Journal of Monetary Economics 118 (2021) 54–86

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Optimal policy with general signal extraction[☆]

Esther Hauk[a], Andrea Lanteri[b], Albert Marcet[c,*]

[a] IAE-CSIC, MOVE, and Barcelona GSE, Institute of Economic Analysis, Spanish Research Council, Campus UAB, Bellaterra 08193, Spain
[b] Duke University and CEPR, Duke University, 213 Social Sciences, Campus Drive, Durham, NC 27708, USA
[c] CREI, ICREA, Barcelona GSE, Universitat Pompeu Fabra, and CEPR, CREI, Universitat Pompeu Fabra, Ramon Trias-Fargas 25–27, Barcelona 08005, Spain

A R T I C L E I N F O

Article history:
Received 3 December 2019
Revised 5 January 2021
Accepted 8 January 2021
Available online 16 January 2021

JEL classification:
D80
E62
H21

Keywords:
Optimal policy
Partial information
Calculus of variations
Fiscal policy

A B S T R A C T

Most available results on optimal decisions under partial information are derived under "separation". But this principle does not always hold. We derive a non-standard first order condition of optimality from first principles when signal extraction and optimal policy must be jointly determined. This allows us to solve a model of optimal fiscal policy where separation does not apply. Tax smoothing prevails in normal times, but taxes respond strongly in recessions. This non-linearity arises because signal extraction interacts differently with optimal policy depending on the value of the observed signals. Existing results based on the "separation principle" follow as special cases.

1. Introduction

Inferring the underlying state of the economy is a key practical difficulty in setting macroeconomic policy, as exemplified by this quote.

> "In the policy world, there is a very strong notion that if we only knew the state of the economy today, it would be a simple matter to decide what the policy should be. The notion is that we do not know the state of the system today,

☆ The authors appreciate helpful comments from the Editor (Urban Jermann), two anonymous referees, Manuel Amador, In-Koo Cho, Wouter den Haan, Cosmin Ilut, Kyle Jurado, Jennifer La'O, Paul Levine, Martí Mestieri, Juan Passadore, Joe Pearlman, Julien Prat, Aleh Tsyvinski, Robert Shimer, Jaume Ventura, Mike Woodford, Pierre Yared, and other seminar participants at LSE Macro Work-in-progress, IAE-CSIC Macro work-in-progress, Midwest Macro Spring Meeting 2014, SED Annual Meeting 2014, University of Cambridge, Columbia University, University of Vienna, University of Nottingham, Duke University, Federal Reserve Board, Barcelona GSE Summer Forum 2016, UCL, Adaptive Learning Workshop Bilbao and MACFINROBODS network meeting. Hauk acknowledges financial support from through CICYT project number ECO2015-67171-P (MINECO / FEDER) and from the Spanish Ministry of Science, Innovation and Universities through project PGC2018-097898-B-100 MCIU/AEI/FEDER. Hauk and Marcet have received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 612796, from the government of Catalonia (2017 SGR 1571) and from the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Program for Centers of Excellence in R&D (CEX2019-000915-S). Lanteri acknowledges support and hospitality from MOVE Barcelona and Institut d'Anàlisi Econòmica-CSIC. Marcet acknowledges support from the Axa Foundation; the Excellence Program of Banco de España; and the European Research Council (ERC) under the European Unions Seventh Framework Programme FP7/2007-2013 GA number 324048-(APMPAL) and Horizon2020 GA number 788547 (APMPAL-HET).

* Corresponding author.
E-mail addresses: esther.hauk@iae.csic.es (E. Hauk), andrea.lanteri@duke.edu (A. Lanteri), albert.marcet@crei.cat (A. Marcet).

https://doi.org/10.1016/j.jmoneco.2021.01.002
0304-3932/

and it is all very uncertain and very hazy whether the economy is improving or getting worse or what is happening. Because of that, the notion goes, we are not sure what the policy setting should be today. [...] In the research world, it is just the opposite. The typical presumption is that one knows the state of the system at a point in time. There is nothing hazy or difficult about inferring the state of the system in most models." (James Bullard, interview on *Review of Economic Dynamics*, November 2013)

This difficulty arises because the observed signals about the state of the economy are, in general, endogenous to policy decisions and depend on multiple shocks hitting the economy. In this paper, we address this issue and solve for optimal policy taking into account the uncertainty ("haziness") about the underlying state, accounting for the endogeneity of the available information.

To illustrate a practical case of endogenous signals in optimal policy, consider the fiscal policy response at the beginning of the financial crisis. In 2008–2009, policy-makers observed a large fall in output and employment, but it was unclear whether the recession was due to a shock to productivity, to a demand shock or some combination of both. Nonetheless, policy-makers had to react to the recession implicitly making a guess about the nature of the shock. In that instance, the G20 opted for "stimulus packages", implicitly saying that they believed the recession to be mainly due to a temporary demand shock. Ex-post, we know that this expansionary policy generated large deficits, leading later to significant fiscal adjustments in some countries. The level of aggregate output observed during this period was, obviously, endogenous to the policy decision that was taken. Automatic stabilizers, e.g., income taxes and unemployment benefits, are other leading examples of policies that respond to aggregate endogenous information such as output, while simultaneously influencing the level of output. This paper addresses the question of how to design such instruments optimally under Partial Information (PI).

Generically, we consider a setup where the government observes a signal s (say, output). The signal is determined in equilibrium by a policy variable τ (say, taxes) and a vector of fundamental unobserved shocks A (say, demand and supply shocks). The government chooses τ as a function of its information s, thus it sets $\tau = \mathcal{R}(s)$ for a certain policy function $\mathcal{R}$. A key difficulty is that the optimal choice for $\mathcal{R}$ depends on the density of A conditional on s, denoted $f_{A|s}$ (the signal extraction, or filtering, problem) while at the same time the density $f_{A|s}$ depends on the policy function $\mathcal{R}$. Therefore, the filtering and optimal choice problem have to be determined simultaneously.

To our knowledge, all existing results in the literature on optimal policy with PI introduce timing assumptions and a structure of the economy guaranteeing that separation holds, that is, $f_{A|s}$ at the optimum can be found independently from $\mathcal{R}$.[1] While this literature has lead to many interesting applications, the assumptions needed for separation typically are not satisfied in modern dynamic models. The absence of general results limits the application of PI in today's literature on policy analysis.

In this paper, we consider the general non-linear case without separation (General Signal Extraction), and address the problem from first principles. We use a calculus-of-variation approach to characterize the optimal policy. Specifically, we consider the effects of deviations from the optimal policy on signals and welfare, and impose that the optimal deviation has to be no deviation, for any observed signal. This method gives a general optimality condition, which is non-standard, as it does not set the average derivative of the objective function to zero. Instead, the derivative at each realization is weighted by an endogenous kernel that accounts for the effect of policy choices on signals. We then provide an easily applicable algorithm that solves this non-linear optimality condition, jointly delivering the optimal policy and the optimal filter at each value of the observed signal. We also show how previous results in the literature arise as special cases of our approach.

We apply this method in a model of optimal fiscal policy and show that lack of separation leads to important non-linearities in optimal policy, especially when a deep recession is possible, and the economy is close to the top of the Laffer curve. Our leading example is a two-period version of the standard fiscal policy model of Lucas and Stokey (1983). We introduce two shocks A (to demand and supply). To make the issue of hidden information relevant we assume incomplete markets as in Aiyagari et al. (2002). We solve for optimal Ramsey taxation τ under the assumption that the government does not observe the realizations of the shocks, but only an endogenous signal, such as output or hours worked, s.

We argue that PI has important consequences for the design of optimal policy and for our understanding of real-world policy decisions. In particular, optimal taxes can be fairly unresponsive to signals in normal times, thus providing a novel rationale for tax smoothing, but highly responsive to signals when the government fears that the worst realizations of the shocks may have materialized. This non-linearity arises because the signal extraction problem interacts in a different way with optimal policy depending on the range of observed signals. In particular, we find that this non-linearity is stronger when government spending is high and the economy is close to the top of the Laffer curve (or government debt is high), as was arguably the case in some European economies in the Great Recession.[2] Several European governments made large

[1] Svensson and Woodford (2004) consider a case where separation fails but they can show that a "certainty equivalence" result holds. This allows them to solve the signal extraction and optimal choice problems sequentially. See Section 2 for more details on this literature. Section 6 derives the result of Svensson and Woodford (2004) from our general optimality condition for the special case of a linear model.

[2] Other sources of non-linearities, such as the zero lower bound in monetary models, are likely to induce similar results in future applications of our method.

fiscal adjustments a few years after the start of the downturn. This delay lead to a large debt accumulation which eventually amplified the fiscal adjustments needed. Our model provides a framework where this delay in increasing taxes could be a feature of optimal decision making.

To highlight our contribution relative to the literature based on separation, we compare our results to a model with exogenous signals, and show that non-linearities are negligible in that case. We then show that General Signal Extraction is present in most dynamic models under rational expectations. To this aim, we present several versions of our model to demonstrate how separation fails generically. In our main example, we model the private sector as a representative private agent with Full Information (FI) about the shocks. Separation fails because the income tax rate τ is determined simultaneously with the signal s, which is the level of hours worked. We then show that separation fails generically in models with heterogeneous agents who have less information than the government and under no simultaneity.

In our second example, we consider heterogeneous agents facing idiosyncratic shocks. Different from our main example, private agents have "less" information than the government. Nonetheless, we show how to cast this model into our general framework. In our third example, taxes are predetermined, and labor supply involves a dynamic choice because of learning by doing. Despite the lack of simultaneity between policy actions and signals, we obtain similar results.

Finally, we build an infinite-horizon model with non-contingent debt. The short-run response of taxes to underlying shocks in most periods is low relative to the case of Full Information. Taxes do not adjust immediately when a negative shock to the tax base. Hence, the high-frequency response of taxes to shocks features tax smoothing. However, a delayed response to shocks may lead to accumulation of large debt over time, eventually leading to larger fiscal adjustments. Therefore, policy implications are different from the case of market completeness and FI (Lucas and Stokey, 1983), since perfect tax smoothing is not possible, and from the case of market incompleteness and FI (Aiyagari et al., 2002), where the delayed response would never occur.

We emphasize that the technique that we develop can be applied to many interesting open questions in several fields beyond macroeconomics. For example, the unobserved shocks A may represent the planner (or a player)'s ignorance about the structure of the economy. Our results would also apply to cases where the government has incomplete information about agents' types or about agents' expectations.

The remainder of the paper is organized as follows. We discuss the related literature is discussed in Section 2. Section 3 contains the main theoretical contribution. Section 4 introduces our two-period model of optimal fiscal policy. In Section 5 we apply our theoretical results to the fiscal policy problem. Section 6 shows how to derive previous results as special cases and considers our fiscal policy model with exogenous signals. Section 7 presents the model with heterogeneous agents. Section introduces a model with predetermined taxes and dynamic labor choice. Section 9 presents the infinite-horizon model with non-contingent debt. Section 10 concludes.

2. Related literature

PI and signal extraction were often present in the early papers on dynamic models with Rational Expectations. Signal extraction with an *exogenous* signal is well understood; it goes as far back as Muth (1960). Typically, it just requires a routine application of the Kalman filter. Because the signal extraction problem is solved independently of policy choices, it is said that a "separation principle" between signal extraction and optimization applies.

A large literature considers competitive agents who use prices or other aggregate endogenous variables as signals of unknown information, as early as in the seminal paper by Lucas (1972), or in Townsend (1983), which analyzes beliefs about other agents' information. For more recent applications, see, for instance, Nimark (2008a), who studies firm pricing decision with imperfect information, or Angeletos and Pavan (2009), Angeletos et al. (2015), and Melosi (2017), who focus on the role of macroeconomic policy. In this literature, the choices of competitive agents do not interact with price determination (in the agents' mind), therefore the filtering problem can be solved using standard techniques.

Few papers have studied optimal policy when signals are endogenous. Pearlman (1992) and Svensson and Woodford (2003) consider linear Gaussian models where the policy-maker and the private sector have the same information set. In other words, information is partial but symmetric. In this case, they show that the "separation principle" continues to hold.[3] Closest to our work is Svensson and Woodford (2004), who consider optimal policy in a non-microfounded linear Rational Expectations model, where the government's information set is a subset of the private sector's information set. They show that, even though the "separation principle" fails because of asymmetric information, there is a suitable modification of the standard Kalman filter that works in the case of linearity and additively separable shocks. Moreover, optimal policy under discretion has the "certainty equivalence" property: under PI the government applies the FI policy to its best estimate of the state.[4]

Our contribution is to consider a fully microfounded optimal policy model and to provide a solution to the General Signal Extraction and optimization problem, when the distribution of the signal depends on government policy (or, more generally,

[3] Baxter et al. (2007) and Baxter et al. (2011) derive an "endogenous Kalman filter" for these cases, which is equivalent to the solution of a standard Kalman filter of a parallel problem where all the states and signals are fully exogenous.

[4] Aoki (2003) applies these results to optimal monetary policy with noisy indicators on output and inflation. Nimark (2008b) applies them to a problem of monetary policy where the central bank uses data from the yield curve knowing that the chosen policy affects the very same data. Relatedly, Morris and Shin (2018) analyze the optimal weight on an endogenous signal in a linear policy rule.

on the actions of a Stackelberg leader). In the general case, separation does not apply and the strict linearity requirements of Svensson and Woodford (2004) do not hold. We show cases where a linear approximation can be misleading.[5] We provide an example that appears to be amenable to linearisation but, as we show, the correct solution is highly non-linear in nature. Similar to Amador and Weill (2010), in our main example we assume that all shocks are "fundamental", i.e., they affect preferences or technology. We also show how previous results based on some shocks being "noise", or measurement error, arise as special cases.

The effect of policy choices on information extraction is also considered in models of learning about unobserved variables or parameters (armed-bandit problems), such as Prescott (1972) and Kiefer and Nyarko (1989). For some applications to dynamic macro policy see Wieland (2000a), Wieland (2000b) and Ellison and Valla (2001). Van Nieuwerburgh and Veldkamp (2006) use a similar learning framework with non-linearities to explain business-cycle asymmetries.[6] These papers avoid the issue of General Signal Extraction by making special assumptions. For instance, signals are predetermined with respect to policy, independent of future policies, because of the absence of endogenous state variables. We argue that in most modern dynamic economic models the issue of General Signal Extraction that we consider would arise generically.[7]

Another related strand of literature is that on robust control. Hansen and Sargent (2012) study Ramsey-optimal policy with ambiguity aversion and find that this leads to violations of "certainty equivalence" even in linear-quadratic setups. In their setup, the optimal solution is found by, first, solving for the optimal policy for each possible value of the state, then the policy is chosen assuming the worst possible state. Therefore, the issue of General Signal Extraction is avoided. Adam (2004) shows that the min-max criterion arises from a sequence of planners that have expected utility and are increasingly risk-averse. Our setup is available for *any* level of risk aversion.[8]

3. Optimal control with general signal extraction

In this section we present the problem of "Optimal Control under General Signal Extraction (GSE)" and derive optimality conditions that can be readily used to find a solution.

Consider a setup where the government observes a signal s. The signal is determined in equilibrium by a policy variable τ and a vector of fundamental unobserved shocks A. The government chooses τ as a function of its information s, thus it sets $\tau = \mathcal{R}(s)$ for a certain policy function $\mathcal{R}$. A key difficulty is that the optimal choice for $\mathcal{R}$ depends on the density of A conditional on s, denoted $f_{A|s}$ (the signal extraction, or filtering, problem), while at the same time the density $f_{A|s}$ depends on the policy function $\mathcal{R}$. Therefore, the filtering and optimal choice problem have to be determined simultaneously. We call this Optimal Control under General Signal Extraction (GSE). Most papers in the literature consider setups where the filtering and optimal choice problem are "separated", but as we show with several applications in later sections (see, in particular, Section and Appendix C.4), separation only holds for very special cases.

3.1. Optimization under GSE

Consider a planner/government who chooses a policy variable $\tau \in \mathcal{T} \subset R$, conditional on observing an endogenous signal $s \in \mathcal{S} \subset R$, where $\mathcal{T}$ and $\mathcal{S}$ denote the set of possible policy actions and signals.[9] The planner's objective is to maximize $E[W(\tau, s, A)]$ for a given payoff function $W : \mathcal{T} \times \mathcal{S} \times \Phi \to R$ and random variables $A \in \Phi \subset R^k$ with a given distribution F_A.[10]

The policy variable τ maps into endogenous signals through the following equation

$$s = h(\tau, A) \text{ a.s. in } F_A \tag{1}$$

for a known function $h : \mathcal{T} \times \Phi \to \mathcal{S}$. The government has to choose τ given an observation on s, without observing the value of A. The government knows the fundamentals W, h, F_A, $\mathcal{T}$, $\mathcal{S}$ and equation (1).

Optimal behavior under uncertainty implies that the government chooses a policy contingent on the observed variable s. Therefore, the government's problem is to set policy actions according to a policy function $\mathcal{R} : \mathcal{S} \to \mathcal{T}$

$$\tau = \mathcal{R}(s) \tag{2}$$

[5] Optimal non-linear policies have been found in the literature but for totally different reasons. Swanson (2006) obtains a non-linear policy when he relaxes the assumption of normality in the linear model with separable shocks. He considers a model where the "separation principle" applies. The non-linearity results entirely from Bayesian updating on the a priori non-Gaussian shocks.

[6] For an application to monopoly behavior see Mirman et al. (1993).

[7] In Appendix C.4 we provide a modification of our model in Section 4 that restores the simplifications in this literature and that may help clarify this issue.

[8] The literature on optimal contracts under private information and incentive compatibility constraints (or the "New Dynamic Public Finance" as in Kocherlakota (2010) is perhaps less related to our work. This literature studies setups where private information is revealed in equilibrium. As we show in Section 3.4, endogenous filtering is not an issue when private information is revealed. On the other hand, this literature assumes that the policy function $\mathcal{R}$ depends on individual choices and, thus, agents act strategically. We leave an application of our results to this interesting case for future work.

[9] Focusing on the univariate case simplifies the analysis, we leave the multivariate case for future work.

[10] F_A represents the governments' perception about A, it may differ from the true distribution of A.

to maximize its objective, in other words the government solves the following problem of **Optimal Control under GSE**

$$\max_{\{\mathcal{R}:\mathcal{S}\to\mathcal{T}\}} E[W(\tau,s,A)] \tag{3}$$
$$\text{s.t.} \quad (1), (2).$$

where it is understood that τ, s in the objective function are evaluated at the values that solve (1) for each A and the chosen $\mathcal{R}$.

The optimality conditions for this problem turn out to be non-standard. To see why this happens, rewrite the objective function as follows

$$\int E[W(\tau,s,A)|s]\; f_s(s)\; ds. \tag{4}$$

where f_s is the density of s. Applying a "standard recipe" to derive first order conditions in stochastic discrete time models we would find the following condition:[11]

$$E[W_\tau + W_s h_\tau |s] = 0 \quad \text{for all } s. \tag{5}$$

Notice that $W_\tau + W_s h_\tau = 0$ is the FI optimality condition for a Stackelberg leader problem that chooses τ taking into account the follower's action s satisfies (1) for a known A. According to the recipe, equation (5) averages this first order condition over values of A that are consistent with the observed s.

