

Fiscal Targeting*

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

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

^(b) Universitat Pompeu Fabra, Barcelona GSE and VU Amsterdam

July 24, 2021

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 this paper, we propose to implement fiscal constraints through a fiscal targeting framework, paralleling central banks' move from monetary rules to inflation targeting. Under fiscal targeting, fiscal policy makers must optimally balance some fiscal objectives (e.g., keeping the deficit below 3%) with their own policy objectives (e.g., stabilizing output at potential). Fiscal targeting can be implemented with minimal assumption on the underlying economic model, and it promises a number of benefits over commonly used fiscal rules: (i) stronger buy-in from policy makers, (ii) higher fiscal discipline, (iii) transparency and ease of monitoring.

JEL classification: C14, C32, E32, E52.

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

*We thank Roel Beetsma, Stephanie Cheng, Davide Debortoli, Luca Fornaro, 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.

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, with constraints on the public deficit, on public expenditures or on the debt level.²

Unfortunately, fiscal rules have had limited success in limiting deficit spending, and fiscal constraints are frequently violated (e.g., Eyraud et al., 2018). This failure can be traced back to two practical limitations facing fiscal rules. First, fiscal rules suffer from a so-called 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. Second, the fiscal rules used in practice are very rigid — generally taking the form of fiscal limits, i.e., hard thresholds on fiscal variables (e.g., Lledó et al., 2017). The simplicity of imposing fiscal limits is appealing, but their rigidity can make them harmful at times, prompting little buy-in from policy makers. Combined with low sanctions, this rigidity can lead to frequent rule violations and low fiscal discipline. While a large academic literature has derived more elaborate (i.e., more flexible) state-contingent rules from specific macroeconomic models, a worry among policy makers is that the assumed model structure may always be too stylized relative to the complexity and unknowns of the economy (e.g. Blanchard, Leandro and Zettelmeyer, 2020).³ In practice, model-based state-contingent rules are seldom used, and decision makers rely on ad-hoc escape clauses to flexibilize fiscal limits, with limited success however.⁴

In this paper, we study how to improve the design and implementation of fiscal constraints under two key practical limitations that have received little attention in the literature. First, any fiscal constraint can be violated as enforcing arbitrarily large sanctions on fiscal policy makers is not possible. Second, there is no universally accepted model of the economy, and it is not possible to rely on a specific economic model to design an “optimal” fiscal rule, i.e.,

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

²See e.g., Eyraud et al. (2018), Schaechter et al. (2012) and Yared (2019).

³See e.g., Galí and Monacelli (2008); Halac and Yared (2014, 2019) for examples of model-based fiscal rules. In the context of monetary policy, many policy makers have noted the practical limitations of following strict model-based policy 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 a fiscal rule that captures ex-ante all relevant contingencies may simply not be possible.

⁴For example, in the European Union (EU) the Stability and Growth Pact (SGP) has undergone a number of reforms (motivated by repeated violations of the 3% deficit and 60% debt limits) that moved the SGP towards higher flexibility but also higher complexity and subjectivity, and this ultimately led to lower credibility, lower compliance and even lower fiscal discipline (e.g., Eyraud et al., 2018; Larch and Santacroce, 2020).

a welfare maximizing state-contingent policy rule.⁵

Given these practical limitations we propose to implement fiscal constraints through a fiscal targeting framework, paralleling central banks’ move from money-growth rules and Taylor type rules to inflation targeting (e.g., Mishkin, 2001). Specifically, instead of imposing hard fiscal limits on policy makers, we propose to provide policy makers with a list of policy objectives to be targeted, just like central banks are provided with a list of targets, e.g., stable inflation and full employment. In the context of fiscal targeting, the targets include the original fiscal objectives, such as keeping the deficit below 3%, but they also include the policy makers’ own objectives, such as stabilizing output at potential. Fiscal targeting then consist in balancing these (often conflicting) objectives. To operationalize fiscal targeting, we require the specification of a high-level auxiliary loss function, as chosen by society, that combines the fiscal objectives with the policy maker’s objectives in any desired way. A policy maker is then deemed fiscally responsible if she minimizes this auxiliary loss function.

Fiscal targeting nests the current practice of imposing hard fiscal limits as a special case —when the auxiliary loss function puts infinite weight on the fiscal objectives—. As such fiscal targeting does not aim to replace fiscal rules with discretion (e.g. Kydland and Prescott, 1977). Instead, it offers a transparent way to add flexibility to existing fiscal rules.

Under our limited enforceability and model uncertainty assumptions we show that fiscal targeting provides a number of important benefits: generality, transparency, flexibility and higher rate of compliance.

First, fiscal targeting can be implemented with minimal assumption on the underlying economic model. This brings both generality and simplicity, as there is no need to agree on one representation of the economy, something that is hard to achieve in practice given the complexity of the underlying economy and the many remaining unknowns.

With fiscal targeting, a policy maker must simply ensure that its fiscal policy is such that the forecasts for the policy objectives “look good”, i.e., best balance the objectives specified in the auxiliary loss function. Importantly, the forecasts for the policy objectives can be constructed by independent agencies even if the specific model cannot be explicitly written down, reflecting the practical approach to macro forecasting where policy makers combine multiple models (statistical and structural), judgment calls, and instinct to incorporate a large amount of information into the forecasting process.

Second, fiscal targeting can be implemented and enforced in an objective, transparent and predictable manner. Once an auxiliary loss function has been agreed upon, we show that assessing whether the policy maker is minimizing that loss function, i.e. complying with

⁵This could be because any given model will be too simple relative to the complexity of the underlying economy —a concern voiced by Blanchard, Leandro and Zettelmeyer (2020)—, or because the different parties involved in the design of a fiscal rule may never agree on the correct theoretical representation of the economy.

fiscal targeting, can be determined by means of a simple statistical test. This “compliance” test provides an objective criteria that is much less prone to discretion, judgment calls and political interference than current procedures.⁶ The test statistic is the gradient of the auxiliary loss function, and its distribution can be determined with minimal assumptions on the underlying economic model. This means that fiscal targeting can be implemented without relying on a specific model.

Intuitively, the compliance test consists in assessing whether a small change to the chosen fiscal policy can lower the auxiliary loss function. If it does —the gradient is non-zero—, we can conclude that fiscal policy makers did not make enough of an effort to satisfy the fiscal target —a case of non-compliance—, as a slightly different policy could have better balanced the macro and fiscal objectives. Importantly, tracing out the effects of small policy changes does not require a fully specified model. Instead, it only requires two sufficient statistics: (i) a set of forecasts for the targets conditional on the proposed policy —the baseline scenario—, and (ii) the causal effect of the policy instruments on the policy objectives —impulse response functions—. These two statistics can be constructed or estimated without relying on any one specific macro model. Macro forecasting routinely combines multiple models, judgment calls, and instinct into the forecasting process. Impulse responses can be transparently identified with minimal model assumptions from a large body of macro-econometric studies, notably instrumental variable methods (Ramey, 2016, 2019).

Third, fiscal targeting automatically incorporates the idea that “fiscal responsibility” is a context-dependent concept as the loss function includes both macro and fiscal objectives. This implies that the gradient of the loss function depends on (i) the economic outlook,⁷ and (ii) the “fiscal technology” of a country at any point in time, i.e., the ability of a country to use fiscal tools to achieve both its macro objectives and its fiscal targets. Importantly, unlike previous proposals to flexibilize rigid fiscal rules, the flexibility embedded in the targeting approach is not an ad-hoc or ex-post adjustment. The adjustment is based on an objective and transparent “loss”-based criterion that can be agreed upon ex-ante, and thus objectively assessed using our test-based procedure.

Fourth, flexible fiscal targeting can alleviate the limited enforceability problem of hard fiscal limits. By being more closely aligned to policy makers’ own objective, fiscal targeting reduces policy makers’ incentives to deviate from the fiscal targets, i.e., it increases compliance. As a result, fiscal targeting can not only improve policy makers’ own objectives by relaxing the hard fiscal limits, but it can also improve overall fiscal discipline: a Pareto improvement. Instead of trying to improve compliance via sanctioning, fiscal targeting en-

⁶See the current excessive deficit procedure of the EU’s Stability and Growth Pact for instance.

⁷For instance, a fiscal constraint would typically receive less weight during a recession: in a deep crisis a country may be fiscally responsible *even* if it runs a large budget deficit, as deficit spending can buoy up aggregate demand and prevent a too large output gap.

hances compliance by cooperation.

Finally, fiscal targeting is to a large extent already practiced by many countries subject to fiscal rules. Indeed, while countries often deviate from hard fiscal limits such as the SGP 3% deficit ceiling, policy makers do try to stay somewhat “close” to the fiscal limits (e.g., Eyraud and Wu, 2015). A fiscal targeting framework allows to formalize what constitutes an appropriate deviation from the fiscal constraint: it provides a means to contract ex-ante on the appropriate balance between satisfying the fiscal constraints and satisfying the policy maker’s own objectives.

To illustrate the workings of fiscal targeting we discuss the results from a pseudo out-of-sample policy evaluation exercise for different countries in the European Union. In particular, we ask whether fiscal policy decisions over 1998-2020 satisfy a fictitious fiscal targeting contract with two objectives: stabilizing GDP growth at potential and keeping the budget deficit below 3%.⁸ To implement the policy evaluation exercise we constructed a new database containing the individual 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 test contract compliance for France and Germany. We find that France is much less fiscally responsible than Germany: even after controlling for the economic outlook, France made less of an effort than Germany in respecting the 3% deficit rule. Looking across all EU countries, we find that fiscal policy can be described by a fiscal targeting contract, but the weight placed on the fiscal objective varies substantially across EU countries. In other words, fiscal responsibility varies greatly across EU members.

