

Optimal Policy with General Signal Extraction*

Esther Hauk[†] Andrea Lanteri[‡] Albert Marcet[§]

November 2019

Abstract

We study optimal policy when the planner has partial information in a general setup where observed signals are endogenous to policy. In this context, signal extraction and policy have to be determined jointly. We derive a general non-standard first order condition of optimality from first principles and we use it to find numerical solutions. This first order condition allows us to identify widely-used special cases in the literature in which the signal extraction and the optimal decision problems can be solved separately, using the well-known separation principle. Our general setup, which does not feature any separation, is relevant for most available dynamic models in macro. We apply our results to a model of fiscal policy and show that optimal taxes are often a very non-linear function of observed hours, calling for tax smoothing in normal times, but for a strong fiscal reaction to output in a deeper recession. This non-linearity arises because signal extraction interacts differently with optimal policy depending on the range of observed signals. The non-linearity is stronger near the top of the Laffer curve or near a debt limit. In a fully dynamic model taxes react with a delay to adverse deficit shocks due to partial information, and this can lead to larger low frequency fluctuations.

JEL codes: C63, D82, E60, H60.

Keywords: Optimal Policy; Partial Information; Separation; Calculus of Variations; Fiscal Policy.

*The authors appreciate helpful comments from 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 CI-CYT project number ECO2015-67171-P (MINECO / 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 (SEV-2015-0563). 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 Union's Seventh Framework Programme FP7/2007-2013 GA number 324048-APMPAL and Horizon2020 GA number 788547 (APMPAL-HET).

[†]Institut d'Anàlisi Econòmica-CSIC, MOVE and Barcelona GSE; email: esther.hauk@iae.csic.es

[‡]Department of Economics, Duke University; email: andrea.lanteri@duke.edu

[§]University College London, Barcelona GSE, CEPR; email: marcet.albert@gmail.com

“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, 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)

1 Introduction

Inferring the underlying state of the economy is a key practical difficulty in setting macroeconomic policy, as exemplified by our opening quote. This is 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 very large deficits, leading later to even larger fiscal adjustments during the worst part of the recession in some countries. The 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.

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 Partial Information 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 led to many interesting applications, the assumptions needed for separation typically are not satisfied in modern dynamic models.

In this paper, we derive an optimality condition from first principles that allows to compute the optimal policy in the general case. This optimality condition is non-standard, as it does not set the average derivative of the objective function to zero. Instead, the derivative at each realisation is weighted by a certain kernel. Previous results in the literature arise as special cases.

We apply this formulation in a model of optimal fiscal policy. 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 to supply). To make the issue of hidden information relevant we assume incomplete insurance 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 some endogenous signal, such as output or hours worked, s . We argue that Partial Information 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 in normal times, but taxes become highly responsive to signals when the government has strong evidence that the worst possible realizations of the shocks materialize.

This non-linearity arises because the signal extraction problem interacts in a different way with optimal policy depending on the range of observed signals. If the signal is close to its worst possible realisation policy-makers are quite sure that the worse combination of shocks A has materialised. In that case they are quite certain that a very low tax revenue is very likely so that they raise taxes today in order to avoid extremely high taxes tomorrow. The non-linearity is stronger with very high government spending, when future taxes could be close to the maximum of the Laffer curve, as was arguably the case in some European economies in the recent crisis. Several European governments in the Great Recession made large fiscal adjustments a few years after the start of the downturn. This delay led 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 and not necessarily of irresponsible politicians, hence the model may

¹Svensson 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 4.7 derives the result of Svensson and Woodford from our general optimality condition.

be useful to distinguish an irresponsible government from one that simply has partial information.

We also show that hidden information can be a driver of tax smoothing, as the contemporaneous response of taxes to underlying shocks is small relative to the case of Full Information. This identifies a reason for tax smoothing across states that is quite different from the standard results in presence of complete markets.

General Signal Extraction is present in most dynamic models under rational expectations. We show two versions of our model to demonstrate how GSE arises generically. In our main example, separation fails because the income tax rate τ , is determined simultaneously with the signal s , which in our example is the level of hours worked. It may seem that one can change the model slightly for separation to hold: if taxes have to be determined one period before they are implemented, then separation would hold in that model. However, we show that if labor supply involves a dynamic decision, e.g. because of learning-by-doing, then we are back in a setup where separation fails and GSE needs to be applied. In our second example with learning-by-doing we get similar results as in the main example.

Finally, we build an infinite-horizon model. 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 after a bad shock to the deficit is realised, thus the high-frequency response of taxes to shocks resembles tax smoothing. This is due to the fact that optimal policy averages over all the possible unobserved contingencies for A , as it occurred in the two-period model. However, this delayed response to shocks can lead to accumulation of large debt over time, eventually leading to large fiscal adjustments and high volatility at lower frequencies. Therefore tax implications are quite different from optimal tax under market completeness and Full Information (Lucas and Stokey, 1983) since perfect tax-smoothing is not found, and different from market incompleteness and Full Information (Aiyagari et al., 2002) where the delayed response would never occur.

It should be noted that the technique we develop can be applied to many interesting open questions in several fields. 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, about the household distribution of income and wealth or about agents' expectations.

The remainder of the paper is organized as follows. The related literature is discussed in Section 2. Section 3 introduces our two-period optimal fiscal policy model with incomplete markets and Partial Information. Section 4 contains the main theoretical contribution. It provides the first order condition for a general static model, it shows how this can be used to compute optimal solutions, it compares the Full Information solution with the Partial Information solution and discusses a case of "invertibility" when the two solutions coincide. We also study when "certainty equivalence" holds. In Section 5 we apply these results to our Ramsey fiscal policy problem.

Section 6 introduces a model with predetermined taxes and dynamic labor choice because of learning-by-doing, showing that our results are robust to changes in timing assumptions. Section 7 presents the infinite-horizon model with non-contingent debt. Section 8 concludes.

2 Related literature

Partial Information 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 wide literature considers competitive agents who use prices or other aggregate endogenous variables as signals of unknown information, as early as Lucas’s (1972) seminal paper, or Townsend’s (1983) contribution about beliefs on other agents’ information. For more recent papers with a focus on macro policy see Angeletos and Pavan (2009) and Angeletos et al. (2015). In this literature, the choices of competitive agents do not interact with price determination (in the agents’ mind), therefore agents’ filtering problem can be solved using standard techniques, and the issue of endogenous signal extraction that we address does not arise.

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. Baxter et al. (2007, 2011) derive an “endogenous Kalman filter” for all 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.

Closest to our work is Svensson and Woodford (2004). They 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 Partial Information the government applies the Full Information policy to its best estimate of the state. Aoki (2003) applies these results to optimal monetary policy with noisy indicators on output and inflation. Nimark (2008) 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.

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, 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.² We provide an example that appears to be amenable to linearisation but where the correct solution is highly non-linear in nature.

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, 2000b) and Ellison and Valla (2001). For an application to monopoly behavior see Mirman et al. (1993). Van Nieuwerburgh and Veldkamp (2006) use a similar learning framework with non-linearities to explain business-cycle asymmetries. In these papers, the density of the signal is exogenous to agents' actions, although the planner can decide how much to learn about this density by its actions. These papers can abstract from the issue of General Signal Extraction because of special assumptions: signals are predetermined with respect to policy, independent of future policies, for instance because of the absence of endogenous state variables. In Section 4.1 we argue that in most modern dynamic economic models this simplification is not available, and the issue of General Signal Extraction that we consider would arise generically.³

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 this approach sidesteps the issue of finding the filter $f_{A|s}$ at the optimal policy \mathcal{R} . Adam (2004) shows that the min-max criterion arises from a sequence of planners that have expected utility and are more and more risk-averse. Our setup is available for *any* level of risk aversion of such a planner.

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 4.5, endogenous filtering is not an issue when private information is revealed, therefore this literature abstracts from the issue of the current paper. On the other hand, this literature

²Optimal 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.

³In Appendix C.4 we provide a modification of our model in Section 3 that restores the simplifications in this literature and that may help clarify this issue.

assumes that the policy function \mathcal{R} depends on individual choices. Hence, agents react strategically to influence the policy action τ . We abstract from this issue by introducing atomistic agents that take the government action τ as given. It would be of interest to apply our results on signal extraction in a setup where agents react to the policy function \mathcal{R} and not only to the planner's action τ . This is left for future research.

3 A simple model of optimal fiscal policy

We now present a simple model of optimal policy with two shocks, to supply and demand respectively. We believe that this model and its extensions in Sections 5, 6 and 7 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 4 in a more general setup could be applied in many fields, also outside of macro.

Specifically, our first 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.

3.1 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

$$E [U (c_1^i, l_1^i, c_2^i, l_2^i; \gamma)] \tag{1}$$

with

$$U (c_1^i, l_1^i, c_2^i, l_2^i; \gamma) = \gamma u (c_1^i) - v (l_1^i) + \beta [u (c_2^i) - v (l_2^i)]$$

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.

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

$$y_t = \theta_t l_t \tag{2}$$

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 $(\tau_1, \tau_2, b^g) : \Phi \rightarrow R^3$ 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 $(c_1, c_2, l_1, l_2, b) : \Phi \rightarrow R^5$ knowing the fiscal policy and the bond price function $q : \Phi \rightarrow R$. 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

$$c_1 + qb = \theta l_1(1 - \tau_1) \tag{3}$$

$$c_2 = \theta_2 l_2(1 - \tau_2) + b \tag{4}$$

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) .

3.2 Competitive equilibrium

Here we provide a definition of competitive equilibrium. The definition is standard in the literature, it is common to both the Full Information (FI) and the Partial Information (PI) equilibria that we analyze.

Definition 1 *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 (1) subject to (3) and (4). In addition, bonds and goods markets clear, so that $\tilde{b}^g = b$ and

$$c_t + g = \theta_t l_t \quad \text{for } t = 1, 2. \quad (5)$$

This definition embeds competitive equilibrium relations insuring 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

$$\frac{v'(l_1)}{u'(c_1)} = \theta \gamma (1 - \tau_1) \quad (6)$$

$$\frac{v'(l_2)}{u'(c_2)} = \theta_2 (1 - \tau_2) \quad (7)$$

$$q = \beta \frac{u'(c_2)}{\gamma u'(c_1)} \quad (8)$$

As anticipated, the demand shock enters the first period labor supply decision described by (6) as well as the bond pricing equation (8). A competitive equilibrium is fully characterized by equations (3) to (8).

3.3 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 best policy for households.

We first give the standard definition when both government and consumers observe the realization of A .⁴

Definition 2 *A **Full Information (FI-) Ramsey equilibrium** is a fiscal policy (τ_1, τ_2, b^g) that achieves the highest utility (1) when allocations are determined in a competitive equilibrium.*

Our interest is in studying optimal taxes under PI. More precisely, 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 . It is important to highlight that the signal depends on aggregate outcomes. In our leading example, we assume $s = l_1 = \int l_1^i di$. The distinction between aggregate and individual choices does not matter for our computations due to the homogeneous agents assumption, but it is important for the interpretation of our equilibrium that the signal only depends on aggregate outcomes, as this justifies that agents take the policy action τ as given.

⁴These definitions take for granted that we only consider tax policies for which a competitive equilibrium exists and is unique.

Definition 3 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

$$\tau_1 = \mathcal{R}(s) \text{ for all } A \in \Phi \quad (9)$$

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

We are interested in the case when (9) prevents the PI-Ramsey equilibrium from achieving the FI outcome.

Note that the definition of competitive equilibrium implies that consumers know the map $\tau_1 : \Phi \rightarrow R$. Hence, in a PI-Ramsey equilibrium, agents' perception of how policy is set is consistent with equation (9). An interpretation is that the government posts its policy \mathcal{R} , agents understand this policy and they know the 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 (9) 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.⁵ 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} .⁶

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 (5), the following implementability condition holds

$$\gamma u'(c_1) c_1 - v'(l_1) l_1 + \beta [u'(c_2) c_2 - v'(l_2) l_2] = 0. \quad (10)$$

The standard approach to find Ramsey policy under FI is to maximize (1) subject to (10) 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 (5) for $t = 1$ to substitute out consumption in (6), we add the constraint

$$\frac{v'(l_1)}{u'(\theta l_1 - g)} = \theta \gamma (1 - \tau_1), \quad (11)$$

⁵This 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*. Therefore in that literature agents internalize the effect that their actions have on policy actions. Using our notation, the government chooses an optimal individual tax $\tau_t^i = \mathcal{R}(l_t^i)$.

⁶We consider the case of predetermined taxes in Section 6

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

$$l = h(\tau, A) \tag{12}$$

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 (5) and (10) 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 (1), welfare for each A can then be written as

$$W(l; A) \equiv U(c_1, l, c_2, l_2; \gamma) \tag{13}$$

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

$$\max_{(\tau, l): \Phi \rightarrow R^2} E[W(l; A)] \tag{14}$$

$$\text{s.t. (12)} \tag{15}$$

Obviously the result is the same as to maximize (1) subject to (10) given A .⁷

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 (9) appears, that $\tau_1 = \tau$ determines $l_1 = l$ in equilibrium through (12) and that the choice is over the policy \mathcal{R} .