When f_s is exogenous to the policy choice this formula is correct. However, in the general case $f_s = \int_\Phi f_{s|A}\, f_A\, dA$ is endogenous to $\mathcal{R}$, (5) ignores this dependence and, therefore, it does *not* hold in the optimal solution. Somehow the fact that f_s depends on $\mathcal{R}$ should play a role when taking derivatives of (4). Economically, the problem is that the policy choice influences the values of the signal that actually materialize and the optimal policy choice should take this into account. In other words, in general, the optimization and filtering problems cannot be separated.

3.2. Assumptions, conditions on $\mathcal{R}$ and outcomes $S(\mathcal{R}+\alpha\delta, A)$.

We now state some assumptions on the fundamentals.

Assumption 1. $W(\cdot, A)$ is continuous, $h(\cdot, A)$ differentiable for all $A \in \Phi$.

Assumption 2. The sets $\mathcal{T}, \mathcal{S}$ are bounded closed intervals, $\mathcal{T} = [\underline{\tau}, \overline{\tau}]$, $\mathcal{S} = [\underline{s}, \overline{s}]$, and Φ is compact.

We will use the following notation. Let $S(\mathcal{R}, A)$ be the set of observable values of s induced by the shock A and a policy $\mathcal{R}$. Formally, for a given A and $\mathcal{R}$, letting H be given by

$$H(s, A; \mathcal{R}) \equiv s - h(\mathcal{R}(s), A), \tag{6}$$

elements $s \in S(\mathcal{R}, A)$ satisfy $H(s, A; \mathcal{R}) = 0$. We refer to $S(\mathcal{R}, A)$ sometimes as "the outcome", as it gives the realisations of the signal for a given policy function $\mathcal{R}$, sometimes we refer to $S(\mathcal{R}, A)$ as "the equilibrium" as the h function is often derived from equilibrium conditions.

The policy variables that are realized for each $\mathcal{R}$ and A are then given by $T(\mathcal{R}, A) = \mathcal{R}(S(\mathcal{R}, A))$. Notice the distinction between the objects S, T and $\mathcal{R}$: the latter is a function of a real variable s while S and T are functionals mapping $\mathcal{R}$ and the realizations of the shocks into $\mathcal{S}$, $\mathcal{T}$.

In several proofs we consider variations of the solution $\mathcal{R}^* + \alpha\delta$ according to a function $\delta : \mathcal{S} \to R$. To guarantee that the outcomes of a given variation $S(\mathcal{R}^* + \alpha\delta, A)$ are well behaved for small α we need the following requirements on h:

Assumption 3. W is differentiable everywhere with respect to (τ, s), with bounded partial derivatives $|W_\tau|, |W_s| < Q$. Also, h_τ is Lipschitz continuous: $\left|\frac{h_\tau(s,A) - h_\tau(s',A)}{s-s'}\right| \leq Q^L$ uniformly on $\mathcal{S} \times \Phi$ for some constant $Q^L < \infty$.[12]

Obviously this implies $|h_\tau| < Q$ for some constant $Q < \infty$.

The first order condition (9) allows for a non-differentiable optimal solution $\mathcal{R}^*$. Indeed, we will see that non-differentiable $\mathcal{R}^*$ arises in our computations. To guarantee that the objective function is well defined for variations $\mathcal{R}^* + \alpha\delta$, we need that non-differentiabilities of $\mathcal{R}^*$ occur with probability zero. We make the following assumption for this purpose.

Assumption 4. Given any pair $(s, \tau) \in \mathcal{S} \times \mathcal{T}$, the set of realizations $\{A : s = h(\tau, A)\}$ has probability zero.

This is easy to ensure if at least one element in the vector of shocks A has a continuous density. We impose two Conditions on $\mathcal{R}$ ensuring that the outcomes $S(\mathcal{R} + \alpha\delta, A)$ are well defined.[13]

[11] More precisely, we take the "standard recipe" to be: *i)* find the derivative of the objective function (or Lagrangean) with respect to a choice variable under certainty, *ii)* take the expectation of this derivative conditional on the information available at the time the variable is chosen and, *iii)* set this conditional expectation to zero.

[12] Obviously Lipschitz continuity of h_τ is guaranteed if $h_{\tau\tau}$ exists everywhere and is uniformly bounded.

[13] Perhaps one can prove that Conditions 1 and 2 hold at the optimum under Assumptions 1–4 or further assumptions. We leave this for future research.

Denote the support of possible outcomes for a certain policy function $\mathcal{R}$ as

$$\mathcal{S}_{\mathcal{R}} = \{s \in \mathcal{S} : s = h(\mathcal{R}(s), A) \text{ for some } A \in \Phi\} = \bigcup_{A \in \Phi} S(\mathcal{R}, A)$$

Condition 1. $\mathcal{R}$ is absolutely continuous in $\mathcal{S}_{\mathcal{R}}$. Also, where $\mathcal{R}'$ exists $\left|\mathcal{R}'(s)\right| < K^L$ for some uniform constant $K^L < \infty$.

One could equivalently state that $\mathcal{R}$ is a Lipschitz function with a Lipschitz constant K^L. Condition 1 will be naturally satisfied at the optimum in most models.

We wish to ignore issues of multiplicity of equilibria in this paper, so we consider a planner constrained to choosing policies for which $S(\mathcal{R}, A)$ is a singleton. This could be justified by appealing to a principle that "good policy" avoids multiple equilibria, or to the principle that a good research strategy starts with the simpler case.[14] For this reason we now state a condition that is nearly equivalent to imposing uniqueness of outcomes.

Condition 2. There is an $\overline{\varepsilon} > 0$ such that, for each $s \in \mathcal{S}_{\mathcal{R}}$ where $\mathcal{R}'$ exists then

$$h_\tau(\mathcal{R}(s), A)\mathcal{R}'(s) \leq 1 - \overline{\varepsilon} \tag{7}$$

at all A such that $s \in S(\mathcal{R}, A)$, $\overline{\varepsilon}$ uniform over s, A.

This Condition guarantees that $H = 0$ has a well-conditioned solution, ensuring uniqueness in the following Lemma. Condition 2 is also used in the optimality condition in Proposition 2 and existence result in Proposition 5. Conditions 1,2 can be easily checked at a computed solution, as we show in Appendix A.6.

The following lemma ensures uniqess and existence of solutions, thus guaranteeing that the objective function in (3) is well defined at the considered policies.

Lemma 1. *Assume Assumptions 1,2,4 and consider an $\mathcal{R}$ that satisfies Conditions 1–2. Then $S(\mathcal{R}, A)$ is a singleton almost surely. Assume in addition Assumption 3 holds and the outcomes of $\mathcal{R}$ are interior, that is $\mathcal{S}_{\mathcal{R}} \subset int(\mathcal{S})$ and $\mathcal{R}(\mathcal{S}_{\mathcal{R}}) \subset int(\mathcal{T})$. Fix a variation $\delta : \mathcal{S} \to R$ uniformly bounded, differentiable everywhere with a uniformly bounded derivative. Then the set of outcomes $S(\mathcal{R} + \alpha\delta, A)$ is a singleton for all $\alpha \in R$ small enough a.s. Furthermore $S(\mathcal{R} + \alpha\delta, A) \to S(\mathcal{R}, A)$ a.s. $\alpha \to 0$*

The "a.s." statement is over the distribution of A.[15] Lemma 1 is proved in detail in Appendix A.1.[16] We now provide a sketch of the proof.

Proof of Lemma 1.

Denote the set of realizations that give rise to non-differentiable outcomes as

$$\Phi_{\mathcal{R}}^{ND} = \{A \in \Phi : \mathcal{R} \text{ non-differentiable at some } s \in S(\mathcal{R}, A)\}.$$

Assumption 4 guarantees that $Prob(\Phi_{\mathcal{R}}^{ND}) = 0$.

Consider an absolutely continuous function $g : X \to X$ for a compact set X. Two properties of such functions are well known: *i)* g has at least one fixed point (Brouwer's fixed point theorem), *ii)* If in addition $g'(x_f) < 1$ at all fixed points $x_f = g(x_f)$ then x_f is unique.[17]

Fix $A \notin \Phi_{\mathcal{R}}^{ND}$ and δ. Consider α small. Clearly $S(\mathcal{R} + \alpha\delta, A)$ is equal to the set of fixed points of the map $g(.) \equiv h((\mathcal{R} + \alpha\delta)(\cdot), A)$. Interiority and Condition 1 guarantee that $h((\mathcal{R} + \alpha\delta)(\cdot), A)$ maps $\mathcal{S}$ into itself, therefore $S(\mathcal{R} + \alpha\delta, A)$ is non-empty due to property *i)*. Condition 2 and property *ii)* guarantee that $S(\mathcal{R}, A)$ is unique. The uniform bounds on h_τ, δ and δ' and Lipschitz continuity of h_τ guarantee that there is an $\varepsilon > 0$ such that $\frac{dh((\mathcal{R}+\alpha\delta)(s),A)}{ds}$ is close to $\frac{dh(\mathcal{R}(S(\mathcal{R},A)),A)}{ds}$ for all $s \in (S(\mathcal{R}, A) - \varepsilon, S(\mathcal{R}, A) + \varepsilon)$ and α small enough.

A bounded δ implies that for small α all fixed points of $h((\mathcal{R} + \alpha\delta)(\cdot), A)$ (that is, all elements in $S(\mathcal{R} + \alpha\delta, A)$) are in the interval $(S(\mathcal{R}, A) - \varepsilon, S(\mathcal{R}, A) + \varepsilon)$. Therefore, Condition 2 and the way ε is chosen implies that $\frac{dh((\mathcal{R}+\alpha\delta)(s),A)}{ds} < 1$ at all $s \in S(\mathcal{R} + \alpha\delta, A)$. Property *ii)* then implies that $S(\mathcal{R} + \alpha\delta, A)$ is a singleton and it converges to $S(\mathcal{R}, A)$ as $\alpha \to 0$. □

3.3. General formulation

Let $\mathcal{F}$ be the value of the objective function for a given $\mathcal{R}$, that is[18]

$$\mathcal{F}(\mathcal{R}) \equiv E[W(T(\mathcal{R}, A),\ S(\mathcal{R}, A),\ A)].$$

---

[14] Multiplicity of equilibria could be dealt with at the cost of having to introduce selection criteria or randomization in the model, we leave this for future research.

[15] Notice that even though the calculus of variations often uses indicator functions δ, these are ruled out in the Lemma by the differentiability requirement on δ. In fact the result in this Lemma would break down with indicator functions because outcomes could fail to exist, that is $S(\mathcal{R} + \alpha\delta, A)$ could be empty with positive probability for any α, continuity of δ ensures that $S(\mathcal{R} + \alpha\delta, A)$ is non-empty. This is why in our proofs and even when we check second order conditions in practice, we always work with continuous δ's.

[16] Although perhaps only a curiosity, we point out that the uniform bound in Condition 2 is stronger than needed for this Lemma, as it also holds if $h_\tau(\mathcal{R}(s), A)\mathcal{R}'(s) < 1$ for all $s \in \mathcal{S}_{\mathcal{R}}$. The stronger Condition 2 is needed in later results only.

[17] Property *ii) is* easy to see with a simple graph. It has been encountered in economics, for example in the old literature on deterministic growth with increasing returns to scale where the law of motion of capital g typically has three steady states, an unstable one with $g' > 1$ and two stable ones with $g' < 1$, the lower of which is a poverty trap. A detailed proof of *ii)* is available upon request.

[18] Notice that $\mathcal{F}$ maps the space of functions into the real line R.

We can now re-state the Optimal Control with GSE problem as

$$\max_{\{\mathcal{R}:\mathcal{S}\to\mathcal{T}\}} \mathcal{F}(\mathcal{R}) \quad \text{s.t. } \mathcal{R}\in\mathcal{E} \tag{8}$$

where $\mathcal{E}$ is a pre-specified set of admissible policy functions. Restriction (1) is left implicit but guaranteed to hold in this formulation, as it is embedded in the definition of $S(\mathcal{R}, A)$.

The expectation defining $\mathcal{F}$ integrates over A. The set $\mathcal{E}$ is meant to be a large set that includes most policy functions so that the optimum is interior, also it should be such that a solution to the optimization problem exists. $\mathcal{E}$ has to be restricted so that $S(\mathcal{R}, A)$ is a singleton a.s. for all $\mathcal{R}\in\mathcal{E}$, guaranteeing that the expectation that defines $\mathcal{F}(\mathcal{R})$ is well defined. Typically $\mathcal{E}$ will exclude policy functions with large derivatives at some points. We discuss later in more detail the choice for $\mathcal{E}$.

We denote the solution to Problem (8) by $\mathcal{R}^*$.

3.4. Apparent partial information: Revelation and invertibility

In some cases the government can still implement the FI policy even if it does not observe all the shocks. This occurs whenever the government can learn the true state of the economy A from observing the signal s. In other words, whenever the government can invert the observed s, τ to find out the realised A.

To formally define Invertibility, let $(\tau^{FI}, s^{FI}) : \Phi \to \mathcal{T}\times\mathcal{S}$ be the FI solution, namely the solution that arises if the government observes A, thus the government chooses $s^{FI}(A), \tau^{FI}(A)$ to maximize $\mathcal{F}$ s.t. (1). Let $\mathcal{S}^{*,FI}$ be the support of s^{FI}.

Definition 1. Invertibility holds if for any signal that occurs in equilibrium $s\in\mathcal{S}^{*,FI}$ there exists a unique $A\in\Phi$ such that $s = s^{FI}(A) = h(\tau^{FI}(A), A)$.

Invertibility will often occur when the dimension of s is the same as the dimension of A. In this case, knowledge of s and of the policy function reveals A. Even if A is high-dimensional, Invertibility holds generically if Φ is a finite set, as only by coincidence would the same equilibrium point (τ^{FI}, s^{FI}) occur for two different realizations of A. Obviously Invertibility is synonymous with the existence of the inverse function $\left(s^{FI}\right)^{-1} : \mathcal{S}^{*,FI}\to\Phi$.

Proposition 1. *Under Invertibility the solution to Optimal Control with GSE is given by* $\mathcal{R}^*(s)\equiv\tau^{FI}\left(\left(s^{FI}\right)^{-1}(s)\right)$ *for all* $s\in\mathcal{S}^{*,FI}$.

In other words, under Invertibility the solution of the PI problem coincides with the solution under FI. This follows from the fact that the partial information (PI) case has the additional constrain (1) relative to the FI problem, therefore the value in the PI case is less than or equal to the FI case, since the value of the FI case is achievable for $\mathcal{R}^*$ as in Proposition 1, this must be the solution for the PI case.

A large part of the literature on partially observed shocks has considered this case. Also all the papers assuming revelation in incentive problems or those justifying the FI assumption by appealing to the fact that signals can be backed out from prices.[19]

3.5. First order conditions under GSE

The case of interest in this paper arises when knowledge of (τ, s) is not sufficient to back out the actual realizations of the shocks from (1), so that Invertibility as in the previous subsection does not obtain. In general, Invertibility fails generically if A has dimension higher than 1, at least one element of A has a continuous distribution and the other elements of A are non-degenerate. Technically, Assumption 4 ensures that Invertibility does not hold.[20]

We can now state our main result, namely a first order optimality condition under GSE. Let $W_\tau^*, W_s^*, h_\tau^*$ denote W_τ, W_s, h_τ evaluated at the optimal solution $s = S(\mathcal{R}^*, A)$ and $\tau = \mathcal{R}^*(S(\mathcal{R}^*, A))$.

Proposition 2. *Assume Assumptions 1–4. Assume that a solution to (8) exists denoted $\mathcal{R}^*$. Assume, $\mathcal{R}^*$ satisfies Conditions 1–2, the interiority conditions stated in Lemma 1 and $\mathcal{R}^*\in int(\mathcal{E})$. Then $\mathcal{R}^*$ satisfies the following necessary first order condition*

$$E\left(\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}}\middle| S(\mathcal{R}^*, A)\right) = 0 \ \text{ a.s. in } A. \tag{9}$$

Given fundamentals W, h, F_A, the interiority requirement on equilibrium (τ, s) can often be satisfied by enlarging the sets $\mathcal{T}, \mathcal{S}$.[21] Interiority of $\mathcal{R}^*$ is likely to hold for $\mathcal{E}$ that impose minimal restrictions on the set of admissible functions.

The proof of Proposition 2 uses variations, as in the literature on calculus of variations. Notice, however, that our approach is not a special case of results in that literature: unlike in the standard case, the choice of $\mathcal{R}$ involves choosing

[19] Invertibility also holds in the supply function model of Klemperer and Meyer (1989), where uncertainty is one-dimensional. Invertibility is not the only case in which the "standard recipe" works and (5) gives the correct solution. We will come back to this in Section 6 where we derive these as special cases of our general solution in Proposition 2.

[20] Even in the case that A is one-dimensional, Invertibility is violated if $s^{FI}(\cdot)$ is non-monotonic. We leave an exploration of this case for future work.

[21] See Appendix A.1

measurability conditions relating the observable signal and the underlying variables of integration. We refer the interested reader to Appendix A.7, where we explain this difference in more detail.

In Section 6 we explain how previous results in the literature on optimal policy under PI can be derived as special cases of Proposition 2. We show why in those cases the "standard recipe" works and equation (5) gives the correct solution.

We state a short version of the proof in the text, for a detailed proof see Appendix A.2.

Proof of Proposition 2.

Fix a deviation δ satisfying the boundedness conditions. Consider the problem

$$\max_{\alpha \in R} \mathcal{F}(\mathcal{R}^* + \alpha\delta) \tag{10}$$

maximizing over small deviations α of the optimal reaction function in the direction determined by δ. Lemma 1 guarantees that the solution $S(\mathcal{R}^*{+}\alpha\delta, A)$ is unique so that $\mathcal{F}(\mathcal{R}^* + \alpha\delta)$ is well defined for small α.

Given the interiority assumption $\mathcal{R}^* + \alpha\delta \in \mathcal{E}$, hence the solution of (10) is attained at $\alpha = 0$ and by the maximum principle

$$\left.\frac{d\mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha}\right|_{\alpha=0} = 0. \tag{11}$$

The appendix shows existence of this and all the derivatives that we use in this proof.

We use the following notation: for each function $G = W, W_\tau, W_s, \delta, \mathcal{R}^*, \delta', \mathcal{R}^{*\prime}, h_\tau, S$ we denote G^δ as the function mapping arguments (α, A) into the value of the corresponding G function evaluated at $S(\mathcal{R}^* + \alpha\delta, A)$ and/or $T(\mathcal{R}^* + \alpha\delta, A)$, namely

$$G^\delta(\alpha, A) \equiv G(T(\mathcal{R}^* + \alpha\delta, A), S(\mathcal{R}^* + \alpha\delta, A), A).$$

For example, we have $W^\delta(\alpha, A) \equiv W(T(\mathcal{R}^* + \alpha\delta, A), S(\mathcal{R}^* + \alpha\delta, A), A)$ and so on.

