Optimal Policy with Endogenous Signal Extraction*

Esther Hauk† Andrea Lanteri‡ Albert Marcet§

September 2015

Preliminary and incomplete

Abstract

This paper studies optimal policy in a setup with hidden information and observed signals that are endogenous to policy chosen in the same period. In this case the signal extraction problem about the state of the economy cannot be separated from the determination of the optimal policy. We derive a non-standard first order condition of optimality from first principles and we use it to find numerical solutions. We show how available results where separation obtains arise as special cases. We use as an example an optimal fiscal policy problem. We find that the policy response can be quite non-linear, calling for tax smoothing in normal times, but for a strong fiscal adjustment during a slump that protracts the downturn. This can provide a rationale for the austerity programs followed by some European countries in the Great Recession. Applying the separation principle and a linear approximation would miss this potential non-linearity and could thus be a poor approximation to the correct solution.

*The authors appreciate helpful comments from Wouter den Haan, Paul Levine, Martí Mestieri, Joe Pearlman, Alex Tsyvinski, Robert Shimer, Jaume Ventura and other participants at the LSE Macro Work-in-progress seminar, IAE-CSIC Macro work-in-progress, Barcelona GSE Winter Workshop 2013, Midwest Macro Spring Meeting 2014, SED Annual Meeting 2014. Hauk and Marcet are grateful for support from Plan Nacional (Spanish Ministry of Science) and the Generalitat de Catalunya. Hauk has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 612796. 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 European Research Council under the EU 7th Framework Programme (FP/2007-2013) under grant agreement 324048-APMPAL.

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

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

§Institut d'Anàlisi Econòmica-CSIC, ICREA, Barcelona GSE, UAB, MOVE; 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

The opening quote states that inferring the underlying state of the economy is a key practical difficulty in setting macro policy. One could say that this is not such an insurmountable problem: a policy-maker should choose the optimal policy taking into account the uncertainty ("haziness") about the underlying state given the information available to him. However, information available to policy-makers is often endogenous to policy decisions. Hence, in general the problem of designing optimal policy contingent on endogenous information cannot be separated from the problem of inferring the state of the world from the observables. In this paper we develop a general solution to the problem of finding optimal policy with signal extraction from endogenous variables and we illustrate our results in a simple model of fiscal policy.

To illustrate the importance of this issue, consider the fiscal policy response to the recent 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 or to a demand shock or some combination of both. Nonetheless, policy-makers had to design a reaction to the recession, based only on a signal extraction on the nature of the shock. Since output itself depends on whether an expansionary fiscal policy or austerity is adopted, the problem of signal extraction is endogenous to the choice of policy.

The G20 decided in its Washington meeting to take aggressive expansionary fiscal policy measures in order to reactivate the economy. Those measures involved spending 2% of world GDP. Policy-makers were uncertain about the fiscal deficits that would ensue. From the relatively optimistic initial estimate many governments have eventually switched to the opposite fiscal policies, implementing tax hikes and austerity programs. In many occasions the switch to austerity has been quite sudden. Austerity had to be even stricter because of the higher spending that the G20 encouraged.[1]

[1] The Spanish Finance Minister at the time, Pedro Solbes, published an article in El País on

The standard literature on Ramsey-optimal fiscal policy assumes that taxes can be a function of the realizations of the shocks hitting the economy. However in the real world these policies can only depend on observable variables such as income, that in turn responds simultaneously to taxes as well as the underlying shocks. Automatic stabilizers, e.g. income taxes and unemployment benefits, are a leading example of policies that respond to aggregate endogenous information. This paper addresses the question of how to design such instruments optimally.

The existing literature on optimal policy with signal extraction has already pointed out that there can be an endogeneity issue when signals respond to policy. However it has worked only with linear models, where a "certainty equivalence" result arises: under Partial Information it is optimal to apply the same policy as under Full Information to the conditional expectation of the underlying shocks. While the signal extraction can in some cases depend on the policy, the optimal policy itself is independent of the information structure and can be derived easily.[2]

A model of fiscal policy is non-linear in nature. For instance, productivity shocks enter multiplicatively with hours worked in the budget constraint. Furthermore, the existence of a Laffer curve creates non-linearities in tax revenue. This paper derives general non-linear methods to solve for optimal policy when there is no separation of any kind between optimization and signal extraction, leading to a violation of "certainty equivalence".

We first analyze optimal policy in a static (or two-period) model of optimal control with multidimensional uncertainty and an endogenous signal. The density of the signal depends on the optimal policy and viceversa. We derive a first order condition (FOC) for the optimal policy relying on first principles. This optimality condition is different from the standard FOC found in dynamic stochastic models: the probabilities of each state of nature need to be weighted by a kernel that depends on the effect of policy on the observed signal. The FOC is derived for a general model so that our results can be widely applied. We show how "certainty equivalence" arises in some special cases as an implication of our general theorem.

Our leading example is a two-period version of the standard fiscal policy model of Lucas and Stokey (1983). We introduce two shocks (to demand and to supply) and to make the issue of hidden information relevant we assume incomplete insurance markets as in Aiyagari et al. (2002). Then we solve for optimal Ramsey taxation under the assumption that the government does not observe the realizations of the shocks, but it only observes some endogenous signal, such as output or hours worked.

This model yields interesting insights about the conduct of fiscal policy. First, hidden information can be a driver of tax smoothing: taxes are less volatile relative to a case with Full Information. This is because one implication of the optimal policy

---

the 10th of November 2013 under the title "Cuando decidí salir del Gobierno" explaining how the outcome of this meeting eventually lead to the current Spanish Debt crisis and near-default in summer 2012.

[2]See Section 2 for a detailed discussion of the literature.

is to average all the possible contingencies using "proper" weights for each contingency that depend on how reactive observables are to policy. This identifies a reason for tax smoothing different from the standard one in Ramsey policy.

When we consider a case with very high government spending, capturing a situation of fiscal crisis, the implied policy reaction can be very non-linear. In particular, in a case when future taxes could be close to the maximum of the Laffer curve, as was arguably the case in some European economies in the current crisis, a government may go very quickly from not reacting to low observed output to increasing taxes very strongly as output goes down. This sudden adjustment occurs once policy-makers realize that a fiscal crisis may materialize under the worst possible realizations of the shocks in the next period, but importantly this inference is endogenous to the output observed and the policy chosen. This may rationalize why some countries first reacted slowly to the crisis and then went for austerity quickly but with a delay. We also show the consequences of using the existing linear methods based on "certainty equivalence" when the actual optimal policy is highly non-linear.

Finally, we build an infinite-horizon model that confirms our main results in a dynamic setup: the Ramsey government under Partial Information may react with a delay to a downturn. The delayed fiscal adjustment induces a longer recession relative to the Full Information benchmark, as has arguably happened in the countries that expanded government spending after November 2008 and eventually introduced austerity measures. It may seem that irresponsible politicians amplified the recession by delaying tax increases, but (in the model) this is a feature of an optimal decision.

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 are main theoretical contribution. It provides the first order condition for a general static model, compares the Full Information solution with the Partial Information solution and discusses a case of "invertibility" when the two solutions coincide. We also study a case where "certainty equivalence" holds and provide a simple algorithm for computing optimal policy. In section 5 we apply these results to our Ramsey fiscal policy problem. Section 6 is dedicated to robustness. Section 7 presents the infinite-horizon model. Section 8 concludes.

2 Related literature

Most of the literature on optimal policy in dynamic models in the last thirty years has disregarded the issue of endogenous signal extraction. But Partial Information and signal extraction were often present in the early papers on dynamic models with Rational Expectations (RE).

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 before deciding on policy it is said that there is a "separation principle" between optimization and estimation.

Few papers have studied optimal policy when signals are endogenous. Pearlman et al. (1986), 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, thanks to linearity and additively separable shocks. Moreover optimal policy has the "certainty equivalence" property: under Partial Information the government applies the Full Information policy to its best estimate of the state. In our setup, because of non-linearity, both the "separation principle" and "certainty equivalence" fail. 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 while at the same time understanding that it is affecting them.

Our contribution is to consider a fully microfounded optimal policy model and to provide a general solution to the endogenous signal extraction and optimization problem when the government (or, more generally, a Stackelberg leader) conditions on a signal simultaneously determined with policy. There is no separation of any kind if the shocks and policy variables are allowed to enter non-linearly in the equilibrium conditions. We find how "certainty equivalence" arises in the linear case using our generic theorem and we show a case where the correct solution is highly non-linear in nature and, therefore, a linear approximation can be misleading in this case.[3]

The effect of policy choices on information extraction about unobserved variables is also considered in the armed-bandit problems of decision theory. For some applications to dynamic macro policy see Kiefer and Nyarko (1989), Wieland (2000a, 2000b) and Ellison and Valla (2001), or to monopolist behavior see Mirmann, Samuelson and Urbano (1993). Van Nieuwerburgh and Veldkamp (2006) use a similar learning framework to explain business-cycle asymmetries. In these papers the planner's decision influences the probability of next period's signal, but the current signal is unaffected

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

by the current policy decision so that separation holds given the past policy decision. It would be interesting to blend the issues in these papers with the one studied in the current paper.

Furthermore, a wide literature considers models where competitive agents use prices as signals of unknown information.[4] Since prices are taken as given by competitive agents the signal extraction problem can be solved with standard filtering techniques and the issue we address does not arise in this literature.

The literature of optimal contracts under private information and incentive compatibility constraints (or the "new dynamic public finance" as in Kocherlakota, 2010) is perhaps less directly related to our work. This literature usually assumes revelation of the private information conditional on the equilibrium actions (the "invertibility" case where we argue below Partial Information is not relevant) but it assumes that agents react strategically to the optimal contingent policy set up by the principal, an issue that we abstract from. In our setup the government conditions on aggregate variables and, since agents are atomistic, the policy function ($\mathcal{R}$ in the text) does not affect agents' decision while the government action (the tax rate τ in our main example) does and it is taken as given. On the other hand, there is an interaction between the signal extraction problem and the optimal policy decision that the literature on incentive compatibility constraints often ignores. Blending the two issues would be of interest but it is left for future research.

3 A simple model of optimal fiscal policy

We first present a simple model of optimal policy. This will serve the purpose of illustrating the problem of endogenous signal extraction and it will be of interest on its own. It is a very simple two-period version of Lucas and Stokey (1983). We introduce incomplete insurance markets to be consistent with the Partial Information story.

3.1 Preferences and technology

The economy lasts two periods $t = 1, 2$. A government needs to finance an exogenous and deterministic stream of expenditure (g_1, g_2), where subscripts indicate time periods, using distortionary income taxes (τ_1, τ_2) and bonds b^g issued 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 utility function

[4] Lucas's (1972) seminal paper, and Guerrieri and Shimer (2013) analyze a competitive market in this setup. For an optimal policy problem see Angeletos and Pavan (2010) where those solving a signal extraction problem are the agents (not the government). The whole controversy about whether asymmetric information RE equilibria could be reached as discussed in Townsend (1983) is also within this framework.

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

where

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

where c_t^i and l_t^i for $t = 1, 2$ are consumption and hours worked respectively, with $u' > 0$, $u'' < 0$, $v' > 0$, $v'' > 0$.

γ is a random variable with distribution F_γ, we refer to it as a "demand shock". When γ is high, agents like first period consumption relatively more than other goods. At the same time high γ makes them willing to work more in their intratemporal labor-consumption decision and also more impatient in their intertemporal allocation of consumption. Viceversa, for low γ. Given that agents are identical, in the following we drop the subscripts i for notational convenience.

The production function in each period is linear in labor and output is given by

$$y_t = \theta_t l_t \; for \; t = 1, 2. \tag{2}$$

The random variable $\theta_1 = \theta$ has distribution F_θ and we will refer to it as the "productivity shock". As far as θ_2 is concerned, we will distinguish two cases, one where $\theta_2 = \theta_1$, in which case the productivity shock is permanent, and one where $\theta_2 = E\theta$, that is, second period productivity is a known constant, equal to the mean of the first period shock, in which case the productivity shock is temporary. Notice that in both cases θ_2 is known given θ_1.

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. γ and θ are assumed to be independent. The distributions F_γ and F_θ represent the government's perceived distribution of the exogenous shocks.

We will consider agents that have rational expectations. To this end, 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 $(\widetilde{\tau}_1, \widetilde{\tau}_2, \widetilde{b}^g) : \Phi \to R^3$ and these are actually the equilibrium values of taxes and government bonds for each A.

Consumers' choices and prices are contingent on the state (γ, θ) observed in period $t = 1$. Agents choose $(c_1, c_2, l_1, l_2, b) : \Phi \to R^5$ knowing the fiscal policy and the bond price function $q : \Phi \to R$. Obviously, the solution of the agents' problem in this setup coincides with the non-stochastic model where γ, θ are 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, so that the period budget constraints of the representative agent are

$$\begin{aligned} c_1 + qb &= \theta l_1 (1 - \widetilde{\tau}_1) && (3) \\ c_2 &= \theta_2 l_2 (1 - \widetilde{\tau}_2) + b && (4) \end{aligned}$$

where q is the price of the governemnt discount bond b. The above budget constraints have to hold for all realizations γ, θ we leave this implicit.

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 $(\widetilde{\tau}_1, \widetilde{\tau}_2, \widetilde{b}^g)$, price q and allocations (c_1, c_2, l_1, l_2, b) such that when agents take $(\widetilde{\tau}_1, \widetilde{\tau}_2, q)$ as given the allocations maximize the agents' utility (1) subject to (3) and (4). In addition, bonds and goods markets clear, so that $b^g = b$ and*

$$c_t + g_t = \theta l_t \; for \; t = 1, 2 \tag{5}$$

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

$$\frac{v'(l_1)}{u'(c_1)} = \theta\gamma(1 - \widetilde{\tau}_1) \tag{6}$$

$$\frac{v'(l_2)}{u'(c_2)} = \theta_2(1 - \widetilde{\tau}_2) \tag{7}$$

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

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 the consumer.