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

$$\max_{\mathcal{R}: R \rightarrow R; (\tau, l): \Phi \rightarrow R^2} E[W(l; A)] \tag{16}$$

$$\text{s.t. (12)}$$

$$\tau = \mathcal{R}(l)$$

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 *general signal extraction*, where the filtering problem $f_{A|l}$ has to be solved *jointly* with the optimal choice for the policy function \mathcal{R} . We solve this problem in Section 4.

⁷Strictly speaking l should be constrained to belong to the set of feasible l 's, we are explicit about this technicality in the next section.

3.4 The economic consequences of PI for taxation policy

Before giving a mathematical solution, it is worthwhile discussing the economic issues raised by PI in the fiscal policy example we use. As is well known, the optimal FI policy is one of tax smoothing over time as the government wants to spread the distortions equally in the two periods. In the case of CRRA preferences

$$u(c) = \frac{c^{1+\alpha_c}}{1+\alpha_c}, \quad v(l) = B \frac{l^{1+\alpha_l}}{1+\alpha_l}$$

for $\alpha_c \leq 0$, $\alpha_l, B > 0$, 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

$$\tau\theta l_1 - g + \beta \frac{u'(c_2)}{\gamma u'(c_1)} (\tau\theta_2 l_2 - g) = 0. \quad (17)$$

It is clear from (17) 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. Clearly, 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 in period 2, the period 1 deficit will materialize, the government 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. The optimal design of these instruments requires a simultaneous determination of the density of taxable income and the policy. The next section studies a generic problem that allows the determination of taxes under limited endogenous information.

4 Optimal Control under GSE

We label maximization problems of the type (16) as “Optimal Control under General Signal Extraction” (GSE). The key difficulty is that the distribution of the exogenous shocks conditional on the signal, $f_{A|s}$, is endogenous to the policy decision \mathcal{R} . At the same time, the optimal choice for \mathcal{R} depends on $f_{A|s}$, therefore separation does not apply. This section provides a formal treatment of Optimal Control under GSE. We first argue that GSE arises generically in most modern dynamic economic models. We then provide results concerning existence of a solution, as well as first and second order optimality conditions that can be used to easily compute the optimal policy.

We can see that the filter $f_{A|s}$ depends on tax policy in the context of the example in Section 3. Combining equations (12) and (9) we see that the shocks $A = (\gamma, \theta)$ can be mapped into a value $l = h(\mathcal{R}(l), A)$ for a given \mathcal{R} . Using the chain rule, it is clear that the conditional density $f_{l|A}$ depends on \mathcal{R} . Since $f_{A|l}$ depends on $f_{l|A}$ (through Bayes’ rule) it depends on \mathcal{R} , so that the filtering and optimization problems can not be separated.

Notice that h and \mathcal{R} can be interpreted as reaction functions. In the above formulation of Optimal Control under GSE there is a hierarchy: h is given, while \mathcal{R} is chosen optimally given h .⁸ This is a natural assumption when, as in Ramsey policy, agents are atomistic and the signal is an aggregate variable. Extending the setup to cases when reaction functions are determined simultaneously is left for future research.⁹

4.1 GSE: a Generic Issue

The literature has, to date, avoided the issue of GSE, often by assuming that policy decisions are predetermined.¹⁰ We show how in some models it is possible to simplify things by assuming a certain timing of policy decisions that restores separation. However, we argue that introducing state variables gives rise to GSE generically, even if policy is predetermined. Therefore, GSE arises naturally under Partial Information in most economic dynamic models.

In the example of Section 3 GSE arises because τ and l are determined simultaneously, as is standard in Ramsey equilibria. A shrewd economist might notice that separation can be restored by assuming that taxes are predetermined, i.e. that τ_t

⁸This may be thought of as a Stackelberg-reaction-function game where the “leader” (in Section 3, the government) chooses its policy (or reaction) function \mathcal{R} taking as given the reaction h of the “follower” (in Section 3, the atomistic consumers), while the follower takes the action of the government (τ) as given.

⁹The literature on supply function equilibria (Klemperer and Meyer, 1989) considers supply functions that are consistent with the other players’ choice. The issue of GSE that we address arise in principal/agent problems and supply function equilibria with multiple sources of uncertainty and no-revelation.

¹⁰One exception is Svensson and Woodford (2004) and subsequent literature, who exploit linearity.

must be chosen in period $t - 1$ after observing l_{t-1} . This might seem like an attractive modelling alternative as policy is still determined under PI (the government chooses taxes without knowing the revenue that a certain tax rate will raise) while avoiding the technicalities of this paper.¹¹

However, consider now the case in which output y_t depends on a factor of production determined in the past, say, physical or human capital. Even if τ_t is predetermined and contingent on y_{t-1} , this tax is likely to influence tomorrow's labor supply l_t , to affect the payoff of factors of production at $t - 1$ and, under Rational Expectations, the choice for τ_t will affect equilibrium y_{t-1} . Therefore we are back to a case where the signal (y_{t-1}) is endogenous to predetermined taxes (τ_t).

Therefore, a predetermined tax may restore separation in the special case of an economy with non-storable goods, but in the presence of an endogenous predetermined factors of production separation breaks down. We demonstrate this in detail in Section 6, where we introduce predetermined taxes and human capital (learning-by-doing) in our leading example and, as we show, issues of GSE determine the solution.

4.2 Optimization under GSE

We now present a generic problem of optimal control under GSE. This adds formal details to the formulation and generalizes (16) without adding any difficulty to the proof. 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 signal.¹² The planner's objective is to maximize $E[W(\tau, s, A)]$ for a given payoff function $W : \mathcal{T} \times \mathcal{S} \times \Phi \rightarrow R$ and random variables $A \in \Phi \subset R^k$ with a given distribution F_A .¹³ This nests the case when endogenous variables other than τ enter the objective function if these can be embedded in W , as we did in the discussion leading to (13).

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

$$s = h(\tau, A) \quad \text{a.s. in } F_A \quad (18)$$

for a known function $h : \mathcal{T} \times \Phi \rightarrow \mathcal{S}$. The government has to choose τ given an observation on s , without observing the value of A . The government knows W , h , F_A , \mathcal{T} , \mathcal{S} and equation (18).

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

¹¹This is the approach commonly used in the armed-bandit literature discussed in Section 2.

¹²Focusing on the univariate case simplifies the analysis, the notation and highlights the main issues. We leave the multivariate case for future work.

¹³ F_A represents the governments' perception about A , it may or may not be equal to the true distribution of A .

policy actions according to a policy function $\mathcal{R} : \mathcal{S} \rightarrow \mathcal{T}$

$$\tau = \mathcal{R}(s) \tag{19}$$

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

$$\begin{aligned} & \max_{\{\mathcal{R}:\mathcal{S}\rightarrow\mathcal{T}\}\{\tau,s:\Phi\rightarrow\mathcal{T}\times\mathcal{S}\}} E [W (\tau, s, A)] \\ \text{s.t.} & \quad (18), (19). \end{aligned} \tag{20}$$

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

$$\int E [W (\tau, s, A) | s] f_s(s) ds. \tag{21}$$

where f_s is the density of s . Applying a “standard recipe” to derive first order conditions in stochastic discrete time models¹⁴ we would find the following condition:

$$E [W_\tau + W_s h_\tau | s] = 0 \quad \text{for all } s. \tag{22}$$

Notice that $W_\tau + W_s h_\tau = 0$ is the FI optimality condition for a Stackelberg leader problem that takes into account the follower’s action s satisfies (18) for a known A . Equation (22) averages this first order condition over values of A that consistent with the observed s .

When $f_s(s)$ is exogenous to the policy choice this formula is correct. However, as argued at the beginning of this section, $f_s = \int_{\Phi} f_{s|A} f_A$ is endogenous to \mathcal{R} , and (22) ignores this dependence. Somehow the fact that f_s depends on \mathcal{R} should play a role when taking derivatives of (21). Economically, the problem is that the policy choice (namely the tax rate τ_1 in the model of Section 3) influences the values of the signal that actually happen (namely labor l_1) and the optimal policy choice should take this into account.

We now provide some further 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

$$H(s, A; \mathcal{R}) \equiv s - h(\mathcal{R}(s), A), \tag{23}$$

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 generic control problem (20), sometimes we refer to $S(\mathcal{R}, A)$ as “the equilibrium” as in a Ramsey problem of Section 3.

¹⁴More 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.

The policy variables that can be realized for each value of the shock A and for a given policy function \mathcal{R} 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 s while S and T are functionals mapping \mathcal{R} and the realizations of the shocks into \mathcal{S} , \mathcal{T} .

4.3 Assumptions, Conditions on \mathcal{R} and outcomes $S(\mathcal{R}, A)$.

We now state some assumptions on the fundamentals and impose some conditions on the admissible set of policy functions \mathcal{R} . These conditions ensure that the outcomes $S(\mathcal{R}, A)$ are well defined when we derive the optimality conditions that are the focus of this section.

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

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

Our aim is to state Conditions on \mathcal{R} that are generically satisfied at an unconstrained maximum and that can be checked at a computed solution. Denote the support of possible outcomes for \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, $|\mathcal{R}'(s)| < K^L$ for some constant $K^L < \infty$ and all $s \in \mathcal{S}_{\mathcal{R}}$ where \mathcal{R}' exists.

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.¹⁵

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.¹⁶

We now state a condition that is generically equivalent to imposing uniqueness of outcomes, that allows us to derive optimality conditions and that can be checked

¹⁵It is possible that under Assumptions 1-2 or further assumptions one can prove that Condition 1 (as well as Condition 2 stated below) hold at the optimum. We leave this for future research.

¹⁶Multiplicity 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.

easily. Let us recall a few well known facts about fixed points. Consider an absolutely continuous function $g : X \rightarrow X$ for a compact set $X \subset \mathbb{R}$. Two facts are well known about fixed points of g : *i)* g has at least one fixed point. Assume in addition that the derivative g' exists at all fixed points $x_f = g(x_f)$ and that $g'(x_f) \neq 0$, then *ii)* a fixed point exists such that $g'(x_f) > 1$ if and only if there are multiple fixed points.¹⁷

Applying these properties to the map $h(\mathcal{R}(\cdot), A)$ establishes that $S(\mathcal{R}, A)$ is non-empty and that imposing that $S(\mathcal{R}, A)$ is a singleton is essentially equivalent to imposing the following

Condition 2 *There is an $\bar{\varepsilon} > 0$ such that, for all $A \in \Phi$ where \mathcal{R}' exists at all elements of $S(\mathcal{R}, A)$ then*

$$h_\tau(\mathcal{R}(s), A)\mathcal{R}'(s) \leq 1 - \bar{\varepsilon} \quad (24)$$

at all $s \in S(\mathcal{R}, A)$.

This Condition guarantees that H is a well-conditioned non-linear equation and this ensures uniqueness in the following Lemma. This condition is also used in the optimality condition in Proposition 2 and existence result in Proposition 5. We show in Section 5.2 that Condition 2 can be checked in a specific computed policy.

To guarantee that the outcomes of a given variation $S(\mathcal{R} + \alpha\delta, A)$ are well defined and they converge to $S(\mathcal{R}, A)$ we need the requirements on h stated in the following

Assumption 3 *W is differentiable everywhere with respect to (τ, s) , with bounded derivatives $|W_\tau|, |W_s| < Q$. Also, $|h_\tau| < Q$ and 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 constants $Q, Q^L < \infty$.¹⁸*

The first order condition (26) allows for a non-differentiable optimal solution \mathcal{R}^* . Indeed, we will see that a non-differentiable \mathcal{R}^* can arise in our computations. To guarantee that the outcomes of variations $S(\mathcal{R} + \alpha\delta, A)$ converge we need that non-differentiability 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.*

¹⁷Obviously *i)* follows from Brouwer's fixed point theorem. Property *ii)* is easy to see with a simple graph, it has been encountered in economics long ago. One example is the old literature on deterministic growth with increasing returns where the law of motion of capital g had multiple steady states, the steady state with $g' > 1$ being unstable while a low stable steady state with $g' < 1$ was a poverty trap. A detailed proof of *ii)* is available upon request.

¹⁸Obviously Lipschitz continuity of h_τ is guaranteed if $h_{\tau\tau}$ exists everywhere and is uniformly bounded. In fact $|h_\tau| < Q$ follows from the Lipschitz condition.

This is easy to ensure if at least one element in the vector of shocks A has a continuous density. We point out in Section 4.6 that this is needed for the case of interest, namely non-invertibility.

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 $\mathcal{S}_{\mathcal{R}} \subset \text{int}(\mathcal{S})$ and $\mathcal{R}(\mathcal{S}_{\mathcal{R}}) \subset \text{int}(\mathcal{T})$. Fix a variation $\delta : \mathcal{S} \rightarrow R$ 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 $S(\mathcal{R} + \alpha\delta, A) \rightarrow S(\mathcal{R}, A)$ as $\alpha \rightarrow 0$ a.s.*

Lemma 1 is proved in Appendix A.1.¹⁹ The “almost surely” statement is over the distribution of A .

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. In our proofs and even when we check second order conditions in practice, we always work with continuous δ 's.