Fix $A \notin \Phi^{ND}_{\mathcal{R}^*}$ (this set has been defined in the proof of Lemma 1). Carefully differentiating W^δ we have

$$\begin{aligned}\frac{dW^\delta(\alpha, A)}{d\alpha} &= \Big[W^\delta_\tau(\alpha, A)(\mathcal{R}^{*\prime\delta}(\alpha, A) + \alpha\delta^{\prime\delta}(\alpha, A)) \\ &\quad + W^\delta_s(\alpha, A)\Big]S^{\delta\prime}(\alpha, A) + W^\delta_\tau(\alpha, A)\delta^\delta(\alpha, A).\end{aligned} \tag{12}$$

Given $\mathcal{R}^*, \delta, A$ consider now

$$s = h((\mathcal{R}^*{+}\alpha\delta)(s), A) \tag{13}$$

as an equation in (s, α). The implicit solution for s is precisely $S^\delta(\alpha, A)$. The implicit function theorem implies

$$S^{\delta\prime}(\alpha, A) = \frac{h^\delta_\tau(\alpha, A)\ \delta^\delta(\alpha, A)}{1 - h^\delta_\tau(\alpha, A)\ \left[\mathcal{R}^{*\prime\delta}(\alpha, A) + \alpha\delta'(S^\delta(\alpha, A))\right]}. \tag{14}$$

Evaluating (12) and (14) at $\alpha = 0$ and rearranging gives

$$\frac{dW^\delta(0, A)}{d\alpha} = \frac{W^*_\tau + W^*_s h^*_\tau}{1 - h^*_\tau \mathcal{R}^{*\prime}}\delta(S(\mathcal{R}^*, \cdot)) \ \text{ for all } A \notin \Phi^{ND}_{\mathcal{R}^*}. \tag{15}$$

Using $Prob(\Phi^{ND}_{\mathcal{R}}) = 0$ and exchanging the order of the derivative and integral signs in the definition of $\mathcal{F}$ (the appendix shows that this is a valid step) we have

$$\int_\Phi \frac{W^*_\tau + W^*_s h^*_\tau}{1 - h^*_\tau \mathcal{R}^{*\prime}}\delta(S(\mathcal{R}^*, \cdot))\ dF_A = 0. \tag{16}$$

Since this equation holds for any bounded δ with a bounded derivative it also holds for any bounded measurable δ. Therefore, the general definition of conditional expectation implies (9). □

As we anticipated the first order condition (9) does not coincide with the "standard recipe" FOC (5). The new term $\frac{1}{1-h^*_\tau \mathcal{R}^{*\prime}}$ acts as a kernel, or measure change relative to the standard case under separation. The reason this term appears is the following: under FI the optimality FOC involves $W^*_s h^*_\tau$ (see our discussion around (5)) because h^*_τ is the marginal change in the signal due to a marginal change in the policy variable τ. However, as shown in the proof, the marginal change in the signal under PI when the policy variable changes (in the direction of the variation δ) is

$$S^{\delta\prime}(0, A) = \frac{h^*_\tau \delta}{1 - h^*_\tau \mathcal{R}^{*\prime}} \tag{17}$$

hence this is what should multiply W^*_s in (15) and eventually gives rise to (9).

At this point, Fig. 1 may provide useful intuition. The illustration is based on our main fiscal policy example, which we introduce in Section 4, with two shocks $A \equiv (\gamma, \theta)$, these are not observed by the government, it only observes the aggregate signal $s \equiv l$. Given a choice for $\mathcal{R}$, upon observing a certain value of the signal $\bar{s} \equiv \bar{l}$, the planner cannot be sure if it was determined by a realization of shocks (γ_1, θ_1) or a realization (γ_2, θ_2), as illustrated by the intersection of a policy $\mathcal{R}$ with two possible reaction functions $h^{-1}(\cdot, A)$. The issue of General Signal Extraction is to realize that the choice of $\mathcal{R}$

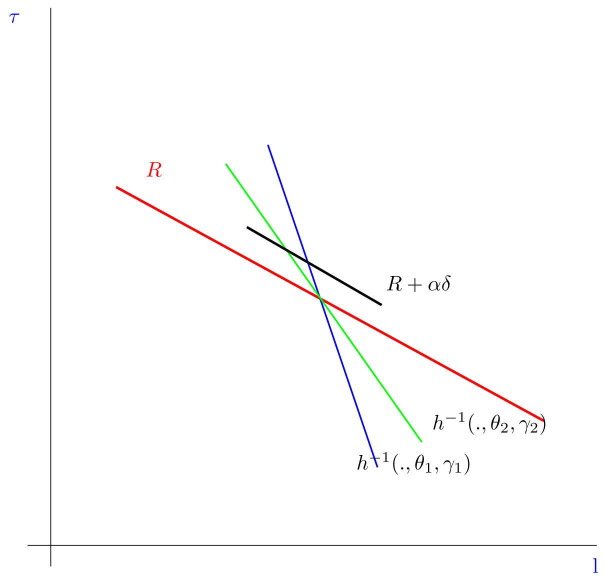

Fig. 1. Optimal policy $\mathcal{R}$ and deviation. The figure refers to the example of Section to illustrate the effect of a deviation $R+\alpha\delta$ on the signal, that is the intersection with two possible reaction functions h associated with different realisations of the shocks θ and γ. The signal is on the x-axis, and the policy on the y-axis.

influences the likelihood of observing $\bar{s}=\bar{l}$. Small changes in the choice of $\mathcal{R}$ impact on s differently for each realization. In particular, consider a small deviation $\mathcal{R}+\alpha\delta$, when δ is the indicator function of an interval around $\bar{s}$.[22] The figure illustrates how the impact of this deviation on the signal will be different if realization (γ_1,θ_1) or (γ_2,θ_2) actually materialize (even though these two realizations give rise to the same outcome for $\mathcal{R}$). The difference in the outcome s for each realization is proportional to the term $\frac{1}{1-h_\tau\mathcal{R}'}$ that appears in (9). Notice that if $\mathcal{R}$ and $h^{-1}(\cdot,A)$ are nearly parallel then $1-h_\tau\mathcal{R}'$ is close to zero and the kernel is large. Hence, the kernel gives more weight in the FOC to realizations such that a small change in the policy has large effects on equilibrium s.

3.5.1. A convenient optimality condition

Condition (9) is useful for comparability with the standard FOC and to point out the role of the kernel $\frac{1}{1-h_\tau R^{*\prime}}$. However, it suggests that computing the solution involves solving a differential equation, since the derivative of the policy function $\mathcal{R}^{*\prime}$ is involved in (9). In that case numerical solutions should ensure that both $\mathcal{R}^*$ and $\mathcal{R}^{*\prime}$ are well approximated. As it turns out, we can simplify things by deriving an optimality condition in which $\mathcal{R}^{*\prime}$ is not present. This happens because $\mathcal{R}^{*\prime}$ is also present in the density $f_{A|s}$ which defines the integral in (9) and the derivative cancels out. A sufficient condition to derive this result is the following assumption. In the following we denote a random variable with an upper case, say X, its density f_X, and we use lower case, say x, for possible values of X.

For a partition of the vector $A=(A_1,A_2)$, denote the support of A_2 conditional on observing s,τ as

$$\Theta_2(s,\tau)=\{a_2 : s=h(\tau,a_1,a_2) \text{ for some } (a_1,a_2)\in\Phi\}.$$

Assumption 5. The vector A can be partitioned as $A=(A_1,A_2)$ for $a_1\in R$ such that, i) the conditional density $f_{A_1|A_2}$ exists with probability one on A, ii) for any $(s,\tau,a_2)\in\mathcal{S}\times\mathcal{T}\times\Theta_2(s,\tau)$ there is a unique a_1 such that $s=h(\tau,a_1,a_2)$, iii) the partial derivative h_{A_1} exists everywhere in $\mathcal{T}\times\Phi$ and it is bounded away from zero.

Part *i)* strengthens Assumption 4. Part *ii)* means that one can recover A_1 from knowledge of the other variables, i.e. there is a well defined function $A_1=\mathcal{A}^*(s,A_2)$ satisfying $s=h(\mathcal{R}^*(s),\mathcal{A}^*(s,a_2),a_2)$ for all $(s,a_2)\in\mathcal{S}_{\mathcal{R}^*}\times\Theta_2(s)$ that occur in the optimal policy. Equivalently, we could define $\mathcal{A}^*(s,a_2)$ as the inverse function of $S(\mathcal{R}^*,\cdot,a_2)$.

Proposition 3. *In addition to the assumptions and conditions needed for Proposition 2, assume Assumption 5. Then, $\mathcal{R}^*$ satisfies the following necessary optimality condition*

$$\int_{\Theta_2(s,\mathcal{R}^*(s))}\frac{W_\tau^*+W_s^*h_\tau^*}{\left|h_{A_1}^*\right|}\,f_{A_1|A_2}(\mathcal{A}^*(s,a_2),a_2)\,dF_{A_2}(a_2)=0 \text{ for almost all } A \tag{18}$$

and for all s where the marginal density of $S(\mathcal{R}^*,A)$ is positive.

[22] This Figure uses an indicator function δ for illustrative purposes and because it is often used in the calculus of variations. However, as pointed out after Lemma 1, in our proofs we require continuous δ's.

The proof is in Appendix A.3.

As promised, this optimality condition does not involve $\mathcal{R}^{*\prime}$. In Section 3.7 we discuss how to set up an algorithm that uses this optimality condition. In addition to simplifying computations this proposition honors the title of the paper, as the proof involves computing the filter $f_{A|s}$ explicitly, highlighting that in general this filter depends on the optimal policy $\mathcal{R}^*$ and its derivative $\mathcal{R}^{*\prime}$.

3.6. Second order conditions and existence

We now state two more results: a second order condition that can be useful in practice when two solutions to (18) are found, and an existence result ensuring that the optimality conditions are meaningful. Proofs are in Appendix A.4 and A.5.

Second Order Optimality Condition

Proposition 2 was obtained using $\left.\frac{\partial \mathcal{F}(\mathcal{R}^*+\alpha\delta)}{\partial\alpha}\right|_{\alpha=0} = 0$. Similarly, the second order condition $\left.\frac{\partial^2 \mathcal{F}(\mathcal{R}^*+\alpha\delta)}{\partial\alpha^2}\right|_{\alpha=0} \leq 0$ has to hold for any δ. For simplicity the proposition is stated for the case where W does not depend on τ.

Assumption 6. W is twice differentiable everywhere with respect to s, h is twice differentiable everywhere with respect to τ and the second derivatives $|W_{ss}|, |h_{\tau\tau}| < Q$ uniformly on $\mathcal{S} \times \mathcal{T} \times \Phi$ for a constant $Q < \infty$.

Condition 3. $\mathcal{R}^*$ is twice differentiable almost everywhere with respect to s, with $\left|\mathcal{R}''\right| < Q$ uniformly on $\mathcal{S}_{\mathcal{R}}$.

Proposition 4. *In addition to the assumptions and conditions on $\mathcal{R}$ needed for Proposition 2, assume Assumption 6 and Condition 3. Then, for all continuously twice differentiable deviations $\delta : \mathcal{S} \to R$, $\mathcal{R}^*$ satisfies the following necessary second order condition*

$$E\left[W_{ss}^*\left(S^{\delta\prime}(0,A)\right)^2 + W_s^* S^{\delta\prime\prime}(0,A)\right] \leq 0 \tag{19}$$

where $S^{\delta\prime}(0,A)$ is given by equation (17) and

$$S^{\delta\prime\prime}(0,A) = \frac{\delta'^* h_\tau^* S^{\delta\prime} + (\delta^*)^2 h_{\tau\tau}^* + \delta^* h_{\tau\tau}^* \mathcal{R}^{*\prime} S^{\delta\prime} - \delta'^* (h_\tau^*)^2 \mathcal{R}^{*\prime} S^{\delta\prime} + \delta^* \delta'^* (h_\tau^*)^2 + \delta^* (h_\tau^*)^2 \mathcal{R}^{*\prime\prime} S^{\delta\prime}}{(1 - h_\tau^* \mathcal{R}^{*\prime})^2} \tag{20}$$

Existence

To guarantee existence of a maximum we can restrict the set of functions $\mathcal{E}$ to a large compact set so that it includes global maxima in most models. The following condition strenghthens Conditions 1–2 so as to guarantee compactness.

Condition 4. Same as Condition 2 but (7) holds at all s such that $|s - h(\mathcal{R}(s), A)| \leq \overline{\delta}$ for some constant $\overline{\delta} > 0$ for some $A \notin \Phi_{\mathcal{R}}^{ND}$.

Proposition 5. *Define $\mathcal{E}$ as the set of all functions $\mathcal{R}$ satisfying Conditions 1 and 4 for some $K^L, \overline{\varepsilon}, \overline{\delta} > 0$. Let Assumptions 1,2,4 hold. Then $\mathcal{E}$ is compact and there exists a solution $\mathcal{R}^*$ attaining the maximum of problem (3).*"

Choosing a large K^L and small $\overline{\varepsilon}, \overline{\delta}$ should include the global optimum in $int(\mathcal{E})$ for most models.

3.7. Putting everything together

The above results give a clear path to find solutions to (3): *i)* Find a solution to $\mathcal{R}^*$ for all (or a fine grid of) s using

Algorithm 1. Discretize the set of possible values for s, i.e., $\mathcal{S}_{\mathcal{R}^*}$. Finding the bounds of this set is easy, given a solution to the FI problem, as we explain in our application in Section 5. For each possible value of $s \in \mathcal{S}_{\mathcal{R}^*}$ and a candidate $\tau = \mathcal{R}^*(s)$, evaluate each function in the integrand of (18) at $(\tau, s, \mathcal{A}^*(s, a_2), a_2)$. Find the integral on the left side of (18) by integrating a_2 over all possible values in $\Theta_2(s, \tau)$ given the candidate τ. This set can be found by checking whether the implied $a_1 = \mathcal{A}^*(s, a_2)$ is within the support of A_1. This maps each possible τ into a value for the integral on the left side of (18). The optimal choice $\mathcal{R}^*(s)$ is found by solving a non-linear equation that maps a candidate τ into an integral as close as possible to zero.

ii) if a unique solution is found check that Conditions 1–2 and interiority conditions are satisfied.

If we find several solutions in *i)* then proceed to: *ii')* check if some solutions can be ruled out with the second order condition.

If still several solutions survive proceed to: *ii")* evaluate the objective function $E(W)$ at each candidate solution and pick the $\mathcal{R}$ with the highest value. [23]

[23] In purity one should still check that the objective function can not be improved by using some $\mathcal{R}$ for which the Lipschitz bounds of Conditions 1–2 and interiority conditions are binding. This would require either to apply Kuhn&Tucker conditions or to prove that the optimal $\mathcal{R}$ can not violate these bounds. We leave this for another paper.

4. A simple model of optimal fiscal policy

We now present a simple model of optimal fiscal policy with two shocks, to supply and demand respectively. We believe that this model and its extensions in Sections 5, and 9 represent an interesting laboratory for our theory and may be useful to gain intuition on the role of endogenous signals for optimal policy. However, we stress that our results developed in Section 3 in a more general setup could be applied in many fields, also outside of macro.

Specifically, our main example is a two-period version of the Lucas and Stokey (1983) model of optimal dynamic taxation under uncertainty. We introduce incomplete markets to be consistent with the presence of Partial Information about the state. For simplicity, we assume that there is a representative agent, and that signals and policy are determined simultaneously. We relax these assumptions in Sections 7 and respectively.

4.1. Preferences and technology

The economy lasts two periods $t = 1, 2$. A government needs to finance an exogenous and constant stream of expenditure $g_1 = g_2 = g$, where subscripts indicate time periods. The government levies distortionary income taxes (τ_1, τ_2) and issues bonds b^g in the first period that promise a repayment in second-period consumption units with certainty.

The economy is populated by a continuum of agents. Each agent $i \in [0, 1]$ has an expected utility function

$$E\left[U\left(c_1^i, l_1^i, c_2^i, l_2^i; \gamma\right)\right] \tag{21}$$

with $U\left(c_1^i, l_1^i, c_2^i, l_2^i; \gamma\right) = \gamma u\left(c_1^i\right) - v\left(l_1^i\right) + \beta\left[u\left(c_2^i\right) - v\left(l_2^i\right)\right]$, where c_t^i and l_t^i for $t = 1, 2$ are consumption and hours worked respectively, and $u' > 0$, $u'' < 0$, $v' > 0$, $v'' > 0$.

We refer to γ, a random variable with distribution F_γ, as a “demand shock”. When γ is high, agents like first period consumption relatively more than other goods. Hence, a high value of γ makes them willing to work more in their intratemporal labor-consumption decision and also makes them more impatient in their intertemporal allocation of consumption. Given that agents are identical, in the remainder we drop the subscripts i for notational convenience, whenever this does not lead to confusion.[24]

The production function in each period is linear in labor. Output is given by

$$y_t = \theta_t l_t \tag{22}$$

for $t = 1, 2$. The random variable $\theta_1 = \theta$ is a “productivity shock” with distribution F_θ. γ and θ are assumed to be independent. As far as θ_2 is concerned, we assume that both agents and government know with certainty that $\theta_2 = E\theta$, that is, the second period productivity is a known constant, equal to the mean of the first period shock.

To summarize, the state of the economy is fully described by a realization of the random variables $A \equiv (\gamma, \theta)$. These variables are observed at the beginning of period $t = 1$ by consumers and firms, but not by the government. The distributions F_γ and F_θ represent the government’s perceived distribution of the exogenous shocks, which may or may not be equal to the true distribution of these variables. Thus this formulation encompasses the case of “true” uncertainty as well as the government’s ignorance about the structure of the economy.

Agents have Rational Expectations: denoting by Φ the space of possible values of A, we assume that agents know that fiscal policy is given by a triplet of functions $(\tau_1, \tau_2, b^g) : \Phi \to R^3$ and these are actually the equilibrium values of taxes and government bonds for each A.

Consumers’ choices and prices are contingent on the state A, which agents observe in period $t = 1$. Agents choose $(c_1, c_2, l_1, l_2, b) : \Phi \to R^5$ knowing the fiscal policy and the bond price function $q : \Phi \to R$. The solution of the agents’ problem in this setup coincides with the non-stochastic model where A is known. Uncertainty will only play a role in the government’s problem, to be specified later.

Firms also observe θ at $t = 1$. Profit maximization implies that agents receive a wage equal to θ_t, observed by agents, so that the period budget constraints of the representative agent are

$$c_1 + qb = \theta l_1 (1 - \tau_1) \tag{23}$$

$$c_2 = \theta_2 l_2 (1 - \tau_2) + b \tag{24}$$

where q is the price of the government bond b. The above budget constraints have to hold for all realizations of A.

The government’s budget constraints are analogous, they restrict the choice of the policy (τ_1, τ_2, b^g).

4.2. Competitive equilibrium

We now provide a definition of competitive equilibrium that is standard in the literature and common to both the Full Information (FI) and the Partial Information (PI) equilibria that we are going to analyze.