The remainder of this paper is organized as follows. We complete the introduction by carefully relating the fiscal targeting approach to the literature. In Section 2 we consider a simple environment that allows us to explain the main ideas that underlie the fiscal targeting approach in an intuitive manner. Sections 3 and 4 then generalize these ideas for a generic macroeconomic environment. The evaluation of compliance with fiscal targeting using hypothesis testing is discussed in Section 5. The general practical implementation of fiscal 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.

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 rules in general (e.g. Claeys, Darvas and Leandro, 2016; Bénassy-

⁸Obviously fiscal targeting was not implemented during the sampling period and therefore our choice for the targets is arbitrary. We note however that the exercise can be repeated using any desired contract and our objective is merely to illustrate the workings of the compliance test.

Quéré et al., 2018; Heinemann, 2018; Constâncio, 2020; Blanchard, Leandro and Zettelmeyer, 2020), but the debate has so far mostly focused on long-run issues such as the appropriate debt ceiling, the optimal level of debt when interest rates fall below the growth rate of GDP or more generally the most appropriate fiscal objectives (e.g., Blanchard, 2019; Furman and Summers, 2020; Blanchard, Leandro and Zettelmeyer, 2020). The present paper focuses on a related, but separate, issue: given a set of fiscal objectives, how can we ensure that countries aim for these fiscal objectives, given that in practice (i) rules cannot incorporate ex-ante all relevant contingencies, and (ii) costly penalty cannot be imposed.

The literature on fiscal rules can be split in two polar approaches. The first approach, prevalent in practical settings, consists in stipulating ad-hoc but simple rules to constrain policy, such as in the SGP in its initial formulation. The simplicity of the approach is appealing, because it is transparent and easy to contract on. A major limitation is the arbitrariness and coarseness of such rules, which ultimately severely limit their usefulness (e.g., Eyraud et al., 2018).

The second approach, more prevalent in the academic literature, consists in starting from a fully-specified macro model and derive fiscal rules (or approximations thereof) that maximize welfare in the presence of deficit bias or fiscal externalities. See for instance, Halac and Yared (2014); Yared (2019); Halac and Yared (2019) at the national level and Beetsma and Jensen (2005); Pappa and Vassilatos (2007); Galí and Monacelli (2008) for monetary unions. This approach can help design rules that capture many relevant contingencies, promising higher welfare than the ad-hoc rules used in practice. A worry however is that the assumed model structure may always be too stylized, relative to the complexity of the economy, to actually deliver these welfare gains (e.g. Portes and Wren-Lewis, 2015; Blanchard, Leandro and Zettelmeyer, 2020). As a result, model-based policy prescriptions may always remain too coarse to adequately constrain policy decisions, and they are seldom used in practice.

Our paper proposes a middle route between these two approaches by trying to strike a balance between precision and robustness. The first route is robust in that it does not depend of any particular model, but also a very coarse means of internalizing or correcting pro-deficit biases. The second route is precise thanks to its micro foundations, but it also very sensitive to model mis-specification. By combining elements from both approaches, fiscal targeting promises benefits from both routes: Pareto improvements relative to simple rules, all the while remaining robust to model specification as well as transparent and simple to contract on and monitor. The insight underlying fiscal targeting is that while policy makers may not be able to explicitly write down and agree on one representation of the economy (because of model complexity and/or model uncertainty), decision makers have a much better sense of their own objectives. Consequently, it is arguably easier to agree on a

set of policy objectives than to agree on a fully specified model.

Finally, fiscal discipline targeting complements the recent call by Blanchard, Leandro and Zettelmeyer (2020) to replace fiscal rules with fiscal standards. By relying on a high-level loss function to define and assess fiscal discipline, our paper offers one way in which such “fiscal standards” could be implemented and enforced in an objective, transparent and predictable manner.

2 Illustrative example

In this section we illustrate our flexible fiscal targeting approach for a simple economy. A general treatment will follow in sections 3 and 4.

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, 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.⁹

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.¹⁰ Whenever the policy maker exceeds the \bar{p} limit, the legislator incurs a loss

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

⁹For instance, we do not assume that p is exogenous as arbitrary correlation may exist between p and ε .

¹⁰Such constraints can be justified by deficit bias or financial externalities in a monetary union (e.g Drazen, 2004; Lledó et al., 2017; Eyraud et al., 2018). For clarity of exposition, the constraint here is directly on the policy maker’s control variable. However, in the more general treatment of section 3, the constraint can be on any variable, for instance the debt-to-gdp ratio, which is not perfectly controlled by the policy maker.

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

While we rely on quadratic loss functions in this illustrative example, this is only for clarity of exposition and our general approach detailed in Section 3 is valid for arbitrary loss functions. More generally, \mathcal{L}^y simply captures the policy maker’s own preference, the loss function that the policy maker would consider in the absence of fiscal constraints, while the loss function \mathcal{L}^x captures any additional loss that the legislator would like the policy maker to internalize. The presence of \mathcal{L}^x could be rationalized from multiple perspectives, notably the existence of a deficit bias due to time inconsistency, political cycles, or bureaucratic behavior (e.g., Drazen, 2004), or the existence of financial externalities in a monetary union. In our context where the specific underlying model is unknown however, the loss function \mathcal{L}^x simply captures the legislator’s preferences.¹¹

The problem

In this paper we study how a contract between the legislator and the policy maker can best ensure that both goals —keeping output close to potential and restraining public spending— are achieved. We study this question under two practical restrictions on the contract:

1. The policy maker can violate any contract at a *finite* sanction cost S
2. $h(w)$ is an unknown function and 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 contract 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.

This restriction contrasts our approach with the traditional academic approach where fiscal contracts are designed *given* a specific model for the economy. In our simple model this would include modeling $h(\cdot)$. This implies that the set of feasible contracts that the policy maker and legislator can agree on is much smaller. In fact, without the function $h(\cdot)$, the contracts can only be defined in terms of the policy objectives $y - y^*$, the policy instrument p and the sanction cost S . The fiscal contracts that we see in practice are defined in terms of these three concepts, never in terms of $h(\cdot)$ (e.g., Lledó et al., 2017).

¹¹This is the crucial difference between our approach and a more traditional model-based approach. Instead of starting with a well-specified model and deriving welfare functions from first-principles, we start with the loss functions, i.e., we take as input the different objectives of the policy maker and the legislator.

Fiscal limits

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

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

Unfortunately a fiscal limit contract has two key limitations.

The first limitation of a fiscal limit is its rigidity. In this simple model, the limit $p \leq \bar{p}$ does not take into consideration the policy maker’s cost —the \mathcal{L}^y -cost— of satisfying the fiscal constraint. With a strictly increasing function \mathcal{L}^y , that 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.

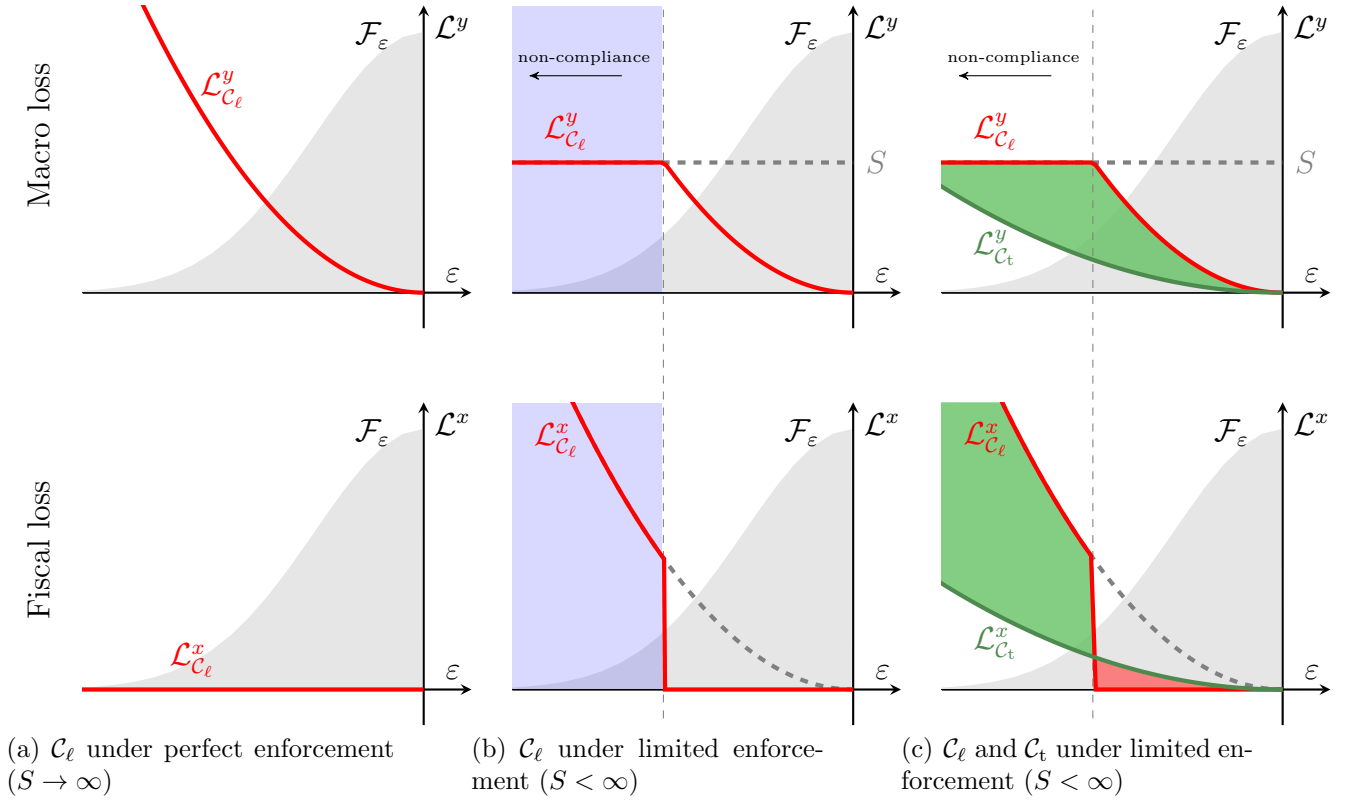
The second limitation is that a fiscal limit contract may end up delivering poor fiscal discipline when the sanction cost S cannot be set at a high enough level, as illustrated in Figure 1b. Since contract compliance is more \mathcal{L}^y -costly in bad times, the policy maker will breach the contract 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 more costly in terms of fiscal discipline, as they lead to larger losses in \mathcal{L}^x (bottom-middle panel). As a result, a fiscal limit contract can deliver poor fiscal discipline if S is too low, i.e., high $\mathbb{E}\mathcal{L}^x$.