4.4 General formulation

Let \mathcal{F} be the value of the objective function for a given \mathcal{R} ²⁰

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

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

$$\begin{aligned} \max_{\{\mathcal{R}:\mathcal{S}\rightarrow\mathcal{T}\}} \mathcal{F}(\mathcal{R}) \\ \text{s.t. } \mathcal{R} \in \mathcal{E} \end{aligned} \tag{25}$$

where, \mathcal{E} is a pre-specified set of admissible policy functions. Note that restriction (18) 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} 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 gives $\mathcal{F}(\mathcal{R})$ is well defined. Also, \mathcal{E} has to be restricted to ensure that a solution to the optimization problem exists. We discuss in more detail the choice for \mathcal{E} in Section 4.9.

We denote the solution to Problem (25) by \mathcal{R}^* .

¹⁹Although 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 $h_{\tau}(\mathcal{R}(s), A)\mathcal{R}'(s) < 1$.

²⁰Notice that \mathcal{F} maps the space of functions into the real line R .

4.5 Apparent Partial Information: Revelation and Invertibility

In some cases the government can still implement the FI policy even if it does not observe the shocks. This occurs whenever the government can learn the true state of the economy A from observing the signal s .²¹

To formally define Invertibility, let $(\tau^{FI}, s^{FI}) : \Phi \rightarrow \mathcal{T} \times \mathcal{S}$ be the FI solution and $\mathcal{S}^{*,FI}$ be the support of s^{FI} .

Definition 4 *Invertibility holds if for any signal that occurs in equilibrium $s \in \mathcal{S}^{*,FI}$ there exists a unique $A \in \Phi$ such that $s^{FI}(A) = h(\tau^{FI}(A), A)$.*

Invertibility will often occur when the dimension of τ is the same as the dimension of A . In this case knowledge of s in equilibrium reveals A . Even if A is high-dimensional, Invertibility holds if Φ is a finite set, as only by coincidence would the same equilibrium point (τ, s) occur for two different realizations of A .

Proposition 1 *Under Invertibility \mathcal{R}^* solves the FI-problem.*

This follows from the fact that the PI case is a constrained FI problem, therefore the value in the PI case is less than or equal to the FI case, and the value of the FI case is achievable under Invertibility.

A large part of the literature has considered this case, in particular 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.

To illustrate a case of Invertibility within the example of Section 3, assume that $\gamma = 1$ with certainty. The government does not observe the random value of θ and it has to choose taxes observing l . This is only apparently a PI problem, because the government can infer θ from observing the labor choice, hence the government can implement the FI policy. In this case the solution would coincide with the FI solution and, therefore, it would show perfect tax smoothing.

4.6 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 (18), so that Invertibility as in the previous subsection does not obtain. In general, Invertibility fails 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. This is related to Assumption 4. For example, in

²¹In the literature of optimal contracts under private information and incentive compatibility constraints this is the standard assumption, which amounts to assuming full revelation. Invertibility also holds in the supply function model of Klemperer and Meyer (1989), where uncertainty is one-dimensional.

the model of Section 3, the values of θ, γ remain hidden even after the choice of τ has been made for a given observed labor l if either θ, γ , (or both) have continuous distributions and the other random variable is non-degenerate.²²

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 $s = S(\mathcal{R}^*, A)$ and $\tau = \mathcal{R}^*(S(\mathcal{R}^*, A))$.

Proposition 2 *Assume Assumptions 1-4. Assume that a solution to (25) exists denoted \mathcal{R}^* . Assume, \mathcal{R}^* satisfies Conditions 1-2, the interiority conditions stated in Lemma 1 hold and $\mathcal{R}^* \in \text{int}(\mathcal{E})$. Then \mathcal{R}^* satisfies the following necessary first order condition*

$$E \left(\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*'}} \Big| S(\mathcal{R}^*, A) \right) = 0 \quad (26)$$

almost surely in A .

We have already discussed the generality of all Assumptions and Conditions involved. 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} .²³ We imposed minimal restrictions on the set of admissible functions \mathcal{E} hence interiority of \mathcal{R}^* is likely to hold.

The proof of Proposition 2 uses an argument similar to the one often used in the literature on calculus of variations. Notice, however, that our approach is not a special case of the standard calculus of variations: unlike in the standard case, the choice of \mathcal{R} involves choosing measurability conditions relating the observable signal and the underlying variables of integration. We refer the interested reader to Section 4.11, where we explain this difference in more detail.

In the following paragraphs we show some of the key steps in the proof and we comment where each assumption is needed. For a detailed proof see Appendix A.2.

Take any function (a variation) $\delta : \mathcal{S} \rightarrow R$ and a constant $\alpha \in R$. Now consider reaction functions of the form $\mathcal{R}^* + \alpha\delta$. For a given δ consider solving the (one-dimensional) maximization problem

$$\max_{\alpha \in R} \mathcal{F}(\mathcal{R}^* + \alpha\delta) \quad (27)$$

so that now we maximize over small deviations of the optimal reaction function in the direction determined by δ . It is clear that $0 \in \arg \max_{\alpha \in \mathbb{R}} \mathcal{F}(\mathcal{R}^* + \alpha\delta)$. Lemma 1 already guaranteed that the solution $S(\mathcal{R}^* + \alpha\delta, A)$ is unique so that $\mathcal{F}(\mathcal{R}^* + \alpha\delta)$ is well defined. Hence, we can take derivatives involved in the function \mathcal{F} . Appendix A.2 shows the first equality in

$$\frac{\partial \mathcal{F}(\mathcal{R}^* + \alpha\delta)}{\partial \alpha} \Big|_{\alpha=0} = E \left(\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*'}} \delta(S(\mathcal{R}^*, A)) \right) = 0 \quad (28)$$

²²Even in the case that A is one-dimensional Invertibility is also violated if $h(\mathcal{R}^{FI*}(\cdot), \cdot)$ is non-monotonic. We leave an exploration of this case for future work.

²³See footnote 42.

and the second equality follows because $\alpha = 0$ is a solution of (27). Since the conditional expectation is a function defining an orthogonal error, this implies (26).²⁴ Furthermore, even if (26) holds “only” almost surely, continuity of \mathcal{R}^* (Condition 1) and the fact that $S(\mathcal{R}^*, A)$ has a density implies that this equation characterizes the whole function \mathcal{R}^* .

As we anticipated the first order condition (26) does not coincide with the “standard recipe” FOC (22). The term $\frac{1}{1-h_\tau^*\mathcal{R}^{*'}} acts as a kernel, or measure change. This is the new term relative to the standard case under separation, when $f_{A|s}$ does not depend on \mathcal{R} . The term $\frac{1}{1-h_\tau^*\mathcal{R}^{*'}$ captures the effect of the choice of \mathcal{R} on the density f_s . More precisely, denoting $S^{\delta'}(\alpha, A) \equiv \frac{dS(\mathcal{R}^*+\alpha\delta, A)}{d\alpha}$, in Appendix A.2 we show that$

$$S^{\delta'}(0, A) = \frac{h_\tau^*\delta}{1-h_\tau^*\mathcal{R}^{*'}} \quad (29)$$


The numerator $h_\tau^*\delta$ captures the direct effect of a small (α) change in policy through (18). The term $\frac{1}{1-h_\tau^*\mathcal{R}^{*'}$ captures the degree to which the signal responds in equilibrium to a marginal α change, near the optimum and along each direction δ .

At this point, Figure 1 may provide useful intuition. The notation for the shocks and signal structure are from the example in Section 3. Given a choice for \mathcal{R} , upon observing a certain value of the signal \bar{s} (in the example, hours worked l), the planner cannot be sure if it was determined by a value of shocks (γ_1, θ_1) or the value (γ_2, θ_2) , as illustrated by the intersection of a policy \mathcal{R} with two possible reaction functions h . The issue of General Signal Extraction is to realise that the choice of \mathcal{R} influences the likelihood of observing \bar{s} . Small changes in the choice of \mathcal{R} impact on s differently for each realisation. In particular, consider a small deviation $\mathcal{R} + \alpha\delta$, when δ is the indicator function of an interval around \bar{s} .²⁵ The figure illustrates how the impact of this deviation on the signal will be different if realisation (γ_1, θ_1) or (γ_2, θ_2) actually materialise (even though these two realisations give rise to the same outcome for \mathcal{R}). The difference in the outcome s for each realisation is proportional to the term $\frac{1}{1-h_\tau\mathcal{R}'}$ that appears in (26). 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 realisations such that a small change in the policy has large effects on equilibrium s .

²⁴More precisely, take random variables X, Y . We know the conditional expectation $E[X | Y] = f(Y)$ for a function f such that $E[(X - f(Y))g(Y)] = 0$ for all functions g . To go from (26) to (28) we take $\frac{W_\tau^* + W_s^* h_\tau^*}{1-h_\tau^*\mathcal{R}^{*'}$ to play the role of X , $S(\mathcal{R}, A)$ to play the role of Y , 0 the role of $f(Y)$ and δ the role of g .

²⁵This Figure uses an indicator function δ for illustrative purposes and because it is often used in the calculus of variations literature. However, as pointed out after Lemma 1, in our proofs we require continuous δ 's.

Figure 1: Optimal policy \mathcal{R} and deviation



The figure refers to the example of Section 3 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.

4.6.1 A convenient optimality condition

Condition (26) 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 (26). 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 (26). 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 .

Denote the support of A_2 's 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*

$$\int_{\Theta_2(s, \mathcal{R}^*(s))} \frac{W_\tau^* + W_s^* h_\tau^*}{|h_{A_1}^*|} f_{A_1|A_2}(\mathcal{A}^*(s, a_2), a_2) dF_{A_2}(a_2) = 0 \text{ for almost all } A \quad (30)$$

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

The proof is in Appendix A.3. As in Proposition 2 the functions $W_\tau^*, W_s^*, h_\tau^*, h_{A_1}^*$ are evaluated at equilibrium values.

As promised, this optimality condition does not involve $\mathcal{R}^{*'}$.

In Section 4.10 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 $\mathcal{R}^{*'}$, as it appears in the Jacobian in (77).

4.7 Relation to previous results on optimal policy with PI

Previous results available in the literature on optimal policy under PI can be derived as special cases of Proposition 2. In those cases the “standard recipe” worked and (22) gave the correct solution.

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 welfare-relevant allocation is a function of τ , the signal s is independent of τ , and given by $s = h(A)$, 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 (26) reduces to the FOC (22) at this s .

The proof is trivial: in all cases the term $\frac{1}{1-h_\tau^* \mathcal{R}^{*'}}$ cancels out because it is known given s so that this term can be pulled outside of the conditional expectation in (26). We now discuss how each case has been used in the literature and how it arises using variations of our example of Section 3.

Case 1 is related to the case discussed in Section 4.5. It generalizes that section as invertibility here can hold only for a subset of $\mathcal{S}_{\mathcal{R}^*}$.

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 3 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 (26) would simplify to equation

$$E [W_l | s] = 0 \tag{31}$$

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 (31). 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, as we already emphasized in Section 4.1, 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 6 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 (26) 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) calls 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.*

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):

$$W(\tau, s) = -\frac{\omega_\tau}{2}\tau^2 - \frac{\omega_s s^2}{2} \quad (32)$$

and the reaction function is

$$s = h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma \quad (33)$$

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

The FOC under PI (26) becomes

$$E[\omega_\tau \tau + \omega_s h_\tau (h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma) | s] = 0 \quad (34)$$

which can be rewritten as

$$\tau = FE[A|s] \quad (35)$$

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 (35) 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.

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

$$E[\theta|s] = E[\theta|h + h_\tau FE[A|s] + h_\theta \theta + h_\gamma \gamma = s] \quad (36)$$

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

4.8 Second order condition

After Proposition 2 we discussed how the result is 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 .

Using this principle, we can obtain the following second order condition that further restricts the solution. This can be useful in practice if we find two solutions to (30), the second order condition may help to select the correct one. The proposition is stated for the case where W does not depend on τ , the case when τ is an argument can be treated similarly to obtain a more involved expression.

First, we make a regularity assumption and impose a new condition that allows us to compute the second derivative of the objective function using a variational argument similar to the one used to prove Proposition 2.

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 $|\mathcal{R}''| < Q$ uniformly on $\mathcal{S}_{\mathcal{R}}$.

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

$$E \left[W_{ss}^* (S^{\delta'}(0, A))^2 + W_s^* S^{\delta''}(0, A) \right] \leq 0 \quad (37)$$

where $S^{\delta'}(0, A)$ is given by equation (29) and

$$S^{\delta''}(0, A) = \frac{\delta'^* h_{\tau}^* S^{\delta'} + (\delta^*)^2 h_{\tau\tau}^* + \delta^* h_{\tau\tau}^* \mathcal{R}' S^{\delta'} - \delta'^* (h_{\tau}^*)^2 \mathcal{R}' S^{\delta'} + \delta^* \delta'^* (h_{\tau}^*)^2 + \delta^* (h_{\tau}^*)^2 \mathcal{R}'' S^{\delta'}}{(1 - h_{\tau}^* \mathcal{R}')^2} \quad (38)$$

Here, W_{ss}^* , h_{τ}^* , $h_{\tau\tau}^*$, δ^* , δ'^* denote the respective functions evaluated at $s = S(\mathcal{R}^*, A)$ and $\tau = \mathcal{R}^*(S(\mathcal{R}^*, A))$. Moreover, $S^{\delta'}$ is evaluated at $\alpha = 0$ as in (37).

The proof is in Appendix A.4.