[24] In Section 7, we generalize this framework to include heterogeneous agents facing idiosyncratic shocks.

Definition 2. A **competitive equilibrium** is a fiscal policy (τ_1, τ_2, b^g), price q and allocations (c_1, c_2, l_1, l_2, b) such that when agents take (τ_1, τ_2, q) as given the allocations maximize the agents' utility (21) subject to (23) and (24). In addition, bonds and goods markets clear, so that $\widetilde{b^g} = b$ and

$$c_t + g = \theta_t l_t \quad \text{for } t = 1, 2. \tag{25}$$

This definition insures that wages are set in equilibrium and that the budget constraint of the government holds in all periods due to Walras' law.

Utility maximization implies for all A

$$\frac{v'(l_1)}{u'(c_1)} = \theta\gamma(1 - \tau_1) \tag{26}$$

$$\frac{v'(l_2)}{u'(c_2)} = \theta_2(1 - \tau_2) \tag{27}$$

$$q = \beta\frac{u'(c_2)}{\gamma u'(c_1)} \tag{28}$$

As anticipated, the demand shock enters the first period labor supply decision described by equation (26) as well as the bond pricing equation (28). A competitive equilibrium is fully characterized by equations (23) to (28).

4.3. Ramsey equilibrium

To describe government behavior, we now provide a definition of Ramsey equilibrium. As is standard, we assume the government has full commitment, perfect knowledge about how taxes map into allocations for a given value of the underlying shocks A, and that it chooses the policy that maximizes household welfare.

We first give the standard definition when both government and consumers observe the realization of A.[25]

Definition 3. A **Full Information (FI-) Ramsey equilibrium** is a fiscal policy (τ_1, τ_2, b^g) that achieves the highest utility (21) when allocations are determined in a competitive equilibrium.

To study optimal taxes under PI, we assume that taxes in the first period have to be set before the shock A is known but after observing a **signal** s that potentially depends on aggregate outcomes observed in period 1, $s = G(c_1, l_1, q, A)$ for a given G. In our leading example, we assume $s = l_1 = \int l_1^i di$. Since in the present example agents are identical, aggregate and individual choices coincide in equilibrium, but it is important for the interpretation of equilibrium that the signal only depends on aggregate outcomes, as this justifies the assumption that agents take the policy action τ as given. Section 7 shows an example in which the distinction between aggregate and individual variables is clearer.

Definition 4. In a **Partial Information (PI-) Ramsey equilibrium**, conditional on observing a signal s, the government chooses the policy that achieves maximum utility, subject to the constraint

$$\tau_1 = \mathcal{R}(s) \text{ for all } A \in \Phi \tag{29}$$

for some measurable function $\mathcal{R} : R \to R$ and allocations are determined in competitive equilibrium.

We are interested in the case when (29) prevents the PI-Ramsey equilibrium from achieving the FI outcome. Therefore we introduce two shocks and one signal so that invertibility does not hold.

Note that the definition of competitive equilibrium implies that consumers' know the map $\tau_1 : \Phi \to R$. Hence, in a PI-Ramsey equilibrium, agents' perception of how policy is set is consistent with equation (29). An interpretation is that the government announces its policy $\mathcal{R}$, agents understand this policy and the associated mapping from A to s. Another equivalent interpretation is that agents simply are endowed with knowledge of $\tau_1(\cdot)$. Even if agents know that (29) holds, they can not exploit this knowledge in their optimization problem as we consider atomistic agents that cannot affect the aggregate signal and, hence, the tax rate.[26] In this model, as is standard in Ramsey equilibria, the tax level τ and the equilibrium allocations (and therefore s) are determined simultaneously as a consequence of the government's choice for $\mathcal{R}$.[27]

FI- Ramsey equilibria

Using the so-called "primal approach" and standard arguments it is easy to show that an allocation is a competitive equilibrium if and only if, in addition to the resource constraints (25), the following implementability condition holds

$$\gamma u'(c_1)c_1 - v'(l_1)l_1 + \beta\left[u'(c_2)c_2 - v'(l_2)l_2\right] = 0. \tag{30}$$

[25] These definitions take for granted that we only consider tax policies for which a competitive equilibrium exists and is unique.

[26] This differs from the setup in the New Dynamic Public Finance, where consumers optimize given a policy function $\mathcal{R}$ that is a function of individual choices, i.e., using our notation, the government chooses an optimal individual tax $\tau_{it} = \mathcal{R}(l_{it})$. Therefore in that literature agents internalize the effect of their actions on taxes.

[27] We consider the case of predetermined taxes in Section

The standard approach to find Ramsey policy under FI is to maximize (21) subject to (30) and the resource constraints. We now deviate slightly from this traditional approach so that the PI problem can be understood as an additional restriction of the FI problem. For this purpose we do not "substitute out" τ_1, instead we keep this tax as an explicit choice variable and using (25) for $t = 1$ to substitute out consumption in (26), we add the constraint

$$\frac{v'(l_1)}{u'(\theta l_1 - g)} = \theta\gamma(1 - \tau_1), \tag{31}$$

Letting h be the function that maps τ_1, θ, γ into the value for l_1 that solves this equation, we can rewrite the above equilibrium condition as

$$l = h(\tau, A) \tag{32}$$

where we have suppressed the time subscript from first period labor and tax rate. This shows how the allocation for labor reacts to a tax choice.

The resource constraints (25) and (30) give three equations defining a map from each possible value of $l = l_1$ into corresponding equilibrium (c_1, c_2, l_2). Plugging this map into the utility function (21), welfare for each A can then be written as

$$W(l; A) \equiv U(c_1, l, c_2, l_2; \gamma) \tag{33}$$

solely as a function of $l = l_1$, embedding in W the equilibrium conditions mentioned above. The FI Ramsey Equilibrium reduces to solving

$$\max_{(\tau,l):\Phi\to R^2} E[W(l; A)] \tag{34}$$

$$\text{s.t. } (32) \tag{35}$$

Clearly, this problem is equivalent to maximizing (21) subject to (30), given A.[28]

PI- Ramsey equilibria

We focus on the case when the signal is just aggregate labor so $s = l_1 = l$. The only difference under PI is that the additional constraint (29) appears, that $\tau_1 = \tau$ determines $l_1 = l$ in equilibrium through (32) and that the choice is over the policy $\mathcal{R}$.

Hence a **PI-Ramsey equilibrium** (given a signal l) solves

$$\max_{\mathcal{R}:R\to R} E[W(l; A)] \tag{36}$$
$$\text{s.t. } (32)$$
$$\tau = \mathcal{R}(l)$$

Notice that, mathematically, the only difference with the FI-problem is the presence of the last constraint.

This gives rise to a non-standard maximization problem under GSE, where the filtering problem $f_{A|l}$ has to be solved *jointly* with the optimal choice for the policy function $\mathcal{R}$.

4.4. The economic consequences of PI for taxation policy

Before presenting our numerical solution, it is worthwhile discussing the economic issues raised by PI in this model. As is well known, the optimal FI policy implies tax smoothing over time, as the government spreads the distortions equally in the two periods. In the case of CRRA preferences $u(c) = \frac{c^{1+\alpha_c}}{1+\alpha_c}$, $v(l) = B\frac{l^{1+\alpha_l}}{1+\alpha_l}$ for $\alpha_c \le 0$, $\alpha_l, B > 0$, tax smoothing will be perfect and Ramsey policy under FI involves setting a tax rate $\tau = \tau_1 = \tau_2$ constant in time, the level of taxes satisfies the intertemporal budget constraint

$$\tau\theta l_1 - g + \beta\frac{u'(c_2)}{\gamma u'(c_1)}(\tau\theta_2 l_2 - g) = 0. \tag{37}$$

It is clear from (37) that the government needs to know the realization of both productivity and demand shocks in order to implement this policy under FI. In particular, the realization of $\theta = \theta_1$ is a crucial piece of information, as it determines the revenue that a given tax rate, together with an observed level of hours worked, is going to raise. The demand shock γ also matters as it affects both the objective function and the interest rate that the government will have to pay on its debt. Furthermore, both shocks clearly contribute to the determination of an allocation (c_1, c_2, l_1, l_2).

Under PI, the government can only condition its policy on l, without knowing what combination of the shocks gives rise to a given observation. Thus, under PI the choice of constant taxes (optimal under FI) is not feasible. The government has to fix τ_1 while it is still uncertain about the revenue that this tax rate will generate and it will enter period 2 with an uncertain amount of debt. Once θ and γ are known, government deficit (or surplus) will materialize, and the government

[28] As explained in the theoretical section l should be constrained to belong to the set of feasible signals.

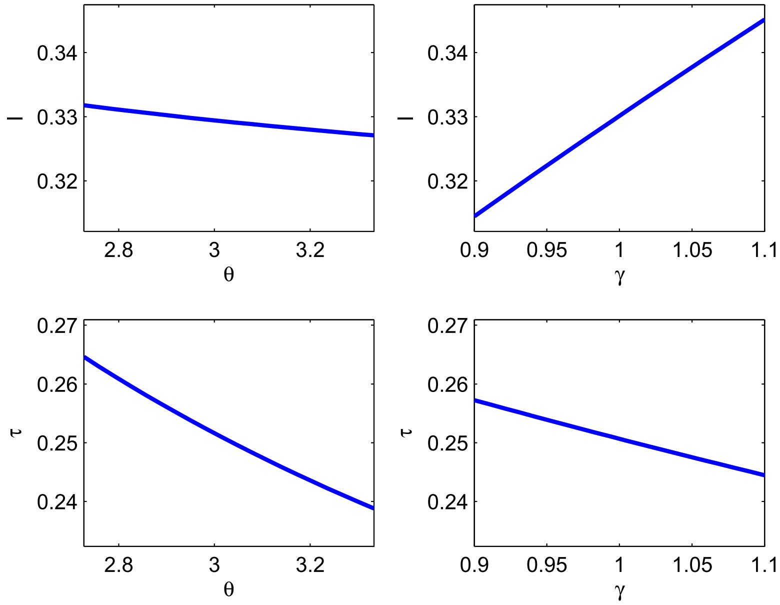

Fig. 2. Hours and taxes with Full Information. Left panels: top, hours under FI as a function of the productivity shock; bottom, tax rate under FI as a function of the productivity shock. Right panels: top, hours under FI as a function of the demand shock; bottom, tax rate under FI as a function of the demand shock.

will have to set τ_2 so as to balance the budget in the second period. Hence, in the eyes of the government, τ_2 is unavoidably a random variable at the time of choosing τ_1.

Arguably, uncertain tax revenue is a crucial feature of actual fiscal policy decisions, and tax rates are decided based on information from equilibrium outcomes that are observed frequently. In this sense, one can interpret this model as a simple model of optimal automatic stabilizers, as these are fiscal instruments that are designed to respond to endogenous outcomes, such as income or unemployment, independently of the source of fluctuations in these variables.

5. Solution of fiscal policy model with GSE

We now apply our theoretical results to solve for optimal policy under PI in the model of Section 4.[29]

5.1. Low government spending and optimal tax smoothing

We parametrize the economy as follows. We assume $u(c) = \log(c)$ and $v(l) = \frac{B}{2}l^2$. We set $\beta = 0.96$. B and the mean of θ are set to normalize average output to 1 and average hours to a third. We let θ be uniformly distributed on a support $[\theta_{\min}, \theta_{\max}]$ and γ uniformly distributed on $[\gamma_{\min}, \gamma_{\max}]$. Government expenditure is constant and equal to 25% of average output. The mean of γ is 1. The supports of both shocks imply a range of $\pm$10% from the mean.

We first present the FI solution in order to illustrate the optimal response of taxes and allocations to the two different shocks we consider. In Figure 2 we show how hours and taxes move with the two different shocks under FI. On the left side of the figure, we keep γ constant and equal to its mean and we show that both hours and taxes are decreasing in the productivity shock. On the right side, we keep θ constant and equal to its mean and show that labor is increasing in γ, while taxes are decreasing. This shows that ideally, if the government sees an increase in l, it would want to react in opposite directions depending on the source of the shock: a tax increase, if driven by low θ, or a tax cut if driven by high γ. However, under PI this is not feasible because θ and γ are not observed.

We now discuss the problem of optimal policy with PI. Notice that the equilibrium condition (26) becomes

$$Bl_1 c_1 = \gamma\theta(1 - \tau_1). \tag{38}$$

After substituting out consumption using the resource constraint, we obtain that labor supply is the positive root of a quadratic equation, so that the reaction function (32) specializes to

$$l = h(\tau, \theta, \gamma) = \frac{Bg\theta^{-1} + \sqrt{(Bg)^2\theta^{-2} + 4B\gamma(1-\tau)}}{2B}. \tag{39}$$

[29] Consistent with the presentation of the model in Section, we focus on productivity and demand shocks that hit the economy only in the first period and assume that the signal is the level of hours worked. In Appendix C, we also consider permanent shocks to productivity (Appendix C.1) and the case in which the signal is output (Appendix C.2).

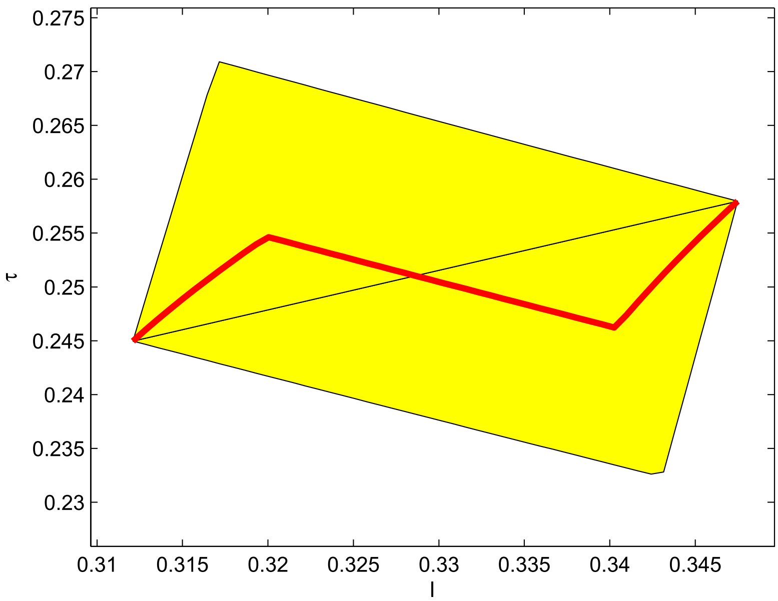

Fig. 3. Optimal policy with low g. Optimal tax rate as function of hours under PI. Thick red line: $\mathcal{R}^*$; yellow region: set of FI pairs (l, τ) for all possible realizations of (θ, γ); thin black line: linear policy connecting the two full revelation points. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

This equation gives the signal observed by the government, namely the level of hours worked, as a function of the tax rate and the two unobserved shocks. While it is clear that hours are unambiguously increasing in the demand shock γ, it is important to note that the productivity shock θ has two opposing effects on l: the substitution effect between leisure and consumption and the wealth effect, that acts in the opposite direction.

By setting the tax rate, the government affects the informativeness of the signal. Notice that the government can actually achieve full revelation of the value of productivity θ, by setting a tax rate of 100%. This would reveal $\theta = \frac{g}{l}$, but imply zero consumption. Hence, this choice would clearly by suboptimal. This illustrates that the government faces a trade-off between learning about the state of the economy and choosing a suitable policy under uncertainty.

To the extent that θ is the most important shock affecting tax smoothing, we could expect that extracting a better signal about θ is valuable to the government. Combining this with the discussion in the previous paragraph suggests that higher taxes in the first period are valuable because they narrow the range of possible θ's.

Furthermore, the extent to which hours are an informative signal about productivity depends importantly on the level of government spending, through its effect on the marginal utility from consumption. If $g = 0$, then substitution and wealth effect exactly offset each other and hours are independent of θ. In the presence of positive government spending, the wealth effect dominates the substitution effect and hence high realizations of θ will lead to low labor, ceteris paribus. The higher g, the higher the marginal utility of consumption, the stronger the wealth effect on labor supply, and the more informative hours become about productivity.[30]

The partial derivatives h_τ and h_θ that we need to derive to solve the optimality condition (18) are easily obtained analytically in this example. In particular,

$$h_\tau(\tau, \theta, \gamma) = \frac{-1}{\sqrt{(Bg)^2(\theta\gamma)^{-2} + \gamma^{-1}4B(1-\tau)}}. \tag{40}$$

It is clear that both the productivity shock and the demand shock affect this slope, therefore the kernel in the optimality condition (18) is not a constant. Hence, we proceed to find a $\mathcal{R}^*$ that satisfies (18) using Algorithm 1 described above. We provide the detailed steps of our solution method for the simple model in Appendix B.1.

Fig. 3 illustrates the optimal policy for this case, plotting the tax rate against observed labor. The red line is $\mathcal{R}^*$, while the yellow region is the set of all equilibrium pairs (l^{FI}, τ^{FI}) that could have been realized under FI.

Let $l = L(\mathcal{R}^*, \theta, \gamma)$ be the signal as a function of government policy and shocks under PI. The boundaries of the feasible set for l are easy to find ex-ante by exploiting the fact that the extreme values for the signal l coincide with the FI allocation, which we denote by $L^{FI}(\theta, \gamma)$. As l is increasing in γ and decreasing in θ, letting $l_{\min}$ and $l_{\max}$ be the extreme values of l

[30] In order to illustrate the non-linear effect of government spending g on the signal extraction, we discuss both a case with low g (in this subsection) and a case with high g (in Section 5.2).

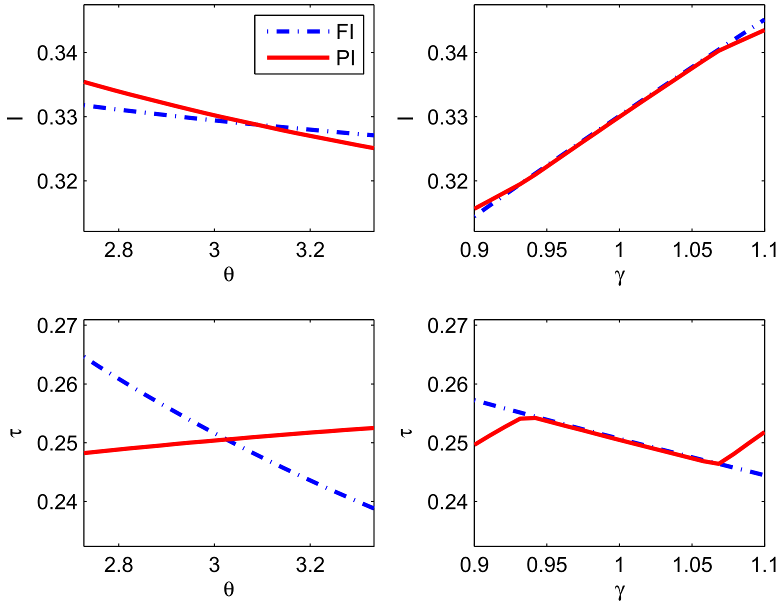

Fig. 4. Hours and taxes with Partial Information and Full Information. Left panels: top, hours as a function of the productivity shock (solid red line for PI, blue dashed-dotted line for FI); bottom, tax rate as a function of the productivity shock. Right panels: top, hours as a function of the demand shock; bottom, tax rate as a function of the demand shock. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

in the PI solution, we have

$$l_{\min} = L(\mathcal{R}^*; \theta_{\max}, \gamma_{\min}) = L^{FI}(\theta_{\max}, \gamma_{\min}); \; l_{\max} = L(\mathcal{R}^*; \theta_{\min}, \gamma_{\max}) = L^{FI}(\theta_{\min}, \gamma_{\max})$$

and the PI solution is in the interval $[l_{\min}, l_{\max}]$.