To avoid these issues, many have called for flexibilizing limits through the addition of cyclical adjustments or escape clauses. Of course, if the underlying model was known and could be contracted upon, one could use the model to design a more elaborate and more appropriate fiscal rule that takes into account these limitations. One would only have to solve the model and devise a rule that can best trade-off fiscal discipline and flexibility to react to shocks. Our starting assumption is that this is not feasible in practice. The function $h(\cdot)$ cannot be contracted upon, either because $h(\cdot)$ is unknown as there is no universally accepted underlying model of the economy, or because the model $h(\cdot)$ is too complex to write down explicitly. The goal of this paper is to propose a fiscal contract that improves the fiscal limits contract, but does not rely on a specific model for the economy.

Fiscal targets

Instead of policy rules, we propose to provide policy makers with policy objectives, i.e., targets. A fiscal targeting contract (\mathcal{C}_t) stipulates an auxiliary loss function that the policy maker should minimize (or face sanction S). The auxiliary loss function is a weighted average

Figure 1: FISCAL LIMIT (\mathcal{C}_ℓ) VS. FLEXIBLE FISCAL TARGETING (\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 cost for the policy maker (high $\mathbb{E}\mathcal{L}_{\mathcal{C}_\ell}^y$) when \mathcal{L}^y is strictly increasing (here convex). Panel (b): Under limited enforcement ($S < \infty$), the fiscal limit 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 (here convex). Panel (c): the flexibility allowed by fiscal 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 limit contract (green area, bottom panel) with smaller (but more frequent) deviations (red area, bottom panel).

of the losses of the policy maker and the legislator, i.e.

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

such that the policy maker's problem becomes

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

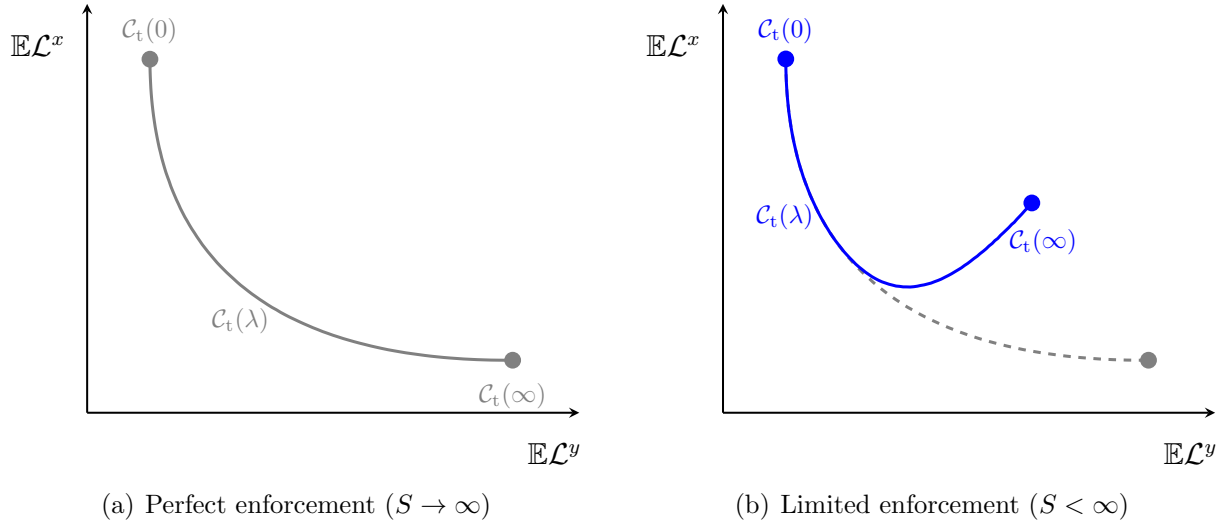
By taking into account the policy maker's own objective fiscal targeting can provide two

benefits: (i) flexibility through constraint relaxation, and (ii) higher fiscal discipline through higher compliance.

First, by explicitly taking into account the policy maker’s own objective, fiscal targeting automatically relaxes the fiscal constraint in recessions, i.e., when \mathcal{L}^y is large. The parameter λ in the auxiliary loss function \mathbf{L} controls this relaxation of the fiscal constraint, allowing both parties to agree ex ante on the desired balance between macro stabilization and fiscal discipline. This is illustrated in Figure 2a, which plots the stabilization–discipline frontier¹² offered by the \mathcal{C}_t contract under perfect contract 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 limit—, and where expectations are taken with respect to \mathcal{F}_ε . Note how fiscal targeting nests fiscal limits as a special case.¹³

Second, under limited enforcement the constraint relaxation can be Pareto improving, reducing both the expected loss of the policy maker *and* the loss of the legislator. This is illustrated in Figure 2b. Starting from the fiscal limit contract ($\mathcal{C}_t(\infty)$) 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 DISCIPLINE–STABILIZATION FRONTIER



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

¹²This frontier is analog to the Taylor curve for monetary policy, see Taylor (1979).

¹³More generally, fiscal targeting can be seen as generalizing fiscal limit contracts. Indeed, the constrained optimization problem implied by a fiscal limit can be represented as the minimization the Lagrangian $L = (y - y^*)^2 + \mu(p - \bar{p})$ where μ is the Lagrange multiplier. Comparing L with the auxiliary loss function \mathbf{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.

Intuitively, flexible targeting can be seen as trading the rare but large and costly deviations from \bar{p} under the fiscal limit contract with smaller (but more frequent) deviations. This is illustrated in Figure 1(c). Faced with a large negative shock, a policy maker under fiscal targeting is allowed to deviate from the fiscal constraint in order to stabilize the economy and avoid large \mathcal{L}^y costs. As a result, the $\mathcal{L}_{\mathcal{C}_t}^y(\varepsilon)$ curve is less steep than $\mathcal{L}_{\mathcal{C}_\ell}^y(\varepsilon)$ curve, and it crosses the non-compliance threshold S later, i.e., for much larger adverse shocks. In other words, non-compliance is less likely under fiscal targeting.

While the policy maker is unambiguously better off under fiscal targeting —lower $\mathbb{E}\mathcal{L}_y$ —, the legislator sees two offsetting effects on its loss function $\mathbb{E}\mathcal{L}_x$. On the one hand, non-compliance is less likely for large adverse shocks and this lowers $\mathbb{E}\mathcal{L}_x$ (green area in the bottom-right panel). On the other hand, fiscal targeting has a cost in terms of flexibility and looser fiscal discipline: deviations from the fiscal constraint are systematic in the face of adverse shocks (red area in the bottom-right panel). Overall, fiscal targeting can offer a Pareto improvement over fiscal limits —lowering both $\mathbb{E}\mathcal{L}^y$ and $\mathbb{E}\mathcal{L}^x$ — if the green area dominates the red area. It turns out that under the assumption that rule violation occur with positive probability, we can always find set of λ 's for which this holds (see Theorem 1 below for a formal proof). Given the high frequency with which fiscal limits have been violated, we think this assumption is a mild one.

A last point to note. While assessing contract compliance may appear more difficult under a \mathcal{C}_t contract given its reliance on an auxiliary loss function, we will show in section 5 that compliance can be assessed using a simple statistical test, thereby preserving the transparency and simplicity of fiscal limits.

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

3 Generic model and assumptions

In this section we outline the generic economic environment in which our study takes place. We allow for arbitrary loss functions and multiple objectives, constraints and fiscal policy instruments.

The policy maker is interested in stabilizing the economy by controlling M_y macroeconomic variables, such as the output gap, inflation, etc. Specifically we impose that at time t the policy maker aims to control the deviation of the variables $y_{i,t+h}$, for $i = 1, \dots, M_y$, from their targets $y_{i,t+h}^*$ over several horizons $h = 0, \dots, H$, where the final horizon H is arbitrary.

The loss that the policy maker incurs is measured by

$$\mathbb{E}_t \mathcal{L}^y(Y_t) , \quad \text{where} \quad Y_t = [y_{i,t+h} - y_{i,t+h}^*]_{i=1,\dots,M_y; h=0,\dots,H} , \quad (1)$$

where $\mathcal{L}^y(\cdot) : \mathbb{R}^{M_y(H+1)} \rightarrow \mathbb{R}^+$ is a strictly increasing function, Y_t is the $M_y(H+1) \times 1$ vector that stacks all target deviations 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 contracts that we propose in this paper do not rely on specific choices for Y_t nor \mathcal{L}^y , and our goal is to provide a general contract structure that has benefits for any specific loss function that the policy maker considers.

To minimize the loss function the policy maker chooses a fiscal policy plan, for instance current and future values of taxes, transfers and spending. We take an abstract approach and simply postulate that the policy has J instruments and the policy plan at time t is defined as

$$p_t = (p_{1,t|t}, \dots, p_{1,t+H|t}, \dots, p_{J,t|t}, \dots, p_{J,t+H|t})' ,$$

where $p_{j,t+h|t}$ is the value of the j th policy instrument announced at time t for period $t+h$. The policy vector p_t has $K = J(H+1)$ entries.

