\boldsymbol{\epsilon}_t}\mathbb{E}_t\mathbf{L}|_{\mathbf{P}_t^{e0}} = 0$ holds. Intuitively, if the policy choice \mathbf{P}_t^{e0} minimizes $\mathbb{E}_t\mathbf{L}$ there should be no deviation from this policy path —including changes in $\boldsymbol{\epsilon}_t$ — that can further lower $\mathbb{E}_t\mathbf{L}$, and the gradient with respect to $\boldsymbol{\epsilon}_t$ should be zero when evaluated at \mathbf{P}_t^{e0} .

¹⁴Expression (10) is simply the structural vector moving-average (VMA) representation of the economy.

Such gradient test, similar to a score test or Lagrange multiplier test, is particularly attractive, because it can be implemented without relying on a specific macro model. Specifically, the gradient test only requires the estimation of the gradient $\nabla_{\epsilon_t} \mathbb{E}_t \mathbf{L}$ under the null, i.e., it only requires the estimation of the gradient at $\mathbf{P}_t^{\epsilon_0}$, which can be done without having to agree on one specific model (consistent with Assumption 2). Two statistics are sufficient to conduct the test: (i) the causal effects of policy shocks on the macro and fiscal objectives, and (ii) the density forecasts for the macro and fiscal objectives conditional on $\mathbf{P}_t^{\epsilon_0}$.

To see this, note that the gradient evaluated at $\mathbf{P}_t^{\epsilon_0}$ can be written as

$$\nabla_{\epsilon_t} \mathbb{E}_t \mathbf{L} |_{\mathbf{P}_t^{\epsilon_0}} = \mathcal{R}^{0y'} \nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon_0}} + \lambda \mathcal{R}^{0x'} \nabla_{\mathbf{X}_t} \mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) |_{\mathbf{P}_t^{\epsilon_0}} \quad (11)$$

where \mathcal{R}^{0y} and \mathcal{R}^{0x} are the causal effects of ϵ_t on \mathbf{Y}_t and \mathbf{X}_t under ϕ^0 and $\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon} = \mathbf{P}_t^{\epsilon_0}}$ and $\nabla_{\mathbf{X}_t} \mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) |_{\mathbf{P}_t^{\epsilon} = \mathbf{P}_t^{\epsilon_0}}$ are functions of the density forecasts $f(\mathbf{Y}_t, \mathbf{X}_t | \mathcal{F}_t, \mathbf{P}_t^{\epsilon_0})$.

As argued in Barnichon and Mesters (2022), the two statistics can be estimated without relying on a specific economic model.

First, \mathcal{R}^{0y} and \mathcal{R}^{0x} are impulse response functions to shocks to ϵ_t^e , and thus can be transparently identified from a large body of macro studies on the propagation of structural shocks: natural experiments, e.g., IV-based methods (Ramey, 2016, 2019), theoretical studies (e.g., Zubairy, 2014; Leeper, Traum and Walker, 2017; Sims and Wolff, 2018) or even macro-econometric models used in fiscal institutions and ministries of finance. Moreover, since the gradient test is based on a necessary condition, it is not necessary to know the *full* maps \mathcal{R}^{0y} and \mathcal{R}^{0x} to construct a test of non-compliance. Any subset of the elements of \mathcal{R}^{0y} and \mathcal{R}^{0x} is enough to construct a gradient test of non-compliance, see Barnichon and Mesters (2022). In theory, the compliance test would be most powerful if it could be based on the full matrices \mathcal{R}^{0y} and \mathcal{R}^{0x} , that is if we could assess the gradient of the loss function \mathbf{L} in all possible directions. In practice however, there will be a trade-off between the number of impulse responses and the power of the test, as impulse responses need to be estimated, and more uncertain impulse response estimates will lead to less powerful tests.

Second, $\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon_0}}$ and $\nabla_{\mathbf{X}_t} \mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) |_{\mathbf{P}_t^{\epsilon_0}}$ can be evaluated from the forecasts densities alone, which do not require a specific model.¹⁵ Indeed, a large forecasting literature has shown how one can construct superior forecasts by combining large and disparate information sources, multiple (imperfect) models and possibly judgment (e.g. Stock

¹⁵To see this, note that as long as we can interchange the differentiation and integration orders we have for \mathbf{Y} (and similarly for \mathbf{X})

$$\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon} = \mathbf{P}_t^{\epsilon_0}} = \int \nabla_{\mathbf{Y}_t} \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon} = \mathbf{P}_t^{\epsilon_0}} dF_{\mathbf{Y}_t | \mathcal{F}_t}$$

where $F_{\mathbf{Y}_t | \mathcal{F}_t}$ denotes the time- t conditional forecast distribution. Thus, if the forecast densities for \mathbf{Y}_t and \mathbf{X}_t given $\mathbf{P}_t^{\epsilon_0}$ are available then $\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}^y(\mathbf{Y}_t) |_{\mathbf{P}_t^{\epsilon_0}}$ and $\nabla_{\mathbf{X}_t} \mathbb{E}_t \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) |_{\mathbf{P}_t^{\epsilon_0}}$ can be easily evaluated.

and Watson, 2002; Geweke and Amisano, 2012; Manganelli, 2009), and fiscal policy makers already rely on that literature to construct such conditional forecasts as part of their decision making procedure.

To summarize, the gradient statistic (11) can be evaluated from (i) the impulse responses to fiscal shocks and (ii) the forecast densities, both of which can be estimated by an independent agency. Importantly, the method to estimate the causal effects and the forecasts can be agreed upon ex-ante, periodically reviewed, and contracted upon transparently. In the appendix, we provide more details on the implementation of the gradient test for quadratic loss functions \mathcal{L}^y and \mathcal{L}^x , as we rely on such specification to empirically illustrate our approach below.