We first give the standard definition when both government and consumers observe the realization of A.[5]

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

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 Partial Information. 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 depends potentially on aggregate outcomes observed in period 1, $s = G(c_1, l_1, q, A)$ for a given G.

Definition 3 *A* ***Partial Information (PI-) Ramsey equilibrium*** *when government observes a signal* s *is a FI-Ramsey equilibrium satisfying*

1. τ_1 *is measurable with respect to* s
2. *fiscal policy* (τ_1, τ_2, b^g) *achieves the highest utility from among all equilibria satisfying 1.*

Restriction 1 can be expressed as the PI-Ramsey equilibrium having to satisfy

$$\tau_1 = \mathcal{R}(s) \ forall A \in \Phi, \ forsome \mathcal{R} : R \rightarrow R \tag{9}$$

We are interested in the case when this restriction prevents the PI-Ramsey equilibrium from achieving the FI version.

Note that consumers may or may not know that (9) holds, in any case they take as given the tax rate that arises from this equation and equilibrium. Even if they knew (9), they would not be able to exploit this knowledge in their optimization problem, because they are atomistic and cannot affect the aggregate signal, and hence the tax rate.[6] We assume that s is an aggregate variable so that as long as consumers are infinitesimal they take the tax level τ as given. 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 $\mathcal{R}$.

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 resource constraints (5), the following implementability condition holds

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

which is derived by combining the representative agent's budget constraints (3) and (4) with his FOCs (6), (7) and (8). The standard approach to find the Ramsey

[6] This differs from the situation in the New Public Finance where consumers optimize given the policy function $\mathcal{R}$ which is a function of individual choices.

policy under FI is to maximize (1) subject to (10). We now slightly deviate from this traditional approach in order to obtain a formulation of the FI problem that is as close as possible to the PI problem.

Implicit in the standard definition of FI Ramsey equilibrium is the assumption that the government knows how the economy reacts to a given tax policy given each value of A. We find it convenient to write out this reaction function explicitly since this reaction function is the natural way to write the problem under Partial Information. Using the first period resource constraint (5) for $t = 1$ to substitute out consumption in the first period consumption-leisure margin (6), we get

$$\frac{v'(l_1)}{u'(\theta l_1 - g_1)} = \theta\gamma(1 - \tau_1), \tag{11}$$

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

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

where we have again suppressed the time subscript from first period labor and tax rate. This shows how the signal l reacts to the tax choice.

Now we write the equilibrium utility function as a function of first period equilibrium allocations and shocks only. Using the resource constraints (5) to substitute out c_t in (10) gives one equation that, for each A, involves only the unknowns l_1, l_2. This defines implicitly a function that maps an equilibrium l_1 into a second period equilibrium labor l_2, call this map $L_2^{imp} : \Phi \times [0,1] \to R$,[7] so that

$$L_2^{imp}(A, l_1)\ for all A \in \Phi \tag{13}$$

solves (10). The welfare of the planner for each A can be written as

$$\begin{aligned} W(l; A) &\equiv U(\theta l - g_1, l, \theta_2 L_2^{imp}(l, A) - g_2, L_2^{imp}(l, A); \gamma) \\ &= U(c_1, l_1, c_2, l_2; \gamma) \end{aligned} \tag{14}$$

Note that the only arguments in W are the observed variable l and the shock γ. The functions U and L_2^{imp} are embedded in the definition of competitive equilibrium and known by the government.

Given the above discussion, the FI Ramsey Equilibrium reduces to solving

$$\begin{aligned} &\max_{\tau:\Phi\to R} E[W(l; A)] \\ &s.t. (12) \end{aligned} \tag{15}$$

[7]Under some specific assumptions on u and v, it will be possible to solve for l_2 as a function of l_1 in closed form. In general, the marginal effect of l_1 on l_2 is easily found by applying the implicit function theorem to (10). See our examples below.

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

The standard approach in Ramsey policy analysis is to not even write the constraint (12) in (15) as it is simply used to substitute out taxes which the government cannot do under PI since A is not observed.

PI- Ramsey equilibria

We focus on the case when the signal is just labor so $s = l_1 = l$ (in the robustness section we consider the case when output $y = \theta l_1$ is observed). The only difference under PI is that the additional constraint (9) appears and that the choice is over a tax contingent on the signal. Hence a **PI-Ramsey equilibrium** (given a signal l) solves

$$\begin{aligned} &\max_{\mathcal{R}:R\to R} E[W(l;A)] \\ &s.t.(12) \\ &\tau = \mathcal{R}(l) \end{aligned} \tag{16}$$

This gives rise to a non-standard maximization problem. We tackle this issue by application of calculus of variations in section 4.

3.4 The Economic Consequences of PI for taxation policy

Before giving a mathematical solution it is worthwhile discussing the economic issues raised by limited information 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

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

for $\alpha_c \leq 0$, $\alpha_l, B > 0$, tax smoothing will be perfect and Ramsey policy under FI involves setting a constant tax rate $\tau = \tau_1 = \tau_2$ that solves the government's intertemporal budget constraint

$$\tau\theta l_1 - g_1 + \beta\frac{u'(c_2)}{\gamma u'(c_1)}(\tau\theta_2 l_2 - g_2) = 0. \tag{17}$$

It is clear from (17) that the government needs to know the realization of both productivity and demand shock 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 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. Depending on the realizations, the government would like to set different tax rates and under some preference assumptions (e.g. log-quadratic, discussed in Section 5) it may even be the case that a certain increase in hours would call for a tax cut if driven by a high realization of γ, but would call for a tax hike if driven instead by low θ. Since the government does not observe these shocks, this makes this model a very interesting framework to study optimal policy with PI.

Clearly, under PI the choice of constant taxes 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 government will have to set τ_2 so as to balance the budget in the second period in order to avoid default, so to 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 decision, 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 Endogenous Signal Extraction (ESE)

Maximization problems such as (16) can be characterized as "optimal control under endogenous signal extraction". The key difficulty in (16) is that the policy choice $\mathcal{R}$ affects both the policy action τ as well as the distribution of the signal $s = l$. The signal provides information about the two unobserved exogenous shocks (γ, θ), so that the optimal policy depends on the conditional density $f_{(\gamma,\theta)|l}$, but this density depends itself on the tax policy $\mathcal{R}$. Therefore the choice of an optimal $\mathcal{R}$ must be consistent with the implied density $f_{A|s}$. To our knowledge this is the first paper to consider such a difficulty. The literature has considered setups where separation holds in the sense that $f_{A|s}$ does not depend on $\mathcal{R}$, see section 2 for a discussion. Here we provide a solution for a general setup.

Many problems in economics have this form. It can be thought of as a Stackelberg-reaction-function game where the "leader" (in section 3, the government) chooses its policy (or reaction) function $\mathcal{R}$ optimally given the reaction of the "follower" h, but h is given and independent of the choice on $\mathcal{R}$. Unlike standard Stackelberg games the actions τ and s are determined jointly in this setup, there is a hierarchy only in the way leader and follower choose the reaction functions $\mathcal{R}, h$. Simultaneity is the standard assumption in Ramsey equilibria, where the equilibrium allocations are influenced only by the policy action τ, not by the whole policy function. This is a natural assumption when, as is standard in Ramsey taxation, followers are atomistic and the signal is an aggregate variable.[8]

We now present a generic problem of optimal control under endogenous signal extraction. This generalizes the PI- Ramsey Problem without adding any difficulty to the proof. The generalization may be useful in other applications.

Consider a planner/government that chooses a policy variable $\tau \in \mathcal{T} \subset R$, observes endogenous signal $s \in \mathcal{S} \subset R$ at the time of choosing τ, when random variables $A \in \Phi \subset R^k$ have a given distribution F_A.[9]

The planner's objective is to maximize $E\left[W\left(\tau, s, A\right)\right]$ for a given payoff function W. The planner knows that a value for the policy variable τ maps into endogenous signals through the following equation

$$y = h(\tau, A) \tag{18}$$

The government is assumed to know W, h, F_A, it does not observe the value of A.

This nests the case when other endogenous variables enter the objective function, as these can be embedded in W. We did this in section 3 through the use of the L_2^{imp} function defined in (13) to substitute out l_2 in the objective function.

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

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

Our interest lies in solving for the case that the exact value of A can not be inferred from the observed s, the chosen τ, knowledge of h and of (18). This happens, generically, when the dimension of A is higher than the dimension of s ($k > 1$) and some of the variables in A have a continuous density. In terms of the fiscal policy example when the signal $s = l$ is observed, the values of θ, γ remain hidden even after the choice of τ has been made for a given observed labor l.

To summarize, we wish to solve the following model of **Optimal Control with Endogenous Signal Extraction:**

$$\begin{aligned} &\max_{\{\mathcal{R}:\mathcal{S}\to\mathcal{T}\}} E\left[W\left(\tau, s, A\right)\right] \\ &s.t. \qquad (18), (19) \end{aligned} \tag{20}$$

[8]Simultaneity also occurs in the literature on supply function equilibria (Klemperer and Meyer, 1989). Here firms simultaneously choose a supply function.

[9]It is possible to generalize the problem to the case of multidimensional policy instruments and signals. However, notation becomes cumbersome, hence we only refer to the univariate case in the following.

To see why a standard first order condition does not apply, we can rewrite the objective function as follows

$$\int E\left[W\left(\tau,\overline{s},A\right)\middle|\,\overline{s}\right]f_s(\overline{s})d\overline{s} \tag{21}$$

Taking derivatives with respect to τ and *if* $f_s(\overline{s})$ *could be taken as given* we would find the following optimality condition

$$E\left[W_\tau+W_s h_\tau\middle|\,\overline{s}\right]=0\ for\ all\ \overline{s}. \tag{22}$$

In most applications in dynamic models this would be correct, but it is not the correct FOC in our case because, in general, $\mathcal{R}$ determines the density f_s. To see this notice that since s is determined implicitly by

$$s=h(\mathcal{R}(s),A)$$

$f_{s|A}$ depends on $\mathcal{R}$. Therefore $f_s=\int_A f_{s|A}\ f_A$ is also endogenous to $\mathcal{R}$. The derivative of f_s with respect to $\mathcal{R}$ should be taken into account in deriving optimality conditions, as we do below.

4.1 General first order conditions

Let us call $S(\mathcal{R},A)$ the observable s induced by the shock A and a policy $\mathcal{R}$. Formally, $S(\mathcal{R},A)$ is defined as follows: define H as

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

then $H(s,A;\mathcal{R})=0$ gives the equilibrium value of s. We consider $\mathcal{R}$ such that $S(\mathcal{R},A)$ is uniquely defined for all $A\in\Phi$.

The policy variable that is realized for each value of the shocks A and for a given policy function $\mathcal{R}$ is then given by

$$T(\mathcal{R},A)=\mathcal{R}\left(S(\mathcal{R},A)\right)$$

Notice the following distinction between the objects S, T and $\mathcal{R}$: the latter is a function of s while S and T are functions of $\mathcal{R}$ and the realizations of the shocks.

Let $\mathcal{F}$ be the value of the objective function for a given choice for $\mathcal{R}$.[10]

$$\mathcal{F}\left(\mathcal{R}\right)\equiv E\left[W\left(T\left(\mathcal{R},A\right),S\left(\mathcal{R};A\right),A\right)\right] \tag{24}$$

We can now re-define the Optimal Control with Endogenous Signal Extraction problem as

$$\max_{\{\mathcal{R}:\mathcal{S}\to\mathcal{T}\}}\mathcal{F}(\mathcal{R}) \tag{25}$$

and denote its solution by $\mathcal{R}^*$.

[10]Notice that $\mathcal{F}$ maps the space of functions into R. The expectation operator integrates over realizations of A using the government's perceived distribution of A, so that the above objective function is mathematically well defined given the above definitions for T, S and under standard boundedness conditions.

4.2 Apparent Private Information: Invertibility

In some cases the government can still implement the FI policy even if it does not observe the shocks. This occurs whenever the information set of the government is invertible, allowing it to learn the true state of the economy A from observing the signal s.[11]

To formally define Invertibility, consider the set of all possible values (τ, s) in the FI case, namely

$$M^{*} \equiv \left\{ (\tau, s) \in R^2 : \tau = \mathcal{R}^{*}_{FI}\left(\overline{A}\right) \ and \ s = h\left(\mathcal{R}^{*}_{FI}\left(\overline{A}\right), \overline{A}\right) \ for\ some\ \overline{A} \in \Phi \right\} \quad (26)$$

Let M^{*}_{τ} (M^{*}_{y}) denote the projection of M^{*} on $\mathcal{T}$ $(\mathcal{S})$

Definition 4 *Invertibility holds if for any $\overline{s} \in M^{*}_{s}$ there exists a unique $\tau \in M^{*}_{\tau}$*

Clearly, Invertibility is satisfied if $h\left(\mathcal{R}^{FI*}(A), A\right) = \overline{s}$ defines implicitly a unique value for A for all $\overline{s}$.

Invertibility will often occur when the dimension of τ is the same as the dimension of A. Even if A is high-dimension, Invertibility also obtains if Φ is a finite set, in this case we can expect to be able to map an equilibrium into the shock since there are finitely many realizations, only by coincidence would the same equilibrium point (τ, l) occur for two different realizations of A.

Proposition 1 *Under Invertibility $\mathcal{R}^{*} = \mathcal{R}^{FI*}$*

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

To illustrate a case of Invertibility, consider 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.

4.3 General case of Limited information

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). Observe that

[11]In the literature of optimal contracts under private information and incentive compatibility constraints this is the standard assumption, which amounts to assuming full revelation, but those papers concentrate on the difficulties raised by the reaction of the agents' to the choice of $\mathcal{R}$, this reaction is the reason that the Full Information solution is not reached in those papers.

Invertibility also holds in the supply function literature (Klemperer and Meyer, 1989) because uncertainty is onedimensional.

Remark 1 *Invertibility is generically violated if*

- *A has higher dimension than s and some element of A has a continuous distribution or*
- *if $h\left(\mathcal{R}^{FI*}(\cdot), \cdot\right)$ is non-monotonic.*

In these cases the solution (20) should take into account that the distribution is endogenous to policy while, at the same time, the policy depends on the optimal filtering of the fundamental shocks A given the observed signal s.