The current and future deviations from target are related to the policy plan via the generic model

$$Y_t = \mathcal{R}^y p_t + \varepsilon_t , \quad \varepsilon_t = h^y(W_t) , \quad (2)$$

where \mathcal{R}^y measures the dynamic causal effects of the policy plan p_t and ε_t captures the effects of all other factors W_t via the function $h^y(W_t)$. We stress that model (2) is generic and most macroeconomic models can be expressed in this way. The linear relationship between p_t and Y_t is made for convenience and can be relaxed.

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

$$x_{i,t+h} \leq \bar{x}_{i,t+h} , \quad i = 1, \dots, M_x , \quad h = 0, \dots, H , \quad (3)$$

where $\bar{x}_{i,t+h}$ is some threshold and there are M_x fiscal variables to control over H horizons.¹⁴ We stack the fiscal variables in $X_t = [x_{i,t+h}]_{i=1,\dots,M_x; h=0,\dots,H}$ and $\bar{X}_t = [\bar{x}_{i,t+h}]_{i=1,\dots,M_x; h=0,\dots,H}$.

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

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

where $\mathcal{L}^x(\cdot) : \mathbb{R}^{M_x(H+1)} \rightarrow \mathbb{R}^+$ is strictly increasing for positive values and $(X_t - \bar{X}_t)_+$ has

¹⁴The thresholds can be time and horizon specific, although in practice fiscal limits are constant across time and horizon, for instance a 3% deficit limit (e.g. Lledó et al., 2017).

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 $\mathcal{L}^x((X_t - \bar{X}_t)_+) = \|X_t - \bar{X}\|_+^\nu$ where ν is capturing the degree of risk aversion of the legislator towards p exceeding \bar{p} . Taking $\nu = 2$ gives the quadratic loss function used in the illustrative example from section 2. Importantly, as with \mathcal{L}^y , the benefits of fiscal targeting do not depend on the specific choice for \mathcal{L}^x as long as it is strictly increasing.

The policy maker affects the fiscal variable X_t with its policy plan p_t through

$$X_t = \mathcal{R}^x p_t + \eta_t, \quad \eta_t = h^x(W_t), \quad (5)$$

where \mathcal{R}^x denotes the causal effects of p_t on X_t and η_t captures other (non-policy) factors. X_t could also include the debt-GDP ratio in which case (5) would capture the law of motion of debt. Equation (5) captures the fact that the policy maker may have only limited control over the fiscal constraints. For instance, η could capture the effect of risk premium shocks that affect the debt servicing cost and thus the debt-GDP ratio. Alternatively, X_t could include the budget deficit, in which case η can capture some mechanical cyclicity of X_t through the automatic stabilizers: in recessions the tax base shrinks and the deficit increases.

In this paper we consider an environment with two limitations on the contract design: (i) limited enforcement, and (ii) model uncertainty:

Assumption 1 (Limited enforcement). *The sanction (in units of $\mathbb{E}_t \mathcal{L}^y$) for non-complying with a fiscal contract is finite and denoted by S .*

Assumption 2 (Model uncertainty). *The functions $h^j(\cdot)$, for $j = x, y$, cannot be written down explicitly and cannot be contracted upon.*

The first assumption allows for limited enforcement of the fiscal contract: if the cost of non-compliance is finite, the policy maker can choose to violate the constraints (3). That assumption captures the fact that in practice it is hard to punish policy makers who choose not to respect a fiscal contract, i.e., the non-compliance sanction S cannot be arbitrarily large.¹⁵

The second limitation captures the fact that in practice the specific model underlying the economy is highly complex and cannot be written down explicitly. As a result, it is difficult to agree on/contract on a specific model structure, and thus to agree on a set of rules provided by a specific model.¹⁶

¹⁵Improvements on the sanction mechanisms are also of great interest but are outside of the scope of this paper.

¹⁶We differ from the Principal-Agent literature (e.g. Bolton and Dewatripont, 2004, Part I) in that there is no asymmetry of information in our model: the policy maker's preferences \mathcal{L}^y and policy choice p_t 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 contract that

With these practical limitations in place we seek to find contracts that are attractive for both the legislator and the policy maker, i.e. contracts that lead to low expected losses for both parties. To formally rank contracts in terms of performance we adopt the following notation. For a given contract \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 contract \mathcal{C} . In general, we define the following criteria for ranking any two contracts.

Definition 1 (Fiscal Discipline). *Given two contracts \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 contracts \mathcal{C}_1 and \mathcal{C}_2 , the \mathcal{C}_2 contract is a Pareto improvement over the \mathcal{C}_1 contract 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 contracts in terms of their ability to induce policy makers to respect the fiscal constraints. Definition 2 allows to rank different contracts in terms of their ability to jointly achieve the policy maker's objectives and the legislator's objective.

4 From fiscal limits to fiscal targets

In this section we present our *flexible fiscal targeting* approach to promote fiscal discipline. We first define formally the fiscal limit contracts commonly used in practice, and we then define our proposed fiscal targeting contracts.

4.1 Fiscal limit contract

The common approach to ensure $X_t \leq \bar{X}_t$ is to directly impose the fiscal limits on the policy maker's program. Formally, we define a Fiscal Limit contract (\mathcal{C}_ℓ) as follows:

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

Under the \mathcal{C}_ℓ contract an optimizing policy maker solves

$$\begin{cases} \min_{p_t} \mathbb{E}_t \mathcal{L}^y(Y_t) & \text{s.t. } \mathbb{E}_t X_t \leq \bar{X}_t & \text{if } \mathbb{E}_t \mathcal{L}^y(Y_t) \leq S \\ \min_{p_t} \mathbb{E}_t \mathcal{L}^y(Y_t) + S & & \text{else} \end{cases} \quad (6)$$

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. The reason that this type of contract is not feasible in our setting is because of Assumption 2: model complexity and the impossibility to contract on $h(\cdot)$. The list of contingencies to take into account would be very long, complex and prone to disagreement, and even likely incomplete because of Knightian uncertainty (Blanchard, Leandro and Zettelmeyer, 2020). In the context of our generic model, this is captured by the fact that the function h cannot be written down and thus cannot be contracted upon.

Clearly, such fiscal limits contract has the benefit of transparency and the vast majority of fiscal rules found in practice can be described by such a contract, 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.

Limitations

Fiscal limit contracts have two related limitations: (i) rigidity which leads to high \mathcal{L}^y -cost (poor macro stabilization), and (ii) poor compliance which leads to high \mathcal{L}^x -cost (poor fiscal discipline). These limitations were discussed in section 2, and they stem from the rigidity of fiscal limits.

Of 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 contract that does not rely on a specific model.

In this practical context, a simple approach would consist in allowing the limit \bar{X}_t to depend on the state of the business cycle Y_t , that is to impose a constraint of the form

$$\mathbb{E}_t X_t \leq \bar{X}_t + \mathcal{A} Y_t \quad \text{with} \quad \mathcal{A} \text{ some constant matrix.} \quad (7)$$

Effectively, this approach amounts to cyclically-adjusting \bar{X}_t , as called by many proposals to reform fiscal rules (e.g. Claey's, Darvas and Leandro, 2016).¹⁷ A limitation of this type of approach however is that it is not clear how one should cyclically-adjust the fiscal limit, i.e., how to choose \mathcal{A} ?

Ultimately, choosing \mathcal{A} is about choosing how to trade-off stabilization vs fiscal discipline. For instance, one could choose \mathcal{A} in order to maximize fiscal discipline, but other choices are possible. In the next section, we will see that flexible fiscal targeting is precisely about contracting on a particular trade-off using an objective and transparent loss-based criterion.

4.2 Fiscal targeting contract

Instead of constraining policy actions by policy rules, we propose to work at a higher-level by appending fiscal objectives to the policy makers' own objectives. To operationalize fiscal targeting in a transparent, quantitative and verifiable way, we rely on an auxiliary loss function that aggregates the different (typically conflicting) objectives.

¹⁷Recent efforts to adjust fiscal rules to the economic outlook and more generally to the country's own context are attempts at flexibilizing fiscal limits along this line and without relying on a specific model. See for instance the reforms of EU Stability and Growth Pact (Eyraud et al., 2018).

Specifically, flexible fiscal targeting stipulates an auxiliary loss function that the policy maker should minimize:¹⁸

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

for some fixed constant $\lambda > 0$, so that a policy maker under fiscal targeting solves

$$\begin{cases} \min_{p_t} \mathbb{E}_t \mathbf{L} & \text{if } \mathbb{E}_t \mathcal{L}^y(Y_t) \leq S \\ \min_{p_t} \mathbb{E}_t \mathcal{L}^y(Y_t) + S & \text{else} \end{cases} . \quad (8)$$

Formally, we define a fiscal targeting contract (\mathcal{C}_t) as follows:

Definition 4 (\mathcal{C}_t contract). *A Fiscal Target contract 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 .*

As we will see below, an important benefit of our approach is that the fiscal targeting contract can be defined for any preferred λ . The next section discusses the attractive properties of the \mathcal{C}_t contract and formally compares the \mathcal{C}_ℓ and \mathcal{C}_t contracts.

4.3 Properties of fiscal targeting

In this section, we discuss two attractive properties of fiscal targeting: (i) the ability to navigate the stabilization–discipline trade-off in a transparent and contractible way and (ii) superior contract performance, i.e., higher compliance in the face of limited enforcement. The third key benefit of fiscal targeting —the ability to transparently assess compliance— is discussed in detail in the next section.