6 Implementing fiscal-macro targeting rules

Implementing fiscal-macro targeting requires the policy maker and the legislator to agree on three elements ex-ante, i.e., at the time of the signing of the fiscal-macro rule: (i) the auxiliary loss function \mathbf{L} —the policy objectives—, (ii) a timeline and evaluation procedure for evaluating compliance, and (iii) the sanction mechanisms. The first two elements have been discussed from economic and statistical perspectives in the previous sections, here we merely discuss some practical considerations that need to be taken into account. Regarding the third element, since the nature of the sanction system is not altered by moving from fiscal ceilings to fiscal targets, we will not discuss this element explicitly. That being said, we note that Theorem 1 implies that fiscal-macro targeting will require less sanctions on average.

6.1 Rule set-up

The first step is to agree on the auxiliary loss function \mathbf{L} , i.e., agree on the vectors of macro objectives \mathbf{Y}_t and fiscal objectives $\mathbf{X}_t - \bar{\mathbf{X}}_t$ along with functional forms for \mathcal{L}^y and \mathcal{L}^x :

$$\mathbf{L} = \mathcal{L}^y(\mathbf{Y}_t) + \lambda \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}})_+) .$$

The key variable to decide upon is the relative weight to assign to the fiscal objectives, that is the parameter λ , the desired macro stabilization - fiscal sustainability trade-off. Indeed the parameter λ captures how the \mathcal{C}_t rule values a marginal gain in \mathcal{L}^x relative to a marginal gain in \mathcal{L}^y , i.e., it captures the marginal rate of substitution between the “macro objective” and the “fiscal objective”. Graphically, picking λ consists in picking a point on the stabilization-discipline frontier depicted in Figure 2.

6.2 Timing and evaluation of compliance

Compliance with a targeting rule can be assessed using statistical methods discussed in Section 5. In practice, the policy maker and legislator will need to agree on (i) how often the test is conducted and (ii) who conducts the test.

In general, it is desirable to let the test be conducted by an independent agency. This ensures that the test is conducted in a transparent and credible way. The agency is then required to construct forecasts, compute impulse responses and implement the test. Since the power of the compliance test depends on the quality (low variance and unbiasedness) of the forecast, it is important to consider an agency with a good track record in terms of forecasting performance.

Interestingly the envisioned role for the independent agency is somewhat similar to that of the Swedish Fiscal Policy Council who has a special responsibility for analyzing how well the Swedish Government achieves its budget policy targets and whether the fiscal policy is sustainable in the long term. Andersson and Jonung (2019) argue that the presence of such agency is one of the components for the success of Swedish fiscal policy over the last three decades.

7 Empirical illustration

In this section we illustrate the workings of a fiscal-macro targeting rule using historical data for the EU and its Stability and Growth Pact (SGP). Obviously, no fiscal-macro targeting rule was officially implemented in the EU, and our exercise illustrates the evaluation of a fictitious fiscal-macro rule at different points in time.

7.1 Fiscal-macro rule setup

The loss function

Since the SGP imposes a 3% ceiling on budget deficits, we consider an auxiliary loss function capturing two objectives: (i) keeping GDP growth y at potential y^* , and (ii) keeping the budget surplus s above $\bar{s} = -3$ percent:

$$\mathbb{E}_t \mathbf{L} = \sum_{h=0}^H \mathbb{E}_t (y_{t+h} - y^*)^2 + \lambda \sum_{h=0}^H \mathbb{E}_t (s_{t+h} - \bar{s})_+^2 \quad (12)$$

Since the SGP requires paths for the next 3 years, we will fix the horizon at $H = 3$ years.¹⁶

¹⁶Naturally, more complicated loss functions are possible —including an additional debt-GDP target for instance.

Testing procedure

At the time of the signing of the treaty, parties must agree on (i) an independent forecasting agency that will create the forecasts (including model uncertainty estimates), and (ii) a set of policy experiments to assess compliance with fiscal-macro targeting, as well as an independent agency in charge of estimating the corresponding impulse responses (including estimation uncertainty).¹⁷

In this example, we use the economic forecasts reported by the individual countries to the EU commission as part of the SGP. Specifically, drawing on SGP records, we constructed a database over 1998-2020 that contains the individual forecasts provided by each union member to the EU commission. The forecasts are conditional on the intended fiscal path. The forecasts for the budget surplus and the real growth rate for France and Germany are shown in Figure 4.

The forecasts for France show a high degree of bias for both GDP growth and the budget surplus. In nearly all periods the forecasts turn out to be over-optimistic about the future path of the economy, especially in the long run as the bottom panel of Figure 4 shows. The bias is also present for Germany albeit less pronounced. We thus first bias-adjust the forecasts and remove the horizon specific trend for each country such that at least unconditionally the forecasts are unbiased.¹⁸ Additionally, it seems important to stress that for any reliable evaluation of fiscal discipline the current forecasting methodology needs to be improved, see also Gilbert and de Jong (2017).

To test compliance, we rely on the set of impulse response estimates from Guajardo, Leigh and Pescatori (2014) that capture the effects of fiscal austerity packages. Given our SGP focus, we only use EU countries in our estimation.

7.2 Assessing rule compliance

We now illustrate the fiscal-macro targeting rule defined by (12) in two ways. First, we illustrate how one would test compliance for France and Germany. Second, we consider a dual use of our framework, whereby we quantify the fiscal discipline of a given country. As we will see, this can allow to compare fiscal discipline across members of a monetary union.

¹⁷Alternatively, the two parties could agree ex-ante on values for the impulse responses (with uncertainty). In terms of timeline, compliance could be evaluated by the legislator (here the European commission) at the time of the signing of the budget.