Let $\mathcal{S}^*$ be the support of the random variable $S(\mathcal{R}^*, A)$. For the general case we will show that

Proposition 2 *Assume there is a solution $\mathcal{R}^*$ to the PI problem (20) and Assumptions 1-6 stated in Appendix A hold. Assume in addition that $S(\mathcal{R}^*; A)$ has a density. Let $\mathcal{S}^*$ be the support of $S(\mathcal{R}^*; A)$. The solution $\mathcal{R}^*$ satisfies the following necessary first order condition*

$$E\left(\left.\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}}\right| S(\mathcal{R}^*; A) = \overline{s}\right) = 0 \tag{27}$$

for $\overline{s} \in \mathcal{S}^*$ in a set of probability one.

To find the optimality condition in Proposition 2 we use a variational argument. In the following paragraphs we show some of the key steps in the proof. For a detailed proof see Appendix A.

Take any function (a variation) $\delta : \mathcal{S}^* \to R$ 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

$$\max_{\alpha \in \Re} \mathcal{F}(\mathcal{R}^* + \alpha\delta) \tag{28}$$

In other words, now we maximize over small deviations of the optimal reaction function in the direction determined by δ. It is clear that

$$0 \in \arg\max_{\alpha \in \Re} \mathcal{F}(\mathcal{R}^* + \alpha\delta) \tag{29}$$

In Appendix A we show that $\frac{\partial \mathcal{F}(\mathcal{R}^* + \alpha\delta)}{\partial \alpha}$ evaluated at $\alpha = 0$ is

$$E\left(\left(W_\tau^* + W_s^* h_\tau^*\right) \frac{\delta}{1 - h_\tau^* \mathcal{R}^{*\prime}}\right) = 0 \tag{30}$$

This can be derived by carefully writing down all the derivatives involved.

The general definition of conditional expectation implies (27).

Notice that, as we anticipated, the first order condition (27) does not coincide with the FOC when separation holds (22). The term $\frac{1}{1-h_\tau^* \mathcal{R}^{*\prime}}$ acts as a kernel, or as a measure change, it is the new term relative to the standard case when $f_{A|s}$ does not depend on $\mathcal{R}$. The term $\frac{1}{1-h_\tau^* \mathcal{R}^{*\prime}}$ captures the effect of the choice of $\mathcal{R}$ on the density f_s that, as we anticipated, had to appear in some way in the optimality conditions.

With Proposition 2 in hand we can derive previous results available in the literature on optimal policy under Partial Information as special cases. The following corollary shows some cases that had been considered in the literature and that gave rise to separation.

Corollary 1 *(separation) Assume for some $\overline{s} \in \mathcal{S}^*$ one of the following hold*

1. *(Invertibility) there is a unique $A \in \Phi$ such that $S(\mathcal{R}^*, A) = \overline{s}$.*
2. *(exogenous signals or linearity) h_τ is given constant for all $A \in \Phi$ such that $S(\mathcal{R}^*, A) = \overline{s}$.*

Then the general FOC (27) reduces to the FOC under separation (22).

The proof is trivial: in both cases $\frac{1}{1-h_\tau^* \mathcal{R}^{*\prime}}$ is known given s so that this terms goes out of the conditional expectation in (27) and it cancels out. Case 1 obtains a generalization of the Invertibility theorem we found before. Case 2 is useful to discuss the various approaches to optimal policy under Partial Information that have been discussed in the literature. Consider the case of exogenous signals, when s is a function of A only, in this case $h_\tau = 0$ so case 2 applies. The armed-bandit problems are also a special case of this, since it is only previously determined policies that influence the signal, so that with respect to the current signal we still have $h_\tau = 0$. Note that the linear models that have been analyzed in the literature (Svensson and Woodford, 2004) arise as a special case of (27) when h_τ is a known constant due to linearity.

The condition (27) is useful for comparability with the standard FOC in the cases of corollary 1, but it turns out to be less convenient for computations. The reason is that (27) contains the derivative of the policy function to be computed $\mathcal{R}^{*\prime}$. An algorithm trying to approximate $\mathcal{R}^*$ numerically will have to ensure that not only $\mathcal{R}^*$ is well approximated but that its derivative is well approximated as well along the iterations. For this purpose it is more convenient to use an envelope condition that delivers the optimality condition that will be stated in Proposition 3.

Note that (27) conditions on s. If in addition we condition on some additional variables in A the remaining variables can be taken as given. More precisely, let us partition $A = (A_1, A_2)$ where $A_1 \in R$ and assume for all possible values of $\left(\overline{y}, \overline{A}_2\right)$ one can back out a unique value of A_1 that is compatible with $\left(\overline{s}, \overline{A}_2\right)$, so that

$$H\left(\overline{s}, A_1, \overline{A}_2; \mathcal{R}^*\right) = 0 \tag{31}$$

holds. This equation defines an implicit function $A_1 = \mathcal{A}^*(s, A_2)$ that, given the information on s, maps a realization of A_2 into the corresponding A_1 for the optimal policy.

Proposition 3 *The optimality condition (27) is equivalent to*

$$\int_{\Theta_2(\overline{s},\mathcal{R})} \frac{W_\tau^* + W_s^* h_\tau^*}{h_{A_1}^*} f_{A_1}(\mathcal{A}^*(\overline{s}, A_2)) f_{A_2}(A_2) dA_2 = 0 \; for\ all\ \overline{s} \tag{32}$$

where $\Theta_2(\overline{s}, \mathcal{R}^*)$ is the support of A_2's conditional on observing $\overline{s}$, stars denote that the partial derivatives and A_1 are evaluated at $\mathcal{A}^*\left(\overline{s}, \overline{A}_2\right)$.

This condition is much easier to use for computations than (27) because it does not involve the derivative of $\mathcal{R}^*$. The only cost is that it has to assume that the function $\mathcal{A}^*$ is uniquely defined. This will not only simplify computations but it also honors the title of the paper, as its derivation involves explicitly $f_{A|s}$ which is precisely the filter that is determined endogenously by the optimal choice of policy.

We now give the main steps of the proof. Using the notation $f(x) + g(x) = (f + g)(x)$ note that

$$E\left(\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}} \middle| S(\mathcal{R}^*, A) = \overline{s}, A_2 = \overline{A}_2 \right) = \frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}} (\overline{\tau}^*, \overline{s}, \mathcal{A}^*\left(\overline{s}, \overline{A}_2\right), \overline{A}_2).$$

By the law of iterated expectations (27) implies

$$E\left(\frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}} (\overline{\tau}^*, \overline{s}, \mathcal{A}^*(\overline{s}, A_2), A_2) \middle| S(\mathcal{R}^*, A) = \overline{s} \right) = 0 \; for\ any\ \overline{s} \in \mathcal{S}^*. \tag{33}$$

so that we can integrate over A_2's only as long as A_1 is substituted by $\mathcal{A}^*(s, A_2)$ whenever these shocks appear as arguments of W_τ^*, W_s^*, h_τ^* and f_{A_1}. Therefore this can be rewritten as

$$\int_{\Theta_2(\overline{s},\mathcal{R}^*)} \frac{W_\tau^* + W_s^* h_\tau^*}{1 - h_\tau^* \mathcal{R}^{*\prime}} (\overline{\tau}^*, \overline{s}, \mathcal{A}^*(\overline{s}, A_2), A_2) f_{A_1}(\mathcal{A}^*(\overline{s}, A_2)) f_{A_2|\overline{s}} dA_2 = 0 \tag{34}$$

To find $f_{A_2|\overline{s}}$ proceed as follows. First we find $f_{s|A_2}$. Notice that given $A_2 = \overline{A}_2$ we have that $s = \mathcal{A}^{*-1}\left(A_1, \overline{A}_2\right)$ where $\mathcal{A}^{*-1}$ is the inverse function of $\mathcal{A}^*(\cdot, \overline{A}_2)$. By the change of variable rule

$$f_{s|A_2}\left(\overline{s}, \overline{A}_2\right) = f_{A_1}\left(\mathcal{A}^*(\overline{s}, \overline{A}_2)\right) \left|\mathcal{A}_s^*(\overline{s}, \overline{A}_2)\right| \tag{35}$$

where $\left|\mathcal{A}_s^*(\overline{s}, \overline{A}_2)\right|$ is the jacobian of $\mathcal{A}^*(\cdot, \overline{A}_2)$. To find this jacobian we apply once again the implicit function theorem to H and get

$$\mathcal{A}_s^*\left(\overline{s}, \overline{A}_2\right) = \left(h_{A_1}^*\right)^{-1} (1 - h_\tau^* \mathcal{R}^{*\prime}) \tag{36}$$

Plugging (36) into (35) gives $f_{s|A_2}$. A standard application of Bayes' rule gives $f_{A_2|\overline{s}}$ in terms of $f_{s|A_2}$ namely

$$f_{A_2|s} = \frac{f_{s|A_2} f_{A_2}}{\int f_{s|\overline{A}_2} f_{\overline{A}_2} d\overline{A}_2} \tag{37}$$

Plugging the above formula for $f_{A_2|\overline{s}}$ in (27) and since the denominator $f_s(s)$ drops out in the first order condition gives (32).

This proof highlights that $f_{A|\overline{s}}$ depends on $\mathcal{R}^{*\prime}$as it appears in the Jacobian in (36) and it also determines $\mathcal{A}^*$.Therefore we have that $f_{A|\overline{y}}$depends on the policy function $\mathcal{R}^*$and Bayes' rule implies that the signal extraction $f_{A_2|s}$ depends on the optimal policy $\mathcal{R}^*$. This justifies that we dub the maximization problem (20) and the title of the paper "endogenous filtering".

4.4 Linearity and "certainty equivalence"

We now show that if we consider a linear model (quadratic objective function and linear reaction function), we obtain the "certainty equivalence" result of Svensson and Woodford (2004).

Assume the objective function is

$$W(\tau, l) = -\frac{\omega_\tau}{2}\tau^2 - \frac{\omega_l l^2}{2} \tag{38}$$

and the reaction function

$$l = h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma \tag{39}$$

where the ω's are positive coefficients and the h's are any real numbers.

Using Svensson and Woddford's (2004) notation, let $X \equiv (1, \theta, \gamma)'$. Then the Full Information optimal policy is

$$\tau = FX$$

where

$$F = \left(-\frac{\omega_l h_\tau h}{\omega_\tau + \omega_l h_\tau^2}, -\frac{\omega_l h_\tau h_\theta}{\omega_\tau + \omega_l h_\tau^2}, -\frac{\omega_l h_\tau h_\gamma}{\omega_\tau + \omega_l h_\tau^2} \right).$$

The FOC under PI (27) becomes

$$E\left[\omega_\tau \tau + \omega_l h_\tau \left(h + h_\tau \tau + h_\theta \theta + h_\gamma \gamma\right) |l\right] = 0 \tag{40}$$

which can be rewritten as

$$\left(\omega_\tau + \omega_l h_\tau^2\right) \tau \omega_l h_\tau h + \omega_l h_\tau h_\theta E\left[\theta | l\right] + \omega_l h_\tau h_\gamma E\left[\gamma | l\right] = 0 \tag{41}$$

which implies that optimal policy is given by

$$\tau = FE\left[X | l\right] \tag{42}$$

where F is the same vector of coefficients found under Full Information, that is, independently of the information structure and distribution of the shocks. This property of optimal policy with PI in linear models is called "certainty equivalence", as the government forms the best estimate of the state and behaves as if this estimate was certainty, or Full Information.

However, the "separation principle" does not hold, because one cannot compute the expectation of the state $E[X|l]$ without knowledge of the policy. To see this, consider for example $E[\theta|l]$, which can be rewritten as

$$E[\theta|l] = E[X|h + h_\tau F E[X|l] + h_\theta \theta + h_\gamma \gamma = l] \tag{43}$$

This shows that optimal policy (42) and signal extraction (43) have to be solved jointly as a system.

4.5 Set of possible values and Transversality Condition

Let us know go back to the general non-linear case. So far we have ignored two issues that often complicate obtaining solutions to dynamic stochastic problems, namely, the ex-ante restriction of the set of possible values of choice variables and the transversality condition.

In many models of optimal choice under uncertainty it is important to know the range of possible values that endogenous variables can take. One reason this is important is that if the true solution is non-linear and approximations to non-linear functions can only be accurate on a compact set, one needs to know, at least roughly, what are the likely values of the equilibrium solution. A second reason is that it is possible that the solution may be simply undefined outside a certain range of the endogenous variable so that the algorithm may break down if it attempts to compute the solution outside that range.[12]

In our case and, in particular, in the two-dimensional uncertainty case posed by the fiscal policy example, the limits of l are easy to find ex-ante by exploiting the fact that the extreme values for l coincide with the FI case. More precisely, assume that l is monotonic in both γ and θ. In particular, assume that l is increasing in γ and decreasing in θ.[13] It is clear that, letting $l_{\min}$ and $l_{\max}$ be the extreme values of l in the PI solution, we have

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

[12] In fact, for applications of the promised utility approach to solving incentive problems, one needs to know the set of possible values to even be able to formulate a consistent recursive problem. This problem does not arise in our case, where the set of possible τ's can be restricted to a large set in the formalization of the maximization problem.

[13] This will be the case in the log-quadratic case shown in Section 5.2. Such monotonicity can be easily proved in the FI case, although it should be checked with a candidate solution for $\mathcal{R}^*$.

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

Another issue is that in most infinite-dimensional optimization problems the first order conditions derived by applying δ variations to the policy function do not determine a unique policy function, one needs to add another condition that holds as a terminal or initial condition, often referred to as a "transversality condition". Since $\mathcal{R}^{*\prime}$ appears in (27) it appears that a transversality condition may be needed as the solution in a given value of s may depend on the solution for other values.

But there are two reasons why this is not a problem in our case. First, because there are two natural end conditions: using the above discussion about extreme values of l it is trivial to establish that at the extremes the optimal tax is given by the FI solution so that

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

give terminal and initial condition for $\mathcal{R}^*$ that are trivially computed.