Two important remarks: *i*) this is only a necessary condition, we leave research for a sufficient second order condition for later, *ii*) different from second order conditions found in the calculus of variations literature, the above condition involves $(\delta^*)^2$ and δ'^* along with δ itself. The literature on the calculus of variations often arrives at simpler expressions by taking indicator functions, indeed these terms would disappear from the above expression in this case, as $(\delta^*)^2 = 0$ or 1 and $\delta'^* = 0$ almost everywhere for δ . However, as pointed out after Lemma 1, in our analysis we need continuous variations, so we stay away from indicator function δ 's so we keep these terms in the expression.

In practice we take δ 's that put almost all weight in a small neighborhood of s . In this way we can check the second order condition “around” many specific points s while maintaining continuity, see Section 5.2.

4.9 Existence of a maximum

Obviously, the optimality conditions can be misleading if a maximum does not exist, as emphasized by some famous examples in the calculus of variations literature.²⁶ One way to guarantee existence is to restrict the set of functions \mathcal{E} to be compact but keeping \mathcal{E} as large as possible so as to include global maxima in most models.

²⁶See page 294 Clarke (2013).

In view of the previous results it may be tempting to specify \mathcal{E} as the set of all functions satisfying Conditions 1-2. It turns out, however, that this set is not compact.²⁷ The following strengthened condition is enough to guarantee compactness.

Condition 4 *Same as Condition 2 but (24) holds at all s where \mathcal{R}' exists, where $|s - h(\mathcal{R}(s), A)| \leq \delta$ for some $\delta > 0$, for some A as in the first line of Condition 2.²⁸*

This condition requires that well-conditioning (24) holds uniformly around all points that are “nearly” fixed points of $h(\mathcal{R}(\cdot), A)$. If \mathcal{R} satisfies Condition 4 then it satisfies Condition 2 as well.

Proposition 5 *Define \mathcal{E} as the set of all functions \mathcal{R} satisfying Conditions 1 and 4 for some $K^L, \bar{\varepsilon}, \delta > 0$. Let Assumptions 1,2,4 hold. Then there exists a solution \mathcal{R}^* attaining the maximum of problem (20).*

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

The proof follows because \mathcal{E} is obviously non-empty, \mathcal{E} a compact set in the sup-norm, because of the continuity of the objective function and boundedness of \mathcal{T}, \mathcal{S} . For details see Appendix A.5

4.10 Putting everything together

The above results give a clear path to find solutions to (20): *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 . 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 (30) at $(\tau, s, \mathcal{A}^*(s, a_2), a_2)$. Find the integral on the left side of (30) 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 (30). 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.

²⁷It is possible to find a sequence of functions f_k each with one fixed point satisfying Condition 2 but converging to a function with two fixed points. Therefore the set of functions satisfying Condition 2 is not closed.

²⁸To clarify, $|s - h(\mathcal{R}(s), A)| \leq \delta$ is only required for A 's such that the derivative of \mathcal{R} exists at all elements of $S(\mathcal{R}, A)$.

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.²⁹

4.11 Relation with Calculus of Variations

In this subsection we clarify the distinction between our contribution and the type of problems addressed with standard Calculus of Variations and explain why we need to go to first principles to derive our optimality condition since our case is not a special case of the standard approach.

It is clear that problem (20) can be reformulated as

$$\max_{\mathcal{R}:S \rightarrow \mathcal{T}} \int_{\Phi} W(\mathcal{R}(S(\mathcal{R}, A)), S(\mathcal{R}, A), A) dF_A. \quad (39)$$

Although related, (39) is not a special case of the standard formulation in the literature on the Calculus of Variations. Consider the following citation:

“The basic problem in the subject that is referred to as *Calculus of Variations* consists of minimizing an integral functional of the type

$$\int \Lambda(x(t), x'(t), t) dt. \quad (40)$$

over a class of functions x ” (p. 287 Clarke, 2013). Optimality conditions for this problem were first formulated by Euler in 1744 and they have been generalized in many directions.³⁰

The two problems are clearly related, they both involve an integral of a known function (W or Λ) and the choice is over a function (\mathcal{R} or x). A notable difference between our problem (39) and (40) is that the object to be chosen in our case (namely \mathcal{R}) is *not* a function of the variable of integration (namely A), while the standard case requires x to be a function of the variable of integration t .³¹

Our problem is to choose the measurability conditions relating observables s and the underlying variable of integration A , while this is not an issue in (40). To our knowledge, this has not been treated in the literature. Hence, we need to derive optimality condition from first principles.

²⁹In 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.

³⁰They have been applied in various economic models including, for example, the neoclassical growth model in continuous time.

³¹Another difference is a simplification: the objective function in the standard case involves the derivative x' , while in our case \mathcal{R}' is not an argument of W .

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 3.³² In Section 6 we extend the analysis to study a dynamic labour supply decision in presence of learning-by-doing and in Section 7 we present an infinite-horizon model.

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 $\beta = .96$. 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 when we introduce PI with only hours being observed, if the government sees an increase in l , it would want to react in opposite directions depending on the source of the shock: this would call for a tax increase, if driven by low θ , or a tax cut if driven by high γ . Hence, this model is particularly interesting to analyze optimal policy with endogenous PI since by observing a certain value l and imposing a tax rate τ , the government cannot infer the value of the shocks.

We now discuss the problem of optimal policy with PI. Notice that the equilibrium condition (6) becomes

$$Bl_1c_1 = \gamma\theta(1 - \tau_1) \quad (41)$$


and 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 (12) specializes to

$$l = h(\tau, \theta, \gamma) = \frac{Bg\theta^{-1} + \sqrt{(Bg)^2\theta^{-2} + 4B\gamma(1 - \tau)}}{2B}. \quad (42)$$

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

³²Consistent with the presentation of the model in Section 3, 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).

Figure 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.

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 be 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 theta 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.³³

The partial derivatives h_τ and h_γ that we need to derive to solve the optimality condition (30) are easily obtained analytically in this example. In particular,

$$h_\tau(\tau, \theta, \gamma) = \frac{-1}{\sqrt{(Bg)^2 (\theta\gamma)^{-2} + \gamma^{-1}4B(1 - \tau)}}. \quad (43)$$

It is clear that both the productivity shock and the demand shock affect this slope, therefore the kernel in the optimality condition (30) is not a constant. Hence, we proceed to find a \mathcal{R}^* that satisfies (30) using Algorithm 1 described above.

Figure 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.

The limits of the signal $l = L(\mathcal{R}^*; \theta_{max}, \gamma_{min})$ 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 in the PI solution, we have

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


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

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 leads to the FI equilibrium for $(\theta_{max}, \gamma_{min})$. Then, taxes start to increase: 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, up the full revelation point $l_{max} = L^{FI}(\theta_{min}, \gamma_{max})$.

³³In 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 5.3).

Figure 3: Optimal policy with low g




Optimal tax rate as function of hours under PI. Red line: \mathcal{R}^* ; yellow region: set of FI pairs (l, τ) for all possible realizations of (θ, γ) ; black line: linear policy connecting the two full revelation points.

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 Figure 4. In all four panels, we reproduce the FI outcomes shown in Figure 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.³⁴ 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, $l = .33$, in Figure 5. The wealth effect of productivity makes

³⁴We will see in Section 5.3 that this property of the solution will change with higher government expenditure.

Figure 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.

this locus an increasing function in the (θ, γ) 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. This can be seen in Figure 6, where the equilibrium cumulative distribution function of tax rates under PI (red line) is contrasted with the one obtained under FI (blue dotted line). 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 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

Figure 5: Set of admissible shocks consistent with observing $l = .33$




Set of combinations of (θ, γ) that have positive density in equilibrium, conditional on a particular realization of the signal, namely $l = .33$. Solid red line for PI, blue dashed-dotted line for FI.

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 very 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 it 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 7.

Figure 6: Equilibrium CDF of tax rates



Cumulative Distribution Function of τ_1 under PI (solid red line) and FI (blue dashed-dotted line).

5.2 Checking Conditions 1-2 and Second Order Condition

In this subsection we briefly discuss how we check Conditions 1 and 2 as well as the Second Order Condition of Proposition 4 in this numerical example. First, notice from our Figures that our computed \mathcal{R}^* is a continuous function and its derivative is never too large, implying that we can easily find a bounding constant $K^L < \infty$ that makes Condition 1 satisfied.

As far as Condition 2 is concerned, we verify numerically that the term $h_\tau \mathcal{R}'$ is bounded above by a number less than 1. Indeed, in this numerical example, its range of values is between -0.34 and 0.08, depending on the realizations of A .

In order to check the Second Order Condition, we implement the following procedure. We parametrize $\delta(s) = \phi(s; \mu, \sigma)$ where ϕ is the probability density function of a Normal distribution with mean μ and standard deviation σ . We then set a very small $\sigma = 0.005$ to put most of the weight around certain μ 's, and we vary μ on a grid of s 's ranging from the minimum value to the maximum value in $\mathcal{S}_{\mathcal{R}^*}$. Next, we check the inequality (37) for all value of μ on this grid.

5.3 Close to the top of the Laffer curve

Let us now look at the case where government expenditure is very high, equal to 60% of average output in both periods.³⁵ We will see that this leads to a very non-linear


³⁵All other assumptions on preferences and shocks are the same as in Section 5.1.

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.

Figure 7 shows optimal policy for this case (red line), again contrasted with the set of tax-labor outcomes under FI (yellow region).

Figure 7: Optimal policy with high g




Optimal tax rate as function of hours under PI. Red line: \mathcal{R}^* ; yellow region: set of FI pairs (l, τ) for all possible realizations of (θ, γ) ; black line: linear policy connecting the two full revelation points.

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 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.

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 Figure 8 and under PI in Figure 9. Consider Figure 8 first. It can be seen that under FI any realization of θ is consistent with an intermediate realization 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 .


Figure 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.

To see this, consider now Figure 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

Figure 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.

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 Figure 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 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 Section 5.1, 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 for agents, then optimal policy calls for being very reactive to observables in order to prevent those cases to materialize.

Figure 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.

This is exemplified by the optimality of increasing taxes steeply in the first period to avoid having to distort the economy too heavily in the second period if realized productivity turn out to be low (and hence the fiscal deficit turns out to be high). This lesson 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 were 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.

5.4 Linear-quadratic approximation


We now compare our solution to existing methods based on linear-quadratic optimization (Svensson and Woodford, 2004). In order to do so, we modify our distributional assumption and we assume that both θ and γ are normally distributed. We then truncate these distributions at three standard deviations from the mean in order to have a bounded support for the shocks in our solution method. The standard deviation of each of the shocks is assumed to be 3% of the mean.

In order to compute the linear approximation, we take a second-order approximation of the objective function and a first-order approximation of the reaction function h around the allocation and policy that arises under FI when the shocks take their

mean value. Then, we compute the “certainty equivalent” policy as described in Section 4.7. Importantly, this policy can be found under FI and then applied to the PI case by simply computing the conditional mean of the shocks for each value of l . Figures 11 and 12 compare the optimal policy and the linear approximation in the case of low g and high g respectively. It can be seen that the approximation is quite accurate for intermediate realizations of labor, but less so for extreme values. This suggests that linear approximations can be misleading when there is endogenous PI and the economy is hit by large shocks.


An important difference between the non-linear solution and the linear approximation is that in the linearized model the government cannot affect the slope of the loci of shocks realizations conditional on the signal. In the Svensson and Woodford (2004) environment, this slope is constant and exogenous to policy, hence there is no scope for a choice of better signal extraction about one shock or the other.

Figure 11: Linear approximation, low g



Optimal tax rate as a function of the signal l , with $g = .25$. Red line: \mathcal{R}^* , blue dashed-dotted line: “certainty equivalent” linear approximation.

Figure 12: Linear approximation, high g



Optimal tax rate as a function of the signal l , with $g = .6$. Red line: \mathcal{R}^* , blue dashed-dotted line: “certainty equivalent” linear approximation.

6 A model with dynamic labor decision and pre-determined 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).³⁶ 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.³⁷ 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.

³⁶Clymo and Lanteri (2019) study optimal fiscal policy under FI with this timing assumption on policy decisions.

³⁷Stantcheva (2015) studies optimal taxation with truthful revelation in a similar model with learning-by-doing and private information about productivity.

6.1 Preferences and technology

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

$$E \sum_{t=1}^3 \beta^{t-1} [\gamma_t u(c_t) - v(l_t)] \quad (44)$$

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

$$c_t + g = \theta_t \chi(l_{t-1}) l_t \quad (45)$$

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

$$c_t + q_t b_t = \theta_t \chi(l_{t-1}) l_t (1 - \tau_t) + b_{t-1} \quad (46)$$

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.

6.2 Government

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

$$\tau_t \theta_t \chi(l_{t-1}) l_t + q_t b_t = g + b_{t-1} \quad (47)$$

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 $t = 3$).

6.3 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 .

$$v'(l_1) - \gamma u'(c_1)\theta\chi(l_0)(1 - \tau_1) - \beta u'(c_2)\bar{\theta}l_2(1 - \tau_2)\chi'(l_1) = 0. \quad (48)$$

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.

6.4 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 $C = (1/3)^{-\psi}$, $\psi = .1$ 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 (44) 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 ($g = .25$) to the solution with high government spending ($g = .60$). These two cases are illustrated in Figures 13 and 14 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 Figures 13 and 14, 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.