Trading fiscal discipline and macro stabilization

The first direct property of flexible fiscal targeting is that it allows both parties to take a stand (and contract) on the desired balance between fiscal discipline and macro stabilization. Through the choice of λ , both parties can contract on the tightness of the fiscal constraint. Indeed, a contract $\mathcal{C}_t(\lambda)$ nests as special cases (i) the fiscal limits contract when $\lambda \rightarrow \infty$ and (ii) the unconstrained solution when $\lambda = 0$. Clearly, there exists a range of rigidity in the fiscal constraint for λ in between these polar cases.

¹⁸Again 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.

Pareto improving relaxation

By relaxing the rigidity of the fiscal constraint, flexible fiscal targeting can improve compliance in the case of finite sanction costs and thereby deliver higher fiscal discipline. In other words, the \mathcal{C}_t contract can deliver the same fiscal discipline as a fiscal limit contract but at a lower cost to the policy maker: a Pareto improvement. The following theorem establishes the result.

Theorem 1. *Let $\xi_t = (\varepsilon'_t, \eta'_t)'$ and suppose that ξ_t takes values in $\Gamma \subseteq \mathbb{R}^{(M_y + M_x)(H+1)}$. Given assumption 1 and that the set $\{\xi_t \in \Gamma : \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_\ell} + \varepsilon_t) > S\}$ is non-empty, where $p_t^{\mathcal{C}_\ell} \in \arg \min_{p_t \in \mathbb{R}^K : X_t \leq \bar{X}_t} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t)$, we have that there exists a $\bar{\lambda}$ such that*

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^y \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y \quad \text{and} \quad \mathbb{E}_t \mathcal{L}_{\mathcal{C}_t}^x \leq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x \quad \text{for all } \lambda \in [\bar{\lambda}, \infty) \quad (9)$$

The theorem states that there exists a range of values for λ that ensure that the \mathcal{C}_t contract Pareto dominates the \mathcal{C}_ℓ contract. The main assumption is that defaults can happen with positive probability (e.g. $\{\xi_t \in \Gamma : \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_\ell} + \varepsilon_t) > S\}$ is non-empty) for the \mathcal{C}_ℓ contract. 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). In fact, fiscal targeting can be seen as a natural extension of fiscal limit: by relaxing the hard fiscal limit in a transparent loss-based fashion, fiscal targeting can offer higher flexibility and higher fiscal discipline.


5 Evaluating compliance with fiscal targeting

This section discusses how compliance with fiscal targeting can be evaluated without assuming a specific structure for the economy. So far we have established that fiscal targeting can improve upon the current fiscal limits contracts under minimal assumptions, e.g. Theorem 1, but showing that it can in fact be implemented under similarly modest assumptions is equally important from a practical perspective.

Assessing compliance with fiscal targeting amounts to verifying that the policy maker's proposed policy plan, say p_t^0 , minimizes the expected auxiliary loss $\mathbb{E}_t \mathbf{L}(\lambda)$. To do so, we can test whether the necessary condition of optimality $\nabla_{p_t} \mathbb{E}_t \mathbf{L}|_{p_t=p_t^0} = 0$ holds, similarly to a score test or Lagrange multiplier test. The idea is illustrated in Figure 3.

Such gradient test is particularly attractive in our setting because it only requires the estimation of the gradient of $\mathbb{E}_t \mathbf{L}$ under the null, i.e., it only requires the estimation of the gradient at p_t^0 , which can be done without having to agree on one specific model (consistent with Assumption 2).

Figure 3: GRADIENT TESTING WITH MINIMAL ASSUMPTIONS



Notes: Displayed is the fiscal targeting loss function $\mathbb{E}_t \mathcal{L}$, where p_t^{opt} is the optimal fiscal policy and p_t^0 is the proposed policy choice of the policy maker. The gradient based test requires the evaluation of the gradient \mathcal{G}_t at p_t^0 (the slope of the dashed red line), and if $\mathcal{G}_t(p_t^0) \neq 0$ as in the figure, we conclude that $p_t^0 \neq p_t^{\text{opt}}$, in which case the policy maker is in violation of the contract.

To see this, note that the gradient evaluated at p_t^0 can be written as

$$\begin{aligned} \mathcal{G}_t(p_t^0) &= \nabla_{p_t} \mathbb{E}_t \mathcal{L} \big|_{p_t=p_t^0} \\ &= \mathcal{R}^{y'} \nabla_{Y_t} \mathbb{E}_t \mathcal{L}^y(Y_t) \big|_{p_t=p_t^0} + \lambda \mathcal{R}^{x'} \nabla_{X_t} \mathbb{E}_t \mathcal{L}^x((X_t - \bar{X}_t)_+) \big|_{p_t=p_t^0} \end{aligned} \quad (10)$$

where \mathcal{R}^y and \mathcal{R}^x are the causal effects of p_t on Y_t and X_t and $\nabla_{Y_t} \mathbb{E}_t \mathcal{L}^y(Y_t) \big|_{p_t=p_t^0}$ and $\nabla_{X_t} \mathbb{E}_t \mathcal{L}^x((X_t - \bar{X}_t)_+) \big|_{p_t=p_t^0}$ are functions of the density forecasts $f(Y_t, X_t | \mathcal{F}_t, p_t^0)$. We note that in the case of quadratic loss functions these functions become equal to the conditional mean forecasts, but for general loss functions other moments of the forecast density are required.

From (10), we can see that the gradient at p_t^0 depends on only two estimable sufficient statistics: (i) the dynamic causal effects of the policy instruments on the policy objectives and (ii) the density forecasts for the objectives conditional on the proposed policy choice.

First, the causal effects of policy instruments is estimable in an objective manner. In fact, \mathcal{R}^y and \mathcal{R}^x are impulse response functions to shocks to p_t^0 , and \mathcal{R}^y and \mathcal{R}^x 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.¹⁹ Crucially, since the

¹⁹Agreeing on a specific set of impulse responses is easier than agreeing on a full economic model, as different models can lead to similar transmissions of fiscal policy but very different predictions in dimensions unrelated to fiscal policy.

gradient test is based on a necessary condition, it is not necessary to know the *full* matrices \mathcal{R}^y and \mathcal{R}^x to construct a test of non-compliance. Any derivative with respect to one of the element of p_t (or with respect to some linear combination of the elements of p_t) is enough to construct a gradient test of non-compliance. In theory, the compliance test would be most powerful if it could be based on on the full matrices \mathcal{R}^y and \mathcal{R}^x , that is if we could assess the gradient of the auxiliary loss function 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, forecasts densities can be constructed even if the specific model cannot be explicitly written down, reflecting the practical approach to macro forecasting where policy makers combine multiple models (statistical and structural), judgment calls, and instinct into the forecasting process.²⁰ Importantly, the method to estimate both 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 Practical implementation of fiscal targeting

Implementing fiscal 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 contract: (i) the auxiliary loss function \mathbf{L} —the policy objectives—, (ii) a timeline and evaluation procedure for evaluating compliance, and (iii) agree on 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 limits to fiscal targets, we will not discuss this element explicitly. That being said, we note that Theorem 1 implies that fiscal targeting will require less sanctions on average.

²⁰To see this, note that as long as we can interchange the differentiation and integration orders we have for Y (and similarly for X)

$$\nabla_{Y_t} \mathbb{E}_t \mathcal{L}^y(Y_t) |_{p_t=p_t^0} = \int \nabla_{Y_t} \mathcal{L}^y(Y_t) |_{p_t=p_t^0} dF_{\xi_t | \mathcal{F}_t}$$

where $F_{Y_t | \mathcal{F}_t}$ denotes the time- t conditional distribution of Y_t . Thus, if the forecast densities for Y_t and X_t are available then $\nabla_{Y_t} \mathbb{E}_t \mathcal{L}^y(Y_t) |_{p_t=p_t^0}$ and $\nabla_{X_t} \mathbb{E}_t \mathcal{L}^x((X_t - \bar{X}_t)_+) |_{p_t=p_t^0}$ can be easily evaluated.

6.1 Contract set-up

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

$$L = \mathcal{L}^y(Y_t) + \lambda \mathcal{L}^x((X_t - \bar{X})_+)$$

as well as the functional forms for \mathcal{L}^y and \mathcal{L}^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. λ captures how the \mathcal{C}_t contract 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.

Naturally, the two parties —the policy maker and the legislator— have conflicting objectives as the policy maker wants to minimize \mathcal{L}^y while the legislator wants to minimize \mathcal{L}^x . In the appendix, we describe a general approach that allows both parties to agree on a λ that best accommodate different objectives, for instance the legislator would like to control the risk that the budget deficit exceeds a certain value, and the policy maker would like to control the risk that the GDP growth gap falls below a certain value.

6.2 Timing and evaluation of compliance

Compliance with fiscal targeting 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 fiscal targeting using historical data for the EU and its Stability and Growth Pact (SGP). Obviously, fiscal targeting was not officially implemented in the EU, and our exercise merely shows the evaluation of a fictitious fiscal targeting contract at different points in time.

7.1 Contract setup

The loss function

The SGP imposes a 3% ceiling on budget deficits, so that 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:

$$L = \sum_{h=0}^H \mathbb{E}_t(y_{t+h} - y^*)^2 + \lambda \sum_{h=0}^H \mathbb{E}_t(s_{t+h} - \bar{s})_+^2 \quad (11)$$

Since the SGP requires plans for the next 3 years, we will take $H = 3$ years.²¹ In effect, this means that the fiscal targeting contract aims at optimally balancing macro stabilization and the fiscal constraint over an horizon of 3 years.

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 thought experiments to assess compliance with fiscal targeting, as well as an independent agency in charge of estimating the corresponding impulse responses (including estimation uncertainty).²²

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 plan. The forecasts for the budget surplus and the real growth rate for France and Germany are shown in Figure 5.

²¹Naturally, more complicated loss functions are possible—including an additional debt-GDP target as well for instance—, but this section is meant to illustrate the workings of the fiscal targeting framework, and not to draw empirical conclusions regarding real contract compliance.