Second, because when we rewrite the optimality condition as (32) the derivative $\mathcal{R}^{*\prime}$ does not appear so the optimal value for each s can be found independently of the solution in other points.

4.6 Algorithm

Given (32) it is easy to calculate the PI solution using the following numerical algorithm. Fix a value for $\overline{s}$. We must be able to compute $\mathcal{A}^*(\overline{s}, A_2)$ for a given candidate of $\overline{\tau} = \mathcal{R}^*(\overline{s})$. Then at a possible candidate $\overline{\tau}$ we can evaluate the integrand in (32) for each possible A_2 since we have the corresponding $A_1 = \mathcal{A}^*(\overline{s}, A_2)$. We compute the integral by running A_2 from the lowest to highest possible value of A_2 given the candidate $\overline{\tau}$. Note that these limits to the possible values of A_2, defining $\Theta_2(\overline{s}, \mathcal{R}^*)$, are endogenous to $\overline{\tau}$ and they have to keep A_1 within the admissible limits of the support of A_1. This operation maps a value of a candidate $\overline{\tau}$ to the left side of (32), the optimal $\tau = \mathcal{R}^*(\overline{s})$ is found by solving this non-linear equation that makes this integral as close as possible to zero.

5 Solution of Fiscal Policy Example with ESE

5.1 Log-quadratic utility

We now study the optimal fiscal policy model introduced above using the tools developed in the previous section. For simplicity we study a special case that allows for an analytical reaction function h. Assume $u(c) = log(c)$ and $v(l) = \frac{B}{2} l^2$. Both the productivity and the demand shock are temporary. In the second period, θ_2 is known to be equal to the mean of θ.

We show numerical examples with the following parameter values and distributional assumptions: let θ be uniformly distributed on a support $[\theta_{min}, \theta_{max}]$, γ uniformly distributed on $[\gamma_{min}, \gamma_{max}]$ and assume $\beta = .96$, B calibrated to get average hours equal to a third and government expenditure constant and equal to 25% of average output. The mean of γ is 1 and the mean of θ is 3. 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. Notice that the equilibrium condition (6) becomes

$$Bl_1 c_1 = \gamma\theta_1(1 - \tau_1) \tag{44}$$

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

$$l = h(\tau, \theta, \gamma) = \frac{Bg_1\theta^{-1} + \sqrt{(Bg_1)^2\theta^{-2} + 4B\gamma(1-\tau)}}{2B}. \tag{45}$$

It is important to note that in general the productivity shock θ has two opposing effects: the substitution effect between leisure and consumption and the wealth effect, that acts in the opposite direction. With log-quadratic preferences, the second effect dominates and hence high realizations of θ will lead to low labor, ceteris paribus. On the other hand it can be seen from equation (45) that hours are increasing in the demand shock γ.

In Figure 1 we illustrate 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 $\bar{l}$ and imposing a tax rate $\bar{\tau}$ the government cannot infer the value of the shocks.

Under PI, for an intermediate value of $\bar{l}$ the government is uncertain whether say both θ and γ are high, or viceversa, and in general there is a continuum of realizations $\left(\bar{\theta}, \widetilde{\gamma}(\bar{\theta}, \bar{l}; \mathcal{R})\right)$ consistent with the observation of $\bar{l}$ and a policy $\mathcal{R}$. Therefore it cannot choose the policy under Full Information (constant taxes) since the realizations of γ and θ enter separately in (17).

The partial derivatives h_τ and h_γ are easily obtained analytically. In particular

$$h_\tau(\tau, \theta, \gamma) = \frac{-1}{\sqrt{(Bg_1)^2\left(\theta\gamma\right)^{-2} + \gamma^{-1}4B(1-\tau)}}. \tag{46}$$

Figure 1: Hours and taxes with Full Information

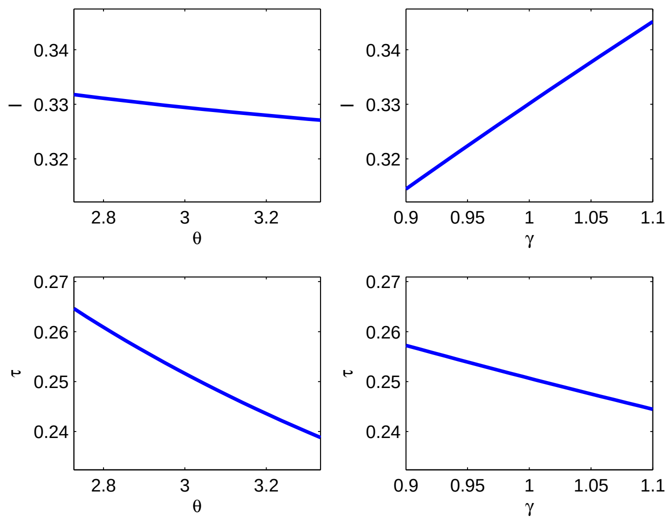

It is clear that both the productivity shock and the demand shock affect this slope, therefore endogenous signal extraction is an issue. Hence we proceed to find a $\mathcal{R}^*$ that satisfies (32) using the algorithm described in Subsection 4.6

Figure 2 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 Full Information.

As explained in Subsection 4.5, the lowest (and highest) labor that is realized under FI is also the lowest (highest) value that can occur under PI and the PI tax is the same as the FI tax. In these extremes 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 (but it is sometimes far from being the average of those tax rates.

For low labor, 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 $l_{\max} = L^{FI}(\theta_{min}, \gamma_{max})$.

Figure 2: Optimal policy with log-quadratic utility

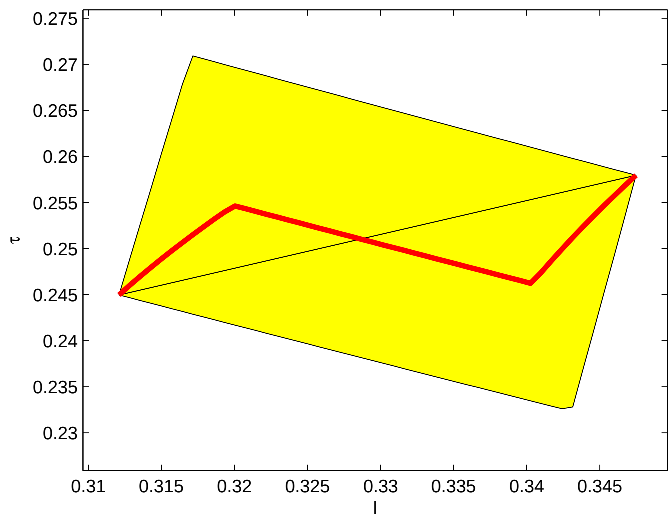

To gain further understanding on the implications of PI for the properties of the model, we plot again hours and taxes as functions of each shock individually in Figure 3. In all four panels, we reproduce the FI outcomes shown in Figure (FI) (blue dashed-dotted lines). The red lines represent the PI outcomes. For instance in the left panels we keep γ equal to its mean and we plot hours and taxes as functions of θ. Of course the government does not observe the values of θ and γ, but only hours. Interestingly, it can be seen that hours become more volatile in response to productivity shocks under PI and 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.[14] 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 fooled 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 realization of shocks for $\bar{l} = .33$ in Figure 4.

[14] We will see in the next subsection that this property of the solution will change with higher government expenditure.

Figure 3: Hours and taxes with Partial Information and Full Information

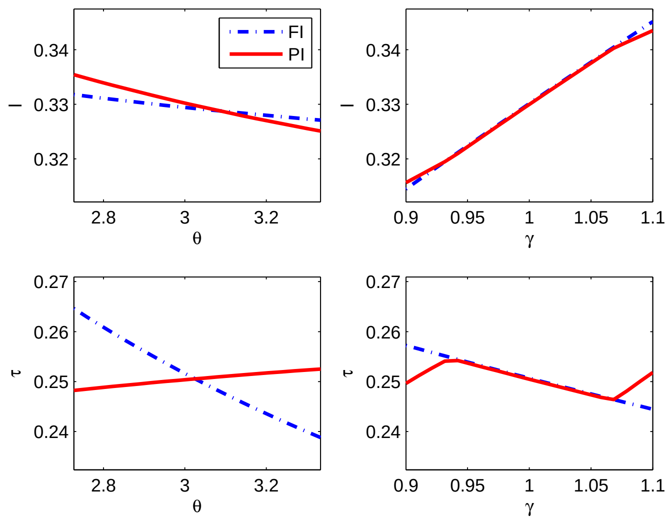

The wealth effect of productivity makes it an increasing function in the (θ, γ) space. Now conditional on l, we can have combinations of high productivity (low wealth effect on labor supply) and high demand or low productivity and low demand.

Figure 4: Set of admissible shocks

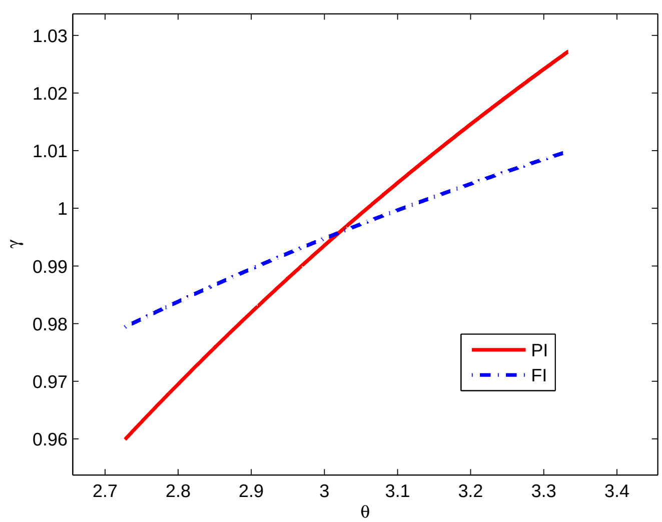

Optimal policy with PI calls for a substantial smoothing of taxes across states. This can be seen in Figure 5, where the equilibrium cumulative distribution function of tax rates under PI (red line) is contrasted with the one obtained under FI (blue dotted line). This result is rather intuitive and it 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.

Figure 5: Equilibrium CDF of tax rates

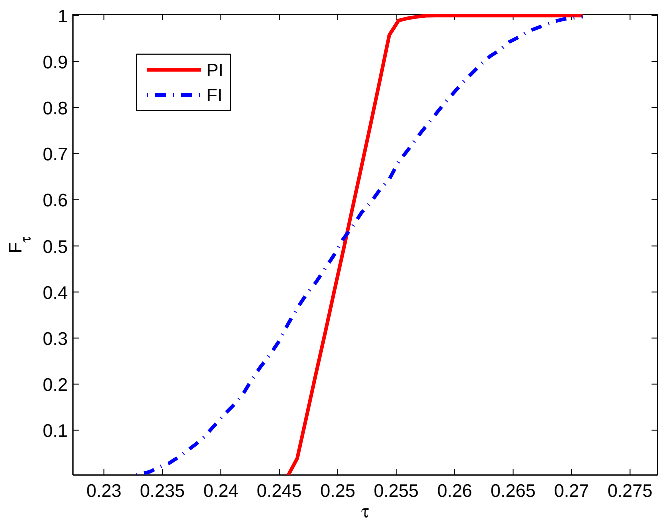

In our model, this smoothing of taxes across states will imply a larger variance of tax rates in the second period with respect to the FI policy. In the second period, all the uncertainty is resolved and the tax rate will be whatever is needed to balance the budget constraint. This is of course taken into account at the time of choosing a policy under uncertainty, so that we could say that optimal policy is very prudent while the source of the observed aggregate variables is not known and then responsive after uncertainty has been resolved. In this sense, this model can rationalize the slow reaction of some governments to big shocks like the current recession. The Spanish example in the latest recession is a case in point. In 2008, it was far from clear how persistent the downturn would be and also whether is was demand-driven or productivity-driven and the government did not adjust its fiscal stance quickly, only to make large adjustments in the subsequent years. We will discuss this further in the paper.

5.2 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.[15] We will see that this leads to a very non-linear optimal policy and to an exception to tax-smoothing across states. This example is of interest for several reasons. From an economic point of view, private information is of more importance here: since the government needs to balance the budget in the second period it is thus now very 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 in this case. 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 and hence consumption is very low. This example will also be of interest because the PI solution has some very different features from the FI outcome.

Figure 6 shows optimal policy for this case (red line), again contrasted with the set of tax-labor outcomes under FI (yellow region).

Figure 6: Optimal policy with high government expenditure

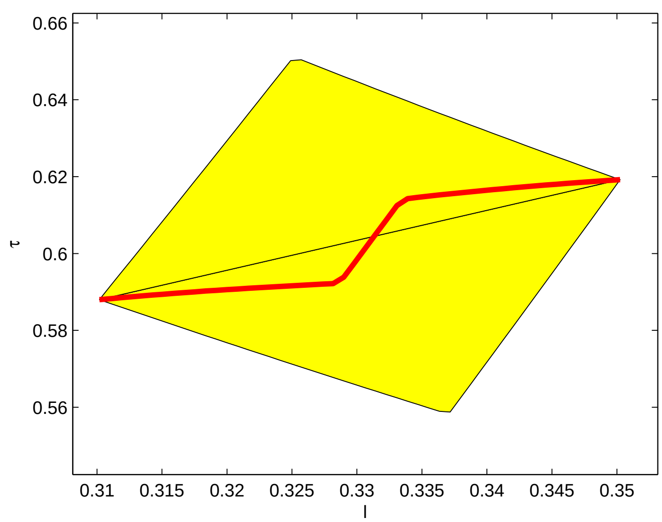

The figure shows that the optimal solution is highly non-linear. The derivative $\mathcal{R}^{*\prime}$ is positive and relatively high in a middle range of levels of l, but both to the left and to the right of this middle range $\mathcal{R}^*$ it is much flatter. Notice that this is the opposite of what happens with a low level of g in the previous subsection. When government expenditure is sufficiently low, the government is very uncertain about

[15] All other assumptions on preferences and shocks are the same as in the previous section.