The intuition for Fig. 3 is as follows. For these extremes values of the signal, there is full revelation, but anywhere between these two extremes the government has to choose a policy without knowing the values of γ, θ that give rise to equilibrium taxes or labor. It can be seen that the optimal policy calls for a tax rate in between the minimum and the maximum FI policies for each observation.

At low realizations of hours, the government learns that productivity must be high, so the tax rate can be rather low. The lowest labor realization $l_{\min}$ coincides with the FI equilibrium for $(\theta_{\max}, \gamma_{\min})$. Then, taxes start to increase with l: higher l's signal lower expected productivity and hence revenue, as the set of admissible θ's is gradually including lower and lower realizations. This goes on up to a point where the set of admissible θ's conditional on l is the whole set $[\theta_{\min}, \theta_{\max}]$. From that point on, the tax rate changes slope and becomes decreasing with respect to l. This is because now, with any θ being possible, increasing l signals an increasing expected revenue, hence allowing lower tax rates on average, up to the point where the highest θ's start being ruled out, at which point the policy becomes increasing again as the government understands that θ is low, up to the full revelation point $l_{\max} = L^{FI}(\theta_{\min}, \gamma_{\max})$.

To gain further understanding on the implications of PI for the properties of the model, we plot again hours and taxes as functions of each shock individually in Fig. 4. In all four panels, we reproduce the FI outcomes shown in Fig. 2 (blue dashed-dotted lines). The red lines represent the PI outcomes. For instance, in the left panels we keep γ equal to its mean and we plot hours and taxes as functions of θ. Interestingly, it can be seen that hours become more volatile in response to productivity shocks under PI, while taxes become smoother and change the sign of their response to θ. This is because under this parametrization the government learns little about the realizations of θ and hence optimally chooses to cut taxes as hours increase.[31]

On the right-hand side, we plot again hours and taxes as functions on γ, keeping θ equal to the mean. For intermediate values of γ, the government is relatively confident about the realization of the demand shock, hence the policies under FI and PI are very close. However, for extreme realizations, the government is uncertain about which shock is driving hours, hence it cuts taxes for very low γ's and increases taxes for very high γ's, believing that changes in productivity are responsible for the observed behavior of hours.

We also plot the locus of admissible realizations of shocks conditional on observing an average level of hours, $l = 0.33$, in Fig. 5. This shows that policy influences the probability of observing a certain outcome, hence optimal policy and filtering problem need to be solved jointly. The wealth effect of productivity makes this locus an increasing function in the (θ, γ)

[31] We will see in Section 5.2 that this property of the solution will change with higher government expenditure.

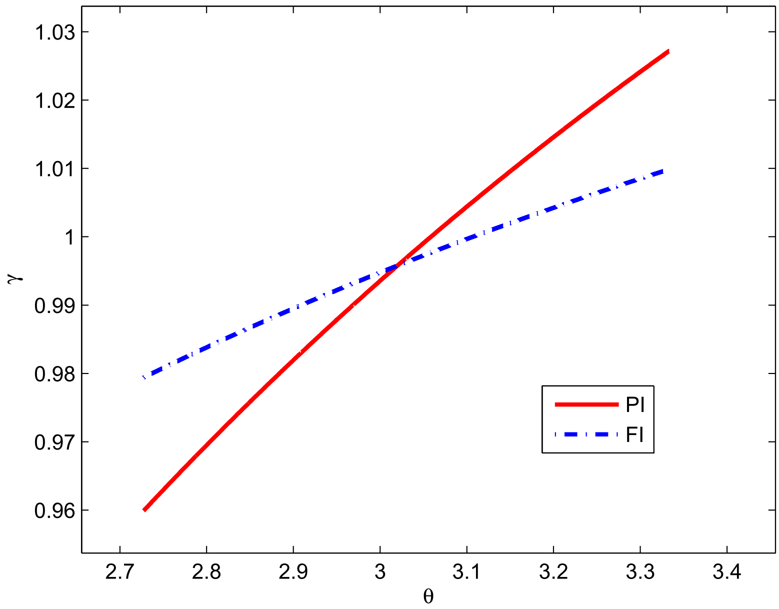

Fig. 5. Set of admissible shocks consistent with observing $l = 0.33$. Set of combinations of (θ, γ) that have positive density in equilibrium, conditional on a particular realization of the signal, namely $l = 0.33$. Solid red line for PI, blue dashed-dotted line for FI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

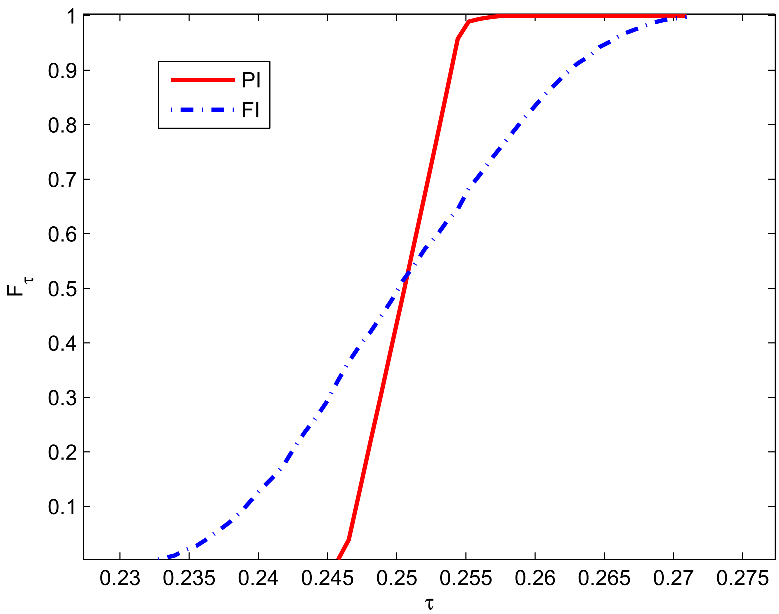

Fig. 6. Equilibrium CDF of tax rates. Cumulative Distribution Function of τ_1 under PI (solid red line) and FI (blue dashed-dotted line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

space. As hours are increasing in γ and decreasing in θ, a given level of hours could be due to combinations of high demand and high productivity, or low demand and low productivity. It should be noted that the locus of shocks realizations conditional on l is endogenous to policy. Importantly, τ affects the slope of this locus, implying that the government can choose to some extent on what shocks the signal extraction will be more precise. To see this, observe that a horizontal locus would imply revelation of the value of γ, while a vertical locus would imply revelation of the value of θ.

Optimal policy with PI calls for a substantial smoothing of taxes across states. We illustrate this result in Fig. 6, where the equilibrium cumulative distribution function of tax rates under PI is contrasted with the one obtained under FI. This result is rather intuitive and it carries a general lesson for optimal fiscal policy decisions under uncertainty: when the government is not sure about what type of disturbance is hitting the economy, it seems sensible to choose a policy that is not too aggressive in any direction and just aims at roughly keeping the budget under control on average.

In our model, this smoothing of taxes across states will imply a larger variance of tax rates in the second period with respect to the FI policy. In the second period, all the uncertainty is resolved and the tax rate will take whatever value is needed to balance the budget constraint. This is of course taken into account at the time of choosing a policy under uncertainty, so that we could say that optimal policy is prudent while the source of the observed aggregate variables is not known and then responsive after uncertainty has been resolved.

This result is related to the question on whether taxes should be smooth across states or over time, depending on the completeness or incompleteness of financial markets. With complete markets and FI, tax smoothing happens across states (Lucas and Stokey, 1983). When markets are incomplete, the FI government substitutes tax smoothing across states with tax smoothing over time (Aiyagari et al., 2002). In our model, with incomplete markets and PI, we find that taxes are smoother across states than over time. This suggests that tax smoothing across states may not necessarily be an indication of market completeness and full insurance on the part of the government, but simply a sign of incomplete information about the state of the economy.

Because of this property, our model can rationalize the slow reaction of some governments to big shocks like the Great Recession. The Spanish example in the latest recession is a case in point. In 2008, it was far from clear how persistent the downturn would be and also whether is was demand-driven or productivity-driven and the government did not adjust its fiscal stance quickly, only to make large adjustments in the subsequent years. We will discuss this effect further in Section 9.

5.2. Close to the top of the Laffer curve

We now consider the case where government expenditure is very high, equal to 60% of average output in both periods.[32] We will see that this leads to a very non-linear optimal policy and to an exception to tax-smoothing across states. This example is of interest for several reasons. From an economic point of view, Partial Information is of higher importance here: since the government needs to balance the budget in the second period it is now concerned about the possibility of a very low level of productivity θ, as in this case tax revenue is low in the first period and a large amount of debt will need to be issued. A high debt, combined with high future expenditure, may call for very high taxes in the future, it could even mean getting the economy closer to the top of the Laffer curve, where taxation is most distortionary.

This example will also be of interest because the PI solution has some very different features from the FI outcome. By increasing marginal utility from consumption, high government spending makes the wealth effect of a productivity shock larger, with the consequence that hours worked become a stronger signal about θ. The government will optimally exploit this in the signal extraction.

Fig. 7 shows optimal policy for this case (red line), again contrasted with the set of possible tax-labor outcomes under FI (yellow region). The figure shows that the optimal solution is highly non-linear. The derivative $\mathcal{R}^{*\prime}$ is positive and relatively high in a middle range of levels of l, but both to the left and to the right of this middle range $\mathcal{R}^*$ it is much flatter. Notice that this is the opposite of what happens with a low level of g in Section 5.1. When government expenditure is sufficiently low, the government is very uncertain about the true realization of θ. Hence higher labor does not allow a more precise signal extraction about productivity. On the other hand, when g is sufficiently high, there is an intermediate region of observables where the government becomes confident about low realizations of θ. In Appendix B.2 we prove this result by illustrating how g affects the slope of the loci of realizations of the shocks through its impact on the wealth effect.

Figs. 8 and 9 show how optimal policy under GSE narrows down the possible values of shocks that are compatible with a given signal. To illustrate how the PI policy involves a relatively precise signal extraction on θ with high government expenditure, compare the sets of possible realizations of θ conditional on hours (on the x-axis) under FI in Fig. 8 and under PI in Fig. 9. Consider Fig. 8 first. It can be seen that under FI any realization of θ is consistent with a large set of intermediate realizations of l, but each of these θ's would call for a different tax rate. However, under PI, there can only be one tax rate for each observed l and the government uses this policy to extract information on θ.

To see this, consider now Fig. 9. The minimum value of l is only consistent with the highest possible θ (and lowest possible γ) because the wealth effect dominates. Under PI, increasing l from this point, the government becomes uncertain and lower realizations of productivity become consistent with the observations. At first, uncertainty is rising with l, but in the intermediate region of l's the government becomes more and more confident about low realizations of productivity. This leads to the sharp increase in the tax rate, which in turn gives rise to feedback effect on the set of possible θ's: high taxes discourage work effort, so higher labor now is an even stronger signal of low θ (high marginal utility from consumption). In this way an optimal policy and a conditional distribution of shocks consistent with it confirm each other in equilibrium.

Consistently with this analysis of the signal extraction, we also plot hours and taxes as functions of each shock individually, and we contrast the PI outcomes with the FI solution in Fig. 10. On the left-hand side we consider productivity shocks only. As illustrated above, in the intermediate region of l's the government has a precise signal about θ, hence PI and FI policies and allocations are very close to each other. However for extreme realizations of θ the government is fooled about the source of the fluctuations and does hardly respond to productivity. On the right-hand side we consider only demand

[32] All other assumptions on preferences and shocks are the same as in Fig. 7.

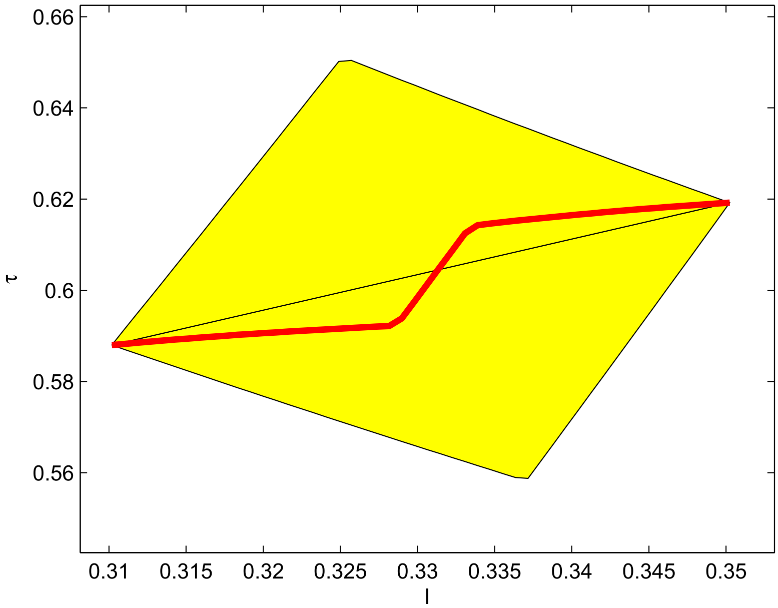

Fig. 7. Optimal policy with high g. Optimal tax rate as function of hours under PI. Thick red line: $\mathcal{R}^*$; yellow region: set of FI pairs (l, τ) for all possible realizations of (θ, γ); thin black line: linear policy connecting the two full revelation points. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

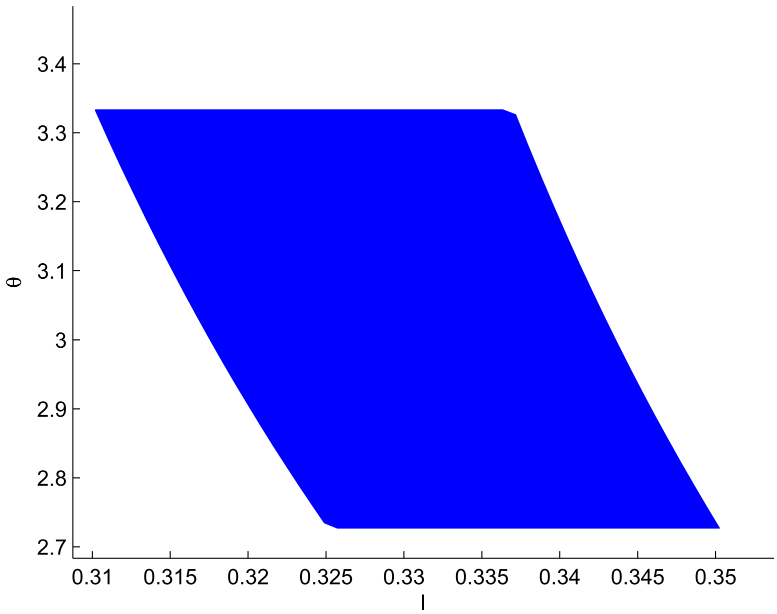

Fig. 8. Set of admissible θ's with FI. The blue region indicates, for each possible value of l, the corresponding set of realizations of θ with positive density in equilibrium under FI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

shocks. It can be seen that the PI government has very imprecise information about γ. Hence it responds to these shocks with the opposite slope with respect to the FI government.

The case of high government expenditure shows that optimal policy with PI can be very non-linear in order to avoid the worst outcomes, e.g. in the model hitting the top of the Laffer curve or, in the real world, a debt crisis. As shown in the previous Proposition 2, when expenditure is low and there are no concerns related to the government budget constraint, policy has to be smooth, but when there are contingencies that are particularly dangerous, then optimal policy calls for being reactive to observables in order to prevent those cases to materialize. This is exemplified by the optimality of increasing taxes steeply in the first period (in the intermediate range of signals l) to avoid having to distort the economy too heavily in the second period if realized productivity turns out to be low (and hence the fiscal deficit turns out to be high). This lesson

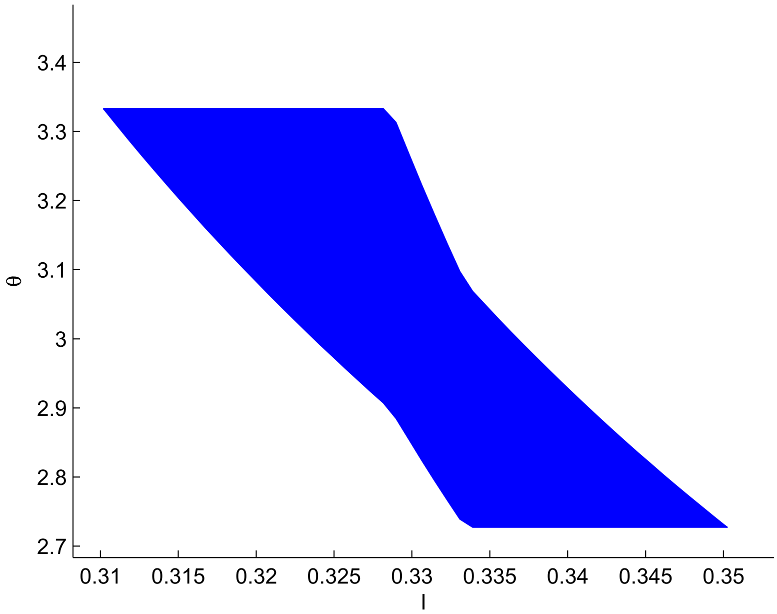

Fig. 9. Set of admissible θ's with PI. The blue region indicates, for each possible value of l, the corresponding set of realizations of θ with positive density in equilibrium under PI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

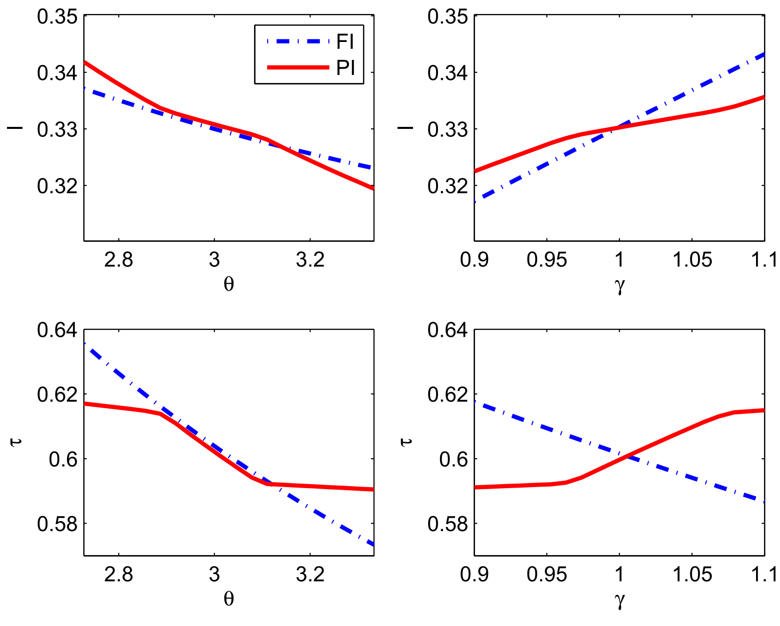

Fig. 10. Hours and taxes with PI and FI: high g. Left panels: top, hours as a function of the productivity shock (solid red line for PI, blue dashed-dotted line for FI); bottom, tax rate as a function of the productivity shock. Right panels: top, hours as a function of the demand shock; bottom, tax rate as a function of the demand shock. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

seems relevant for the understanding of the fiscal policy reaction to the financial crisis in 2008 and afterward, especially in countries like Spain and Italy, that arguably where in danger of getting close to the top of the Laffer curve, as testified by the fact that significant increases in taxes after 2009 did not raise the amount of revenue as much as it was desired by these governments.