²²Alternatively, 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.

The forecasts for France are 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 5 shows. The bias is also present for Germany albeit somewhat less pronounced. In order to illustrate our fiscal targeting approach we first bias-adjust the forecasts and remove the horizon specific trend for each country such that at least unconditionally the forecasts are unbiased.²³ 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 Two illustrations

We now illustrate the flexible fiscal targeting contract defined by (11) in two ways. First, we illustrate how one would test non-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.

Testing compliance: France vs. Germany

Figure 4 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 targeting approach to fiscal discipline. Indeed, under fiscal 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 fiscal targeting. In fact, we can characterize the evolution of Germany’s budget surplus over 1998-2020 in terms of a (fictitious) flexible fiscal contract $\mathcal{C}_t(\lambda^{\text{DE}})$, where we compute λ^{DE} by minimizing the sum-of-squared gradient statistics over 1998-2020.²⁴ Based on our estimated λ^{DE} , Figure 7(a) plots the Gradient statistic for Germany over 1998-2020 and shows that we can never reject that Germany was complying with the fictitious fiscal targeting contract $\mathcal{C}_f(\lambda^{\text{DE}})$.

The situation of Germany contrasts with that of France. While the two surpluses moved

²³Clearly more advanced bias adjustment methods can be considered, but for our purpose of illustrating fiscal targeting the simple bias adjustment is sufficient.

²⁴Specifically, $\lambda^{\text{DE}} = \arg \min \sum_t \mathcal{G}_t^{\text{DE}}(p_t^0; \lambda)^2$ where p_t^0 is the policy implemented by Germany at time t .

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.²⁵ In a flexible fiscal targeting contract, equality of treatment across members of a monetary union imply that the same auxiliary loss function L should apply to all countries, i.e., with the same weight λ on the fiscal constraint. Thus, we can reformulate the question as follows: given Germany's fictitious fiscal targeting contract $\mathcal{C}_t(\lambda^{DE})$, can we reject that France was complying with fiscal targeting? If we can, it would mean that France made less of a fiscal effort in respecting the SGP.

Figure 7(b) plots the Gradient statistic for France over 1998-2020. We can see that France violated $\mathcal{C}_t(\lambda^{DE})$ numerous times when Germany did not, meaning that France was doing less of an effort than Germany in satisfying the SGP.²⁶ As shown in Figure 5, 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 fiscal targeting, there are a number of instances where France's violation of the 3% ceiling are tolerated by the $\mathcal{C}_t(\lambda^{DE})$ contract. Most notably, fiscal targeting 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.

A dual viewpoint: measuring fiscal discipline across countries

Once we reject that France complied with a virtual $\mathcal{C}_t(\lambda^{DE})$ contract describing Germany as a fiscal targeter, the dual question to ask is "Which $\mathcal{C}_t(\lambda^{FR})$ contract, i.e., which parameter λ^{FR} , best describes France as a fiscal targeter?". Using again a minimum sum-of-squares criterion, we estimate $\lambda^{FR} = 0.3$ smaller than our estimate $\lambda^{DE} = 2$ 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 as a fiscal targeter according

²⁵This 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>.

²⁶Interestingly, the official surplus forecasts from France suggest otherwise with systematic reduction the deficit (Figure 5, middle-left panel). The reality has been very different however, and the France forecasts have systematically over-predicted the reduction in the deficit. In other words, the 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 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 fiscal targeting.

to (11). 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 8 plots the resulting estimates, ranking countries from lowest fiscal discipline (lowest λ^i) to highest discipline (highest 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 objectives, but France and Belgium fares just as poorly in terms of fiscal discipline, indicating that France and Belgium are not any more fiscally responsible than the southern European countries once the superior economic outlook of France and Belgium is taken into account. 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).

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 rules like the EU Stability and Growth Pact face a difficult trade-off between flexibility and enforceability.

In this paper, we propose to implement fiscal constraints through a flexible fiscal targeting framework, paralleling the flexible inflation targeting in place in most leading central banks. Instead of constraining policy makers with rigid rules, we propose to provide policy makers with a list of mandates—for instance macro stabilization and a low budget deficit—, similar in spirit to the Fed’s dual mandate of price stability and full employment.

Through the use of an auxiliary loss function, fiscal discipline can be enforced in a transparent and objective manner, as assessing compliance 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 inflation targeting. While the design of the SGP was inspired by monetary rules like the 4.5% growth rate for the monetary base (Thygesen et al., 2019), central banks replaced these ad-hoc, rigid and rarely followed monetary rules with forecast targeting. Our paper follows the same idea, as we propose to replace fiscal rules with a forecast-targeting approach to assessing fiscal discipline.

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.
- Beetsma, Roel M.W.J., and Henrik Jensen. 2005. “Monetary and fiscal policy interactions in a micro-founded model of a monetary union.” *Journal of International Economics*, 67(2): 320 – 352.
- 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.
- Chernozhukov, Victor, Han Hong, and Elie Tamer. 2007. “Estimation and Confidence Regions for Parameter Sets in Econometric Models.” *Econometrica*, 75(5): 1243–1284.
- 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.
- 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.
- Kydland, Finn E., and Edward C. Prescott.** 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy*, 85(3): 473–491.
- 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.
- Mishkin, Frederic S.** 2001. *From monetary targeting to inflation targeting*. World Bank Publications.
- Pappa, Evi, and Vangelis Vassilatos.** 2007. “The unbearable tightness of being in a monetary union: Fiscal restrictions and regional stability.” *European Economic Review*, 51(6): 1492 – 1513.
- Portes, Jonathan, and Simon Wren-Lewis.** 2015. “Issues in the Design of Fiscal Policy Rules.” *The Manchester School*, 83(S3): 56–86.
- 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.

- Ramey, Valerie A., and Sarah Zubairy.** 2018. “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data.” *Journal of Political Economy*, 126.
- Romer, Christina D., and David H. Romer.** 2010. “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” *American Economic Review*, 100(3): 763–801.
- Schaechter, Andrea, Tidiane Kinda, Nina T Budina, and Anke Weber.** 2012. “Fiscal rules in response to the crisis-toward the ‘next-generation’ rules: A new dataset.”
- 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.
- Taylor, John B.** 1979. “Estimation and Control of a Macroeconomic Model with Rational Expectations.” *Econometrica*, 47(5): 1267–1286.
- Thygesen, N, R Beetsma, M Bordinon, S Duchêne, and M Szczurek.** 2019. “Assessment of EU fiscal rules with a focus on the six and two-pack legislation.”
- Yared, Pierre.** 2019. “Rising Government Debt: Causes and Solutions for a Decades-Old Trend.” *Journal of Economic Perspectives*, 33(2): 115–40.
- Zubairy, Sarah.** 2014. “On fiscal multipliers: Estimates from a medium scale DSGE model.” *International Economic Review*, 55(1): 169–195.

Appendix

A1: main proofs

Proof of Theorem 1. We observe that for any fixed ε_t and η_t , and $\lambda > 0$ we have

$$\begin{aligned}
\mathcal{L}^y(\mathcal{R}^y p_t^{*,\mathcal{C}_\ell} + \varepsilon_t) &\in \min_{p_t \in \mathbb{R}^K : X_t \leq \bar{X}_t} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t) \\
&= \min_{p_t \in \mathbb{R}^K : X_t \leq \bar{X}_t} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t) + \lambda \mathcal{L}^x((\mathcal{R}^x p_t + \eta_t - \bar{X}_t)_+) \\
&\geq \min_{p_t \in \mathbb{R}^K} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t) + \lambda \mathcal{L}^x((\mathcal{R}^x p_t + \eta_t - \bar{X}_t)_+) \\
&\supseteq \mathcal{L}^y(\mathcal{R}^y p_t^{*,\mathcal{C}_t} + \varepsilon_t) + \lambda \mathcal{L}^x((\mathcal{R}^x p_t^{*,\mathcal{C}_t} + \eta_t - \bar{X}_t)_+) \\
&\geq \mathcal{L}^y(\mathcal{R}^y p_t^{*,\mathcal{C}_t} + \varepsilon_t)
\end{aligned} \tag{12}$$

where $p_t^{*,\mathcal{C}_\ell} \equiv p_t^{*,\mathcal{C}_\ell}(\varepsilon_t, \eta_t) \in \arg \min_{p_t \in \mathbb{R}^K : X_t \leq \bar{X}_t} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t)$ and $p_t^{*,\mathcal{C}_t} \equiv p_t^{*,\mathcal{C}_t}(\varepsilon_t, \eta_t) \in \arg \min_{p_t \in \mathbb{R}^K} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t) + \lambda \mathcal{L}^x((\mathcal{R}^x p_t + \eta_t - \bar{X}_t)_+)$. Next, let $\xi_t = (\varepsilon'_t, \eta'_t)'$ which takes values in $\Gamma \subseteq \mathbb{R}^{(M_y + M_x)(H+1)}$. For the fiscal limits contract we define

$$\mathcal{S}_{\mathcal{C}_\ell} = \{\xi_t \in \Gamma : \mathcal{L}^y(\mathcal{R}^y p_t^{*,\mathcal{C}_\ell} + \varepsilon_t) \leq S\}$$

and

$$p_t^{\mathcal{C}_\ell} = \begin{cases} p_t^{*,\mathcal{C}_\ell} & \text{if } \xi_t \in \mathcal{S}_{\mathcal{C}_\ell} \\ p_t^* & \text{else} \end{cases}$$

where $p_t^* \in \arg \min_{p_t \in \mathbb{R}^K} \mathcal{L}^y(\mathcal{R}^y p_t + \varepsilon_t)$ which is the unconstrained minimizer. Similarly, for the fiscal targeting contract we define

$$\mathcal{S}_{\mathcal{C}_t} = \{\xi_t \in \Gamma : \mathcal{L}^y(\mathcal{R}^y p_t^{*,\mathcal{C}_t} + \varepsilon_t) \leq S\}$$

and

$$p_t^{\mathcal{C}_t} = \begin{cases} p_t^{*,\mathcal{C}_t} & \text{if } \xi_t \in \mathcal{S}_{\mathcal{C}_t} \\ p_t^* & \text{else} \end{cases}$$