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.

To illustrate how the PI policy involves a relatively precise signal extraction on θ with high government expenditure, consider the sets of possible realizations of θ under FI in Figure 7 and under PI in Figure 8.

Figure 7: Set of admissible θ's with FI

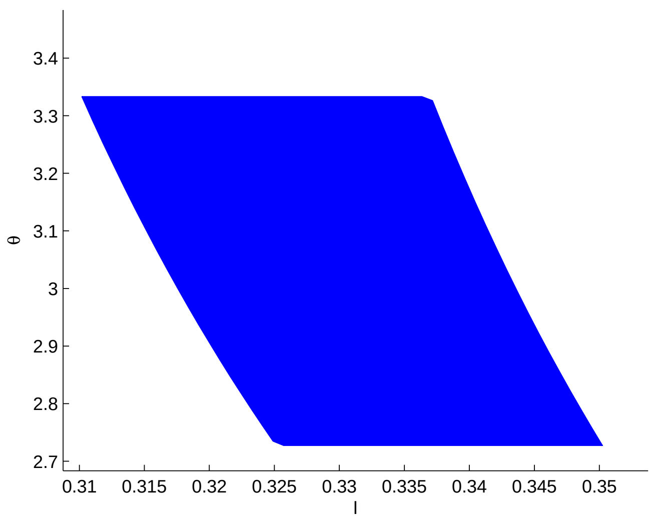

Consider Figure 7 first. It can be seen that under FI any realization of θ is consistent with an intermediate realization of l, but each of these θ's would call for a different tax rate. However, under PI, there can only be one tax rate for each observed l and the government uses this policy to extract information on θ.

To see this, consider now Figure 8. The minimum value of l is only consistent with the highest possible θ (and lowest possible γ) because the wealth effect dominates. Under PI, increasing l from this point, the government becomes uncertain and lower realizations of productivity become consistent with the observations. At first, uncertainty is rising with l, but in the intermediate region of l's the government becomes more and more confident about low realizations of productivity. This leads to the sharp increase in the tax rate, which in turn gives rise to feedback effect on the set of possible θ's: high taxes discourage work effort, so higher labor now is an even stronger signal of low θ (high marginal utility from consumption). In this way an optimal policy and a conditional distribution of shocks consistent with it confirm each other in equilibrium.

Figure 8: Set of admissible θ's with PI

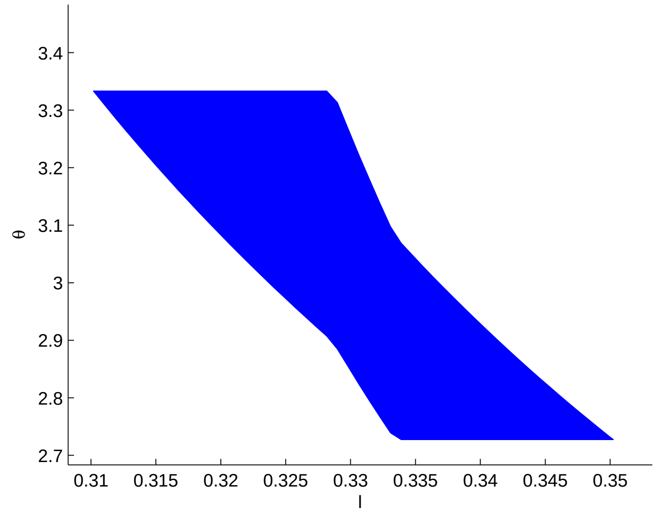

Consistently with this analysis of the signal extraction, we also plot hours and taxes as functions of each shock individually, and we contrast the PI outcomes with the FI solution in Figure 9. On the left-hand side we consider productivity shocks only. As illustrated above, in the intermediate region 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 subsection, 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. 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 afterwards, especially in countries like Spain and Italy, that arguably where in danger of getting close to the top of the Laffer curve, as testified

Figure 9: Hours and taxes with PI and FI: high g

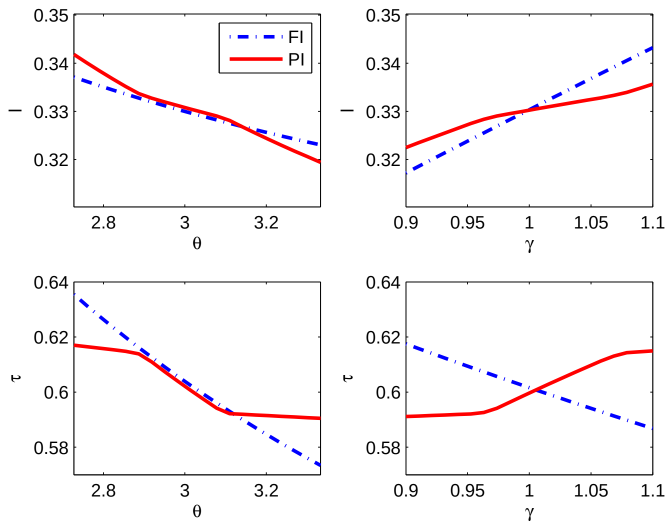

by the fact that significant increases in taxes after 2009 did not raise the amount of revenue as much as it was desired by these governments.

5.3 Linear 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 Full Information when the shocks take their mean value. Then, we compute the certainty equivalent policy as described in subsection 4.4. Importantly, this policy can be found under Full Information and then applied to the Partial Information case by simply computing the conditional mean of the shocks for each value of $\bar{l}$. Figures 10 and 11 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 Partial Information and the economy is hit aby large shocks.

Figure 10: Linear approximation, low g

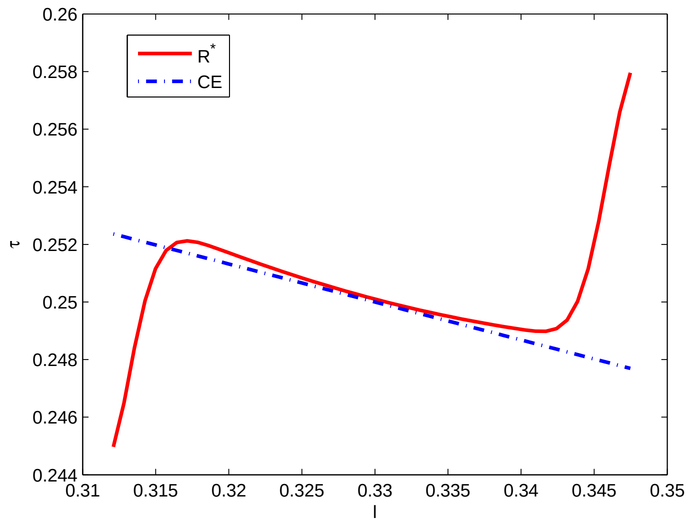

6 Robustness

6.1 Risk aversion and precautionary fiscal adjustments

As we have seen in the previous section, when the economy is close to the top of the Laffer curve, optimal policy is very non-linear in the observable variable, creating a region of sharp fiscal adjustments for intermediate realized values of labor. The government raises taxes dramatically in the first period, in order to prevent the worst scenarios with low productivity, high taxes and low consumption in the second period. In this section, we consider a general CRRA utility function from consumption

$$u(c) = \frac{c^{1+\alpha_c}}{1+\alpha_c}$$

and investigate how optimal policy changes with different degrees of the risk-aversion parameter α_c, while keeping quadratic disutility from labor effort as in our baseline parametrization.

Figure 12 illustrates the optimal policy for $\alpha_c = -1$ (log utility, that is the baseline case), $\alpha_c = -1.5$ and $\alpha_c = -2$. It can be seen that as risk aversion increases, the area where policy is more reactive of observables becomes wider, and the government reacts strongly even for weaker signals of a recession. Intuitively, this is because the government wants to avoid contingencies with high debt that would lead to high taxes in the second period. The more risk averse the agent is, the more painful it is to be in those states, where consumption has to be cut substantially. However, this larger

Figure 11: Linear approximation, high g

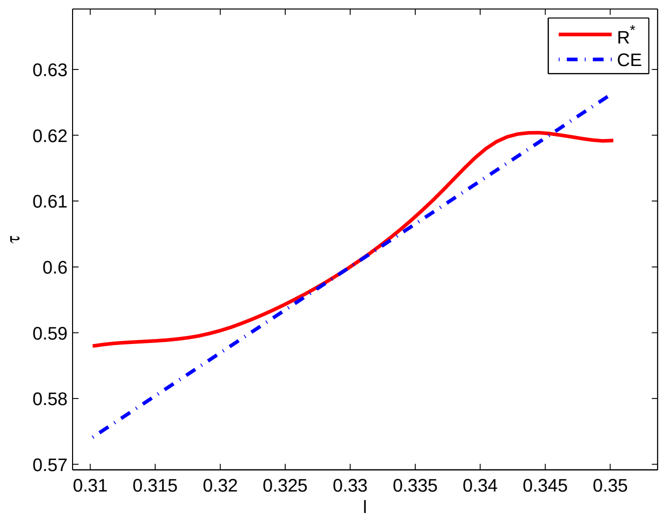

region of reaction also makes the policy function less steep, as can be seen from the picture.

6.2 Permanenent shocks and linear-quadratic utility

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

$$l = h(\tau, \theta, \gamma) = \frac{\gamma\theta}{B}(1 - \tau), \tag{47}$$

The optimal policy is illustrated in Figure 13 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. While the optimal policy is not quite linear, in this example a linear approximation would not be too wrong.

Figure 12: Changing risk aversion

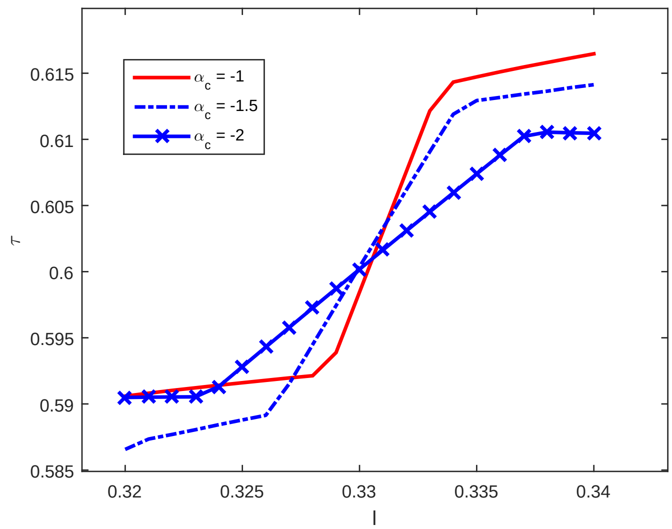

Figure 13: Optimal policy with linear-quadratic utility

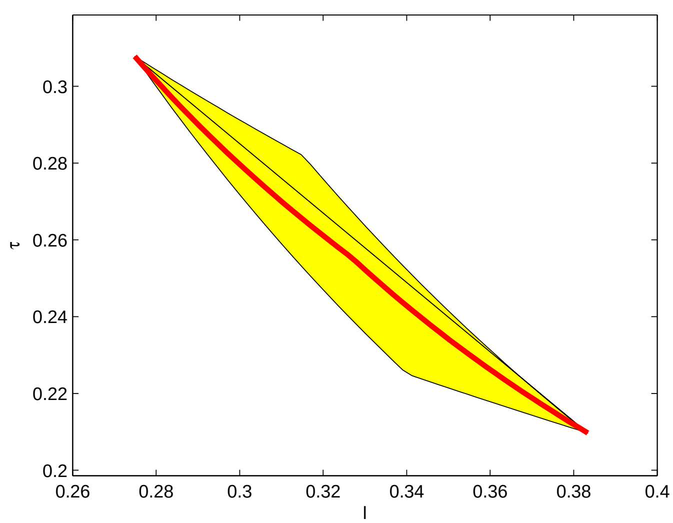

6.3 Observable output

Our previous example assumed that labor was observable. This made notation and presentation easier as the shocks are not involved in the signal directly. We now consider the case when the signal observed by the government is the total product, so $s = y = \theta l$. The government cannot sort out the values of the θ and l independently, hence there still is a problem of signal extraction as the government can not back out the values of (θ, γ). In addition we assume again linear-quadratic utility and a permanent productivity shock.

Figure 14 illustrates the optimal policy for this case. It can be seen that the result is remarkably similar to that obtained in the previous subsection with observable labor. However, in this case the government has a lot more information than in the previous case. This is because 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.[16] As we have seen, uncertainty about debt is key to get large fiscal adjustments as in section 5.2.

Figure 14: Optimal Policy with output observed

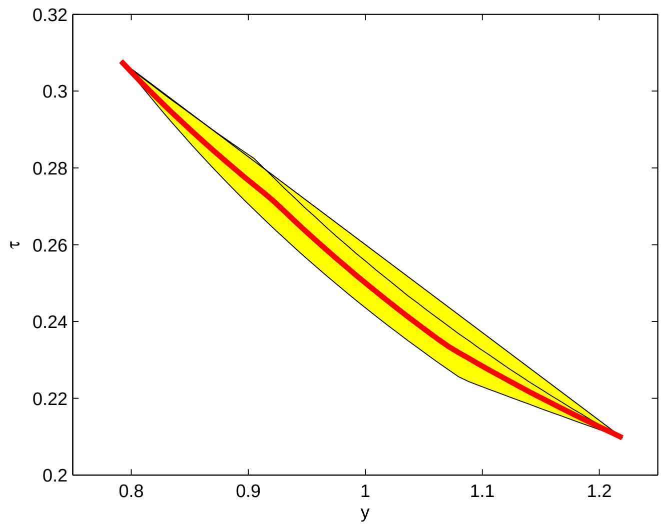

7 An infinite-horizon model with debt

In this section we present an infinite horizon version of the optimal fiscal policy model we have considered. We will see that some key intuitions developed in the two-period model are still present and they lead to interesting dynamics. In particular, under PI the government sometimes reacts slowly to recessions and as a consequence needs to raise taxes for a longer time endogenously prolonging slumps. In this case the slow reaction and the ensuing deepening of the recession is fully desirable.