6. Relation to previous results on optimal policy with PI

We now show how previous results in the literature on optimal policy under PI can be derived as special cases of our results. We then use our model of fiscal policy as a laboratory to illustrate the differences between these approaches and our results.

6.1. Existing results as special cases

We begin this analysis by formulating the following Corollary to Proposition 2 in the general formulation of Section 3.

Corollary 1. *Assume for some $s \in \mathcal{S}_{\mathcal{R}^*}$ one of the following Cases holds*

1. *(Invertibility) there is a unique $A \in \Phi$ such that $S(\mathcal{R}^*, A) = s$.*
2. *(exogenous signals) the signal s is independent of τ, $s = h(A)$, hence $h_\tau = 0$.*
3. *(linearity of h) h_τ is constant for all $A \in \Phi$ such that $S(\mathcal{R}^*, A) = s$.*

Then the general FOC (9) reduces to the FOC (5) at this s.

The proof is trivial: in all cases the term $\frac{1}{1-h_\tau^* \mathcal{R}^{*\prime}}$ cancels out because it is known given s so that this term can be pulled outside of the conditional expectation in (9). We now discuss how each case has been used in the literature and how it arises using variations of our example of Section.

Case 1 is related to the case discussed in Section 3.4. It generalizes that section as invertibility here can hold only for a subset of $\mathcal{S}_{\mathcal{R}^*}$. Existing results on optimal policy with FI apply directly to this case.

In Case 2, signals are exogenous. There is indeed PI, but the optimal policy can be found under separation. Some references for this case are in the first two paragraphs of the literature review (Section 2). For a concrete example, modify our fiscal policy example of Section so the signal is given by $s = \theta + \psi$, where ψ is noise (measurement error) and the state is $A = (\theta, \gamma, \psi)$. The associated first order condition (9) would simplify to equation

$$E[W_l|s] = 0 \tag{41}$$

as there would be no kernel related to the endogeneity of the distribution of the signal. Therefore, in this case there is separation: one can compute $f_{A|s}$ without knowledge of the optimal solution (for example, using the Kalman filter). This density is then applied to find the optimal solution with (41). In this case, the government would have to choose τ also under PI without knowing the revenue generated by this tax, so that perfect tax smoothing is not possible, but the signal extraction problem is exogenous to the optimal solution.

We discussed in Section 2 how the literature on learning and experimentation (or armed-bandit problems) could abstract from issues of GSE because it considered models where Case 2 can be applied. This is because this literature assumes a sequential structure between signals and actions, implying that signals are predetermined with respect to policy, and can thus be treated as exogenous in the optimization. We illustrate this in the context of our model of fiscal policy in Appendix C.4. However, in models with state-variables and signals generated by forward-looking behaviour, Case 2 is violated and our general result applies. We show in Section a modification of our main example that illustrates this point.

Case 3 highlights that for the kernel to be relevant, h_τ has to assume different values for different possible realizations of the shocks, conditional on the signal. It applies if h is linear in τ as in Svensson and Woodford (2004). Since this paper is the closest to ours, we now derive their results as a special case of the FOC (9) in the Corollary below.

First of all, it is helpful to distinguish two different roles played by the function h in our setup. First, h maps the shocks into a part of the allocation that enters the objective function W (in our fiscal policy example, hours). Second, it is what the PI literature (e.g., Svensson and Woodford, 2004) refers to as the "measurement equation", mapping the shocks (and the policy) into an observed signal.

Corollary 2. *If W is a quadratic function of τ and s and h is linear in τ, then optimal policy has the "certainty equivalence" property, that is, optimal policy under PI calls for applying the FI policy function to the conditional expectation of the shocks, satisfying (5).*

Here, a linear h means that h_τ does not depend on A. Let $A = (1, \theta, \gamma)$ and for simplicity of notation, consider the case where the objective function depends only on the policy τ and an endogenous signal s, and it involves only quadratic terms (i.e., the linear terms are zero):

$$W(\tau, s) = -\frac{\omega_\tau}{2}\tau^2 - \frac{\omega_s s^2}{2} \tag{42}$$

and the reaction function is

$$s = h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma \tag{43}$$

where the ω's are positive coefficients and the h's are constants.

The FOC under PI (9) becomes

$$E\left[\omega_\tau \tau + \omega_s h_\tau \left(h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma\right)|s\right] = 0 \tag{44}$$

which can be rewritten as

$$\tau = FE[A|s] \tag{45}$$

for a vector of coefficients $F \equiv \left(-\frac{\omega_s h_\tau h}{\omega_\tau + \omega_s h_\tau^2}, -\frac{\omega_s h_\tau h_\theta}{\omega_\tau + \omega_s h_\tau^2}, -\frac{\omega_s h_\tau h_\gamma}{\omega_\tau + \omega_s h_\tau^2}\right)$. It is easy to check that the FI solution satisfies $\tau = FA$. Comparing this FI solution with (45) shows that there is "certainty equivalence", that is, the government forms the best estimate of the state and behaves as if this estimate was certainty, or FI. Therefore, (5) holds.

However, as noted by Svensson and Woodford (2004), the "separation principle" does not hold. To see this, notice for example that $E[\theta|s]$ satisfies

$$E[\theta|s] = E\left[\theta|h + h_\tau FE[A|s] + h_\theta\theta + h_\gamma\gamma = s\right] \tag{46}$$

so that the density of $f_{\theta|s}$ depends on the optimal choice determined by F.

6.2. Exogenous signals in the fiscal policy model

In Section 5, we have considered optimal policy in the presence of an endogenous signal, and compared the equilibrium outcomes with the Full Information benchmark. To disentangle the role of *endogeneity* of information with respect to policy in our model from the more standard effects of *incompleteness* of information, we now consider a case in which information is incomplete, but exogenous (Case 2 of Corollary 1).

Specifically, we assume that the government receives an exogenous signal about the fundamental shocks, but this signal is contaminated by noise, or measurement error, thus leading to uncertainty about the true realization of the fundamental shocks. We denote the fundamental shocks $A^f \equiv (\theta, \gamma)$, and assume they are uniformly distributed as in our baseline example. We assume measurement error ϵ is also uniformly distributed, with range $\left[-\frac{\omega_i}{2}, \frac{\omega_i}{2}\right]$, where $i = 1, 2$ denotes the elements of ϵ, and ω_i is an inverse measure of precision of the signal. Thus, overall, the vector of unobservables is $A = (A^f, \epsilon)$.

The government receives a signal $s = A^f + \epsilon$. This implies a uniform posterior distribution for A_i^f in the range $\left[\max\left\{A^f_{i,\min}, s_i - \frac{\omega_i}{2}\right\}, \min\left\{A^f_{i,\max}, s_i + \frac{\omega_i}{2}\right\}\right]$ Clearly, the FI benchmark is nested as a the special case with $\omega_i = 0$ for $i = 1, 2$.

The government chooses its policy optimally subject to the measurability restriction $\tau = \mathcal{R}(s)$. Denoting by $\hat{f}(A^f, s)$ the posterior density of A^f, we can apply Corollary 1 and express the first order condition with respect to the tax rate in the first period as Eq. 5, where the expectation is taken using the posterior $\hat{f}(A^f, s)$. Relative to our general first order condition, there are two main differences. First, the kernel $\frac{1}{1-h_\tau \mathcal{R}'}$ drops out of the equation. Second, the posterior distribution used to integrate the marginal effects of taxes on welfare is exogenous.

To illustrate our results, we focus on the case in which measurement error applies only θ, and we use the same parameterization of the model as in Section 5.2, with $g = 0.6$.[33] In Fig. 11, we plot the optimal tax rate as a function of the posterior mean of θ, and we compare the outcome with two levels of informativeness of the signal ($\omega = 0.3$ and $\omega = 0.6$), as well as in the Full Information case (dashed line).[34] We choose these two values for ω because when $\omega = 0.6$, intermediate values of the signal are almost entirely uninformative, as the posterior dispersion of θ approximately equals the prior dispersion. Full Information coincides with a perfectly informative signal ($\omega = 0$) and, thus, $\omega = 0.3$ represents an intermediate case between these two extremes.

First, we notice that despite significant uncertainty about the true realization of θ, the government chooses a similar policy to the one that would be optimal if the signal was infinitely precise. To understand this result, notice that if the objective function was quadratic and the constraints were linear (as in Svensson and Woodford (2004), these two policies would exactly coincide, because of the certainty-equivalence property that is often used in the literature. Our utility function is log in consumption, thus violating certainty equivalence. However, the Full Information policy remains a good approximation to the optimal policy with measurement error.

With incomplete information, the government optimally chooses a slightly higher tax under Partial Information, and more so when the signal is less informative, to partially insure the low-productivity realizations that would lead to high deficit and high distortions in the second period. For this reason, the distance between the two lines is larger, in the central part of the figure, where the posterior range of θ with positive density is wider. In contrast, extreme signals are more precise because they allow the government to rule out many realizations in the range $[\theta_{\min}, \theta_{\max}]$. Overall, we find that the optimal policy is approximately linear and the deviations from Full Information are quantitatively small. In Fig. 12, we also display the counterpart of Fig. 5 for this version of the model, i.e., the set of admissible realizations of shocks conditional on a signal. In this case, this set is a straight line, and, importantly, exogenous with respect to policy.

From this analysis, we derive two conclusions. First, the sharp non-linearity in our general framework (see Fig. 7) arises because of endogeneity of signals, and not because of large non-linearities in the model environment. The fact that the signal is endogenous implies that, in different points in the state space, the government learns more or less about one or the other shock. On the contrary, how much is learned with exogenous signal is, clearly, exogenous. If the measurement

[33] We obtain similar results when we consider measurement error about γ

[34] As explained above, with our assumptions on the information structure, the posterior coincides with the signal in the intermediate region of the figure.

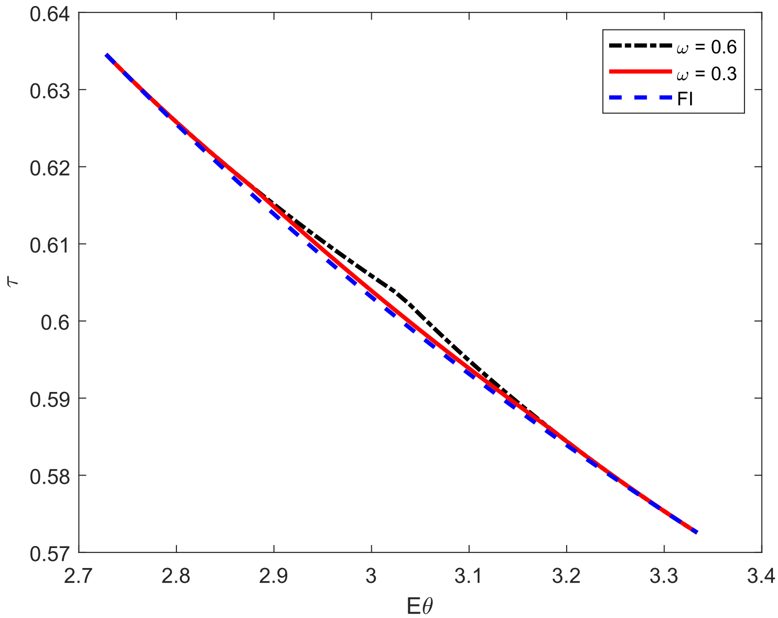

Fig. 11. Optimal policy with exogenous information. The figure illustrates the optimal tax rate with exogenous incomplete information (solid line for $\omega = 0.3$ and dashed-dotted line for $\omega = 0.6$) and with Full Information (dashed line). The posterior mean of θ is on the x-axis, and the policy on the y-axis.

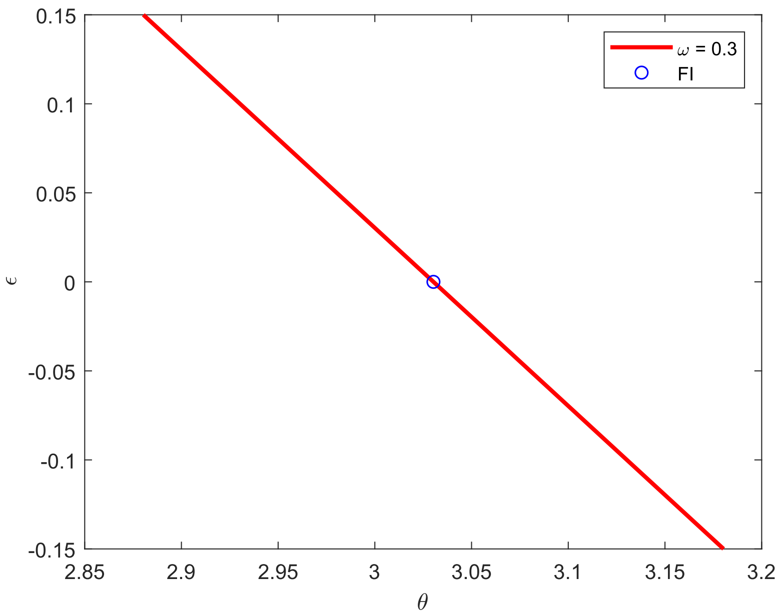

Fig. 12. Set of admissible shocks conditional on $s = E\theta$. The figure illustrates the set of admissible realizations of θ (x-axis) and measurement error ϵ (y-axis) conditional on a signal equal to the prior mean of θ. The red line refers to the case of measurement error with $\omega = 0.3$, the blue circle denotes the corresponding point under Full Information. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

error has constant variance, its informativeness does not change with government actions or different realizations of the fundamental shocks.

Second, our choice to use the Full Information policy as benchmark to compare the outcomes in our application is robust, in the sense that it also approximates well the equilibrium outcomes in the presence of exogenous incomplete information.

Furthermore, signal extraction does not change with the level of g. Therefore, models with exogenous signals ignore a number of issues related to the interaction between signal extraction and economic variables. These issues are instead highlighted under GSE.

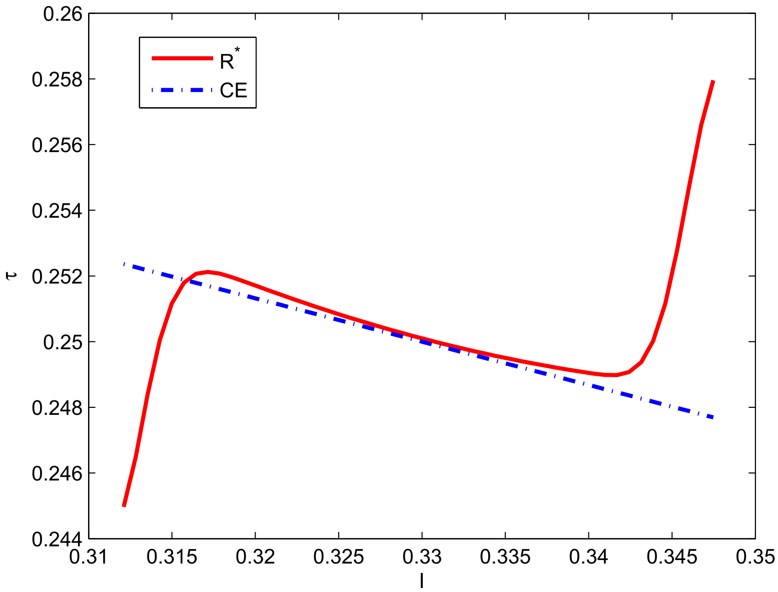

Fig. 13. Linear approximation, low g. Optimal tax rate as a function of the signal l, with $g = 0.25$. Red line: $\mathcal{R}^*$, blue dashed-dotted line: "certainty equivalent" linear approximation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

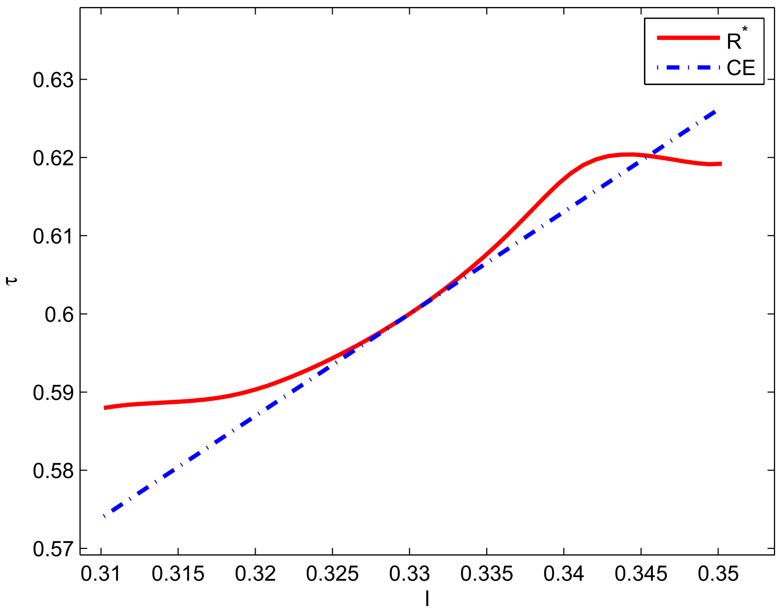

Fig. 14. Linear approximation, high g. Optimal tax rate as a function of the signal l, with $g = 0.6$. Red line: $\mathcal{R}^*$, blue dashed-dotted line: "certainty equivalent" linear approximation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

6.3. Linear-quadratic approximation in the fiscal policy model

We also solve for optimal policy in a linear-quadratic approximation (Case 3 of Corollary 1) of the fiscal policy model of Section 4, assuming the distribution of the shocks is a truncated normal. When we compare the solution to our fully non-linear results, we find that the approximation is accurate for intermediate realizations of the shocks, but inaccurate for large realizations of the shocks, as illustrated in Figs. 13 and 14 for the case of low and high government spending respectively. Details of this analysis are in Appendix C.5.