Note that for any finite λ (12) implies 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 $\mathcal{O} = \mathcal{S}_{\mathcal{C}_\ell}^\perp \cap \mathcal{S}_{\mathcal{C}_t}$. Next, the loss of the policy maker under the \mathcal{C}_ℓ contract is given by

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y = \int_{\xi_t : \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_\ell} + \varepsilon_t) dF_{\xi_t|\mathcal{F}_t} + \int_{\xi_t : \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} (\mathcal{L}^y(\mathcal{R}^y p_t^* + \varepsilon_t) + S) dF_{\xi_t|\mathcal{F}_t}$$

The loss of the policy maker under the \mathcal{C}_f contract is given by

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^y = \int_{\xi_t : \xi_t \in \mathcal{S}_{\mathcal{C}_t}} \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_t} + \varepsilon_t) dF_{\xi_t|\mathcal{F}_t} + \int_{\xi_t : \xi_t \in \mathcal{S}_{\mathcal{C}_t}^\perp} (\mathcal{L}^y(\mathcal{R}^y p_t^* + \varepsilon_t) + S) dF_{\xi_t|\mathcal{F}_t}.$$

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

$$\begin{aligned}
\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^y &= \int_{\xi_t : \xi_t \in \mathcal{O}} (\mathcal{L}^y(\mathcal{R}^y p_t^* + \varepsilon_t) + S) - \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_t} + \varepsilon_t) dF_{\xi_t|\mathcal{F}_t} \\
&\quad + \int_{\xi_t : \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_\ell} + \varepsilon_t) - \mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_t} + \varepsilon_t) dF_{\xi_t|\mathcal{F}_t}
\end{aligned}$$

The first integral is non-negative as over \mathcal{O} the \mathcal{C}_t contract does not default and hence $\mathcal{L}^y(\mathcal{R}^y p_t^{\mathcal{C}_t} + \varepsilon_t) \leq (\mathcal{L}^y(\mathcal{R}^y p_t^* + \varepsilon_t) + S)$. The second term is also positive by (12). Hence, we have $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^y \geq \mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^y$. Next, for the loss of the legislator we note that under the \mathcal{C}_ℓ contract we have

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x = \int_{\xi_t: \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} \mathcal{L}^x((\mathcal{R}^x p_t^* + \eta_t - \bar{X}_t)_+) dF_{\xi_t|\mathcal{F}_t}$$

and under the \mathcal{C}_t contract

$$\mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^x = \int_{\xi_t: \xi_t \in \mathcal{S}_{\mathcal{C}_t}} \mathcal{L}^x((\mathcal{R}^x p_t^{\mathcal{C}_t} + \eta_t - \bar{X}_t)_+) dF_{\xi_t|\mathcal{F}_t} + \int_{\xi_t: \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}^\perp} \mathcal{L}^x((\mathcal{R}^x p_t^* + \eta_t - \bar{X}_t)_+) dF_{\xi_t|\mathcal{F}_t}$$

Subtracting the losses gives

$$\begin{aligned} \mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^x &= \int_{\xi_t \in \mathcal{O}} \mathcal{L}^x((\mathcal{R}^x p_t^* + \eta_t - \bar{X}_t)_+) - \mathcal{L}^x((\mathcal{R}^x p_t^{\mathcal{C}_t} + \eta_t - \bar{X}_t)_+) dF_{\xi_t|\mathcal{F}_t} \\ &\quad - \int_{\xi_t: \xi_t \in \mathcal{S}_{\mathcal{C}_\ell}} \mathcal{L}^x((\mathcal{R}^x p_t^{\mathcal{C}_t} + \eta_t - \bar{X}_t)_+) dF_{\xi_t|\mathcal{F}_t} \end{aligned}$$

Note that for $\lambda \rightarrow \infty$ we have $\mathcal{O} \rightarrow \emptyset$ and $p_t^{\mathcal{C}_t} \rightarrow p_t^{\mathcal{C}_\ell}$ and thus $\mathbb{E}_t \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E}_t \mathcal{L}_{\mathcal{C}_f}^x \rightarrow 0$. Also, for $\lambda = 0$ we have that $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x \leq 0$ as $p_t^{\mathcal{C}_t} = p_t^*$. So if the gradient is negative for $\lambda \rightarrow \infty$ (e.g. $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x$ approaches zero from above) we know that there is at least one $\bar{\lambda}$ for which $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x \geq 0$ as $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x$ 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 $\mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x < 0$ as increasing λ places more weight on the fiscal objective, hence reducing $\mathcal{L}^x((\mathcal{R}^x p_t^{\mathcal{C}_t} + \eta_t - \bar{X}_t)_+)$. Together, this implies that $\nabla_\lambda (\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x) > 0$. Finally, since $\mathcal{L}^x((\mathcal{R}^x p_t^{\mathcal{C}_t} + \eta_t - \bar{X}_t)_+)$ is continuously decreasing as $\lambda \rightarrow \infty$ we have that $\mathbb{E} \mathcal{L}_{\mathcal{C}_\ell}^x - \mathbb{E} \mathcal{L}_{\mathcal{C}_f}^x \geq 0$ 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.²⁷

The loss functions become

$$\mathbb{E}_t \mathcal{L}_t^y = \mathbb{E}_t \sum_{m=1}^{M_y} \omega_m^y \sum_{h=0}^H \beta_h^y (y_{m,t+h} - y_m^*)^2 \quad \mathbb{E}_t \mathcal{L}_t^x = \mathbb{E}_t \sum_{m=1}^{M_x} \omega_m^x \sum_{h=0}^H \beta_h^x (x_{m,t+h} - x_m^*)_+^2, \quad (13)$$

where the parameters $\omega^j = (\omega_1^j, \dots, \omega_{M_j}^j)'$ and $\beta^j = (\beta_1^j, \dots, \beta_{M_j}^j)'$, for $j = x, y$, allow for different weights on the different macro and fiscal targets. We can conveniently express the fiscal targeting loss function as

$$\begin{aligned} \mathbb{E}_t \mathbf{L} &= \mathbb{E}_t \mathcal{L}_t^y + \lambda \mathbb{E}_t \mathcal{L}_t^x \\ &= \mathbb{E}_t \|\mathcal{W}_y^{1/2} Y_t\|^2 + \lambda \mathbb{E}_t \|\mathcal{W}_x^{1/2} X_t\|_+^2 \end{aligned} \quad (14)$$

²⁷In fact, in the empirical application we rely on the forecasts of the euro area countries conditional on their proposed policy plans. Unfortunately the European commission only provides point forecasts limiting the implementation of the general gradient test.

where $\mathcal{W}_y = \text{diag}(\beta^y \otimes \omega^y)$ and $\mathcal{W}_x = \text{diag}(\beta^x \otimes \omega^x)$. We have defined $\|c\|_+^2 = \sum_{i=1}^{M_x} (c_i)_+^2$ for any $c \in \mathbb{R}^{M_x}$.

To verify whether p_t^0 satisfies the gradient condition, we note that the gradient evaluated at p_t^0 is given by

$$G_t^0 = \mathcal{R}_y' \mathcal{W}_y \mathbb{E}_t Y_t^0 - \lambda \mathcal{R}_x' \mathcal{W}_x \mathbb{E}_t X_t^0, \quad (15)$$

To construct a test statistic based on G_t^0 we need to (a) estimate the dynamic causal effects \mathcal{R}_y and \mathcal{R}_x , and (b) approximate the oracle forecasts $\mathbb{E}_t Y_t^0$ and $\mathbb{E}_t X_t^0$.

For the estimation of dynamic causal effects many methods have been developed in macroeconomics. For instance, recently narrative instrument sequences that extract variations in policy that are unrelated to the variables of interest have become popular, see for instance Romer and Romer (2010) and Ramey and Zubairy (2018). Such quasi-experimental approach is attractive in our setting as it requires minimal modeling assumptions (Ramey, 2019). In our empirical work below we provide the details for such approach, but we stress that any agreed upon set of dynamic causal effects estimates can be used.

In what follows we assume that the dynamic causal effects \mathcal{R}_y and \mathcal{R}_x 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

$$\hat{r} \stackrel{a}{\sim} N(r, \Omega) \quad (16)$$

where $r = (\text{vec}(\mathcal{R}_y)', \text{vec}(\mathcal{R}_x)')'$ 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 (16) implies that we can recover the distribution of the dynamic causal effects using using (16).

Next, we approximate the distribution of the oracle forecasts $\mathbb{E}_t Y_t^0$ and $\mathbb{E}_t X_t^0$. In practice, forecasters typically produce point estimates, say \hat{Y}_t and \hat{X}_t , for the macro and fiscal variables. To test whether the policy maker is fiscally responsible we need the distribution of $\hat{Y}_t - \mathbb{E}_t Y_t^0$ and $\hat{X}_t - \mathbb{E}_t X_t^0$, 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

$$\widehat{W}_t - \mathbb{E}_t W_t^0 \sim F_{W_0}, \quad \text{where} \quad \widehat{W}_t = \begin{bmatrix} \hat{Y}_t \\ \hat{X}_t \end{bmatrix}, \quad \mathbb{E}_t W_t^0 = \begin{bmatrix} \mathbb{E}_t Y_t^0 \\ \mathbb{E}_t X_t^0 \end{bmatrix}. \quad (17)$$

The distributions $\hat{r}_t \stackrel{a}{\sim} N(r, \Omega)$ and F_{W_0} are used to construct tests for fiscal responsibility.