[16]Note that if we had assumed, as in section 5.1, that θ_2 is not random the model would not be interesting, since in that case there would be no uncertainty about future revenue.

As is well known, in the case of non-linear utility current interest rate is determined by future taxes, hence the optimal policy under full commitment would be time inconsistent, leading to some complications in the solution of optimal policy under full commitment.[17] In order to avoid these difficulties we assume a linear utility of consumption.

We also assume that the shocks are i.i.d. over time as this reduces the number of state variables and it abstracts from the issues of government experimentation that has been studied in the armed-bandit literature of optimal policy that we discussed in section 2.[18] In this way we are left with the simplest infinite horizon model of fiscal policy where endogenous signal extraction plays a role.[19]

7.1 Full Information

Our model under Full Information 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:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\gamma_t c_t - v(l_t)\right] \tag{48}$$

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

The period t budget constraint of the representative agent is

$$c_t + q_t b_t = \theta_t l_t (1 - \tau_t) + b_{t-1} \tag{49}$$

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 maximisation are

$$\frac{v'(l_t)}{\gamma_t} = \theta_t (1 - \tau_t) \tag{50}$$

and

$$q_t = \beta \frac{\bar{\gamma}}{\gamma_t}. \tag{51}$$

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

[17] Under Full Information and uncertainty this issue was first addressed in Aiyagari et al (2002).

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

[19] Combining this with issues of commitment and optimal experimentation is of interest but we leave it for future work.

constraint $c_t + g = \theta_t l_t$. Under FI, the government can choose a sequence of taxes conditional on a sequence of shocks $A^t = (A_t, A_{t-1}, ..., A_0)$, where $A_t = (\theta_t, \gamma_t)$.

The period-t implementability constraint is

$$b_{t-1} = c_t - \frac{v'(l_t)}{\gamma_t} l_t + \beta \frac{\bar{\gamma}}{\gamma_t} b_t. \tag{52}$$

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 very high number in order to mimic a model with an occasionally binding borrowing constraint while still retaining differentiability of the problem.

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

$$\gamma_t \theta_t - v'(l_t) + \lambda_t \theta_t - \frac{\lambda_t}{\gamma_t} \left[v'(l_t) + v''(l_t) l_t \right] = 0 \tag{53}$$

and

$$\lambda_t \frac{\bar{\gamma}}{\gamma_t} = E_t \lambda_{t+1} + \chi (b_t - b^{max}) I_{[b^{\max}, \infty)}(b_t). \tag{54}$$

where λ_t is the Lagrange multiplier of constraint (49) 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 $W^{FI}(b_{t-1}, A_t)$. 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 W as the value of the utility (48) at the optimal choice for given 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 W satisfies the following Bellman equation

$$W(b) = \max_{\mathcal{R}\colon \Re^2 \to \Re_+} E\left[\gamma(\theta l - g) - v(l) + \beta W\left(\frac{(b + g - \theta l + \frac{v'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}}\right) + \right.$$

$$\left. - \beta\frac{\chi}{2}\left(\frac{(b + g - \theta l + \frac{v'(l)l}{\gamma})\gamma}{\beta\bar{\gamma}} - b^{max}\right)^2\right] \tag{55}$$

where l satisfies $l = h(\mathcal{R}(b, l), \theta, \gamma)$ where $h(\tau, \theta, \gamma)$ is the labor that satisfies (50).

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.

At this point it is important to pause the maths for a second and discuss 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 θ and, therefore, a low tax revenue, occurred with a delay. The reason that b is an argument in $\mathcal{R}$ is that in the presence of incomplete markets, debt piles up after a few bad shocks, much 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 (55) we have substituted future debt using the budget constraint (52). 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 (55).

For a given guess for the value function, this is just a function of observed labor to which we can apply the main theorem of the paper (Proposition 2) and obtain the general first order condition with PI.[20]

This first order condition involves the derivative $W'(b)$. 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 C we show that an

[20]The FOC is explicitly shown in Appendix C.

analogous envelope condition holds under our PI model so that

$$W'(b) = E\left[\frac{\gamma}{\bar{\gamma}} W'(b')\right] - \chi(b' - b^{max}) I_{[b^{\max},\infty)}(b'). \tag{56}$$

Hence by solving the first order condition using (56) and iterating on the Bellman equation (55), 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 can nonetheless rationalize important features of the fiscal response to the Great Recession, with slow and large fiscal adjustments inducing protracted slumps.

7.3 Numerical results

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

Figure 15 illustrates a key property of optimal policy by showing the tax policy as a function of outstanding debt, for two different realization of the shocks: in both panels we keep γ equal to its means, but while in the upper panel θ is also equal to the mean, in the lower panel θ is equal to its minimum value. We contrast optimal policy under PI with the FI counterpart (dashed-dotted line). In the upper panel, we can observe a non-linearity of optimal policy with respect to debt. When the economy gets closer to the debt limit, the PI policy calls for a significantly larger increase in taxes than the FI when the shocks are equal to their mean. This is because observed labor takes an intermediate value for this realization of the shocks. Hence, the government is uncertain about what combination of the shocks has been realized and needs to avoid exceeding the debt constraint for all possible realizations of θ's. On the other hand the FI government knows that a lower tax rate is sufficient to avoid a fiscal crisis for this combination of the shocks. In the lower panel, on the other hand, signal extraction under PI is more precise, as the economy is hit by a large negative productivity shock and both PI and FI call for similar tax responses, independently of the debt level.

We now show impulse response functions for our economy.[21] In particular, we engineer two different scenarios that give rise to a 1% fall in observed hours. With linear utility from consumption, hours are increasing in θ, differently from the two-period model. Hence a fall in hours can be driven by either low θ and high γ or viceversa. In the first case (Figure 16), the shocks hitting the economy are a negative θ shock combined with a (smaller) positive γ shock. In the second case, we consider the opposite combination of θ and γ.

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

Figure 15: Tax policy as a function of debt

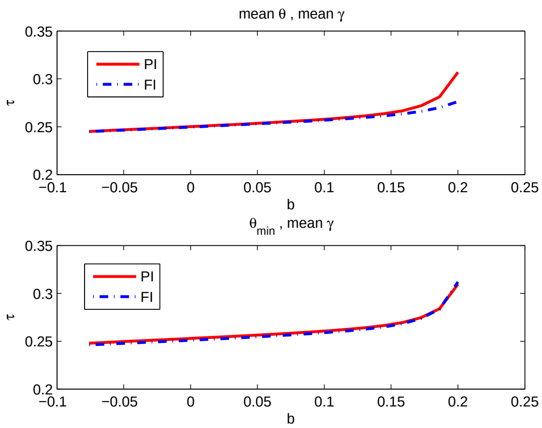

Figure 16: Impulse response function: low θ, high γ

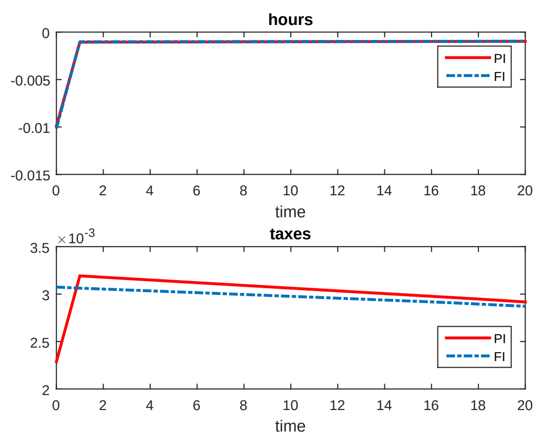

When the true nature of the recession is productivity, the tax response under PI is smaller then under FI. Hence, the government needs to respond by more in the second period, after the nature of the shock becomes known. Eventually taxes increase by more than under FI and stay high for a long time, inducing a (slightly) more persistent fall in hours and output. This behavior of the economy is qualitatively similar to what happened in some European countries after the financial crisis (e.g. Spain), where an initial slow reaction, or even an expansionary policy, has been followed by a necessary large fiscal adjustment and the recovery has so far been very slow and weak.

On the other hand, when the true nature of the fall in hours is a demand shock (Figure 17), the PI government reacts in the wrong direction, increasing taxes, while

Figure 17: Impulse response function: high θ, low γ

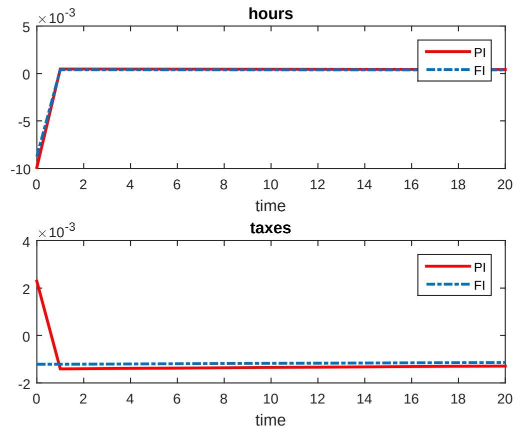

the FI government cuts them. This policy is reversed in the second period, after the past realization of the shocks become known.

While the PI policy is optimal in our setup where only current income can be taxed, the above findings suggest that allowing for retrospective taxation could improve welfare. Society would be better off if the government could adjust taxes on past income after observing the realization of past shocks and consumers knew of this possibility. However, retrospective taxation might not be easily implementable in the real world due to time-consistency issues, as ex-post surprising taxes on past income are non-distortionary.

Figure 18 illustrates a long stochastic simulation of the model. It is easy to see that taxes are very responsive to debt. One interesting question is whether taxes are smoother or more volatile under PI with respect to FI. Intuitively, there seem to be two opposing forces. On the one hand, the PI government does not observe the shocks, and hence smooths its policy across states for a given debt level. However, this policy induces necessary fiscal adjustments following the dynamics of debt, so that this pushes towards higher volatility under PI. The results from long simulations is that this second effect seems to dominate (although slightly) and the FI government is more successful than the PI government at smoothing tax rates. It can be seen that often when debt gets close to the borrowing limit (20% of mean output) the PI government imposes larger fiscal adjustments. This can be rationalized in analogy with the example of the two-period economy close to the top of the Laffer curve. Fear of future large required adjustments in the event of low θ lead the PI government to raises taxes significantly.

Figure 18: Simulation

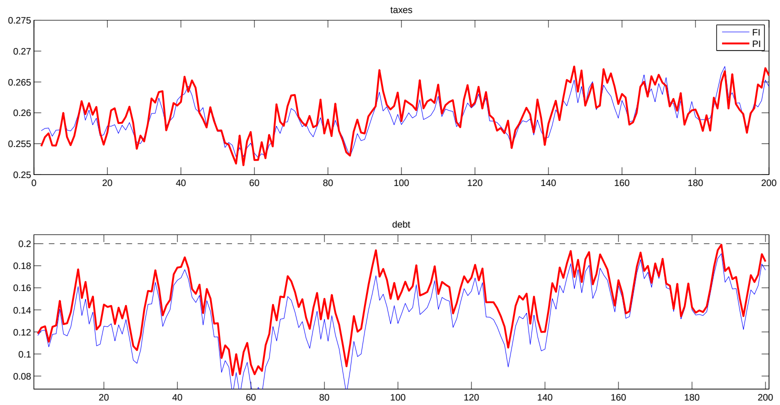

8 Conclusion

We derive a method to solve models of optimal policy with partial information without any separation assumption between the optimization and signal extraction problem. In our model the optimal decision influences the distribution of the shocks conditional on the observed endogenous signal, therefore the signal extraction and optimization problem need to be solved consistently and simultaneously. The method works in general and we show algorithms that solve these problems using standard ideas for solving dynamic models. 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 even if the model is not highly non-linear.

Optimal fiscal policy under endogenous 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 it becomes more consistent with a slump in productivity the government becomes certain that a lower productivity has occurred and this certainty accelerates.

Clearly, while we have illustrated the technique in a model of optimal fiscal policy, the methodology can be easily extended to other dynamic models, for example in the analysis of optimal monetary policy in sticky price models (e.g. Clarida et al. 1999) under the assumption of Partial Information. Our optimal policy smoothing result is likely to extend to that setup, potentially leading to a microfoundation for smooth nominal interest rates.

References

[1] 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.

[2] Angeletos, George-Marios and Alessandro Pavan (2009) ”Policy with Dispersed Information” *Journal of the European Economic Association* 7(1):11–60

[3] Aoki, K. (2003) “On the Optimal Monetary Policy Response to Noisy Indicators” *Journal of Monetary Economics*, Elsevier, vol. 50(3), pages 501-523, April.

[4] Baxter,Brad, Graham,Liam and Stephen Wright (2007) ”The Endogenous Kalman Filter”, Birkbeck School of Economics, Mathematics and Statistics working paper BWPEF 0719

[5] Baxter,Brad, Graham,Liam and Stephen Wright (2011) ”Invertible and Non-invertible Information Sets in Linear Rational Expectations Models” *Journal of Economic Dynamics & Control* 35: 295–311

[6] Clarida, Richard, Jordi Galí and Mark Gertler (1999) ”The Science of Monetary Policy: A New Keynesian Perspective”, *Journal of Economic Literature* Vol. XXXVII:1661–1707

[7] Ellison, Martin and Natacha 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.

[8] Guerrieri, Veronica and Robert Shimer (2013) ”Markets with Multidimensional Private Information”, Society for Economic Dynamics Meeting Papers series 2013 n. 210.

[9] Kiefer, Nicholas and Yaw 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.

[10] Klemperer, Paul D. and Margaret A. Meyer (1989) ”Supply Function Equilibria in Oligopoly under Uncertainty”, Econometrica, Vol. 57, No. 6, pp. 1243-1277

[11] Kocherlakota, N.J. (2010), *The New Dynamic Public Finance,* Princeton University Press.

[12] Lucas Jr, Robert E., Jr. (1972) ”Expectations and the Neutrality of Money”, J*ournal of Economic Theory*, 4(2), 103-124.