Figure 13: Optimal policy in the 3-period model with low g



Optimal tax rate τ_2 as function of hours l_1 under PI, with $g = .25$. Red line: \mathcal{R}^* ; yellow region: set of FI pairs (l_1, τ_2) for all possible realizations of (θ, γ) .

7 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.


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.³⁸ In order to avoid these difficulties we assume linear utility from consumption.

We also assume that the shocks are i.i.d. over time as this reduces the number of state variables and allows us to abstract from the issues of government experimentation that have been studied in the armed-bandit literature of optimal policy that we discussed in Section 2.³⁹ In this way we are left with the simplest infinite horizon model of fiscal policy where endogenous signals play a role.

³⁸Under FI, this issue was first addressed in Aiyagari et al. (2002).

³⁹For example Wieland (2000a, 2000b), Kiefer and Nyarko (1989), Ellison and Valla (2001)

Figure 14: Optimal policy in the 3-period model with high g



Optimal tax rate τ_2 as function of hours l_1 under PI, with $g = .6$. Red line: \mathcal{R}^* ; yellow region: set of FI pairs (l_1, τ_2) for all possible realizations of (θ, γ) .

7.1 Full Information

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. Preferences of the representative agent are given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t [\gamma_t c_t - v(l_t)] \quad (49)$$

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

The period t budget constraint of the representative agent is

$$c_t + q_t b_t = \theta_t l_t (1 - \tau_t) + b_{t-1} \quad (50)$$

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

$$\frac{v'(l_t)}{\gamma_t} = \theta_t (1 - \tau_t) \quad (51)$$

and

$$q_t = \beta \frac{\bar{\gamma}}{\gamma_t}. \quad (52)$$

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 $A^t = (A_t, A_{t-1}, \dots, A_0)$, where $A_t = (\theta_t, \gamma_t)$.

The period- t implementability constraint is

$$b_{t-1} = c_t - \frac{v'(l_t)}{\gamma_t} l_t + \beta \frac{\bar{\gamma}}{\gamma_t} b_t. \quad (53)$$

We now introduce an upper bound on debt, b_{max} . We will 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:

$$\gamma_t \theta_t - v'(l_t) + \lambda_t \theta_t - \frac{\lambda_t}{\gamma_t} [v'(l_t) + v''(l_t) l_t] = 0 \quad (54)$$

and

$$\lambda_t \frac{\bar{\gamma}}{\gamma_t} = E_t \lambda_{t+1} + \chi (b_t - b^{max}) I_{[b^{max}, \infty)}(b_t). \quad (55)$$

where λ_t is the Lagrange multiplier of constraint (50) 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 $\tau_t = \mathcal{R}^{FI}(b_{t-1}, A_t)$

7.2 Partial Information

We start the description of the PI problem by specifying its timing. At the beginning of each period t , the Ramsey 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) .

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. In other words, the government cares about past realizations of the exogenous shocks

only to the extent that they affect the level of current outstanding debt. 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 the utility (49) 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

$$V(b) = \max_{\mathcal{R}: \mathbb{R}^2 \rightarrow \mathbb{R}_+} E \left[\gamma(\theta l - g) - v(l) + \beta V \left(\frac{(b + g - \theta l + \frac{v'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}} \right) + \right. \\ \left. - \beta \frac{\chi}{2} \left(\frac{(b + g - \theta l + \frac{v'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}} - b^{max} \right)^2 I_{[b^{max}, \infty)} \left(\frac{(b + g - \theta l + \frac{v'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}} \right) \right] \quad (56)$$

where l satisfies $l = h(\mathcal{R}(b, l), \theta, \gamma)$ and $h(\tau, \theta, \gamma)$ is the level of hours that satisfies (51).

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

$$l = v'^{-1}(\theta\gamma(1 - \tau)), \quad (57)$$

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 pausing the maths and 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 materialised 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 (56) we have substituted future debt using the budget constraint (53). 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.

In order 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 (56).

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.⁴⁰

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

$$V'(b) = E \left[\frac{\gamma}{\bar{\gamma}} V'(b') \right] - \chi \frac{\gamma}{\bar{\gamma}} (b' - b^{max}) I_{[b^{max}, \infty)}(b'). \quad (58)$$

Hence by solving the first order condition using (58) and iterating on the Bellman equation (56), we can approximate the optimal policy. In the next subsection, we show some numerical results obtained after parametrizing the economy. 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 limits.

7.3 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 $\pm 5\%$ 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. Figure 15 illustrates this property by plotting taxes against labor, as given by the function $\mathcal{R}^*(b, \cdot)$ evaluated at different values of debt. Notice that taxes are in general decreasing in labor, because, with linear-quadratic utility, hours are increasing in θ , since the only effect of this shock is

⁴⁰The FOC is explicitly shown in Appendix D.

the substitution effect. It can be seen that for lower levels of debt (namely for $b = 0$ and $b = .1$), taxes are relatively flat with respect to the signal. However, close to the debt limit ($b=.18$), they become highly responsive, with a higher slope in the middle region of realizations of hours.

Figure 15: Tax policy for different levels of debt



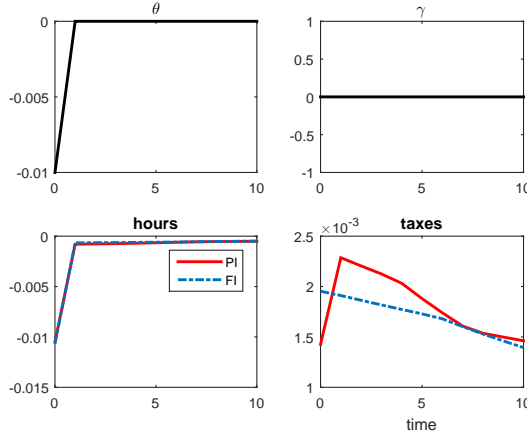
Optimal tax rate as a function of the signal l , for different values of outstanding debt b . Red solid line: $b = .18$; black dashed-dotted line: $b = .10$; blue dashed line: $b = 0$.

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 Figure 16.⁴¹

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).

⁴¹These non-linear impulse response functions are computed as percentage deviations from the path that would arise absent all shocks.

Figure 16: 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.

8 Conclusion

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. These 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 illus-

trated 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.

References

- [1] Adam, K. (2004), “On the Relation between Bayesian and Robust Decision Making”, *Journal of Economic Dynamics and Control*, vol. 28(10) 2105-2117
- [2] Aiyagari, S.R., A. Marcet, T.J. Sargent and J. Seppala (2002) “Optimal Taxation without State-Contingent Debt”, *Journal of Political Economy*, University of Chicago Press, vol. 110(6), pages 1220-1254, December.
- [3] Angeletos, G.M. and A. Pavan (2009), “Policy with Dispersed Information”, *Journal of the European Economic Association* 7(1):11–60
- [4] Angeletos, G.M., L. Iovino and J. La’O (2015), “Efficiency and Policy with Endogenous Learning”, Working Paper
- [5] Aoki, K. (2003), “On the Optimal Monetary Policy Response to Noisy Indicators” *Journal of Monetary Economics*, Elsevier, vol. 50(3), pages 501-523, April.
- [6] Baxter, B., G. Liam and S. Wright (2007), “The Endogenous Kalman Filter”, Birkbeck School of Economics, Mathematics and Statistics working paper BW-PEF 0719
- [7] Baxter, B., G. Liam and S. Wright (2011), “Invertible and Non-invertible Information Sets in Linear Rational Expectations Models”, *Journal of Economic Dynamics & Control* 35: 295–311
- [8] Clarke, F. (2013), *Functional Analysis, Calculus of Variations and Optimal Control*, Springer Graduate Texts in Mathematics
- [9] Clymo, A. and A. Lanteri (2019), “Fiscal Policy with Limited-Time Commitment”, *The Economic Journal*, forthcoming
- [10] Ellison, M. and N. Valla (2001), “Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions”, *Journal of Monetary Economics*, Elsevier, vol. 48(1), pages 153-171, August.
- [11] Hansen, L.P. and T.J. Sargent (2012), Three types of ambiguity, *Journal of Monetary Economics*, Elsevier, 59(5), 422-445.
- [12] Kiefer, N. and Y. Nyarko, (1989), “Optimal Control of an Unknown Linear Process with Learning”, *International Economic Review*, Dept. of Economics, U. of Pennsylvania and Osaka U. ISERA, vol. 30(3), pages 571-86, August.
- [13] Klemperer, P.D. and M.A. Meyer (1989), “Supply Function Equilibria in Oligopoly under Uncertainty”, *Econometrica*, Vol. 57, No. 6, pp. 1243-1277

- [14] Kocherlakota, N.J. (2010), *The New Dynamic Public Finance*, Princeton University Press.
- [15] Lucas, R.E., Jr. (1972), “Expectations and the Neutrality of Money”, *Journal of Economic Theory*, 4(2), 103-124.
- [16] Lucas, R.E., Jr. and N.L. Stokey (1983), “Optimal Fiscal and Monetary Policy in an Economy without Capital”, *Journal of Monetary Economics* , 12, 55-93
- [17] Mirman, L. J., Samuelson, L., and Urbano, A. (1993), “Monopoly Experimentation”, *International Economic Review*, 549-563.
- [18] Muth, J. F. (1960), “Optimal Properties of Exponentially Weighted Forecasts”, *Journal of the American Statistical Association*, Volume 55, Issue 290
- [19] Nimark, K. (2008), “Monetary Policy with Signal Extraction from the Bond Market”, *Journal of Monetary Economics* 55, 1389–1400
- [20] Pearlman, J. (1992), “Reputational and Nonreputational Policies under Partial Information”, *Journal of Economic Dynamics and Control* 16(2): 339-357
- [21] Prescott, E.C. (1972), The Multi-Period Control Problem Under Uncertainty, *Econometrica*, 40(6), 1043-1058
- [22] Stantcheva, S. (2015), “Learning and (or) Doing: Human Capital Investments and Optimal Taxation”, *NBER Working Paper 21381*
- [23] Svensson, L.E.O. and M. Woodford (2003), “Indicator Variables for Optimal Policy”, *Journal of Monetary Economics* 50, 691–720
- [24] Svensson, L.E.O. and M. Woodford (2004), “Indicator Variables for Optimal Policy under Asymmetric Information”, *Journal of Economic Dynamics & Control* 28, 661 – 690
- [25] Swanson, E.T. (2006), “Optimal Nonlinear Policy: Signal Extraction with a Non-normal Prior”, *Journal of Economic Dynamics and Control* 30: 185-203
- [26] Townsend, R.M. (1983), “Forecasting the Forecasts of Others”, *Journal of Political Economy*, 91(4), 546-588
- [27] Van Nieuwerburgh, S., Veldkamp, L. (2006), “Learning Asymmetries in Real Business Cycles” *Journal of Monetary Economics*, Elsevier, vol. 53(4), pages 753-772, May.
- [28] Wieland, V. (2000a), “Monetary Policy, Parameter Uncertainty and Optimal Learning”, *Journal of Monetary Economics* 46: 199-228

- [29] Wieland, V. (2000b), “Learning by Doing and the Value of Optimal Experimentation”, *Journal of Economic Dynamics & Control* 24: 501-534

A Proofs

A.1 Proof of Lemma 1

To prove Lemma 1, we first state a generic result guaranteeing that for a sequence of functions $f, f^k : X \rightarrow R$ such that $f^k \rightarrow f$, the solutions χ, χ^k satisfying $f(\chi) = f^k(\chi^k) = 0$, are well defined and $\chi^k \rightarrow \chi$. To obtain convergence we need to ensure that the non-linear equations $f^k(\cdot) = 0$ are well-conditioned.⁴²

We make the following **Assumptions**

- **L1** : X is compact. f and f^k absolutely continuous.
- **L2** : f^k converge uniformly to f : $d(f, f^k) \rightarrow 0$ as $k \rightarrow \infty$.

where d is distance implied by the sup-norm.

- **L3** : there is a unique solution to $f(\chi) = 0$, satisfying $\chi \in \text{int}(X)$.

The interiority condition holds generically, if it did not, it can often be restored by expanding X .⁴³

Let $X^d \subset X$ be the set containing all points where the derivatives f' and $f^{k'}$ exist for all k . Because of absolute continuity the set X^d has full Lebesgue measure in X .

- **L4** The derivatives of f^k converge uniformly on X^d :

$$\sup_{X^d} |f'(x) - f^{k'}(x)| \rightarrow 0 \text{ as } k \rightarrow \infty$$

- **L5** : The derivative of f exists at the solution and $f'(\chi) \neq 0$.

L5 ensures that the non-linear equation $f(x) = 0$ is well-conditioned. L5 strengthens mildly L3 in the sense that, loosely speaking, L3 and L5 are generically either both satisfied or both violated for a given f .⁴⁴

Lemma 2 Under Assumptions L1-L5, for k sufficiently high there exists a unique solution to $f^k(\cdot) = 0$, this solution satisfies $\chi^k \in \text{int}(X)$, and $\chi^k \rightarrow \chi$ as $k \rightarrow \infty$.

⁴²Note that the notation for this Lemma differs from the rest of paper, here f denotes a function and X denotes the admissible set of variables x .

⁴³To see this, consider this example: $f(x) = x$, for $X = [0, 1]$. The interiority assumption is violated because $\chi = 0$. In this case one can find $f^k \rightarrow f$ where the Lemma would indeed fail, for example $f^k(x) = -1/k + x$, where it is clear that χ^k is empty. But just enlarging X , for example to $X = [-1, 2]$, would restore the interiority condition of χ and then $\chi^k = 1/k$ exists.