¹⁸Clearly more advanced bias adjustment methods can be considered, but for our purpose the simple bias adjustment is sufficient. More generally, there is still considerable room for improvement in the quality of the forecasts reported to the EU commission, see also Gilbert and de Jong (2017).

France vs. Germany

Figure 3 contrasts the evolution of the budget surpluses of France and Germany over the past 20 years.

Germany occasionally deviated from the 3 percent deficit ceiling, but the breaches are short and in fact close in spirit to a fiscal-macro targeting rule. Indeed, under fiscal-macro targeting, deviations from a 3 percent ceiling are allowed, but these allowed deviations depend on the economic outlook. In the case of Germany, all 3% breaches occurred in the early stages of recessions, consistent with the prescription of a fiscal-macro targeting rule that balances the macro and the fiscal objectives.

The situation of Germany contrasts with that of France. While the two surpluses moved in tandem until 2004, since then France has done little fiscal consolidation and has since consistently breached the 3% limit.

A natural question is then whether the economic situation in France was so much worse than the one in Germany to justify the much larger budget deficits of France? Equivalently did France make less of an effort than Germany in respecting the SGP.¹⁹ With a fiscal-macro rule, equality of treatment across members of a monetary union imply that the same auxiliary loss function L should apply to all countries, i.e., the same weight λ on the fiscal objectives. Thus, we can reformulate the question as follows: given some fictitious fiscal-macro targeting rule $\mathcal{C}_f(\lambda^{\text{DE}})$ describing Germany, can we reject that France was complying with that rule? If we can, it would mean that France made less of a fiscal effort in respecting the SGP.

To characterize the evolution of Germany’s budget surplus over 1998-2020 in terms of a (fictitious) fiscal-macro targeting rule $\mathcal{C}_t(\lambda^{\text{DE}})$, we compute the value λ^{DE} that minimizes the sum-of-squared gradient statistics over 1998-2020.²⁰ Intuitively, λ^{DE} is chosen to make the gradient for Germany as small as possible on average over 1998-2020 (in an L^2 -norm sense), that is λ is chosen to minimize our ability to reject that Germany complied with the $\mathcal{C}_t(\lambda^{\text{DE}})$ rule. In other words, λ^{DE} is the parameter for which the rule $\mathcal{C}_t(\lambda^{\text{DE}})$ “best” describes Germany can be as a fiscal-macro targeter.

Based on our estimated λ^{DE} , Figure 6(a) plots $\nabla_{\epsilon_t} \mathbb{E}_t L |_{\mathbf{P}_t^{e_0}}$, the gradient statistic for Germany over 1998-2020, and shows that we can never reject that Germany was complying with the fictitious fiscal-macro rule $\mathcal{C}_t(\lambda^{\text{DE}})$.

Next, Figure 6(b) plots the gradient statistic for France over 1998-2020. We can see that France violated $\mathcal{C}_t(\lambda^{\text{DE}})$ numerous times when Germany did not, meaning that France was

¹⁹This is a common suspicion in Germany. See for instance some German reactions to a recent French proposals to reform the SGP: *France in preelection push to soften the eurozone’s budget rules* DW, May 2021 <https://www.dw.com/en/france-in-preelection-push-to-soften-the-eurozones-budget-rules>.

²⁰Specifically, $\lambda^{\text{DE}} = \arg \min \sum_t \mathcal{G}_t^{\text{DE}}(P_t^0; \lambda)^2$ where P_t^0 is the policy implemented by Germany at time t .

doing less of an effort than Germany in satisfying the SGP.²¹ As shown in Figure 4, the economic outlook was indeed similar in France and Germany and thus cannot justify the laxer fiscal stance of France.

That being said, thanks to the flexibility incorporated in a fiscal-macro rule, there are a number of instances where France’s violation of the 3% ceiling are tolerated by the $\mathcal{C}_t(\lambda^{\text{DE}})$ rule. Most notably, the fiscal-macro rule automatically relaxes the fiscal constraint during the COVID pandemic: despite the large increase in the deficit, the gradient statistic is close to zero, because of the large drop in GDP growth. In other words, there is no need for additional ad-hoc escape clauses. Even for an unforeseeable event like COVID, the targeting rule automatically incorporates the macro stabilization–fiscal discipline trade-off at play.

Fiscal discipline across EU countries

Once we reject that France complied with a virtual $\mathcal{C}_t(\lambda^{\text{DE}})$ fiscal-macro rule describing Germany, the dual question to ask is “Which $\mathcal{C}_t(\lambda^{\text{FR}})$ rule, i.e., which parameter λ^{FR} , best describes France as balancing fiscal discipline with macro stabilization objectives?”. Using again a minimum sum-of-squares criterion, we estimate $\lambda^{\text{FR}} = 0.3$ smaller than our estimate $\lambda^{\text{DE}} = 2$ and confirming the looser fiscal discipline of France.

More generally, we can repeat the procedure for each EU member country (denoted by i) and compute the parameter λ^i that best describes country i ’s implicit fiscal-macro targeting rule according to (12). In other words, given a list of policy objectives, λ^i can provide a metric to compare the level of fiscal discipline across countries.

Figure 7 plots the resulting estimates, ranking countries from lowest fiscal discipline (lowest weight λ^i on the fiscal objectives) to highest discipline (highest weight λ^i). Two separate groups clearly stand out in terms of fiscal discipline. The southern countries (Greece, Portugal and Spain) put the least weight on fiscal discipline relative to macro stabilization, with France and Belgium putting almost just as little weight on keeping the budget deficit under the 3% ceiling.²² In contrast, the northern countries (Holland, Germany, Denmark, Finland and Sweden) form a second group that puts much more weight on fiscal discipline (again, taking the economic outlook into account).