In particular, to test whether the fiscal plan p_t^0 is responsible we compute by simulation methods.

$$\{\widehat{G}_t^{(1)}, \dots, \widehat{G}_t^{(B)}\} \quad \text{where} \quad \widehat{G}_t^{(j)} = \widehat{\mathcal{R}}_y^{(j)'} \mathcal{W}_y \widehat{Y}_t^{(j)} - \lambda \widehat{\mathcal{R}}_x^{(j)'} \mathcal{W}_x \widehat{X}_t^{(j)} \quad \forall j = 1, \dots, B. \quad (18)$$

A simulated draw of the gradient $\widehat{G}_t^{(j)}$ is obtained from the previously defined distributions.

Formally,

$$\hat{r}^{(j)} \stackrel{iid}{\sim} N(\hat{r}, \Omega) \quad \widehat{W}_t^{(j)} = \widehat{W}_t + U^{(j)} \quad U^{(j)} \stackrel{iid}{\sim} F_{W_0} \quad (19)$$

Under the null hypothesis that the proposed fiscal policy plan p_t^0 is optimal, we should have that $G_t^0 = 0$ and we can assess, based on the simulated distribution $\{\widehat{S}_t^{(1)}, \dots, \widehat{S}_t^{(B)}\}$, whether this hypothesis is supported. Specifically, we can study for any desired level of confidence whether the simulated set includes the zero vector. Moreover, we can evaluate whether specific policies, e.g. $p_{k,t}^0$, are optimally set by evaluating whether $\{G_{k,t}^{(j)}, j = 1, \dots, B\}$ includes zero.

A3: A general method to eliciting λ

The thought experiment that underlies our approach is as follows. Suppose that over the last n periods, the policy maker and legislator could have adjusted the policy plan p_t^0 by a fiscal targeting contract. What would be the “optimal” fiscal targeting contract in terms of λ that would meet their objectives.

To set this up, suppose that the policy maker wants to ensure $\mathbb{E}\kappa_y(Y_t) = c_y$, where κ_y is some function of the macro deviations Y_t and $c_y \geq 0$ is a pre-defined constant. The function κ_y may be taken equal to \mathcal{L}^y , but this is not necessary as any other, or possibly multiple criteria can be specified. For instance $P(Y_t < a) = \mathbb{E}\mathbf{1}(Y_t < a) = c_y$ is also possible.²⁸ Similarly, suppose that the policy maker is interested in $\mathbb{E}\kappa_x(X_t - \bar{X}_t) = c_x$. Again for arbitrary, possibly vector valued functions κ_x .

Now suppose that the policy choice p_t^0 for some period $t = 1, \dots, n$ could have been adjusted by $\delta_t(\lambda)$, where $\delta_t(\lambda)$ is based on the fiscal targeting contract. The optimal choice implied by the \mathcal{C}_t contract is given by

$$\delta_t(\lambda) = \arg \min_{\delta_t \in \mathbb{R}^K} \mathbb{E}_t \mathcal{L}^y(Y_t^0 + \mathcal{R}^y \delta_t) + \lambda \mathbb{E}_t \mathcal{L}^x((X_t^0 + \mathcal{R}^x \delta_t - \bar{X}_t)_+)$$

Evaluating $\delta_t(\lambda)$ in practice implies in general that the distribution of $(Y_t^{0'}, X_t^{0'})' | \mathcal{F}_t$ needs to be estimated. For quadratic loss functions only the $\mathbb{E}[(Y_t^{0'}, X_t^{0'})' | \mathcal{F}_t]$ is required. The latter is often publicly available, at least for the euro area.

We stack their objectives in the moment vector

$$g(\lambda) = \mathbb{E} \begin{bmatrix} \mathbb{E}_t \kappa_y(Y_t^0 + \mathcal{R}^y \delta_t(\lambda)) \\ \mathbb{E}_t \kappa_x(X_t^0 + \mathcal{R}^x \delta_t(\lambda) - \bar{X}_t) \end{bmatrix} - \begin{bmatrix} c_y \\ c_x \end{bmatrix}$$

which has empirical counterpart

$$\hat{g}(\lambda) = \frac{1}{n} \sum_{t=1}^n \begin{bmatrix} \mathbb{E}_t \kappa_y(Y_t^0 + \mathcal{R}^y \delta_t(\lambda)) \\ \mathbb{E}_t \kappa_x(X_t^0 + \mathcal{R}^x \delta_t(\lambda) - \bar{X}_t) \end{bmatrix} - \begin{bmatrix} c_y \\ c_x \end{bmatrix}$$


and the minimizing λ is obtained by solving

$$\hat{\lambda} = \arg \min_{\lambda \in \mathbb{R}^+} \|\hat{g}(\lambda)\|^2$$

²⁸Inequality constraints can also be considered, but this will require the set estimation methods for λ , see for instance Chernozhukov, Hong and Tamer (2007).

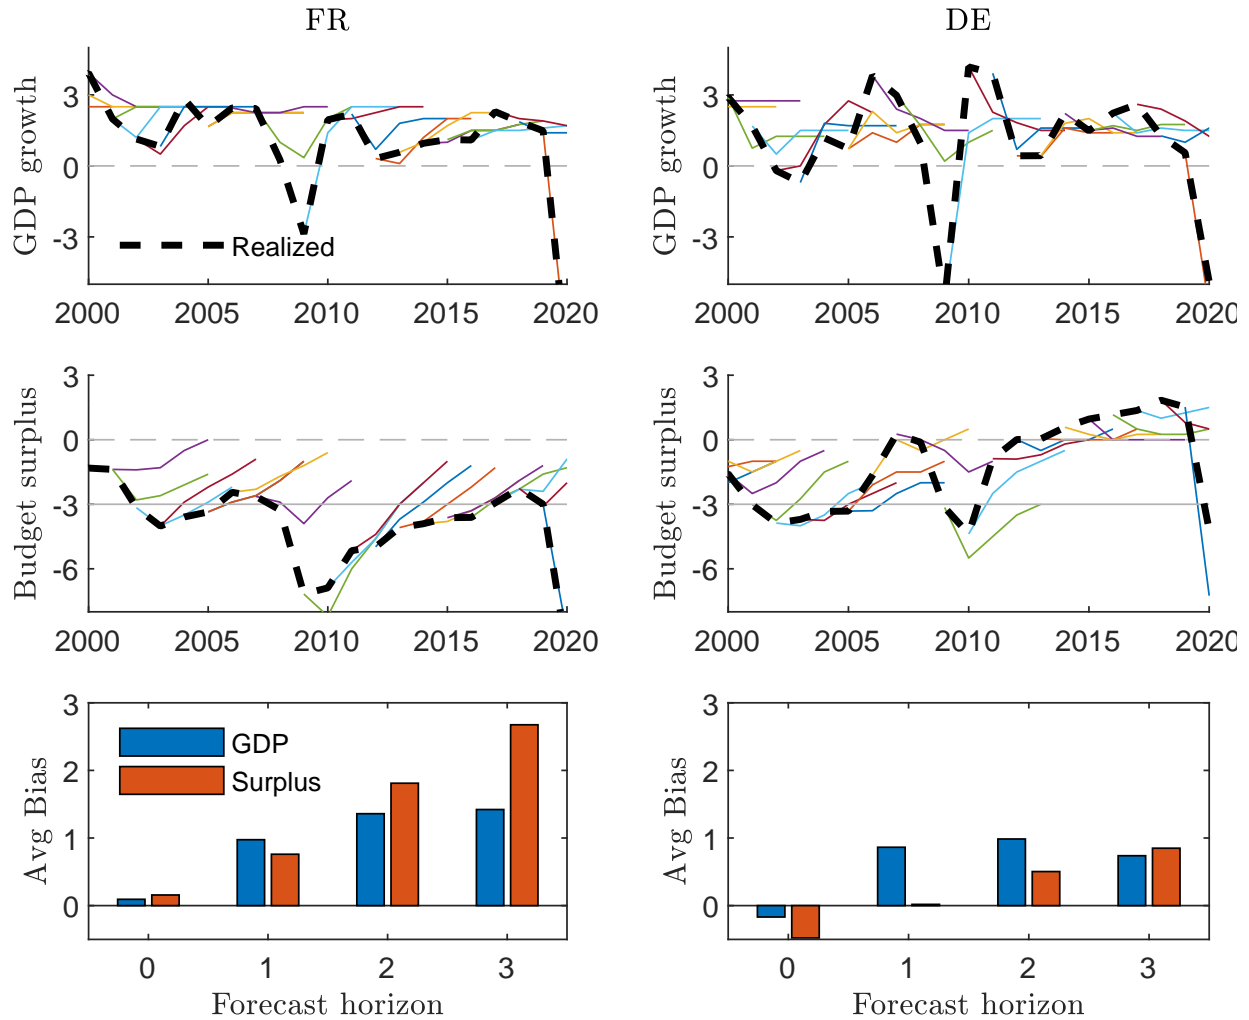
The estimator $\hat{\lambda}$ chooses the λ , based on the \mathcal{C}_f contract, that would have minimized a list of general objectives specified by the legislator and the policy maker.

Figure 4: 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 5: 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 6: 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 7: Gradient test, 1998-2020



Notes: Gradient statistic with 95 confidence band based on $\lambda^{DE} = 1.2$ for France in panel (a) and Germany in panel (b). A non-zero value for the Gradient test indicates non-compliance with the flexible fiscal contract $C_f(\lambda^{DE})$.

Figure 8: FISCAL DISCIPLINE ACROSS THE EU



Notes: Implied fiscal targeting contract for different EU countries. Each bar depicts the preference parameter λ estimated to minimize the sum-of-squares of the gradient statistic implied by the loss function (11).