⁴⁴Consider for example $f(x) = yx + x^3$ for $y \geq 0$; both L3 and L5 hold for all $y > 0$; the exception is only for $y = 0$, where L3 holds but L5 fails. Or, consider $f(x) = |x| + y$; both L3 and L5 fail for all $y \neq 0$, the exception is for $y = 0$ where L3 holds but L5 fails. Assumption L5 just ensures that the cases $y = 0$ are excluded.

Proof of Lemma 2

Without loss of generality we consider $f'(\chi) > 0$. The statements in the proof are taken to hold for sufficiently high k .

We first prove that the solution to $f^k(\cdot) = 0$ exists and is interior. Given L3 and L5 there is an $\varepsilon > 0$ sufficiently small so that $(\chi - \varepsilon, \chi + \varepsilon) \subset \text{int}(X)$ and, for some $K > 0$,

$$f'(x) \text{ exists and } f'(x) > K \text{ for all } x \in (\chi - \varepsilon, \chi + \varepsilon). \quad (59)$$

L1, $f(\chi) = 0$ and the fundamental theorem of calculus imply the equality in

$$f(\chi + \frac{\varepsilon}{2}) = \int_{\chi}^{\chi + \varepsilon/2} f' > K \frac{\varepsilon}{2}, \quad (60)$$

the inequality follows from (59). Similarly, integrating over $[\chi - \varepsilon/2, \chi]$ we find

$$-\frac{\varepsilon}{2}K > f(\chi - \frac{\varepsilon}{2})$$

L2 implies that there is a k sufficiently high $d(f^k, f) < \frac{\varepsilon}{4}K$. Together with the last two inequalities we have

$$f^k(\chi - \frac{\varepsilon}{2}) < -\frac{\varepsilon}{4}K < 0 < \frac{\varepsilon}{4}K < f^k(\chi + \frac{\varepsilon}{2})$$

Therefore $f^k(\cdot)$ takes a positive and a negative value in $(\chi - \varepsilon, \chi + \varepsilon)$. Continuity and the intermediate value theorem imply that a solution to $f^k(\cdot) = 0$ exists in $(\chi - \varepsilon, \chi + \varepsilon)$ and, therefore, that this solution is interior.

Take for now χ^k to be the set of solutions, so far we have proved χ^k is non-empty and one of its elements is in $(\chi - \varepsilon, \chi + \varepsilon)$. We now prove χ^k is a singleton. Let $V \equiv \inf_{\{x: |x-\chi| \geq \varepsilon\}} |f(x)|$. Compactness of $\{x \in X : |x - \chi| \geq \varepsilon\}$ and continuity of f imply that the infimum is attained at some $x^* \in X$ with $|x^* - \chi| \geq \varepsilon$. L3 implies that x^* can not solve $f(\cdot) = 0$, so we have the second inequality in

$$|f(x)| \geq |f(x^*)| = V > 0 \text{ for all } x, |x - \chi| \geq \varepsilon. \quad (61)$$

This and uniform convergence implies that for large k

$$|f^k(x)| > V/2 > 0 \text{ for all } x, |x - \chi| \geq \varepsilon. \quad (62)$$

Consider two possibly different solutions $x', x'' \in \chi^k$. Applying a similar argument as the one below equation (60), denoting $x' \geq x''$ then

$$\int_{x''}^{x'} f^{k'} = f^k(x') - f^k(x'') = 0. \quad (63)$$

Now, (62) implies $x', x'' \in (\chi - \varepsilon, \chi + \varepsilon)$. Therefore (59) holds in the interval $[x'', x']$ we have

$$\int_{x''}^{x'} f^{k'} > \frac{K}{2}(x' - x'').$$

Since $K > 0$, combining this with (63) implies $x' = x''$ so that χ^k is a singleton.

To prove convergence of χ^k , using similar arguments as before

$$|f(\chi) - f^k(\chi)| = |f^k(\chi)| = \left| \int_{\chi}^{\chi^k} f^{k'} \right| > \frac{K}{2} |\chi^k - \chi|,$$

L2 implies that the left side goes to zero so that $\chi^k \rightarrow \chi$ as $k \rightarrow \infty$. ■

Proof of Lemma 1

Fix $A \in \Phi$ such that \mathcal{R} is differentiable at all $s \in S(\mathcal{R}, A)$. Denote the bounds $K^\delta, K^{\delta'} < \infty$ for $|\delta(s)| \leq K^\delta$ and $|\delta'(s)| \leq K^{\delta'}$ for all $s \in \mathcal{S}$.

The discussion in the text prior to the statement of Condition 2 shows that $S(\mathcal{R}, A)$ is a singleton.

Given $A \in \Phi$ as above, the proof of convergence applies previous Lemma 2 when $-s$ in the text takes the role of x in Lemma 2.

$-f$ and f^k in Lemma 2 are given by $f(\cdot) \equiv H(\cdot, A; \mathcal{R})$ and $f^k(\cdot) \equiv H(\cdot, A; \mathcal{R} + \alpha_k \delta)$ for a sequence $\alpha_k \rightarrow 0$. The interiority assumption ensures that $\mathcal{R} + \alpha_k \delta$ takes values in \mathcal{T} for k large, hence $H(\cdot, A; \mathcal{R} + \alpha_k \delta)$ is well defined.

$-S(\mathcal{R}, A), S(\mathcal{R} + \alpha_k \delta, A)$ take the role of χ, χ^k .

We now check assumptions L1-L5.

Assumptions 1, 2, and Condition 1 imply L1.

Assumption 3 implies

$$|H(s, A; \mathcal{R}) - H(s, A; \mathcal{R} + \alpha \delta)| < Q |\alpha| K^\delta.$$

Since $Q, K^\delta < \infty$ we have $H(\cdot, A; \mathcal{R} + \alpha_k \delta)$ converge uniformly to $H(\cdot; \mathcal{R}^*)$ so that L2 holds.

L3 follows because $S(\mathcal{R}, A)$ is a singleton.

Using $H_s(s, A; \mathcal{R}) = 1 - h_\tau(\mathcal{R}(s), A) \mathcal{R}'(s)$ we have

$$\begin{aligned} & |H_s(s, A; \mathcal{R}) - H_s(s, A; \mathcal{R} + \alpha \delta)| = \\ & | [h_\tau((\mathcal{R} + \alpha \delta)(s), A) - h_\tau(\mathcal{R}(s), A)] (\mathcal{R}' + \alpha \delta')(s) + h_\tau(\mathcal{R}(s), A) \alpha \delta'(s) | \leq \\ & Q^L |\alpha| K^\delta (K^R + |\alpha| K^{\delta'}) + |\alpha| K^{\delta'} Q \end{aligned}$$

the equality follows from adding and subtracting $h_\tau(\mathcal{R}(s), A) \alpha \delta'(s)$, where Q^L, Q as in Assumption 3 and K^R as in Condition 1. Letting $\alpha \rightarrow 0$ we have that $H_s(\cdot, A; \mathcal{R} + \alpha_k \delta) \rightarrow H_s(\cdot, A; \mathcal{R})$ uniformly as $k \rightarrow \infty$ so that L4 holds.

L5 is given by Condition 2.

All that is left to show is that $A \in \Phi$ such that \mathcal{R} is differentiable at all $s \in S(\mathcal{R}, A)$ occurs with probability 1. Notice that Assumption 4 will only be needed in this part of the proof.

Denote the set of non-differentiable points of \mathcal{R}

$$\mathcal{S}_{\mathcal{R}}^{ND} = \{s \in \mathcal{S}_{\mathcal{R}} : \mathcal{R} \text{ is non-differentiable at } s\}.$$

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

$$\Phi_{\mathcal{R}}^{ND} = \{A \in \Phi : \text{for some } s \in S(\mathcal{R}, A) \text{ we have } s \in \mathcal{S}_{\mathcal{R}}^{ND}\}.$$

We need to prove that $Prob(\Phi_{\mathcal{R}}^{ND}) = 0$. Clearly

$$\Phi_{\mathcal{R}}^{ND} \subset_{s \in \mathcal{S}_{\mathcal{R}}^{ND}} \{A : s = h(\mathcal{R}(s), A)\}. \quad (64)$$

For any $s \in \mathcal{S}_{\mathcal{R}}^{ND}$, $Prob\{A : s = h(\mathcal{R}(s), A)\} = 0$ by Assumption 4. Since by Condition 1 $\mathcal{S}_{\mathcal{R}}^{ND}$ has Lebesgue measure zero, this implies $Prob\left(\bigcup_{s \in \mathcal{S}_{\mathcal{R}}^{ND}} \{A : s = h(\mathcal{R}(s), A)\}\right) = 0$. Then (64) implies that $Prob(\Phi_{\mathcal{R}}^{ND}) = 0$. ■

A.2 Proof of Proposition 2

Consider the problem defined by (27). Fix δ . Given the boundedness assumptions on δ and interiority it is clear that $\mathcal{R}^* + \alpha\delta \in \mathcal{E}$ for α small enough and that $\mathcal{R}^* + \alpha\delta$ satisfies Conditions 1-2. This implies that the solution of (27) is attained at $\alpha = 0$ and by the maximum principle

$$\left. \frac{d\mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha} \right|_{\alpha=0} = 0 \quad (65)$$

if this derivative exists. We first prove that this derivative exists and that (65) implies (26).

We use the following notation: for each function $G = W, W_{\tau}, W_s, \delta, \mathcal{R}^*, \delta', \mathcal{R}^*, h_{\tau}, S$ we denote G^{δ} a 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

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

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

Let

$$M(\alpha, A) \equiv \frac{W^{\delta}(\alpha, A) - W^{\delta}(0, A)}{\alpha}. \quad (66)$$

Take any sequence $\alpha_k \rightarrow 0$. We now show that the limit and integration operators can be exchanged as follows

$$\lim_{k \rightarrow \infty} \int_{\Phi} M(\alpha_k, A) dF_A(A) = \int_{\Phi} \lim_{k \rightarrow \infty} M(\alpha_k, A) dF_A(A). \quad (67)$$

Fix $A \notin \Phi_{\mathcal{R}^*}^{ND}$ where, as in the proof of Lemma 1, $\Phi_{\mathcal{R}^*}^{ND}$ denotes the set of realizations where \mathcal{R}^* is non-differentiable at $S(\mathcal{R}^*, A)$.

It is clear that as long as

$$S^{\delta'}(\alpha, A) \equiv \frac{dS^\delta(\alpha, A)}{d\alpha} = \frac{dS(\mathcal{R}^* + \alpha\delta, A)}{d\alpha} \quad (68)$$

exists we have

$$\begin{aligned} \frac{dW^\delta(\alpha, A)}{d\alpha} &= [W_\tau^\delta(\alpha, A)(\mathcal{R}^{*\delta}(\alpha, A) + \alpha\delta'(\alpha, A)) \\ &\quad + W_s^\delta(\alpha, A)] S^{\delta'}(\alpha, A) + W_\tau^\delta(\alpha, A)\delta^\delta(\alpha, A). \end{aligned} \quad (69)$$

To prove existence of $S^{\delta'}(\alpha, A)$, for given A, \mathcal{R}^*, δ interpret

$$s = h((\mathcal{R}^* + \alpha\delta)(s), A) \quad (70)$$

as a an equation involving s and α . The implicit solution of this equation for each given (small) α is $S^\delta(\alpha, A)$. Lemma 1 implies that the map $\alpha \rightarrow S^\delta(\alpha, A)$. The implicit function theorem implies that $S^{\delta'}$ in (68) exists. Furthermore, carefully differentiating (70) with respect to α we have

$$S^{\delta'}(\alpha, A) = \frac{h_\tau^\delta(\alpha, A) \delta^\delta(\alpha, A)}{1 - h_\tau^\delta(\alpha, A) \mathcal{R}^{*\delta}(\alpha, A)} \quad (71)$$

Notice how the numerator of this expression is non-zero by Condition 2.

This implies that $\frac{dW^\delta(0, A)}{d\alpha}$ exists, that

$$\lim_{k \rightarrow \infty} M(\alpha_k, A) = \frac{dW^\delta(0, A)}{d\alpha} \quad (72)$$

and $\frac{dW^\delta(0, A)}{d\alpha}$ given by (69) and (71) for all $A \notin \Phi_{\mathcal{R}^*}^{ND}$.

At the end of the proof of Lemma 1 we showed that $Prob(\Phi_{\mathcal{R}^*}^{ND}) = 0$. Therefore (72) holds almost surely in A .

Collecting the bounds implied by various assumptions and conditions

$$|M(\alpha_k, A)| \leq Q \left([K^L + \alpha_k K^{\delta'} + 1] \frac{Q}{\varepsilon} + K^\delta \right) \quad (73)$$

for all k , where $K^\delta, K^{\delta'}$ were defined at the beginning of the proof of Lemma 1. With this uniform bound we can apply Lebesgue Dominated Convergence to conclude (67).

Given that

$$\frac{\mathcal{F}(\mathcal{R}^* + \alpha_k\delta) - \mathcal{F}(\mathcal{R}^*)}{\alpha_k} = \int_{\Phi} M(\alpha_k, A) dF_A(A)$$

(67) implies that the derivative $\left. \frac{d\mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha} \right|_{\alpha=0}$ exists and it is equal to $\int_{\Phi} \frac{dW^\delta(0, A)}{d\alpha} dF_A$.