7. Heterogeneous households: Idiosyncratic shocks and aggregate signal

In our main example, households have better information than the government about the economy, in particular, they know the wage θ and the aggregate labor supply parameter γ, but the government does not. This ranking of information may seem undesirable, as governments arguably have better and more up to date information than individuals about wages and other macroeconomic variables. Accordingly, the literature often considers situations where heterogeneous agents have "dispersed information" about economic variables (Angeletos and Pavan, 2009), or where policy reveals superior information on the part of the government (Melosi, 2017).

This feature of our main example is for simplicity. Generically, in a heterogeneous agent setup individuals may have "less" information than government about aggregate variables and yet the government needs to solve a problem of GSE. More precisely, consider an economy with a continuum of individual agents indexed by $i \in [0,1]$. Their decisions depend on some underlying shocks x_i. Assume the government knows the distribution $x_i \sim \Phi_x(\cdot, A)$, but it ignores the exact value of some parameters A.[35] The value of A is policy-relevant as it potentially determines aggregate outcomes and welfare. If the government observes an endogenous signal s which provides valuable information about A, then optimal policy must be determined under GSE by forming the filter $f_{A|s}$, but since this information is beyond the reach of individuals they have "less" information about A and aggregate variables.

To be more precise, we now introduce idiosyncratic uncertainty in our main example and we show that the policy problem is a *special case* of our framework for optimal control with GSE for some specific h and W functions. This heterogeneous-agents setup clarifies that (i) our assumptions on the (asymmetric) information sets of government and private agents are natural in the context of a model of optimal macro policy in competitive equilibrium, and (ii) in models with heterogeneous agents and dispersed information, our GSE analysis is necessary to solve for the optimal policy.

Households indexed $i \in [0,1]$ are now heterogeneous with utility function

$$U_i \equiv u(c_{i1}) + \frac{1}{\gamma_i} v(l_{i1}) + \beta[u(c_{i2}) + v(l_{i2})] \tag{47}$$

where γ_i is an idiosyncratic preference shock.[36] Agent i's productivity is given by θ_i. Household i observes γ_i, θ_i as well as prices and taxes.

As in the main example, there is no uncertainty at date $t = 2$; households and the government can trade a riskless bond at price q; θ_i is the wage for household i in period 1; the (homogeneous) level of productivity is normalized to one in the second period; the government raises proportional income taxes at rates τ_1 and τ_2. The government observes aggregate hours, i.e. the endogenous signal is $s = \mathrm{E}_i(l_{i1})$, where the expectation operator E_i integrates over realizations of the idiosyncratic shocks. The resource constraints of the economy and budget constraint of the government are

$$\mathrm{E}_i c_{i1} + g = \mathrm{E}_i \theta_i l_{i1} \tag{48}$$

$$\mathrm{E}_i c_{i2} + g = \mathrm{E}_i l_{i2} \tag{49}$$

$$\tau_1 \mathrm{E}_i \theta_i l_{i1} + q\tau_2 \mathrm{E}_i l_{i2} = g(1+q) \tag{50}$$

The shocks (γ_i, θ_i) have a distribution $\Phi_{\gamma,\theta}(\cdot, A)$,. The government does not observe the means $A \equiv (\mu_\gamma, \mu_\theta)$, but it has a prior distribution F_A about these means and knows everything else about the distribution $\Phi_{\gamma,\theta}(\cdot, A)$,.

The government maximizes expected Benthamite utility, therefore its objective is to maximize $\mathrm{E}_A(\mathrm{E}_i(U_i))$.

Now we show that this is a special case of Optimal control under GSE as described in Section 3. Optimality conditions analogous to (26), (27) and (28) for the heterogeneous-agent case and the individual budget constraint give four equations to determine individual equilibrium allocations $(l_{i1}, c_{i1}, l_{i2}, c_{i2}) = G(\tau_1, \tau_2, q; \theta_i, \gamma_i)$, where G is the same across all agents. It is clear that each of the aggregate variables $\mathrm{E}_i(\cdot)$ that appear in (48)-(49)-(50) is a function of $(\tau_1, \tau_2, q; A)$. To see this, note for example that aggregate consumption in $t = 1$ is given by

$$\mathrm{E}_i(c_{i1}) = \int G_{c_1}(\tau_1, \tau_2, q, \cdot)\ d\Phi_{\gamma,\theta}(\cdot; A)$$

where G_{c_1} denotes the corresponding coordinate of G.

Therefore, given any feasible combination $(\tau_1; A)$, any pair of equations in (48)-(49)-(50) can be used to solve for the corresponding equilibrium $(\tau_2, q) = \mathcal{F}(\tau_1; \mu_\theta, \mu_\gamma)$ (the remaining third equation among (48)-(49)-(50) is satisfied due to Walras' law).

Thus, the *reaction function h* is given by

$$\mathrm{E}_i(l_{i1}) = s = h(\tau_1; A) \equiv \int G_{l_1}(\tau_1, \mathcal{F}(\tau_1; A), \cdot)\ d\Phi_{\gamma,\theta}(\cdot; A) \tag{51}$$

[35] A necessary condition for the government to be unaware of the exact value of A is that the government does not observe $\{x_i\}_{i \in [0,1]}$.

[36] We now assume the preference shock hits the labor disutility term instead of the consumption term to abstract from heterogeneity in intertemporal marginal rates of substitution, thereby simplifying the derivations.

and the *welfare function* W is given by

$$W(\tau_1; A) = \int w(\tau_1, \mathcal{F}(\tau_1; A), \cdot)\ d\Phi_{\gamma,\theta}(\cdot, A) \tag{52}$$

for utility function $w(\cdot) = u(G_{c_1}(\cdot)) + \frac{1}{\gamma_i} v(G_{l_1}(\cdot)) + \beta[\ u(G_{c_2}(\cdot)) + v(G_{l_2}(\cdot))\ ]$.

Hence, the model in this section is a special case of Optimal Policy under GSE for the above W, h: The government needs to use observed aggregate labor to update its information about the unknown means $A = (\mu_\theta, \mu_\gamma)$, taking into account the endogeneity of labor with respect to taxes.

To derive explicit expressions for functions h, W, we now specialize the example to the linear-quadratic case when $u(c) = c$ and $v(l) = \frac{Bl^2}{2}$ and assume that $\Phi_{\gamma,\theta}$ is log-normal with (known) variances $\sigma_\gamma^2, \sigma_\theta^2$ and γ_i, θ_i independent. Using (26), (27) and (28) we see that the reaction function is given by

$$s = B^{-1}(1-\tau_1)\mathrm{E_i}\theta_i\gamma_i = B^{-1}(1-\tau_1)e^{\mu_\theta+\mu_\gamma+\frac{\sigma_\theta+\sigma_\gamma}{2}} = h(\tau_1, A) \tag{53}$$

To find W, we express the government objective $\mathrm{E}_i(U_i)$ as a function of policy, signal, and aggregate shocks. By combining the aggregate resource constraint and the reaction function, and using log-normality of the idiosyncratic shocks, we get aggregate utility from consumption and disutility from labor at $t = 1$ as follows

$$\mathrm{E}_i c_{i1} = B^{-1}(1-\tau_1)\mathrm{E_i}\theta_i^2\gamma_i - g = se^{\mu_\theta+\frac{3}{2}\sigma_\theta^2} - g \tag{54}$$

$$-\frac{B}{2}\mathrm{E_i}\frac{l_{i1}^2}{\gamma_i} = -\frac{(1-\tau_1)^2}{2B}\mathrm{E_i}\theta_i^2\gamma_i = -\frac{(1-\tau_1)s}{2}e^{\mu_\theta+\frac{3}{2}\sigma_\theta^2} \tag{55}$$

while for $t = 2$ we have

$$\begin{aligned} \mathrm{E_i}c_{i2} &= B^{-1}(1-\tau_2) - g \\ -\frac{B}{2}\mathrm{E_i}l_{i2}^2 &= -\frac{(1-\tau_2)^2}{2B} \end{aligned} \tag{56}$$

Furthermore, combining (50) with the household optimality conditions and the definition of the signal we have

$$\tau_1 se^{\mu_\theta+\frac{3}{2}\sigma_\theta^2} + \beta B^{-1}\tau_2(1-\tau_2) = g(1+\beta) \tag{57}$$

Equation (57) is a quadratic equation in τ_2 and can thus be solved to express τ_2 as a function of (s, τ_1, A).[37]

An explicit formula for $W(\tau_1, s; A) \equiv \mathrm{E}_i\left[c_{i1} - \frac{B}{2\gamma_i}l_{i1}^2 + \beta\left(c_{i2} - \frac{B}{2}l_{i2}^2\right)\right]$ can be found plugging in (54), (55), (56) for τ_2 that solves (57).

Notice that in this example the objective function depends on the policy instrument τ also beyond its effect on the signal. This justifies our choice to derive our key Propositions for a general formulation in which both policy and signal enter the objective function independently.

8. Dynamic labor decision and predetermined taxes

We now describe a version of the model with a dynamic labor supply choice and a different timing assumption for taxes and show that our results are robust to this modification and do not depend on taxes and signals happening simultaneously. In this model, the policy instrument is predetermined, as a future tax rate is decided conditional on a currently realized signal (hours).[38] Because of learning-by-doing, the labor supply decision is dynamic and the signal function h is implicitly defined by an Euler equation for labor supply.[39] Agents have Rational Expectations and choose current hours depending, among other things, on taxes paid next period, making the signal endogenous to future policy. This shows that (i) the general issue of endogeneity of signals to policy does not depend on simultaneity of policy and signals and (ii) our results can be easily applied to models where the signal is determined by a dynamic equation, as opposed to an intratemporal condition as in the previous section.

8.1. Preferences and technology

The economy lasts for three periods $t = 1, 2, 3$. The representative agent's utility is

$$E\sum_{t=1}^{3}\beta^{t-1}[\gamma_t u(c_t) - v(l_t)] \tag{58}$$

[37] The fact that there are two solutions only reflects a standard Laffer curve effect, therefore we pick the τ_2 that delivers the highest aggregate welfare (typically it will be the lower τ_2).

[38] Clymo and Lanteri (2020) study optimal fiscal policy under FI with this timing assumption on policy decisions.

[39] Stantcheva (2015) studies optimal taxation with truthful revelation in a similar model with learning-by-doing and private information about productivity.

with $\gamma_1 = \gamma$, a random variable with distribution F_γ (demand shock), while $\gamma_2 = \gamma_3 = 1$.

The resource constraints in the three periods are

$$c_t + g = \theta_t \chi(l_{t-1}) l_t \tag{59}$$

for $t = 1, 2, 3$, given l_0, with $\theta_1 = \theta$, a random variable with distribution F_θ (productivity shock). Instead, $\theta_2 = \theta_3 = \bar{\theta}$, a known number. χ is a "learning-by-doing" function that maps work experience in period t into productivity in period $t+1$, with $\chi(l) > 0$, $\chi' > 0$, $\chi'' < 0$.

The household's budget constraints are

$$c_t + q_t b_t = \theta_t \chi(l_{t-1}) l_t (1 - \tau_t) + b_{t-1} \tag{60}$$

for $t = 1, 2, 3$, starting from zero initial debt b_0 and with $b_3 = 0$. We assume that learning-by-doing is fully internalized by individual agents. This makes the labor choice truly "dynamic", which is what we need for future taxes to affect current behavior.

8.2. Government

The government finances expenditure g using income taxes and debt. The government's budget constraints are

$$\tau_t \theta_t \chi(l_{t-1}) l_t + q_t b_t = g + b_{t-1} \tag{61}$$

for $t = 1, 2, 3$, with $b_3 = 0$.

Taxes are announced one period in advance. Hence, the tax rate τ_1 is a given of the problem and cannot be modified by the government. This captures the idea that a new government inherits the current fiscal policy stance from its predecessor. The tax rate τ_2 is announced by the government at $t = 1$, conditional on the realization of l_1, and τ_3 is announced at $t = 2$ (and will take the value that balances the government budget at $t = 3$).

8.3. Competitive equilibrium and signal determination

The household's first order conditions with respect to l_1 and c_1 can be combined to give the intertemporal optimality condition that defines the "reaction function" h.

$$v'(l_1) - \gamma u'(c_1) \theta \chi(l_0)(1 - \tau_1) - \beta u'(c_2) \bar{\theta} l_2 (1 - \tau_2) \chi'(l_1) = 0. \tag{62}$$

The last term of this expression is the key new element, and represents the discounted marginal benefit of experience through its future productivity effect, net of future taxes. Hours not only depend on current taxes through the standard substitution effect, but also on future taxes, because they affect the returns from learning-by-doing. The implicit function h maps the *expected* policy τ_2 into the signal l_1, for a given realization of the shocks and associated future level of hours l_2.

The other optimality conditions, as well as a brief discussion of the FI-Ramsey problem, are in Appendix C.3.

8.4. Optimal policy

We assume $u(c) = \log(c)$ and $v(l) = B\frac{l^\eta}{\eta}$. In the numerical example below, we set $\eta = 1$ and set the value of B such that average hours equal 1/3. The learning-by-doing function is given by $\chi(l) = Cl^\psi$ and we set $C = (1/3)^{-\psi}$, $\psi = 0.1$ and $l_0 = 1/3$. Assumptions on the distributions of the shocks are as in the previous sections.

The government chooses a function $\mathcal{R}$ such that $\tau_2 = \mathcal{R}(l_1)$ in order to maximize (58) subject to competitive equilibrium conditions. In order to apply the optimality condition derived in Proposition 3, two steps are required.

First, we need to compute the marginal effect of hours in the first period on hours in the second and third periods; this is done by using total differentiation of the first order conditions with respect to l_1, l_2 and l_3. Second, we need to compute the derivatives of the reaction function with respect to future taxes (h_{τ_2}) and with respect to the demand shock (h_γ); to do this, we totally differentiate the system of FOCs with respect to l_1, l_2, l_3, τ_2 and γ_1 and solve the obtained system for the derivatives of interest. This new additional step is required because the signal is determined through a dynamic condition that involves l_2, differently from the baseline model without learning-by-doing.

We compare the solution with low government spending ($g = 0.25$) to the solution with high government spending ($g = 0.60$). These two cases are illustrated in Figs. 15 and 16 respectively. As can be seen in the two figures, the results are qualitatively similar to those arising in the two-period model with simultaneity. In particular, both the presence of non-linearities and the different slope of $\mathcal{R}$ in the central region of the signal are evident in Figs. 15 and 16, where we plot the optimal tax function $\mathcal{R}$ (red line) as well as the set of Full Information couples (l_1, τ_2) (yellow region).

The results of this model show that for the endogeneity of signals to arise, it is not necessary that signals and policy apply in the same period. As agents are forward-looking, their actions respond to future taxes, leading to GSE issues in a dynamic context.

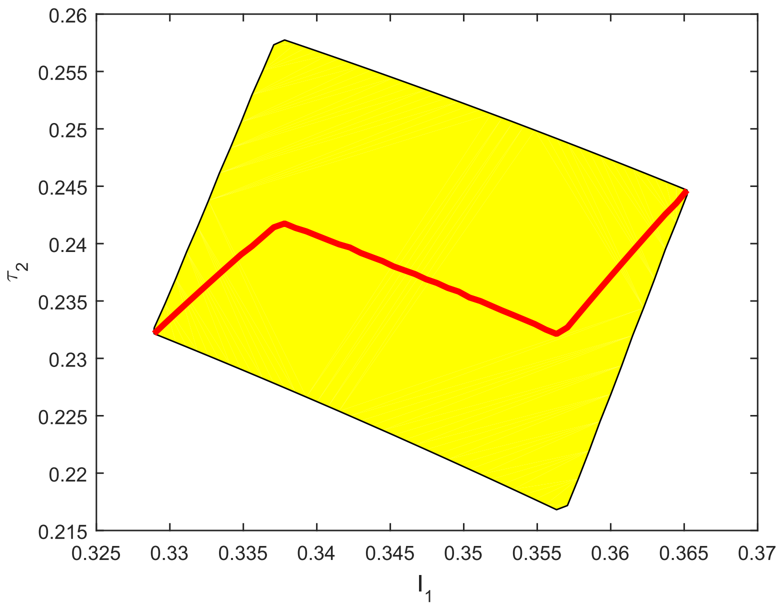

Fig. 15. Optimal policy in the 3-period model with low g. Optimal tax rate τ_2 as function of hours l_1 under PI, with $g = 0.25$. Red line: $\mathcal{R}^*$; yellow region: set of FI pairs (l_1, τ_2) for all possible realizations of (θ, γ). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

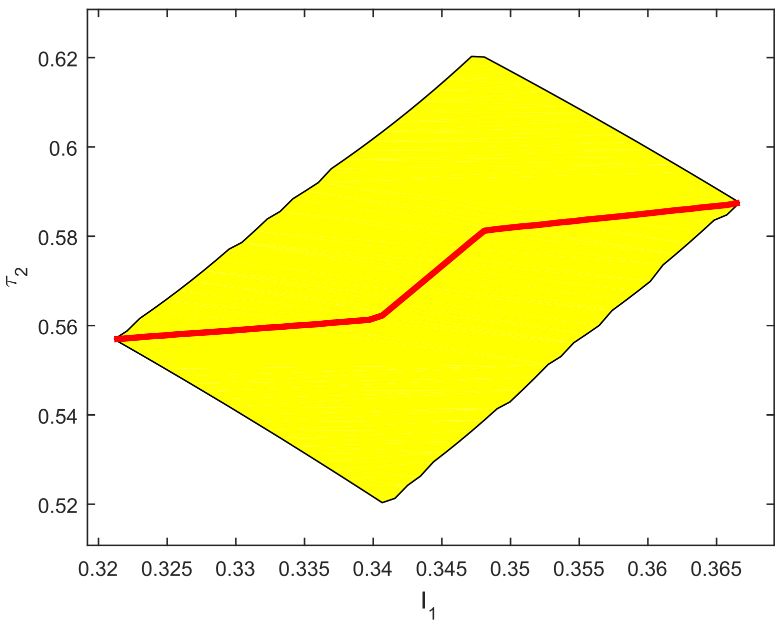

Fig. 16. Optimal policy in the 3-period model with high g. Optimal tax rate τ_2 as function of hours l_1 under PI, with $g = 0.6$. Red line: $\mathcal{R}^*$; yellow region: set of FI pairs (l_1, τ_2) for all possible realizations of (θ, γ). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

9. An infinite-horizon model with debt

In this section, we present an infinite-horizon version of the optimal fiscal policy model we have considered. We show that the key intuitions developed in the two-period model are still present and lead to interesting dynamics. In particular, optimal policy is more responsive to signals when the government is close to a debt limit. Moreover, the government sometimes reacts slowly to recessions and, as a consequence, needs to raise taxes by more and for a long time.

Our objective is to formulate the simplest infinite horizon model of fiscal policy where endogenous signals play a role. In order to avoid difficulties related to time inconsistency, we assume linear utility from consumption.[40] Moreover, we assume that shocks are i.i.d. over time, as this reduces the number of state variables and allows us to abstract from issues of optimal experimentation.[41]

9.1. Environment and full information

Preferences of the representative agent are given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t [\gamma_t c_t - v(l_t)] \tag{63}$$

where γ_t is a demand shock, i.i.d. over time.