[13] Lucas, Robert E., Jr. and Nancy L. Stokey (1983) "Optimal Fiscal and Monetary Policy in an Economy without Capital", *Journal of Monetary Economics* 12, 55-93

[14] Mehra, R.and E.C. Prescott.(1980)."Recursive Competitive Equilibrium: the Case of Homogeneous Households", *Econometrica* 48(6),1365–1379

[15] Mirman, L. J., Samuelson, L., and Urbano, A. (1993), "Monopoly Experimentation", *International Economic Review*, 549-563.

[16] Muth, J. F. (1960), "Optimal Properties of Exponentially Weighted Forecasts" *Journal of the American Statistical Association*, Volume 55, Issue 290

[17] Nimark, Kristoffer (2008) "Monetary Policy with Signal Extraction from the Bond Market", *Journal of Monetary Economics* 55, 1389–1400

[18] Pearlman, J., D. Currie and P.Levine (1986) "Rational Expectation Models with Private Information", *Economic Modelling* 3(2): 90-105

[19] Pearlman, Joseph (1992) "Reputational and Nonreputational Policies under Partial Information", *Journal of Economic Dynamics and Control* 16(2): 339-357

[20] Svensson, Lars E.O. and Woodford, Michael (2003) "Indicator Variables for Optimal Policy", *Journal of Monetary Economics* 50, 691–720

[21] Svensson, Lars E.O. and Woodford, Michael (2004) "Indicator Variables for Optimal Policy under Asymmetric Information" *Journal of Economic Dynamics & Control* 28, 661 – 690

[22] Swanson, Eric T. (2006) "Optimal Nonlinear Policy: Signal Extraction with a Non-normal Prior", *Journal of Economic Dynamics and Control* 30: 185-203

[23] Townsend, Robert M. (1983) "Forecasting the Forecasts of Others", *Journal of Political Economy* 91(4), 546-588

[24] Van Nieuwerburgh, S., Veldkamp, L. (2006) "Learning Asymmetries in Real Business Cycles" *Journal of Monetary Economics*, Elsevier, vol. 53(4), pages 753-772, May.

[25] Wieland, Volker (2000a) "Monetary Policy, Parameter Uncertainty and Optimal Learning", *Journal of Monetary Economics* 46: 199-228

[26] Wieland, Volker (2000b) "Learning by Doing and the Value of Optimal Experimentation", *Journal of Economic Dynamics & Control* 24: 501-534

Appendix A: Proof of Proposition 2

To arrive at Proposition 2 we first need to show that for any variation δ and small α the equilibrium $S(\mathcal{R}^* + \alpha\delta, A)$ is well defined with probability one. This will be stated in Lemma 2. We first state a generic result on the existence and uniqueness of solutions.

On the sets $X \subset R$, $Y \subset R^n$ we consider functions $f, f^k : X \times Y \to R$ for $k = 1, 2, ...$

Let $d(f, g)$ be the sup norm

$$d(f,g) = \sup_{X \times Y} |f(x,y) - g(x,y)|$$

We make the following assumptions

- **Assumption L1:** X and Y are compact, $f(\cdot, y)$ and $f^k(\cdot, y)$ are absolutely continuous for each y.

- **Assumption L2:** f^k converge uniformly to f, i.e. $d(f, f^k) \to 0$ as $k \to \infty$

- **Assumption L3:** for each $y \in Y$ there is a unique solution to $f(\cdot, y) = 0$. This solution lies in the interior of X.

  Hence, there is a well-defined mapping $\chi : Y \to X$ with the property $f(\chi(y), y) = 0$ and $\chi(Y) \subset int(X)$.

  Let $X^d(y) \subset X$ be the set containing all points where the derivatives $f_x(\cdot, y)$ and $f_x^k(\cdot, y)$ exist for all k. Note that because of absolute continuity $f_x(\cdot, y)$ and $f_x^k(\cdot, y)$ are differentiable almost everywhere so $X^d(y)$ has measure zero.

- **Assumption L4:** The partial derivatives of f^k converge as follows

$$\sup_{x \in X^d(y), y} \left| f_x(x,y) - f_x^k(x,y) \right| \to 0$$

- **Assumption L5:** f_x is uniformly bounded away from zero near the zeroes. Formally, there is a constant $K > 0$ and an $\varepsilon > 0$ such that for all $y \in Y$ and all x such that $|x - \chi(y)| < \varepsilon$, then $x \in X$ and if $f_x(x, y)$ exists then

$$|f_x(x,y)| > K$$

  Furthermore, for all y, the sign of $f_x(x, y)$ is the same for any x such that $|x - \chi(y)| < \varepsilon$ where f_x exists.

Note that once we have uniqueness as in Assumption L3 the last assumption is easily guaranteed, it just requires that the partial f_x is not close to 0 near the zeroes, for example if $f(x, y) = \phi x + x^3$ it excludes $\phi = 0$ but any other ϕ works. The sign restriction excludes, for example, $f(x, y) = |x|$.

Lemma 1 ***1 (Existence and uniqueness of solutions)*** *Under Assumptions L1-L5 there exists a unique solution to $f^k(\cdot, y) = 0$ for all y for k is sufficiently high. Let $\chi^k(y)$ denote this solution. Furthermore, $d(\chi^k, \chi) \to 0$ as $k \to \infty$.*

Proof

Define the neighborhood of the zeros as the set

$$XY^\varepsilon \equiv \{(x, y) : for\, all\, y \in Y\, and\, x\, such\, that\, |x - \chi(y)| < \varepsilon\, for\, some\, y \in Y\}$$

where ε is as in assumption L5.

We first prove that $f^k(\cdot, y)$ has a solution for all y for sufficiently high k.

The fundamental theorem of calculus gives

$$\int_{\chi(y)}^{x} f_x(\overline{x}, y) d\overline{x} = f(x, y) \tag{57}$$

for all (x, y).[22] Consider a y such that the sign of f_x near $\chi(y)$ is positive. Using assumption L5, if $|x - \chi(y)| < \varepsilon$ and if $x > \chi(y)$ we have

$$\int_{\chi(y)}^{x} f_x(\cdot, y) > K(x - \chi(y)) \tag{58}$$

Plugging $x = \chi(y) \pm \varepsilon/2$ in (57) we have

$$\begin{aligned} \frac{\varepsilon}{2}K &< f(\chi(y) + \frac{\varepsilon}{2}, y) \\ \frac{\varepsilon}{2}K &> f(\chi(y) - \frac{\varepsilon}{2}, y) \end{aligned}$$

A similar argument gives the reverse inequalities for y's such that the sign of f_x is negative near $\chi(y)$.

This all implies that $f(\chi(y) \pm \frac{\varepsilon}{2}, y)$ are bounded away from zero, and of oposite signs for all y.

Note that here we have used $\chi(Y) \subset int(X)$ in guaranteeing that f is well defined at $x = \chi(y) \pm \frac{\varepsilon}{2}$.

Applying uniform convergence of f^k there is k sufficiently high so that $d(f^k, f) < \frac{\varepsilon}{4}K$. The last two equations imply that if $K > 0$

$$\begin{aligned} \frac{\varepsilon}{4}K &< f^k(\chi(y) + \frac{\varepsilon}{2}, y) \\ \frac{\varepsilon}{4}K &> f^k(\chi(y) - \frac{\varepsilon}{2}, y) \end{aligned}$$

Therefore $f^k(\cdot, y)$ takes a positive and a negative value. This implies by the intermediate value theorem that a solution to $f^k(\cdot, y) = 0$ exists for all y and for k high

[22] It is understood that for $x < \chi(y)$ upper and lower limit of the integral need to be exchanged.

enough. As promised, we have shown that $f^k(\cdot, y)$ has a solution for all y for sufficiently high k. Let us call this solution $\chi^k(y)$, since we have not proved uniqueness all we know thus far is that $\chi^k(y)$ is a non-empty set.

We now prove that $\chi^k(y) \subset XY^\varepsilon$. We first show that

$$|f(x,y)| \geq V forall (x,y) \notin XY^\varepsilon \tag{59}$$

for a constant $V > 0$. To see this consider the infimum of $|f|$ outside XY^ε

$$V \equiv \inf_{X \times Y - XY^\varepsilon} |f(x,y)|$$

Compactness of $X \times Y - XY^\varepsilon$ and continuity of f implies the infimum is attained at some $(x^*, y^*) \in X \times Y - XY^\varepsilon$ such that $f(x^*, y^*) = V$. If $V = 0$ this implies that $x^* = \chi(y^*)$ and it contradicts the fact that all the zeroes of f are contained in XY^ε. Therefore $V > 0$.

Uniform convergence together with (59) implies

$$\left|f^k(x,y)\right| > V/2 > 0 \; \forall (x,y) \in X \times Y - XY^\varepsilon \; for\, k\, large\, enough$$

so that $f^k(x,y) = 0$ can only happen for $(x,y) \in XY^\varepsilon$ for k large.

This proves that all solutions of f^k (namely (x,y) where $x \in \chi^k(y)$) must lie in XY^ε for k large enough.

All that is left to show is that $\chi^k(y)$ has only one element for k large enough. Applying (57) and using assumption L5 in a similar way as we did above implies that given any y for which f_x is positive in assumption L5, for any elements $x', x'' \in \chi^k(y)$ with $x' \geq x''$ then

$$\int_{x''}^{x'} f_x(\overline{x}, y) d\overline{x} = f(x', y) - f(x'', y) = 0$$

which is impossible since

$$\int_{x''}^{x'} f_x(\overline{x}, y) d\overline{x} > K(x' - x'') > 0$$

Therefore there exists a unique solution of $f^k(\cdot, y)$ for k large enough for all y or, equivalently, $\chi^k(y)$ has one element for all y.

To prove convergence of the solutions to $\chi^k(y)$ note that, given y, for any convergence subsequence $\left\{\chi^{k_j}(y)\right\}_{j=1}^{\infty}$ where the limit is denote $\lim_{j\to\infty} \chi^{k_j}(y) = L(y)$, uniform convergence of f^k implies that $f(L(y), y) = 0$, so that $L(y) = \chi(y)$ and $\chi^k(y) \to \chi(y)$ as $k \to \infty$.$\blacksquare$

We apply Lemma 1 to show that $S(\mathcal{R}^* + \alpha\delta, A)$ is well defined. We take h and W as defined on sets for signals $s \in \mathcal{S}$ and policies $\tau \in \mathcal{T}$.

We need the following assumptions on our model. Recall H has been defined in (23).

Assumptions

1. The sets $\mathcal{S}$,$\mathcal{T}$ are both compact. The set of realizations Φ is also compact
2. $h(\cdot;A)$ is differentiable everywhere, $|h_\tau| < Q$ uniformly on (s, A) for a constant $Q < \infty$ and h_τ is Lipschitz continuous with respect to s for $s = S(\mathcal{R}^*, A)$ for almost all A.
3. For each A there is a unique solution $S(\mathcal{R}^*, A) \in int(\mathcal{S})$ setting $H(S(\mathcal{R}^*, A), A; \mathcal{R}^*) = 0$
4. $\mathcal{R}^*$ is continuous in s. Furthermore, $|\mathcal{R}^{*\prime}| < K^R$, for a constant $K^R < \infty$ for $s = S(\mathcal{R}^*, A)$ for almost all A.
5. Either $h_\tau(\mathcal{R}^*(s); A)\mathcal{R}^{*\prime}(s) < 1 - K$ or $h_\tau(\mathcal{R}^*(s); A)\mathcal{R}^{*\prime}(s) > 1 + K$ for all $s = S(\mathcal{R}^*; A)$ where the derivatives exist and for some constant $K > 0$.

Assumptions 1-2 can be imposed before knowing the solution to the PI problem. Assumptions 3-5 depend on the solution $\mathcal{R}^*$ which is not known ahead of time, but they can be checked ex-post with a candidate solution. Typically $\mathcal{R}^*$ will have a bounded derivative almost everywhere and since we consider cases where $S(\mathcal{R}^*, A)$ has a density these assumptions are easily verified.

Lemma 2 *2 Consider a bounded function (a "variation") $\delta : \mathcal{S} \to R$ differentiable everywhere and with a uniformly bounded derivative. Under assumptions 1-5 $S(\mathcal{R}^* + \alpha\delta, A)$ exists and is unique for α small enough.*

Proof

We apply Lemma 1 when we take a sequence $\alpha_k \to 0$. The objects in our model maps into the notation of Lemma 1 as follows

-s takes the role of x and A takes the role of y in Lemma 1.

-$f(\cdot) \equiv H(\cdot; \mathcal{R}^*)$

-$f^k(\cdot) \equiv H(\cdot; \mathcal{R}^* + \alpha_k\delta)$

We now have to check that assumptions 1-5 imply assumptions L1-L5.

Assumptions 1,2,4 imply Assumption L1.

Absolute continuity of h implies

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

where $Q < \infty$ is the uniform bound on $|h_\tau|$. Since δ is bounded this implies that $H(.; \mathcal{R}^* + \alpha_k\delta)$ convergence uniformly to $H(\cdot; \mathcal{R}^*)$ as in assumption L2.

Assumption L3 and 3 are equivalent.

Consider s, A where the derivatives $h_\tau(\cdot, A)$ and $\mathcal{R}^{*\prime}$ exist. We have

$$H_s(s, A; \mathcal{R}^* + \alpha_k\delta) = 1 - h_\tau((\mathcal{R}^* + \alpha\delta)(s), A)(\mathcal{R}^{*\prime} + \alpha\delta')(s)$$

Therefore

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

the second equality follows from adding and substracting $h_\tau(\mathcal{R}^*(s),A)(\mathcal{R}^{*\prime}+\alpha\delta')(s)$ and where $K^{\delta'}$ is the bound on δ', K^δ the bound on δ, and Q^L the Lipschitz constant for h_τ.

This guarantees Assumption L4.

Assumption L5 is given by assumption 5.

Therefore, Lemma 1 implies that $S(\mathcal{R}^*+\alpha\delta,A)$ is well defined for α small enough.■

To prove Proposition 2 we now need to add

Assumption 6: W is continuously differentiable with respect to (τ,s) for almost all A and when the derivative exists it is uniformly bounded.

