Self-fulfilling deflations

Roberto Piazza *

March 26, 2015

Abstract

What types of monetary and fiscal policy rules give rise to self-fulfilling deflationary paths that are monotonic and empirically relevant? The paper provides simple theoretical conditions that guarantee the existence of these paths in a general equilibrium model with sticky prices. These sufficient conditions turn out to be weak enough to be satisfied by most monetary and fiscal policy rules. A quantification of the model which combines a real shock à la Hayashi and Prescott (2002) with a simultaneous sunspot that dis-anchors inflation expectations matches the main empirical features of the Japanese deflationary process during the "lost decade". The results also highlight the key role of the assumption about the anchoring of inflation expectations for the size of fiscal multipliers and, in general, for any policy analysis.

1 Introduction

Since 2008, aggressive monetary policy has brought the short-term interest rate at, or close to, the zero lower bound in major advanced economies. Time and time again "deflation scares" and the occurrence of liquidity traps have been a source of concern for policy-makers around the world. While liquidity traps are an interesting empirical phenomenon and an important topic of policy discussions, the theoretical debate on their nature is still very much open. The paper contributes to this debate by presenting new theoretical and quantitative results on the properties of self-fulfilling liquidity traps.

Self-fulfilling liquidity traps may arise because, under rational expectations, monetary and fiscal feedback rules typically give rise to multiple equilibria. This fact has been recognized at least since Sargent and Wallace (1975). When the monetary authority pursues an inflation target, additional equilibria coexist, where inflation expectations are "dis-anchored" from the target. A sunspot that dis-anchors inflation expectations may therefore be the reason why an economy is driven into a deflationary liquidity trap. The historical record tells us that this is more than a mere theoretical possibility. The power of inflation expectations in shaping liquidity traps and in driving policymakers to extreme measures that appear, at the same time, bold and desperate, emerges clearly in the words of Bank of Japan's Governor Kuroda:

*I would like to thank James Bullard, John Cochrane and Stephanie Schmitt-Grohé for their useful comments on a previous version of the paper. Moreover, I would also like to thank Kozo Ueda for providing me with his original dataset on inflation expectations. The views expressed are those of the author and do not necessarily reflect those of the Bank of Italy. Contact information: Bank of Italy, via Nazionale 91, 00184, Roma, e-mail: roberto.piazza@bancaditalia.it

«Japan is different from countries like the United States, which has inflation expectations anchored at 2 percent. [...] There was a risk that despite having made steady progress, we could face a delay in *eradicating the public's deflation mindset* [...]. It's important for the BOJ [...] to *get its price target firmly embedded in people's mindset.*»[1]

What types of policy rules can give rise to self-fulfilling deflationary paths that are interesting from an empirical standpoint? The literature has not provided a general answer to this question. The main reason is that the dynamic equation defining the evolution of dis-anchored inflation expectations is, as Fernandez-Villaverde (2014) puts it, «hard to characterize» in general terms. Thus, while a case-by-case approach focusing on specific functional forms for the policy rules has led to important seminal results (Benhabib, Schmitt-Grohé and Uribe [2001]), a more comprehensive answer is still lacking. The paper addresses this issue.

First, the paper proves that any monetary and fiscal policy rules that jointly satisfy certain easy-to-check sufficient conditions are guaranteed to give rise to "reasonably looking" self-fulfilling deflationary paths. In my definition, "reasonably looking" is equivalent to monotonicity. Focusing on monotonic deflationary paths is attractive because, as I will argue, they represent empirically relevant situations. By selecting paths that feature simple dynamics I am thus deliberately leaving aside complicated or chaotic expectational dynamics (Benhabib, Schmitt-Grohé and Uribe [2002b]). Such equilibria may be more a curious mathematical object rather than an empirically relevant representation of how people's expectations actually evolve. Second, the paper shows that the aforementioned sufficient conditions are satisfied by a very wide set of potential policy rules. This set includes any static or forward looking rule that prescribes some minimal form of stimulus when inflation, actual and/or expected, drifts below the target. The stimulus may even be minimal in the sense that, for instance, the sufficient conditions are satisfied also for monetary rules that call for an *increase* in policy interest rate in some deflationary states.