²¹The France forecasts are highly-biased (much more so than the German forecasts), as shown in the bottom panel, and it is only once we account for this bias that the lesser fiscal discipline of France becomes clear. In contrast, the biases for GDP growth are roughly comparable across countries. More generally, this finding reinforces the importance of relying on *independent* forecast agencies to assess compliance with a targeting rule.

²²In other words, once we take into account the superior economic outlook of France and Belgium relative to the southern EU countries, France and Belgium are no more fiscally responsible than the southern EU countries.

8 Conclusion

Fiscal constraints are essential to limit policy makers’ pro-deficit bias, but designing efficient yet flexible fiscal constraints has proved a formidable task. Most notoriously, fiscal ceilings—the main form of fiscal rules used in practice— have had limited success at restraining debt financing as rule violations are very frequent (e.g., Eyraud et al., 2018).

In this paper, we present the attractive properties of fiscal-macro targeting rules, which consist in providing policy makers with a loss function to be minimized, and where the loss function includes both fiscal objectives and macro stabilization objectives. Compared to fiscal ceilings, the flexibility offered by fiscal-macro targeting can not only improve policy makers’ own objectives of macro stabilization, but it can also improve rule compliance and thereby improve overall fiscal discipline. Monitoring compliance is transparent and objective, as it amounts to a statistical test.

We conclude by noting a strong parallel between our paper and the way central banks replaced the use of rigid rules with forecast targeting. While the design of the SGP was inspired by monetary rules like the 4.5% growth rate ceiling for the monetary base (Thygesen et al., 2019), central banks replaced these ad-hoc, rigid and rarely followed monetary ceiling with forecast targeting; a promise to set policy in order to meet a list of (possibly conflicting) objectives at a given horizon.²³ Our paper follows the same idea, as we propose to replace fiscal ceilings with a targeting approach to enforcing fiscal discipline.

²³See for instance Bernanke (2015)’s description of the Fed policy rule: “The Fed has a rule. The Fed’s rule is that we will go for a 2percent inflation rate; we will go for the natural rate of unemployment; we put equal weight on those two things; we will give you information about our projections, our interest rate. That is a rule and that is a framework that should clarify exactly what the Fed is doing.”

References

- Andersson, Fredrik N. G., and Lars Jonung.** 2019. “The Swedish Fiscal Framework The Most Successful One in the EU?” Lund University, Department of Economics Working Paper 2019:006.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub.** 2021. “Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models.” *Econometrica*, 89(5): 2375–2408.
- Barnichon, Regis, and Geert Mesters.** 2022. “A Sufficient Statistics Approach to Macro Policy Evaluation.” working paper.
- Bénassy-Quéré, A, M. Brunnermeier, H. Enderlein, E. Farhi, M. Fratzscher, C. Fuest, P-O Gourinchas, P. Martin, J. Pisani-Ferry, H. Rey, I. Schnabel, N. Veron, B. Weder di Mauro, and J. Zettelmeyer.** 2018. “Reconciling risk sharing with market discipline: A constructive approach to euro area reform.” CEPR Policy Insight No 91.
- Bernanke, Ben S.** 2015. “The Taylor Rule: A benchmark for monetary policy?” *Ben Bernankes Blog*, 28.
- Blanchard, Olivier.** 2019. “Public debt and low interest rates.” *American Economic Review*, 109(4): 1197–1229.
- Blanchard, Olivier, Alvaro Leandro, and Jeromin Zettelmeyer.** 2020. “Redesigning EU fiscal rules: From rules to standards.” *Economic Policy*.
- Bolton, Patrick, and Mathias Dewatripont.** 2004. *Contract Theory*. Mit press.
- Chahrour, Ryan, and Kyle Jurado.** 2018. “News or Noise? The Missing Link.” *American Economic Review*, 108(7): 1702–36.
- Claeys, G., Z. Darvas, and Á. Leandro.** 2016. “A Proposal to Revive the European Fiscal Framework.” Bruegel Policy Contribution.
- Constâncio, Vítor.** 2020. “The Return of Fiscal Policy and the Euro Area Fiscal Rule.” *Comparative Economic Studies*, 62: 358372.
- Drazen, Allan.** 2004. “Fiscal rules from a political economy perspective.” In *Rules-based fiscal policy in emerging markets*. 15–29. Springer.
- Eyraud, Luc, and Tao Wu.** 2015. *Playing by the rules: Reforming fiscal governance in Europe*. International Monetary Fund.
- Eyraud, Luc, Xavier Debrun, Andrew Hodge, Victor Lledó, and Catherine Pattillo.** 2018. *Second-generation fiscal rules: Balancing simplicity, flexibility, and enforceability*. Vol. 2018, International Monetary Fund.
- Furman, Jason, and Lawrence Summers.** 2020. “A Reconsideration of Fiscal Policy in the Era of Low Interest Rates.” *Discussion Draft, Harvard University*.