Proof of Proposition 2

Consider a variation $\delta:\mathcal{S}\to R$ such that δ is continuous, differentiable, $|\delta(s)|\leq K^\delta$ and $|\delta'(s)|\leq K^{\delta\prime}$ for all s for some constants $K^\delta, K^{\delta'}<\infty$.

Consider the problem defined by (28). For α small enough $S(\mathcal{R}^*+\alpha\delta;A)$ is well defined by Lemma 2, and so is $\mathcal{F}$. It is clear that since $\mathcal{R}^*+\alpha\delta$ is a feasible policy function in the PI problem the solution of (28) is attained at $\alpha=0$. Since (28) is a standard one-dimensional maximization problem this implies

$$\left.\frac{d\mathcal{F}(\mathcal{R}^*+\alpha\delta)}{d\alpha}\right|_{\alpha=0} = 0 \tag{60}$$

if this derivative exists. We now prove that this derivative exists and that (60) implies (27). We consider here the one-dimensional case for τ and s, namely $m=n=1$, the generalization to multivariate case is left for the final version of the paper.

The assumptions on W and h imply

$$|W_\tau|,|W_s|,|h_\tau| < K^{Wh}$$

for some finite constant K^{Wh} whenever the derivatives exist.

Take any sequence $\alpha_k\to 0$. For each k we can write

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

for

$$M_k(A) \equiv \frac{W\left(T(\mathcal{R}^*{+}\alpha_k\delta,A),S(\mathcal{R}^*{+}\alpha_k\delta,A),A\right) - W\left(T(\mathcal{R}^*,A),S(\mathcal{R}^*,A),A\right)}{\alpha_k}$$

At all points A where $\mathcal{R}^*, W, h$ are differentiable at $s = S(\mathcal{R}^*, A)$ we have that

$$M_k(A) \to \left.\frac{dW\left(\left(\mathcal{R}^* + \alpha\delta\right)\left(S(\mathcal{R}^*{+}\alpha\delta; A)\right), S\left(\mathcal{R}^*{+}\alpha\delta, A\right), A\right)}{d\alpha}\right|_{\alpha=0} \tag{61}$$

$$= W_\tau^{*\prime}(A)\left\{\left[\mathcal{R}^{*\prime}(A)\right] S_\delta^{*\prime}(A) + \delta^*(A)\right\} + W_s^{*\prime}(A) S_\delta^{*\prime}(A) \tag{62}$$

where

$$\begin{aligned}
\mathcal{R}^{*\prime}(A) &= \mathcal{R}^{*\prime}\left(S(\mathcal{R}^*; A)\right) \qquad (63)\\
\delta^{*\prime}(A) &= \delta'\left(S(\mathcal{R}^*, A)\right)\\
\delta^*(A) &= \delta\left(S\left(\mathcal{R}^*, A\right)\right)\\
W_x^*(A) &= W_x\left(\left(\mathcal{R}^*\left(S(\mathcal{R}^*, A)\right), S\left(\mathcal{R}^*, A\right)\right)\right) \ for\ x = \tau, s\\
S_\delta^{*\prime}(A) &= \left.\frac{dS(\mathcal{R}^*{+}\alpha\delta, A)}{d\alpha}\right|_{\alpha=0} \qquad (64)
\end{aligned}$$

(note there is a slight abuse of notation on $\mathcal{R}^{*\prime}$ and $\delta^{*\prime}$ as we use the symbol for the function of s as well as for the function of A).

The only non-obvious term is the derivative $S_\delta^{*\prime}$, furthermore, this is the only term that depends on δ at $\alpha = 0$. The term $S_\delta^{*\prime}$ can be found by applying the implicit function theorem to

$$S(\mathcal{R}^*{+}\alpha\delta; A) = h((\mathcal{R}^*{+}\alpha\delta)(S(\mathcal{R}^*{+}\alpha\delta; A)), A)$$

Carefully differentiating with respect to α we have

$$\begin{aligned}
\frac{dS(\mathcal{R}^*{+}\alpha\delta; A)}{d\alpha} &= h_\tau((\mathcal{R}^*{+}\alpha\delta)(S(\mathcal{R}^*{+}\alpha\delta; A)), A)\left[(\mathcal{R}^{*\prime}{+}\alpha\delta')(S(\mathcal{R}^*{+}\alpha\delta; A))\right] \cdot\\
&\quad \cdot \frac{dS(\mathcal{R}^*{+}\alpha\delta; A)}{d\alpha} + \delta(S(\mathcal{R}^*{+}\alpha\delta; A))
\end{aligned}$$

So that

$$S_\delta^{*\prime}(A) \equiv \left.\frac{dS(\mathcal{R}^*{+}\alpha\delta, A)}{d\alpha}\right|_{\alpha=0} = \frac{h_\tau^*(A)\delta^*(A)}{1 - h_\tau^*(A)\mathcal{R}^{*\prime}(A)}$$

Using the boundedness assumptions it is clear that

$$|M_k(A)| \leq K^{Wh}\left\{\left[K^R + \alpha K^{\delta\partial}\right]\frac{K^{Wh}}{K} + K^\delta\right\} + K^{Wh}\frac{K^{Wh}}{K}$$

uniformly on k for almost all A.

The assumption that $S(\mathcal{R}^*, A)$ has a density implies that the limit in (62) occurs with probability one in A. Since M_k is uniformly bounded Lebesgue dominated convergence implies

$$\frac{\mathcal{F}(\mathcal{R}^* + \alpha_k\delta) - \mathcal{F}(\mathcal{R}^*)}{\alpha_k} \to \int_\Phi \left(\left[W_\tau^*\mathcal{R}^{*\prime} + W_s^*\right] S_\delta^{*\prime} + W_\tau^*\delta^*\right) dF_A$$

Since this holds for any sequence $\alpha_k \to 0$ it proves that the derivative of $\mathcal{F}$ with respect to α exists for any variation δ and from (60) we have

$$\int_{\Phi} \left(\left[W_\tau^* \mathcal{R}^{*\prime} + W_s^* \right] S_\delta^{*\prime} + W_\tau^* \delta^* \right) dF_A = 0 \tag{65}$$

Using the formula for $S_\delta^{*\prime}(A)$ and rearranging, we conclude that for any variation δ

$$\int_{\Phi} \left(W_\tau^* + W_s^* h_\tau^* \right) \frac{\delta^*}{1 - h_\tau^* \mathcal{R}^{*\prime}} dF_A = 0 \tag{66}$$

Since (66) holds for any bounded δ with bounded derivative it also holds when δ is any bounded function measurable with respect to s. Therefore, the general definition of conditional expectation implies (27).

Appendix 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 observables.

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

$$h(\mathcal{R}(\bar{l}), \theta, \tilde{\gamma}(\mathcal{R}(\bar{l}), \theta, \bar{l})) - \bar{l} = 0. \tag{67}$$

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

$$\zeta = -\frac{h_\theta}{h_\gamma} = \frac{g\sqrt{(Bg)^2 + 4B\theta^2\gamma(1-\tau)} + Bg^2}{2\theta^3(1-\tau)} \tag{68}$$

Now we differentiate ζ with respect to g and get

$$\frac{d\zeta}{dg} = \frac{\partial \zeta}{\partial g} + \frac{\partial \zeta}{\partial \tau}\frac{\partial \tau}{\partial g} \tag{69}$$

where all these partial derivatives 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 observables. For this purpose we will take a first order approximation of the loci $\tilde{\gamma}$.[23]

Consider Figures 19 and 20. When g is sufficiently low, the map of loci (solid blue lines) moving in the direction of increasing l's looks like in Figure 19. Starting at l_{min}

[23] In our computed examples these loci are very close to linear.

(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 $\gamma's$. 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 $\mathcal{R}^*$ in section 5.1 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 20, 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 in section 5.2 with high government expenditure.

Figure 19: Loci of shocks realizations with low g

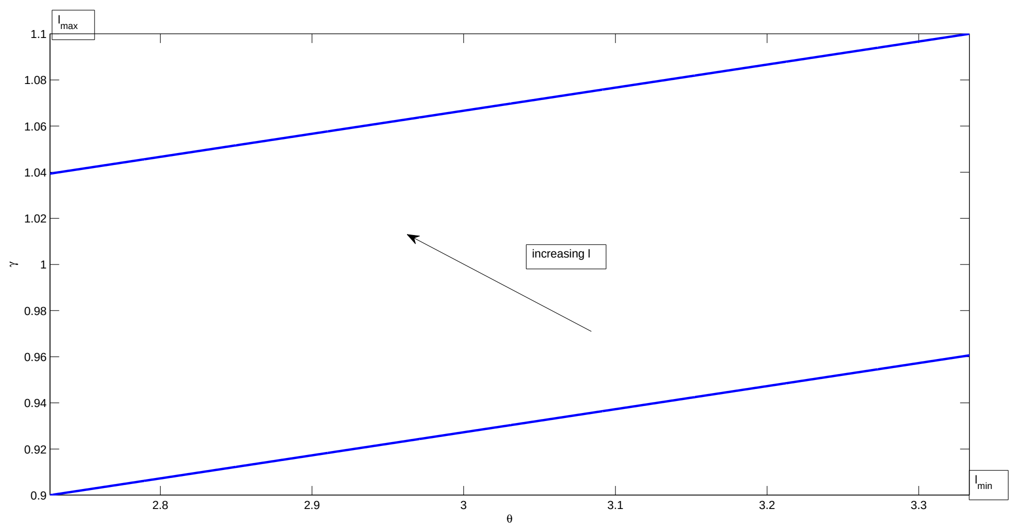

Appendix C: Derivation of the Envelope Condition (56)

In this Appendix we derive the Envelope Condition (56). 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

$$H(l, A, \mathcal{R}) \equiv l - h(\mathcal{R}(b, l), A), \tag{70}$$

Figure 20: Loci of shocks realizations with high g

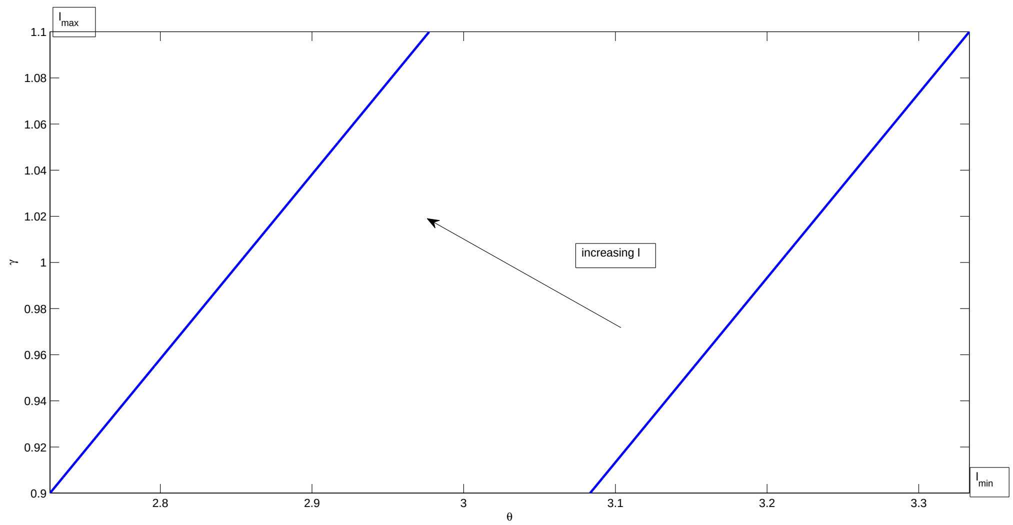

in analogy with the two-period model. By total differentiation of (70), the partial derivative of labor with respect to debt, L_b, is given by

$$L_b(\mathcal{R}, b, A) = \frac{\gamma\theta\mathcal{R}_b(b,l)}{v''(l) + \gamma\theta\mathcal{R}_L(b,l)}. \tag{71}$$

Now, for simplicity consider a case without borrowing penalty. In order to derive the envelope condition, we differentiate (55) with respect to b and get

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

where

$$\begin{aligned} b_L^{*\prime} &= \frac{-\theta\gamma + [v''(L(\mathcal{R}^*, A, b))L(\mathcal{R}^*, A, b) + v'(L(\mathcal{R}^*, A, b))]}{\beta\bar{\gamma}} \\ l^* &= L(\mathcal{R}^*; b, A) \\ L_b^* &= L_b(\mathcal{R}^*; b, A). \end{aligned}$$

Using (71) we can write

$$W'(b) = E\left[\left(\gamma\theta - v'(l^*) + \beta W'\left(b^{*\prime}\right)b_L^{*\prime}\right)\frac{-\gamma\theta\mathcal{R}_b^*(l,b)}{v''(l) + \gamma\theta\mathcal{R}_L^*(l,b)} + W'\left(b^{*\prime}\right)\frac{\gamma}{\bar{\gamma}}\right]. \tag{72}$$

Using Proposition 2, the FOC of PI Ramsey problem is

$$E\left[\left(\theta\gamma - v'(l^*) + \beta W'(b'^*)b_L^{*'}\right)\frac{h_\tau^*}{1-h_\tau^*\mathcal{R}_L^*}|\bar{l}\right] = 0 \tag{73}$$

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 (72) we get

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

Now, applying the law of iterated expectations, using the fact that $\mathcal{R}_b(l,b)$ is known given L, b and using (73), we obtain

$$\begin{aligned} W'(b) &= E\left[E\left(\left(\gamma\theta - v'(L^*) + \beta W'\left(b^{*'}\right)b_L^{*'}\right)\frac{h_\tau^*\mathcal{R}_b^*(L,b)}{1-h_\tau^*\mathcal{R}_L^*(L,b)}\middle| L\right) + W'\left(b^{*'}\right)\frac{\gamma}{\bar{\gamma}}\right] && (74) \\ &= E\left[0 + W'\left(b^{*'}\right)\frac{\gamma}{\bar{\gamma}}\right] && (75) \end{aligned}$$

Finally, adding the marginal cost of excessive debt this becomes

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