To find $\frac{dW^\delta(0, A)}{d\alpha}$ we evaluate (69) at $\alpha = 0$ and

$$\left. \frac{d\mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha} \right|_{\alpha=0} = \int_{\Phi} \{ [W_\tau^* \mathcal{R}^{*\delta} + W_s^*] S^{\delta'}(0, \cdot) + W_\tau^* \delta(S(\mathcal{R}^*, \cdot)) \} dF_A$$

Finally, (65), (71) and rearranging gives

$$\int_{\Phi} \frac{W_{\tau}^* + W_s^* h_{\tau}^*}{1 - h_{\tau}^* \mathcal{R}^{*'}} \delta(S(\mathcal{R}^*, \cdot)) dF_A = 0. \quad (74)$$

Since (74) 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 (26) (see footnote 24). ■

A.3 Proof of Proposition 3

Take random variables X, Y, Z . Assume there is a function g such that $X = g(Y, Z)$ a.s.; clearly for any function ϕ

$$E(\phi(X, Y) | Z) = E(\phi(g(Y, Z), Y) | Z). \quad (75)$$

Letting A_1 play the role of X , $S(\mathcal{R}^*, A)$ the role of Z , A_2 the role of Y and \mathcal{A}^* the role of g we have the second equality in

$$\begin{aligned} & E \left(\frac{W_{\tau}^* + W_s^* h_{\tau}^*}{1 - h_{\tau}^* \mathcal{R}^{*'}} \middle| S(\mathcal{R}^*, A) = s \right) = \\ & E \left(\frac{W_{\tau} + W_s h_{\tau}}{1 - h_{\tau} \mathcal{R}^{*'}} (\tau^*, s, A_1, A_2) \middle| S(\mathcal{R}^*, A) = s \right) = \\ & E \left(\frac{W_{\tau} + W_s h_{\tau}}{1 - h_{\tau} \mathcal{R}^{*'}} (\tau^*, s, \mathcal{A}^*(s, A_2), A_2) \middle| S(\mathcal{R}^*, A) = s \right) = \\ & \int_{\Theta_2(s, \mathcal{R}^*(s))} \frac{W_{\tau} + W_s h_{\tau}}{1 - h_{\tau} \mathcal{R}^{*'}} (\tau^*, s, \mathcal{A}^*(s, a_2), a_2) dF_{A_2|S(\mathcal{R}^*, A)}(a_2, s) \end{aligned} \quad (76)$$

For simplicity in the last equality and the rest of the proof we consider the case where A_2 has a density, if A_2 is a discrete variable the argument can be readily adapted.

We apply Bayes' formula to find $dF_{A_2|S(\mathcal{R}^*, A)}$. First we find $f_{S(\mathcal{R}^*, A)|A_2}$. Since $\mathcal{A}^*(s, A_2)$ is the inverse function of $S(\mathcal{R}^*, \cdot, A_2)$, by the change of variable rule the density $f_{S(\mathcal{R}^*, A)|A_2}$ exists and is given by

$$f_{S(\mathcal{R}^*, A)|A_2}(s, a_2) = f_{A_1|A_2}(\mathcal{A}^*(s, a_2), a_2) |\mathcal{A}_s^*(s, a_2)| \quad (77)$$

where $\mathcal{A}_s^* = \frac{\partial \mathcal{A}^*}{\partial s}$. To find this derivative we apply once again the implicit function theorem to H to find

$$\mathcal{A}_s^* = \frac{1 - h_{\tau}^* \mathcal{R}^{*'}}{h_{A_1}^*}, \quad (78)$$

given Assumption 5 $h_{A_1}^* \neq 0$ hence this derivative is well defined. Plugging (78) into (77) gives $f_{S(\mathcal{R}^*, A)|A_2}$. Applying Bayes' rule, for any point with a positive density $f_{S(\mathcal{R}^*, A)} = \int f_{S(\mathcal{R}^*, A)|A_2} f_{A_2} da_2$ we have

$$dF_{A_2|S(\mathcal{R}^*, A)} = \frac{f_{S(\mathcal{R}^*, A)|A_2} dF_{A_2}}{f_{S(\mathcal{R}^*, A)}}. \quad (79)$$

Plugging the above formula for $dF_{A_2|S(\mathcal{R}^*, A)}$ we see the last expression in (76) becomes

$$\int_{\Theta_2(s, \tau^*)} (W_\tau + W_s h_\tau) \frac{f_{A_1|A_2}}{f_{S(\mathcal{R}^*, A)} |h_{A_1}^*|} (\tau^*, s, \mathcal{A}^*(s, a_2), a_2) dF_{A_2}(a_2)$$

where we cancelled out $1 - h_\tau^* \mathcal{R}^{*'} > 0$ in the numerator and denominator. Using (26) we set this equal to zero, the term $f_{S(\mathcal{R}^*, A)}(s)$ drops out. This gives (30). ■

A.4 Proof of Proposition 4

We consider a variation $\mathcal{R}^* + \alpha\delta$ as in the proof of Proposition 2. To prove Proposition 4, we need to compute the second derivative $\frac{d^2 \mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha^2} \Big|_{\alpha=0}$. A necessary condition for a maximum is that this second derivative is non-positive.

Thanks to Assumption 6, we can use Lebesgue Dominated Convergence to pass the derivative inside the integral over the realizations of the shocks A (as we did in the proof of Proposition 2).

Hence, we get

$$\frac{d^2 \mathcal{F}(\mathcal{R}^* + \alpha\delta)}{d\alpha^2} = E \frac{d^2 W^\delta(\alpha, A)}{d\alpha^2} \quad (80)$$

Next, we compute the derivative inside the expectation term in (80). This is given by

$$\frac{d^2 W^\delta(\alpha, A)}{d\alpha^2} = \frac{d(W_{ss}^\delta(\alpha, A) (S^{\delta'}(\alpha, A)))}{d\alpha} = W_{ss}^\delta(\alpha, A) (S^{\delta'}(\alpha, A))^2 + W_s^\delta(\alpha, A) S^{\delta''}(\alpha, A) \quad (81)$$

where the first derivative of the signal with respect to α , $S^{\delta'}(\alpha, A)$, is given by equation (71) and the second derivative, $S^{\delta''}(\alpha, A)$, is

$$S^{\delta''}(\alpha, A) = \frac{B_1^\alpha [1 - h_\tau(\mathcal{R}^{*'} + \alpha\delta')] + B_2^\alpha \delta h_\tau}{[1 - h_\tau(\mathcal{R}^{*'} + \alpha\delta')]^2} \quad (82)$$

with

$$B_1^\alpha \equiv \delta' h_\tau S^{\delta'} + \delta^2 h_{\tau\tau} + \delta h_{\tau\tau} (\mathcal{R}^{*'} + \alpha\delta') S^{\delta'}$$

$$B_2^\alpha \equiv h_{\tau\tau} (\mathcal{R}^{*'} + \alpha\delta')^2 S^{\delta'} + \delta h_{\tau\tau} (\mathcal{R}^{*'} + \alpha\delta') + h_\tau \delta' + h_\tau (\mathcal{R}^{*''} + \alpha\delta'') S^{\delta'}$$

where we dropped the functions arguments to avoid excessively heavy notation, but it is understood that $S^{\delta'}$ is evaluated at (α, A) .

Evaluating this expression at $\alpha = 0$, we get

$$S^{\delta''}(0, A) = \frac{B_1 (1 - h_\tau^* \mathcal{R}^{*'}) + B_2 \delta h_\tau^*}{(1 - h_\tau^* \mathcal{R}^{*'})^2} \quad (83)$$

with

$$B_1 \equiv \delta' h_\tau^* S^{\delta'} + \delta^2 h_{\tau\tau}^* + \delta h_{\tau\tau}^* \mathcal{R}^{*'} S^{\delta'}$$

$$B_2 \equiv h_{\tau\tau}^* (\mathcal{R}^{*'})^2 S^{\delta'} + \delta h_{\tau\tau}^* \mathcal{R}^{*'} + h_{\tau}^* \delta' + h_{\tau}^* \mathcal{R}^{*''} S^{\delta'}$$

where $S^{\delta'}$ is now evaluated at $(0, A)$.

Doing some simple algebra, this expression simplifies to (38). Hence, evaluating (81) at $\alpha = 0$, integrating it over the distribution of the shocks A and imposing that this integral is non-positive gives condition (37). ■

A.5 Proof of Proposition 5

It is clear that \mathcal{E} is non-empty. For example, a constant policy $\mathcal{R}(s) = \tilde{\tau}$ for some $\tilde{\tau} \in \mathcal{T}$ satisfies Conditions 1 and 4.

To prove that \mathcal{E} is compact take a sequence $\mathcal{R}_k \in \mathcal{E}$ for $k = 1, 2, \dots$. Condition 1 is equivalent to Lipschitz continuity, the Arzelà-Ascoli theorem guarantees that there is a subsequence \mathcal{R}_{k_i} such that $\mathcal{R}_{k_i} \rightarrow \mathcal{R}^{\text{lim}}$ uniformly in \mathcal{S} as $i \rightarrow \infty$ and \mathcal{R}^{lim} is Lipschitz, therefore \mathcal{R}^{lim} satisfies Condition 1.

Take A such that the derivative of \mathcal{R}^{lim} exists at all elements of $S(\mathcal{R}^{\text{lim}}, A)$ and take s such that $|s - h(\mathcal{R}^{\text{lim}}(s), A)| < \delta$. Fix any $\varepsilon > 0$. For any s' such that $|s - h(\mathcal{R}^{\text{lim}}(s), A)| < \delta$ we have

$$\begin{aligned} \frac{h(\mathcal{R}^{\text{lim}}(s), A) - h(\mathcal{R}^{\text{lim}}(s'), A)}{s - s'} &= \frac{h(\mathcal{R}^{\text{lim}}(s), A) - h(\mathcal{R}_{k_i}(s), A)}{s - s'} \\ &+ \frac{h(\mathcal{R}_{k_i}(s'), A) - h(\mathcal{R}^{\text{lim}}(s'), A)}{s - s'} + \frac{h(\mathcal{R}_{k_i}(s), A) - h(\mathcal{R}_{k_i}(s'), A)}{s - s'} \end{aligned}$$

Since $\mathcal{R}_{k_i} \rightarrow \mathcal{R}^{\text{lim}}$ the first two numerators to the right of the equality are eventually less than ε . Applying (yet again) the fundamental theorem of calculus, the last fraction is less than the a.e. derivative of $h(\mathcal{R}_{k_i}(s), A)$ and, since all \mathcal{R}_{k_i} 's satisfy Condition 4, the last fraction above is less than $1 - \bar{\varepsilon}$ for all i . Hence, taking limit as $i \rightarrow \infty$, for all $\varepsilon > 0$ we have

$$\frac{h(\mathcal{R}^{\text{lim}}(s), A) - h(\mathcal{R}^{\text{lim}}(s'), A)}{s - s'} \leq \frac{2\varepsilon}{s - s'} + 1 - \bar{\varepsilon}.$$

Since this holds for all $\varepsilon > 0$, it also holds for $\varepsilon = 0$ so that $\frac{h(\mathcal{R}^{\text{lim}}(s), A) - h(\mathcal{R}^{\text{lim}}(s'), A)}{s - s'} \leq 1 - \bar{\varepsilon}$ and it means that \mathcal{R}^{lim} satisfies (24). Therefore \mathcal{R}^{lim} satisfies Condition 4.

This proves that any sequence $R_k \in \mathcal{E}$ has a subsequence converging uniformly to a limit \mathcal{R}^{lim} that satisfies Conditions 1 and 4 therefore \mathcal{E} is compact. Since the objective function is continuous and it takes values on bounded sets $\mathcal{S} \times \mathcal{T}$, a maximum exists. ■

B The effect of government expenditure on the signal extraction

In this Appendix we show that in the log-quadratic case an increase in g can change the sign of optimal policy in the intermediate region of the signal.

First, for a given \bar{l} consider the locus $\tilde{\gamma}(\mathcal{R}(\bar{l}), \theta, \bar{l})$ which is the inverse of the reaction function h with respect to γ , implicitly defined by

$$h(\mathcal{R}(\bar{l}), \theta, \tilde{\gamma}(\mathcal{R}(\bar{l}), \theta, \bar{l})) - \bar{l} = 0. \quad (84)$$

Let ζ be the derivative of this function with respect to θ . This can be found using the implicit function theorem, which gives the positive slope

$$\zeta = -\frac{h_\theta}{h_\gamma} = \frac{g\sqrt{(Bg)^2 + 4B\theta^2\gamma(1-\tau)} + Bg^2}{2\theta^3(1-\tau)} \quad (85)$$

Now we differentiate ζ with respect to g and get

$$\frac{d\zeta}{dg} = \frac{\partial\zeta}{\partial g} + \frac{\partial\zeta}{\partial\tau} \frac{\partial\tau}{\partial g}. \quad (86)$$

All the partial derivatives in the right-hand side of (86) are positive. Hence, higher government expenditure makes the loci of realizations of the shocks steeper.