- Galí, Jordi, and Tommaso Monacelli.** 2008. “Optimal monetary and fiscal policy in a currency union.” *Journal of International Economics*, 76(1): 116 – 132.
- Geweke, John, and Gianni Amisano.** 2012. “Prediction with Misspecified Models.” *American Economic Review*, 102(3): 482–86.
- Giavazzi, Francesco, Veronica Guerrieri, Guido Lorenzoni, and Charles-Henri Weymuller.** 2021. “Revising the European Fiscal Framework.”
- Gilbert, N., and J. de Jong.** 2017. “Do European fiscal rules induce a bias in fiscal forecasts? Evidence from the Stability and Growth Pact.” *Public Choice*, 170: 1–32.
- Guajardo, Jaime, Daniel Leigh, and Andrea Pescatori.** 2014. “Expansionary austerity? International evidence.” *Journal of the European Economic Association*, 12(4): 949–968.
- Halac, Marina, and Pierre Yared.** 2014. “Fiscal rules and discretion under persistent shocks.” *Econometrica*, 82(5): 1557–1614.
- Halac, Marina, and Pierre Yared.** 2019. “Fiscal rules and discretion under limited enforcement.” National Bureau of Economic Research.
- Heinemann, Friedrich.** 2018. “How could the Stability and Growth Pact be simplified?” Report written at the request of the European Parliaments Committee on Economic and Monetary Affairs.
- Killian, Lutz, and Simone Manganelli.** 2008. “The Central Banker as a Risk Manager: Estimating the Federal Reserve’s Preferences under Greenspan.” *Journal of Money, Credit and Banking*, 40(6): 1103–1129.
- Larch, Martin, and Stefano Santacroce.** 2020. “Numerical compliance with EU fiscal rules: The compliance database of the Secretariat of the European Fiscal Board.” European Commission.
- Leeper, Eric M, Nora Traum, and Todd B Walker.** 2017. “Clearing up the fiscal multiplier morass.” *American Economic Review*, 107(8): 2409–54.
- Lledó, Victor, Sungwook Yoon, Xiangming Fang, Samba Mbaye, and Young Kim.** 2017. “Fiscal Rules at a Glance.” International Monetary Fund.
- Manganelli, Simone.** 2009. “Forecasting with judgment.” *Journal of Business & Economic Statistics*, 27(4): 553–563.
- Martin, Philippe, Jean Pisani-Ferry, and Xavier Ragot.** 2021. “Reforming the European fiscal framework.” *Les notes du conseil danalyse économique*, 63.
- Ramey, Valerie.** 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*, ed. J. B. Taylor and H. Uhlig. Amsterdam, North Holland:Elsevier.
- Ramey, Valerie A.** 2019. “Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?” *Journal of Economic Perspectives*, 33(2): 89–114.

- Rogoff, Kenneth.** 1985. “The optimal degree of commitment to an intermediate monetary target.” *The quarterly journal of economics*, 100(4): 1169–1189.
- Sims, Eric, and Jonathan Wolff.** 2018. “The output and welfare effects of government spending shocks over the business cycle.” *International Economic Review*, 59(3): 1403–1435.
- Smets, Frank, and Rafael Wouters.** 2007. “Shocks and frictions in US business cycles: A Bayesian DSGE approach.” *American economic review*, 97(3): 586–606.
- Stock, James H, and Mark W Watson.** 2002. “Macroeconomic Forecasting Using Diffusion Indexes.” *Journal of Business & Economic Statistics*, 20(2): 147–162.
- Svensson, Lars E. O.** 1997. “Inflation forecast targeting: Implementing and monitoring inflation targets.” *European Economic Review*, 41(6): 1111–1146.
- Thygesen, N, R Beetsma, M Bordignon, S Duchêne, and M Szczurek.** 2019. “Assessment of EU fiscal rules with a focus on the six and two-pack legislation.”
- Walsh, Carl E.** 2017. *Monetary theory and policy*. MIT press.
- Woodford, Michael.** 2003. *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.
- Zubairy, Sarah.** 2014. “On fiscal multipliers: Estimates from a medium scale DSGE model.” *International Economic Review*, 55(1): 169–195.

Appendix

A1: Proofs

Proof of Theorem 1. For convenience let $\mathbf{Z}_t = (\mathbf{Y}'_t, \mathbf{X}'_t, \mathbf{P}'_t, \mathbf{W}'_t)'$ and define for a given realization of Ξ_t the deterministic system

$$\begin{cases} \mathcal{A}_{yy} \mathbf{Y}_t - \mathcal{A}_{yx} \mathbf{X}_t - \mathcal{A}_{yp} \mathbf{P}_t + \mathcal{A}_{yw} \mathbf{W}_t = \mathcal{B}_{y\xi} \Xi_t \\ \mathcal{A}_{xx} \mathbf{X}_t - \mathcal{A}_{xy} \mathbf{Y}_t - \mathcal{A}_{xp} \mathbf{P}_t + \mathcal{A}_{xw} \mathbf{W}_t = \mathcal{B}_{x\xi} \Xi_t \\ \mathcal{A}_{ww} \mathbf{W}_t - \mathcal{A}_{wy} \mathbf{Y}_t - \mathcal{A}_{wx} \mathbf{X}_t - \mathcal{A}_{wp} \mathbf{P}_t = \mathcal{B}_{w\xi} \Xi_t \end{cases} . \quad (13)$$

The optimal solutions under the ceilings are fiscal-macro rules are given by

$$\mathbf{Z}_t^{\mathcal{C}_\ell} \in \arg \min_{\mathbf{Z}_t} \mathcal{L}_t(\mathbf{Y}_t) \quad \text{s.t.} \quad (13) \quad \text{and} \quad \mathbf{X}_t \leq \bar{\mathbf{X}}_t ,$$

and

$$\mathbf{Z}_t^{\mathcal{C}_t} \in \arg \min_{\mathbf{Z}_t} \mathcal{L}_t(\mathbf{Y}_t) \quad \text{s.t.} \quad (13) .$$

We have for any $\lambda \geq 0$ that

$$\begin{aligned} \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_\ell}) &= \min_{\mathbf{Z}_t} \mathcal{L}^y(\mathbf{Y}_t) \quad \text{s.t.} \quad (13) \quad \text{and} \quad \mathbf{X}_t \leq \bar{\mathbf{X}}_t \\ &= \min_{\mathbf{Z}_t} \mathcal{L}^y(\mathbf{Y}_t) + \lambda \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) \quad \text{s.t.} \quad (13) \quad \text{and} \quad \mathbf{X}_t \leq \bar{\mathbf{X}}_t \\ &\geq \min_{\mathbf{Z}_t} \mathcal{L}^y(\mathbf{Y}_t) + \lambda \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) \quad \text{s.t.} \quad (13) \\ &= \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) + \lambda \mathcal{L}^x((\mathbf{X}_t - \bar{\mathbf{X}}_t)_+) \\ &\geq \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) , \end{aligned} \quad (14)$$

which holds for any realization of Ξ_t .