The period t budget constraint of the representative agent is

$$c_t + q_t b_t = \theta_t l_t (1 - \tau_t) + b_{t-1} \tag{64}$$

where θ_t is an i.i.d. productivity shock. Note that the government can only issue real riskless bonds b_t.

The standard first order conditions for utility maximization are

$$\frac{v'(l_t)}{\gamma_t} = \theta_t (1 - \tau_t) \tag{65}$$

and

$$q_t = \beta \frac{\bar{\gamma}}{\gamma_t}. \tag{66}$$

where $\bar{\gamma}$ is the unconditional expectation of the demand shock γ.

The Ramsey government finances a constant stream of expenditure $g_t = g \, \forall t$ and chooses taxes and non-contingent one-period debt in order to maximize utility of the agent subject to the above competitive equilibrium conditions as well as the resource constraint $c_t + g = \theta_t l_t$. Under FI, the government can choose a sequence of taxes conditional on a sequence of shocks $A^t = (A_t, A_{t-1}, \ldots, A_0)$, where $A_t = (\theta_t, \gamma_t)$.

The period-t implementability constraint is

$$b_{t-1} = c_t - \frac{v'(l_t)}{\gamma_t} l_t + \beta \frac{\bar{\gamma}}{\gamma_t} b_t. \tag{67}$$

We introduce an upper bound on debt, $b^{\max}$. We assume that whenever debt goes above this threshold, the government pays a quadratic utility cost $\beta \frac{\chi}{2} (b_t - b^{\max})^2$ and we will set the parameter χ to be a large number in order to mimic a model with an occasionally binding borrowing constraint while still retaining differentiability of the problem.

The first order conditions for Ramsey allocations with respect to hours and debt are:

$$\gamma_t \theta_t - v'(l_t) + \lambda_t \theta_t - \frac{\lambda_t}{\gamma_t} \left[v'(l_t) + v''(l_t) l_t \right] = 0 \tag{68}$$

and

$$\lambda_t \frac{\bar{\gamma}}{\gamma_t} = E_t \lambda_{t+1} + \chi (b_t - b^{\max}) I_{[b^{\max}, \infty)}(b_t). \tag{69}$$

where λ_t is the Lagrange multiplier of constraint (64) and we denote by $I_{[b^{\max},\infty)}(b)$ the indicator function for the event $b > b^{\max}$.

Thanks to the assumption of linear utility from consumption, the Ramsey policy is time-consistent and allocations satisfy a Bellman equation that defines a value function $V^{FI}(b_{t-1}, A_t)$. Thus, optimal taxes are given by a time-invariant policy function $\tau_t = \mathcal{R}^{FI}(b_{t-1}, A_t)$

9.2. Partial information

At the beginning of each period t, the government observes the realization of the exogenous shocks of last period A_{t-1}, the value of its outstanding debt b_{t-1} and the realization of current labor l_t. Based on this information, but before knowing the value of A_t, it sets the tax rate τ_t. Formally, the choice of taxes at time t is contingent on - i.e. a function of - (A^{t-1}, l_t).

[40] Our model under FI is a small variation of Example 2 of Aiyagari et al. (2002), with linear utility from consumption and standard convex disutility from labor effort. As is well known, in the case of non-linear utility, the current bond price depends on future taxes, hence the optimal policy under full commitment would be time inconsistent, leading to some complications in the solution of optimal policy. Under FI, this issue was first addressed in Aiyagari et al. (2002).

[41] These issues are analyzed in the armed-bandit literature that we discussed in Section 2. For example Wieland (2000a), Wieland (2000b), Kiefer and Nyarko (1989), Ellison and Valla (2001).

Note that because of the i.i.d. assumption on the shocks A_t, information about outstanding debt summarizes all the information about past realizations that is relevant in terms of the objective function and the constraints of the Ramsey problem. As a consequence, debt is a sufficient state variable in addition to the current observed signal l_t. Hence the optimal policy has a recursive structure and taxes are given by a policy function $\tau_t = \mathcal{R}(b_{t-1}, l_t)$.

Define V as the value of utility (63) at the optimal choice for given initial debt, before seeing the realization of l_0. By a standard argument, the choice from period 1 onwards is feasible from period 0 onwards given the same level of debt. Therefore V satisfies the following Bellman equation

$$V(b) = \max_{\mathcal{R}:\Re^2 \to \Re_+} E[\gamma(\theta l - g) - \nu(l) + \beta V(b') - \beta \frac{\chi}{2}(b' - b^{\max})^2 I_{[b^{\max},\infty]}(b')] \tag{70}$$

where $b' = \frac{(b+g-\theta l + \frac{\nu'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}}$, $l = h(\mathcal{R}(b,l), \theta, \gamma)$ and $h(\tau, \theta, \gamma)$ is the level of hours that satisfies (65).

The only difference with respect to the reaction function in the two-period model is that now the government should recognize that debt affects labor indirectly through the tax rate.

With linear utility from consumption, the reaction function is given by

$$l = \nu'^{-1}(\theta\gamma(1-\tau)), \tag{71}$$

hence the locus of shocks realization conditional on l is increasing, that is, a certain level of hours can be generated by combinations of high productivity and low demand or vice versa. High θ and low γ is good news for the government for two reasons: revenue is high and the interest rate is low. On the other hand, low θ and high γ put pressure on the government budget constraint both by inducing low revenue and a high interest rate on the newly issued debt.

At this point it may be worth discussing how shocks and PI influence the optimal choice of taxes. Under incomplete markets, a sequence of adverse shocks (low θ) will lead to an increase in debt. This will be more so under PI than under FI, because under PI the government only learns that a low tax revenue materialized with a delay. The reason why b is an argument in $\mathcal{R}$ is that, in the presence of incomplete markets, debt grows after a few bad shocks, more so than under FI, therefore the government will have to increase the level of taxes for a given l_t to avoid debt from becoming unsustainable.

Note that in (70) we have substituted future debt using the budget constraint (67). It is important to highlight a key difference with respect to the FI problem: while in that case a choice of τ_t implied a choice of b_t, now, b_t is a random variable even for a given choice of τ_t. In other words, just like in the two-period model, the government is uncertain about how much debt will need to be issued and in particular must take into account that bad realizations of productivity may lead to a debt level above $b^{\max}$, if taxes are not sufficiently high.

To solve the model, we exploit its recursive structure, by solving for the PI first order condition at each point on a grid for debt and iterating on the value function of the problem. To see how this works, consider the objective function defined by the right-hand side of (70). For a given guess for the value function, this is just a function of observed labor to which we can apply Proposition 3 and obtain the general first order condition with PI.[42]

This first order condition involves the derivative $V'(b)$. In standard dynamic programming it is well known that an envelope condition applies that allows the simplification of the derivative of the value function. In Appendix D we show that an analogous envelope condition holds under our PI model so that

$$V'(b) = E\left[\frac{\gamma}{\bar{\gamma}} V'(b')\right] - \chi \frac{\gamma}{\bar{\gamma}}(b' - b^{\max}) I_{[b^{\max},\infty)}(b'). \tag{72}$$

Hence by solving the first order condition using (72) and iterating on the Bellman equation (70), we can approximate the optimal policy. In the next subsection, we show some numerical results. While the model is not meant to be a quantitative model of fiscal policy, it is helpful to illustrate the key properties of optimal policy with endogenous signals in a dynamic model with debt.

9.3. Numerical results: Non-linearities and delayed adjustments

In order to parametrize the economy, we assume quadratic disutility from labor, and the other parameters are as in the two-period model. The shocks are uniformly distributed on a support of ± 5% from their means, implying a standard deviation of 2.89%. The debt limit is set at 20% of average output.

We now illustrate two interesting properties of optimal policy in this model. First, the response of taxes to endogenous signals is quite non-linear and it depends on the level of debt. When debt is low, optimal policy calls for smooth taxes regardless of the realization of labor. When debt is close to the limit, taxes become highly responsive to the signal, in particular in the central region of the signal, where there is higher uncertainty on the state of the economy. Fig. 17 illustrates this property by plotting taxes against labor, as given by the function $\mathcal{R}^*(b, \cdot)$ evaluated at different values of debt.[43] It can be seen that for lower levels of debt (namely for $b = 0$ and $b = 0.1$), taxes are relatively flat with respect to the signal.

[42] The FOC is explicitly shown in Appendix D.

[43] Notice that taxes are in general decreasing in labor, different from our two-period model, because, with linear-quadratic utility, hours are increasing in θ, since the only effect of this shock is the substitution effect.

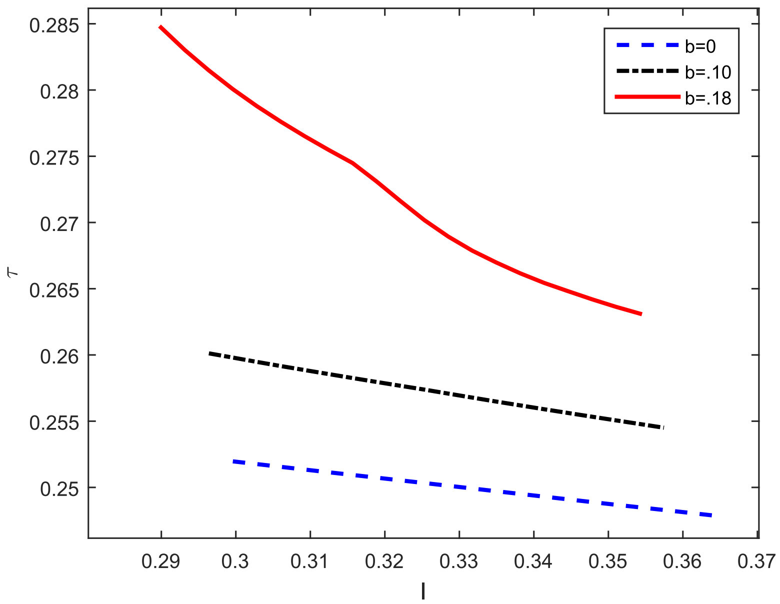

Fig. 17. Tax policy for different levels of debt. Optimal tax rate as a function of the signal l, for different values of outstanding debt b. Red solid line: $b = 0.18$; black dashed-dotted line: $b = 0.10$; blue dashed line: $b = 0$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

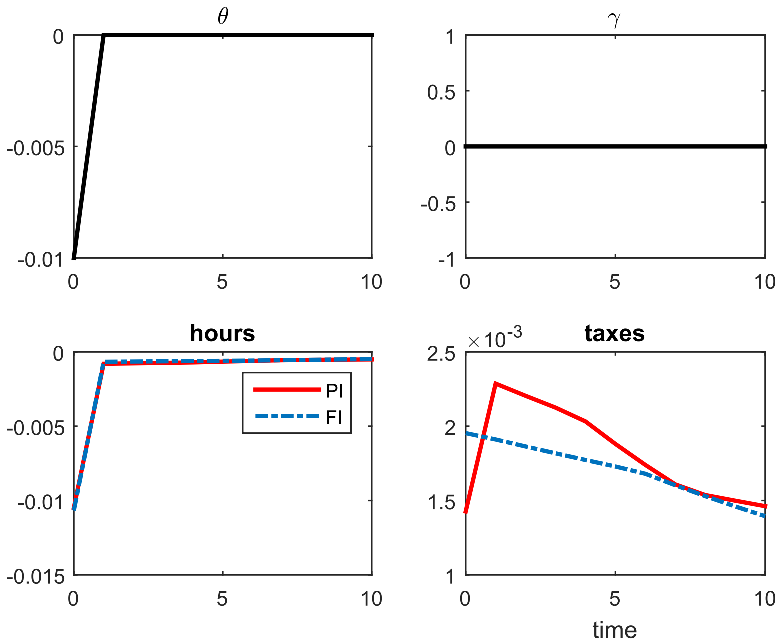

Fig. 18. Impulse response function: θ shock. Impulse response function to a negative 1% productivity shock. Top left panel: productivity shock; top right panel: demand shock; bottom left panel: hours; bottom right panel: tax rate. Solid red line: PI; blue dashed-dotted line: FI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

However, close to the debt limit (b=0.18), they become highly responsive, and with a higher slope in the middle region of realizations of hours.

Second, optimal policy with PI can rationalize delayed fiscal adjustments following negative productivity shocks. While optimal policy with FI can respond contemporaneously to such shocks, under PI the government can only respond with a lag, after observing a higher than expected level of debt, leading to a delayed, persistent tax adjustment. This feature of the model is illustrated in the Impulse Response Function in Fig. 18.[44]

We initialize both the FI and the PI economy at a level of debt equal to.12. When the productivity shock hits, the FI government is successful at making a tax correction to keep the economy sufficiently far from the borrowing constraint. On the other hand, the initial tax increase under PI is small, and the PI government imposes a fiscal adjustment only with a delay, after it learns that debt had increased. As a consequence of this delay, the postponed tax hike needs to be larger. Notice that because of market incompleteness, both the tax change under FI and PI are highly persistent, even though the shock is transitory as the government smooths taxes over time (Aiyagari et al., 2002).

10. Conclusion

Optimal policy with Partial Information is usually derived under a separation result, implying that the signal extraction problem can be solved ahead of optimization. However, separation fails to hold generically and this limits the scope of policy analysis with Partial Information.

We derive a method to solve models of optimal policy with Partial Information generically. In our setup, there is no separation between the optimization and signal extraction problem, as the optimal decision influences the distribution of the shocks conditional on the observed endogenous signal. Therefore, the signal extraction and optimization problem are solved consistently and simultaneously.

The method works in general and we show an easily applicable numerical algorithm to find a solution. We also show that Partial Information on endogenous variables matters as some revealing non-linearities appear in very simple models and the signal extraction problem is endogenous to economic variables. The non-linearities are due to the fact that in different regions of the observed signal the information revealed about the underlying state changes in a non-linear fashion.

Optimal fiscal policy under General Signal Extraction calls for smooth tax rates across states when the government budget is under control, and for regions of large response to aggregate data when the economy is close to the top of the Laffer curve or to a borrowing limit. Uncertainty about the state of the economy helps to understand the slow reaction of some European governments to the Great Recession, followed by sharp fiscal adjustments and prolonged downturns: as the signal worsens and debt increases, the government becomes more certain about lower productivity. The tax choice can be highly non-linear as a function of the signal even in setups that are essentially linear.

Our methodology allows to study the effects of Partial Information about the state of the economy in models that feature important non-linearities. While we have illustrated the technique in a simple model of optimal fiscal policy, the methodology can be easily extended to address many other questions in optimal policy, such as models with heterogeneity and unobserved idiosyncratic shocks and models with unobserved belief fluctuations.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2021.01.002.

References

Adam, K., 2004. On the relation between bayesian and robust decision making. Journal of Economic Dynamics and Control 28 (10), 2105–2117.
Aiyagari, S., Marcet, A., Sargent, T., Seppala, J., 2002. Optimal taxation without state-contingent debt. Journal of Political Economy 110 (6), 1220–1254.
Amador, M., Weill, P.-O., 2010. Learning from prices: public communication and welfare. Journal of Political Economy 118 (5), 866–907.
Angeletos, G., Iovino, L., La'O, J., 2015. Efficiency and Policy with Endogenous Learning. Technical Report.
Angeletos, G., Pavan, A., 2009. Policy with dispersed information. J. Eur. Econ. Assoc. 7 (1), 11–60.
Aoki, K., 2003. On the optimal monetary policy response to noisy indicators. J. Monet. Econ. 50 (3), 501–523.
Baxter, B., Liam, G., Wright, S., 2007. The Endogenous Kalman Filter. Technical Report. Birbeck School of Economics, Mathematics and Statistics.
Baxter, B., Liam, G., Wright, S., 2011. Invertible and non-invertible information sets in linear rational expectations models. Journal of Economic Dynamics and Control 35, 295–311.
Clymo, A., Lanteri, A., 2020. Fiscal policy with limited-time commitment. The Economic Journal 130 (627), 623–652.
Ellison, M., Valla, N., 2001. Learning, uncertainty and central bank activism in an economy with strategic interactions. J. Monet. Econ. 48 (1), 153–171.
Hansen, L., Sargent, T., 2012. Three types of ambiguity. J. Monet. Econ. 59 (5), 422–445.
Kiefer, N., Nyarko, Y., 1989. Optimal control of an unknown linear process with learning. Int Econ Rev (Philadelphia) 30 (3), 571–586.
Kocherlakota, N., 2010. The new dynamic public finance. Princeton University Press.
Lucas, R.J., 1972. Expectations and the neutrality of money. J. Econ. Theory 4 (2), 103–124.
Lucas, R.J., Stokey, N., 1983. Optimal fiscal and monetary policy in an economy without capital. J. Monet. Econ. 12 (1), 55–93.
Melosi, L., 2017. Signalling effects of monetary policy. Review of Economic Studies 84 (2), 853–884.
Mirman, L., Samuelson, L., Urbano, A., 1993. Monopoly experimentation. Int. Econ. Rev. (Philadelphia) 34 (3), 549–563.
Morris, S., Shin, H., 2018. Central bank forward guidance and the signal value of market prices. American Economic Association Papers and Proceedings 108, 572–577.
Muth, J., 1960. Optimal properties of exponentially weighted forecasts. J. Am. Stat. Assoc. 55 (290), 299–306.
Nimark, K., 2008. Dynamic pricing and imperfect common knowledge. J. Monet. Econ. 55, 365–382.
Nimark, K., 2008. Monetary policy with signal extraction from the bond market. J. Monet. Econ. 55, 1389–1400.
Pearlman, J., 1992. Reputational and nonreputational policies under partial information. Journal of Economic Dynamics and Control 16 (2), 339–357.

[44] These non-linear impulse response functions are computed as percentage deviations from the path that would arise absent all shocks.

Prescott, E., 1972. The multi-period control problem under uncertainty. Econometrica 40 (6), 1043–1058.
Stantcheva, S., 2015. Learning and (or) Doing: Human Capital Investments and Optimal Taxation. Technical Report. NBER WP 21381.
Svensson, L., Woodford, M., 2003. Indicator variables for optimal policy. J. Monet. Econ. 50 (3), 691–720.
Svensson, L., Woodford, M., 2004. Indicator variables for optimal policy under asymmetric information. Journal of Economic Dynamics and Control 28 (4), 661–690.
Swanson, E., 2006. Optimal nonlinear policy: signal extraction with a non-normal prior. Journal of Economic Dynamics and Control 30 (2), 185–203.
Townsend, R., 1983. Forecasting the forecasts of others. Journal of Political Economy 91 (4), 546–588.
Van Nieuwerburgh, S., Veldkamp, L., 2006. Learning asymmetries in real business cycles. J. Monet. Econ. 53 (4), 753–772.
Wieland, V., 2000. Learning by doing and the value of optimal experimentation. Journal of Economic Dynamics and Control 24 (4), 501–534.
Wieland, V., 2000. Monetary policy, parameter uncertainty and optimal learning. J. Monet. Econ. 46 (1), 199–228.