Now we illustrate how this effect changes the nature of the signal extraction on the shocks and hence the slope of optimal policy in the intermediate region of the signal. For this purpose, we will take a first order approximation of the loci $\tilde{\gamma}$.⁴⁵

Consider Figures 17 and 18. When g is sufficiently low, the map of loci (solid blue lines) moving in the direction of increasing l 's looks like in Figure 17. Starting at l_{min} (bottom right corner) and increasing l the loci first hit the bottom-left corner, where the lowest θ becomes possible, and then the top-right corner, where the highest values for θ start to be inconsistent with the observed l 's. Hence in the intermediate region of l 's all θ 's are possible, but clearly not all γ 's. In this region, the government learns little about productivity. All the government learns is that the agent is working more as l increases so expected output is higher and taxes can be lower. This gives the negative slope of \mathcal{R}^* with low government expenditure.

On the other hand, when g is sufficiently high, the slope of the loci becomes higher. Hence, as illustrated in Figure 18, in the intermediate region of l 's the government learns that only a relatively small set of θ 's is possible, whereas any γ is consistent with the observations. This leads to the positive slope of the optimal tax rate with high government expenditure.

⁴⁵In our computed examples these loci are very close to linear.

Figure 17: Loci of shocks realizations with low g



X-axis: θ , Y-axis: γ . This figure illustrates how the set of (θ, γ) with positive density changes as the value of the signal l increases, starting from the bottom-right corner for l_{min} and moving up and left until reaching the top-left corner for l_{max} . With low g , these loci of shocks hit the bottom-left corner at a lower value of l than the top-right corner.


C Additional models

In this appendix we present two additional versions of our two-period model. In the first additional model, we assume that the productivity shock is permanent and we maintain labor as the signal. In the second model, we consider output as a signal, instead of labor. Next, we provide some additional derivations of the model with learning-by-doing introduced in Section 6. Finally, we present a modification of our model that helps clarifying the relation between GSE and issues of learning and experimentation studied in the previous literature.

C.1 Permanent productivity shock

We now consider a permanent productivity shock, that is $\theta_2 = \theta_1 = \theta$, while we maintain the assumption of transitory demand shock γ . This setup would be uninteresting with log-quadratic preferences because hours would become independent of θ as wealth and substitution effects would cancel out. Therefore we assume linear utility from consumption (as in the infinite-h orizon model) $u(c) = c$ and quadratic disutility from labor $v(l) = \frac{B}{2}l^2$, which give another case with an analytical solution

Figure 18: Loci of shocks realizations with high g



X-axis: θ , Y-axis: γ . This figure illustrates how the set of (θ, γ) with positive density changes as the value of the signal l increases, starting from the bottom-right corner for l_{min} and moving up and left until reaching the top-left corner for l_{max} . With high g , these loci of shocks hit the top-right corner at a lower value of l than the bottom-left corner.

for the reaction function h and its derivatives.

In particular, it is easy to see from the first order condition (6) that the reaction function (12) specializes to


$$l = h(\tau, \theta, \gamma) = \frac{\gamma\theta}{B}(1 - \tau), \quad (87)$$

The optimal policy is illustrated in Figure 19 and is compared with the set of FI tax rates conditional on l . In general, \mathcal{R}^* is decreasing as higher observed labor suggests higher conditional expectation for productivity, hence allowing to balance the intertemporal budget constraint with a lower distortionary tax. The figure also compares the optimal policy with a linear policy obtained connecting the two full revelation points with a straight line.

C.2 Observable output

Our benchmark example assumed that labor was observable. As a further exercise, we now consider the case in which the signal observed by the government is output, so $s = y = \theta l$. The government cannot sort out the values of the θ and l independently,

Figure 19: Optimal policy with permanent productivity shock




Optimal tax rate as a function of the signal l when the productivity shock is permanent. Red line: \mathcal{R}^* ; yellow region: set of equilibrium (l, τ) under FI.

hence there still is a problem of signal extraction as the government can not back out the values of (θ, γ) . In addition, to make the problem of PI more relevant, we assume that the productivity shock is permanent, that is $\theta_1 = \theta_2 = \theta$ as in the subsection above. We consider again preferences that are linear in consumption and quadratic in labour.

Figure 20 illustrates the optimal policy for this case. When the signal is output, the government has a lot more information than in the case studied in the main text. This can be seen by observing that the yellow region is thinner than in Figure 19 (for the same size of the shocks). The reason for this is that current revenue $\tau\theta l$ is known, and hence there is no uncertainty about the amount of debt that needs to be issued. The only uncertainty is about the amount of revenue that will be collected in the future, as the value of (permanent) productivity is unknown.⁴⁶ As we have seen, uncertainty about debt is key to get fiscal adjustments as in our benchmark example and this motivates our choice to focus on hours as the signal in our analysis.

⁴⁶Note that if we had assumed that θ_2 is not random the model would not be interesting, since in that case there would be no uncertainty about either current or future revenue.

Figure 20: Optimal Policy with output observed



Optimal tax rate as a function of the signal y (output). Red line: \mathcal{R}^* ; yellow region: set of equilibrium (l, τ) under FI.

C.3 Additional derivations of the model with learning-by-doing in Section 6

We report here the optimality conditions omitted in Section 6 and define the FI-Ramsey problem in this model.

First order conditions with respect to bonds, and with respect to hours and consumption in periods 2 and 3 give

$$q_t = \beta \frac{\gamma_{t+1} u'(c_{t+1})}{\gamma_t u'(c_t)} \quad (88)$$

for $t = 1, 2$,

$$v'(l_2) - u'(c_2) \bar{\theta} \chi(l_1) (1 - \tau_2) - \beta u'(c_3) \bar{\theta} l_3 (1 - \tau_3) \chi'(l_2) = 0 \quad (89)$$

and

$$v'(l_3) - u'(c_3) \bar{\theta} (1 - \tau_3) = 0. \quad (90)$$

By substituting all household FOCs into the intertemporal government budget constraint, and using our assumed functional forms, we can obtain the implementability constraint of the FI-Ramsey problem:

$$\gamma - B l_1^{1+\eta} + \beta [1 - B l_2^{1+\eta} (1 - \psi)] + \beta^2 [1 - B l_3^{1+\eta} (1 - \psi (1 - \psi))] = 0 \quad (91)$$

The FI-Ramsey problem is to choose the allocation that maximizes (44) subject to the implementability constraint (91) and subject to the additional constraint

$$Bl_1^{1+\eta} - \frac{\gamma\theta(1-\tau_1)l_1}{c_1} - \beta\psi(Bl_2^{1+\eta} - \beta\psi Bl_3^{1+\eta}) = 0, \quad (92)$$

given τ_1 .

This additional, non-standard, constraint is obtained by combining the labor FOCs in periods 1 and 2 and reflects the fact that the government is not free to choose any allocation that satisfies implementability, because the tax rate in period 1 is predetermined. Hence, the Ramsey allocation has to satisfy the FOC for labor supply in period 1 at the given value of τ_1 .

C.4 Learning, experimentation and GSE

We now present a modification of our model that may help to clarify the distinction between General Signal Extraction and the related issues that are the focus of other papers in which posterior beliefs about the uncertain state of the economy are endogenous to policy decisions. For instance, Prescott (1972), Kiefer and Nyarko (1989) and Wieland (2000a, 2000b) study optimal experimentation in linear models, whereas Van Nieuwerburgh and Veldkamp (2006) consider a model of endogenous learning with non-linearities, but abstract from incentives to active experimentation. Our example shows that by assuming that policy actions are predetermined with respect to the determination of signals, these papers avoid the issue of GSE, which arises in more general setups and is the focus of our paper.

In order to illustrate this point, we modify our simple model of Section 3 to obtain endogenous posterior beliefs, but in absence of endogeneity of signals to the contingent policy function, because taxes are predetermined. The economy lasts for three periods $t = 1, 2, 3$.

The utility function is

$$\gamma u(c_1) - v(l_1) + \beta [u(c_2) - v(l_2)] + \beta^2 [u(c_3) - v(l_3)] \quad (93)$$

As in our baseline model, the preference shock is temporary and only hits in the first period. However, in order to generate incentives to learn, we assume the productivity shock is permanent, that is $\theta_1 = \theta_2 = \theta_3 = \theta$.

Taxes in the first two periods are predetermined with respect to allocations. Specifically, the government chooses τ_1 before observing any signal (it is a non-contingent policy), then chooses τ_2 contingent on the signal l_1 , that is $\tau_2 = \mathcal{R}(l_1)$, without observing the realizations of the shocks, nor the rest of the allocations and prices in period 1 and finally, in period 3, observes the realizations of the shocks and the allocation and sets τ_3 to the value that balances the budget.

In this model, the government understands that a choice of τ_1 will affect the distribution of the signal l_1 , making the posterior distributions of the shocks at $t = 2$

endogenous to policy in period 1. This endogeneity of posteriors generates an incentive to experiment, and violates the “separation principle”, as discussed by Wieland (2000a, 2000b). However, importantly, the signal l_1 does not depend on \mathcal{R} , because of the absence of any physical state variable linking the labor supply decision at $t = 1$ to taxes at $t = 2$.

To see this, notice that in equilibrium, hours in the first period are given by

$$\frac{v'(l_1)}{u'(\theta l_1 - g)} = \theta\gamma(1 - \tau_1) \quad (94)$$

and hence do not depend on the function \mathcal{R} . In this model, the only variable in period 1 that responds to expectations about future taxes is the bond price q_1 .

In a more general model with endogenous state variables, such as the model of learning-by-doing of Section 6, or in a model where policy is contingent on contemporaneous signals, such as the simple model of Section 3, the distribution of the signal is a function of the policy \mathcal{R} , hence the issue of GSE, which is the focus of our paper, arises.

D Derivation of the Envelope Condition (58)

In this Appendix we derive the Envelope Condition (58). First of all let us introduce the necessary notation. A tax policy is a function of debt and labor $\mathcal{R}(b, l)$ and labor is a function of a policy \mathcal{R} , outstanding debt and the exogenous shock, $L(\mathcal{R}; b, A)$ defined by the zero of

$$H(l, A, \mathcal{R}) \equiv l - h(\mathcal{R}(b, l), A), \quad (95)$$

in analogy with the two-period model. By total differentiation of (95), the partial derivative of labor with respect to debt, L_b , is given by

$$L_b(\mathcal{R}, b, A) = \frac{\gamma\theta\mathcal{R}_b(b, l)}{v''(l) + \gamma\theta\mathcal{R}_L(b, l)}. \quad (96)$$

Now, for simplicity consider a case without borrowing penalty. In order to derive the envelope condition, we differentiate (56) with respect to b and get

$$V'(b) = E \left[(\gamma\theta - v'(l^*)) L_b^* + V'(b^*) \left(\frac{\gamma}{\bar{\gamma}} + \beta b_L^* L_b^* \right) \right]$$

where

$$b_L^* = \frac{-\theta\gamma + [v''(L(\mathcal{R}^*, A, b))L(\mathcal{R}^*, A, b) + v'(L(\mathcal{R}^*, A, b))]}{\beta\bar{\gamma}}$$

$$l^* = L(\mathcal{R}^*; b, A)$$

$$L_b^* = L_b(\mathcal{R}^*; b, A).$$

Using (96) we can write

$$V'(b) = E \left[\left(\gamma\theta - v'(l^*) + \beta V'(b^*) b_L^* \right) \frac{-\gamma\theta \mathcal{R}_b^*(l, b)}{v''(l) + \gamma\theta \mathcal{R}_L^*(l, b)} + V'(b^*) \frac{\gamma}{\bar{\gamma}} \right]. \quad (97)$$

Using Proposition 3, the FOC of PI Ramsey problem is

$$E \left[\left(\theta\gamma - v'(l^*) + \beta V'(b^*) b_L^* \right) \frac{h_\tau^*}{1 - h_\tau^* \mathcal{R}_L^*} \Big| \bar{l} \right] = 0 \quad (98)$$

for all \bar{l} . Furthermore, we have that the partial derivative of the reaction function h with respect to taxes is

$$h_\tau = \frac{-\gamma\theta}{v''(l)}.$$

So from (97) we get

$$V'(b) = E \left[\left(\gamma\theta - v'(l^*) + \beta V'(b^*) b_L^* \right) \frac{h_\tau^* \mathcal{R}_b^*(l, b)}{1 - h_\tau^* \mathcal{R}_L^*(L, b)} + V'(b^*) \frac{\gamma}{\bar{\gamma}} \right].$$

Now, applying the law of iterated expectations, using the fact that $\mathcal{R}_b(l, b)$ is known given L, b and using (98), we obtain

$$\begin{aligned} V'(b) &= E \left[E \left(\left(\gamma\theta - v'(L^*) + \beta V'(b^*) b_L^* \right) \frac{h_\tau^* \mathcal{R}_b^*(L, b)}{1 - h_\tau^* \mathcal{R}_L^*(L, b)} \Big| L \right) + V'(b^*) \frac{\gamma}{\bar{\gamma}} \right] \\ &= E \left[0 + V'(b^*) \frac{\gamma}{\bar{\gamma}} \right] \end{aligned} \quad (100)$$

Finally, adding the marginal cost of excessive debt this becomes

$$V'(b) = E \frac{\gamma V'(b^*)}{\bar{\gamma}} - \frac{\gamma}{\bar{\gamma}} \chi(b' - b^{\max}) I_{[b^{\max}, \infty)}(b').$$