Next, define the optimal policy that ignores all fiscal considerations by

$$\mathbf{Z}_t^{\mathcal{S}} \in \arg \min_{\mathbf{Z}_t} \mathcal{L}_t^y(\mathbf{Y}_t) \quad \text{s.t.} \quad (13)$$

and the sets

$$\mathcal{S}_{\mathcal{C}_\ell} = \left\{ \Xi_t \in \Gamma : \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_\ell}) \leq S \right\} \quad \text{and} \quad \mathcal{S}_{\mathcal{C}_t} = \left\{ \Xi_t \in \Gamma : \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) \leq S \right\}$$

These sets define the realizations for Ξ_t under which no sanction costs are incurred under the different rules. Note that for any finite λ (14) we have that $\mathcal{S}_{\mathcal{C}_\ell} \subset \mathcal{S}_{\mathcal{C}_t}$ and for $\lambda \rightarrow \infty$ we have $\mathcal{S}_{\mathcal{C}_t} \rightarrow \mathcal{S}_{\mathcal{C}_\ell}$. Define $\mathcal{O} = \mathcal{S}_{\mathcal{C}_\ell} \cap \mathcal{S}_{\mathcal{C}_t}^\perp$, where $\mathcal{S}_{\mathcal{C}_\ell}^\perp$ is the complement of $\mathcal{S}_{\mathcal{C}_\ell}$.

Next, the expected loss of the policy maker under the \mathcal{C}_ℓ rule is given by

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y = \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_\ell}) dF_{\Xi_t | \mathcal{F}_t} + \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} (\mathcal{L}^y(\mathbf{Y}_t^{\mathcal{S}}) + S) dF_{\Xi_t | \mathcal{F}_t} ,$$

where $F_{\Xi_t | \mathcal{F}_t}$ is the distribution of the shocks conditional on the information set \mathcal{F}_t . The loss of the policy maker under the \mathcal{C}_t rule is given by

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^y = \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_t}} \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) dF_{\Xi_t | \mathcal{F}_t} + \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_t}^\perp} (\mathcal{L}^y(\mathbf{Y}_t^{\mathcal{S}}) + S) dF_{\Xi_t | \mathcal{F}_t} .$$

Subtracting the two losses using $\mathcal{S}_{\mathcal{C}_\ell} \subset \mathcal{S}_{\mathcal{C}_t}$ and $\mathcal{O} = \mathcal{S}_{\mathcal{C}_\ell}^\perp \cap \mathcal{S}_{\mathcal{C}_t}$ gives

$$\begin{aligned} \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^y &= \int_{\Xi_t \in \mathcal{O}} (\mathcal{L}^y(\mathbf{Y}_t^s) + S) - \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) dF_{\Xi_t | \mathcal{F}_t} \\ &\quad + \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_\ell}) - \mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) dF_{\Xi_t | \mathcal{F}_t} \end{aligned}$$

The first integral is non-negative as over \mathcal{O} the \mathcal{C}_t rule does not default and hence $\mathcal{L}^y(\mathbf{Y}_t^{\mathcal{C}_t}) \leq (\mathcal{L}^y(\mathbf{Y}_t^s) + S)$. The second term is also positive by (14). Hence, we have $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y \geq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^y$.

Next, the expected loss of the legislator under the \mathcal{C}_ℓ rule is given by

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x = \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} \mathcal{L}^x((\mathbf{X}_t^s - \bar{\mathbf{X}}_t)_+) dF_{\Xi_t | \mathcal{F}_t}$$

and under the \mathcal{C}_t rule

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^x = \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_t}} \mathcal{L}^x((\mathbf{X}_t^{\mathcal{C}_t} - \bar{\mathbf{X}}_t)_+) dF_{\Xi_t | \mathcal{F}_t} + \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} \mathcal{L}^x((\mathbf{X}_t^s - \bar{\mathbf{X}}_t)_+) dF_{\Xi_t | \mathcal{F}_t}$$

Subtracting the losses gives

$$\begin{aligned} \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^x &= \int_{\Xi_t \in \mathcal{O}} \mathcal{L}^x((\mathbf{X}_t^s - \bar{\mathbf{X}}_t)_+) - \mathcal{L}^x((\mathbf{X}_t^{\mathcal{C}_t} - \bar{\mathbf{X}}_t)_+) dF_{\Xi_t | \mathcal{F}_t} \\ &\quad - \int_{\Xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^x((\mathbf{X}_t^{\mathcal{C}_t} - \bar{\mathbf{X}}_t)_+) dF_{\Xi_t | \mathcal{F}_t} \end{aligned}$$

Note that for $\lambda \rightarrow \infty$ we have $\mathcal{O} \rightarrow \emptyset$ and $\mathbf{Z}_t^{\mathcal{C}_t} \rightarrow \mathbf{Z}_t^{\mathcal{C}_\ell}$ and thus $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^x \rightarrow 0$. Also, for $\lambda = 0$ we have that $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x \leq 0$ as $\mathbf{Z}_t^{\mathcal{C}_t} = \mathbf{Z}_t^s$. So if the gradient is negative for $\lambda \rightarrow \infty$ (e.g. $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x$ approaches zero from above) we know that there is at least one $\bar{\lambda}$ for which $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x \geq 0$ as $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x$ must cross zero. To see that this is indeed the case, note that $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x > 0$ if $\mathcal{S}_{\mathcal{C}_\ell}^\perp \neq \emptyset$ and $\nabla_\lambda \mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x = 0$, but $\mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x < 0$ as increasing λ places more weight on the fiscal objective, hence reducing $\mathcal{L}^x((\mathbf{Y}_t^{\mathcal{C}_t} - \bar{\mathbf{X}}_t)_+)$. Together, this implies that $\nabla_\lambda (\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x) > 0$. Finally, since $\mathcal{L}^x((\mathbf{Y}_t^{\mathcal{C}_t} - \bar{\mathbf{X}}_t)_+)$ is continuously decreasing as $\lambda \rightarrow \infty$ we have that $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_t}^x \geq 0$ for all $\lambda \in [\bar{\lambda}, \infty)$. \square

A2: Gradient test implementation for quadratic loss functions

We will discuss the implementation of the gradient test for quadratic loss functions as we rely on such specification to empirically illustrate our approach below.²⁴

The loss functions become

$$\mathbb{E}_t \mathcal{L}_t^y = \mathbb{E}_t \mathbf{Y}_t' \mathcal{W}_y \mathbf{Y}_t \quad \mathbb{E}_t \mathcal{L}_t^x = \mathbb{E}_t (\mathbf{X}_t - \bar{\mathbf{X}}_t)'_+ \mathcal{W}_x (\mathbf{X}_t - \bar{\mathbf{X}}_t)_+ , \quad (15)$$

²⁴In fact, in the empirical application we rely on the forecasts of the euro area countries conditional on their proposed policy paths. Unfortunately the European commission only provides point forecasts limiting the implementation of the general gradient test.

where the diagonal maps \mathcal{W}_y and \mathcal{W}_x , allow for discounting and different weights on the different macro and fiscal targets.

To verify whether $\mathbf{P}_t^{e_0}$ satisfies the gradient condition, we note that the gradient evaluated at $\mathbf{P}_t^{e_0}$ is given by

$$\nabla_{\epsilon_t} \mathbb{E}_t \mathbf{L} |_{\mathbf{P}_t^{e_0}} = \mathcal{R}^{0y'} \mathcal{W}_y \mathbb{E}_t \mathbf{Y}_t^0 - \lambda \mathcal{R}^{0x'} \mathcal{W}_x \mathbb{E}_t (\mathbf{X}_t^0 - \bar{\mathbf{X}}_t)_+ , \quad (16)$$

To construct a test statistic based on $\nabla_{\epsilon_t} \mathbb{E}_t \mathbf{L} |_{\mathbf{P}_t^{e_0}}$ we need to (a) estimate the dynamic causal effects \mathcal{R}^{0y} and \mathcal{R}^{0x} , and (b) approximate the oracle forecasts $\mathbb{E}_t \mathbf{Y}_t^0$ and $\mathbb{E}_t \mathbf{X}_t^0$. Note how for quadratic loss functions only conditional point forecasts (given $\mathbf{P}_t^{e_0}$) are required to evaluate the gradient.

Now in practice, estimating the entire causal effect maps \mathcal{R}^{0y} and \mathcal{R}^{0x} is complicated as it requires identifying all policy news shocks ϵ_t^e . Fortunately to evaluate whether a given policy decision is obtain we can rely on any subset of the gradient, which forms a necessary condition for the optimality of policy. In other words, we can leverage existing identification methods to identify any subset or linear combination of the structural shocks, say $\epsilon_{a,t}^e$, and use these shocks to evaluate policy.

Formally, let \mathcal{R}_a^{0y} and \mathcal{R}_a^{0x} denote the causal effects that pertain to the subset of shocks $\epsilon_{a,t}^e$ that can be identified. The subset gradient becomes

$$\nabla_{\epsilon_{a,t}} \mathbb{E}_t \mathbf{L} |_{\mathbf{P}_t^{e_0}} = \mathcal{R}_a^{0y'} \mathcal{W}_y \mathbb{E}_t \mathbf{Y}_t^0 - \lambda \mathcal{R}_a^{0x'} \mathcal{W}_x \mathbb{E}_t (\mathbf{X}_t^0 - \bar{\mathbf{X}}_t)_+ , \quad (17)$$

and $\nabla_{\epsilon_{a,t}} \mathbb{E}_t \mathbf{L} |_{\mathbf{P}_t^{e_0}} = 0$ is a necessary condition for optimality.

In what follows we assume that the dynamic causal effects \mathcal{R}_a^{0y} and \mathcal{R}_a^{0x} can be estimated by the researcher and that confidence bands can be obtained. More specifically, we assume that the researcher is able to obtain estimates $\hat{r} = \left(\text{vec}(\hat{\mathcal{R}}^y)', \text{vec}(\hat{\mathcal{R}}^x)' \right)'$ that satisfy

$$\hat{r} \stackrel{a}{\sim} N(r, \Omega) \quad (18)$$

where $r = (\text{vec}(\mathcal{R}^{0y})', \text{vec}(\mathcal{R}^{0x})')'$ and Ω is the variance matrix of all impulse responses: across horizons and instruments. We assume that the variance matrix can be consistently estimated and we denote the estimate by $\hat{\Omega}$. The distribution (18) implies that we can recover the distribution of the dynamic causal effects using using (18).

Next, we approximate the distribution of the oracle forecasts $\mathbb{E}_t \mathbf{Y}_t^0$ and $\mathbb{E}_t \mathbf{X}_t^0$. In practice, forecasters typically produce point estimates, say $\hat{\mathbf{Y}}_t$ and $\hat{\mathbf{X}}_t$, for the macro and fiscal variables. We need the distribution of $\hat{\mathbf{Y}}_t - \mathbb{E}_t \mathbf{Y}_t^0$ and $\hat{\mathbf{X}}_t - \mathbb{E}_t \mathbf{X}_t^0$, i.e. the model mis-specification distribution. In practice, this distribution can be assessed by carefully analyzing the forecasting model and past forecasting performance. In our empirical work below we rely on historical forecasting performance, but alternative approaches can also be considered.


In general, we postulate that the forecast misspecification distribution can be approximated by

$$\hat{\mathbf{V}}_t - \mathbb{E}_t \mathbf{X}_t^0 \sim F_{\mathbf{V}_0} , \quad \text{where} \quad \hat{\mathbf{V}}_t = \begin{bmatrix} \hat{\mathbf{Y}}_t \\ \hat{\mathbf{X}}_t \end{bmatrix} , \quad \mathbb{E}_t \mathbf{V}_t^0 = \begin{bmatrix} \mathbb{E}_t \mathbf{Y}_t^0 \\ \mathbb{E}_t \mathbf{X}_t^0 \end{bmatrix} . \quad (19)$$

The distributions $\hat{r} \stackrel{a}{\sim} N(r, \Omega)$ and $F_{\mathbf{V}_0}$ are used to compute the distribution of the subset gradient.


In particular, we simulate B independent draws from $\hat{r} \stackrel{a}{\sim} N(r, \Omega)$ and $F_{\mathbf{V}_0}$, and for each draw we compute the gradient. From this simulated sample of draws we typically report the mean and upper and lower quantiles.

Figure 3: BUDGET SURPLUS: FRANCE VS. GERMANY




Notes: Top panel: government budget balance in percent of GDP (“budget surplus”) for France (FR) and Germany (DE) over 1995-2020. The bottom panel reports the difference between the two series.

Figure 4: SGP FORECASTS: FRANCE VS. GERMANY




Notes: The top two panels report the realized values (dashed-thick lines) for GDP growth and the budget surplus for France (left column) and Germany (right column), along with the forecasts successively reported to the EU commission (colored lines). The bottom row reports the average bias of these forecasts by forecast horizon.

Figure 5: IMPULSE RESPONSES TO A FISCAL AUSTERITY SHOCK



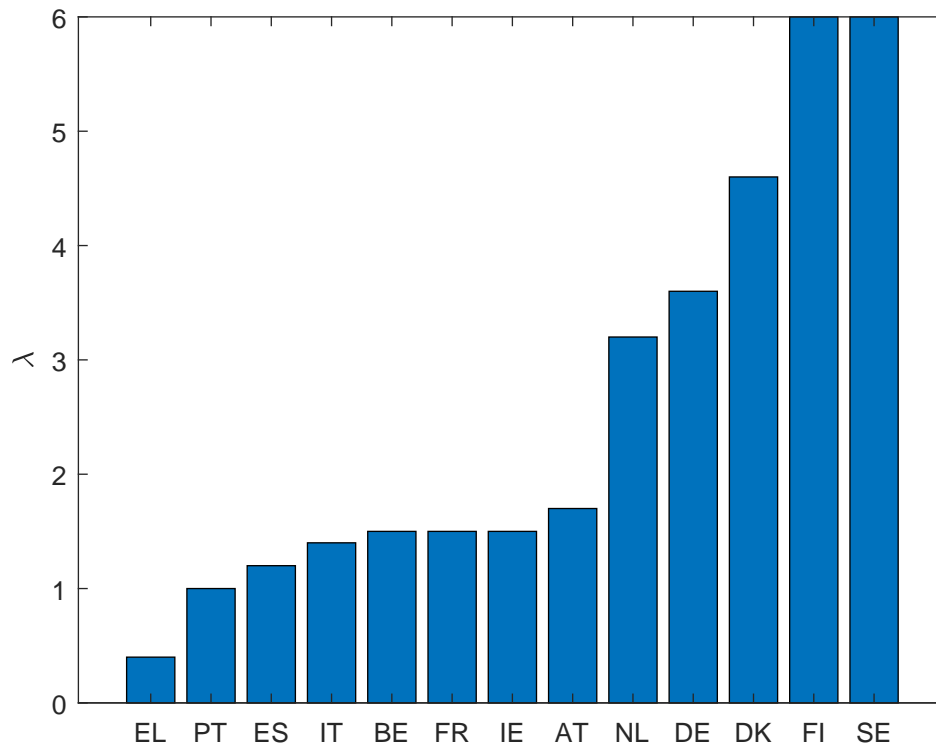
Notes: Impulse responses to fiscal austerity shock, estimation based on Guajardo, Leigh and Pescatori (2014) narratively identified shocks.

Figure 6: Gradient test, 1998-2020



Notes: Gradient statistic $\nabla_{P_t} \mathbb{E}_t L |_{P_t^0}$ with 95 confidence band for the fiscal-macro targeting rule $\mathcal{C}_f(\lambda^{\text{DE}})$ for Germany (panel a.) and France (panel b.). A non-zero value for the gradient test indicates non-compliance with the $\mathcal{C}_f(\lambda^{\text{DE}})$ rule.

Figure 7: FISCAL DISCIPLINE ACROSS THE EU



Notes: Implied fiscal-macro targeting rule for different EU countries. Each bar depicts the parameter λ —the weight on the 3% budget deficit ceiling—estimated to minimize the sum-of-squares of the gradient statistic implied by the loss function (12).