The findings above indicate that the existence of "reasonably looking" deflationary equilibria is a pervasive phenomenon in sticky-price models with rule-based policies. However, while all minimally stimulative rules are subject to self-fulfilling and monotonic deflationary paths, it is nonetheless true that rules that are *more* stimulative have some advantages. In particular, the paper proves that a very stimulative rule requires a sizable initial dis-anchoring of inflation expectations for the self-fulfilling deflationary path to form (*Type II* equilibrium). When the rule, instead, is not very stimulative the deflationary path forms even for an initial arbitrarily small dis-anchoring of expectations (*Type I* equilibrium).

How well can a basic sticky price model match the empirical counterpart of a deflationary process driven by a dis-anchoring of expectations? I choose to focus, as my empirical counterpart, on the Japanese deflation during the "lost decade" 1992-2002, for a number of reasons. First, the Japanese deflation is arguably the most studied and, up to recently, unique instance of a liquidity trap in the post war era. Second, the Japanese deflationary process was monotonic. Third, there is a widespread perception, as also Japanese policymakers have repeatedly stated, that the dis-anchoring of inflation expectations is an indication of the self-fulfilling nature of the Japanese liquidity trap. This idea finds further support, as I argue, in the suggestive evidence that the slow process of dis-anchoring occurred at a very early stage of the lost decade. In fact, during the lost decade the Japanese inflation process can be decomposed into two parts. One is a series of *unexpected* and temporary negative shocks to

[1] Reuters Fri Oct 31, 2014 6:27am EDT.

inflation which, at the *global* level, is common across major advanced countries. The other is an *underlying and fully anticipated* (self-fulfilling, in my interpretation) deflationary process, which was specific only to Japan and slowly led to the liquidity trap.

An appropriately calibrated general equilibrium sticky price model, albeit very simple in its elements, does a good job in matching the salient features of the Japanese deflationary process during the lost decade. This result is based on three crucial ingredients. The first is a sunspot that dis-anchors inflation expectations. The second is a fall in the natural real rate induced, during the Japanese lost decade, by a reduction in potential growth. For this part of the calibration, I rely on the estimates of Hayashi and Prescott (2002) on the evolution of market real rates and on TFP growth in Japan. The third ingredient is the calibration of the monetary policy rule, which is assumed to have a Taylor form. Once calibrated to match some initial conditions on the nominal interest rate in Japan, the inflation parameter of the rule equals 1.38, which is below the typical value of 1.5 reported by Taylor (1993). Using our earlier discussion on the theoretical results, such a rule might be deemed "not very stimulative" (for the case of Japan, this point was made also by Bernanke and Gertler [2000]). Overall, the quantitative results can then be summarized as follows: if a real growth shock à la Hayashi-Prescott is complemented with a contemporaneous sunspot shock to inflation expectations and with a "not very stimulative" monetary rule, then a basic general equilibrium sticky price model accounts well for the evolution of real an nominal variables during the Japanese lost decade. On a similar vein, it is worth noticing that the possible coexistence between real shocks and sunspots shock to inflation expectations has been recently considered by Aruoba and Schorfheide (2012) as a possible explanation of low inflation periods in the USA.

A final set of results concerns the effects of fiscal expansions. As for the case of stimulative monetary policies, the paper proves that anticipated expansionary fiscal policies cannot by themselves prevent the insurgence of monotonic self-fulfilling deflationary paths. Of particular interest is the effect of a fiscal expansion at the zero lower bound for the nominal rate. The paper shows that, if inflation expectations are dis-anchored, short-term fiscal multipliers are small and equal to a standard value of 0.5 (longer term multipliers are even smaller), a result that echoes Mertens and Ravn (2012). This result is not in contrast with the findings that fiscal multipliers at the zero lower bound can be greater than 1 (Eggertsson [2010], Christiano, Eichenbaum, Rebelo [2011]). In fact, the two results are obtained under two different assumptions about the equilibrium selection issue (Christiano, Eichenbaum [2012]). Large long-run fiscal multipliers are obtained if inflation expectations are assumed to be anchored. If, instead, inflation expectations are dis-anchored, then fiscal multipliers have about the same size as outside the zero lower bound. Therefore equilibrium selection choices, which often are only implicitly made, are far from harmless for policy analysis in standard sticky price models (Cochrane [2011], Cochrane [2013]).

The rest of the paper proceeds as follows. The next section provides a brief review of the theoretical literature on liquidity traps and describes some stylized facts, including some novel empirical evidence, about inflation expectations in Japan. Section 3 presents the main theoretical results using a simple version of the aggregate equilibrium equations of the model. Section 4 provides a detailed microfoundation for the aggregate equilibrium equations. Section 5 explores the quantitative properties of the calibrated model and Section 6 concludes.

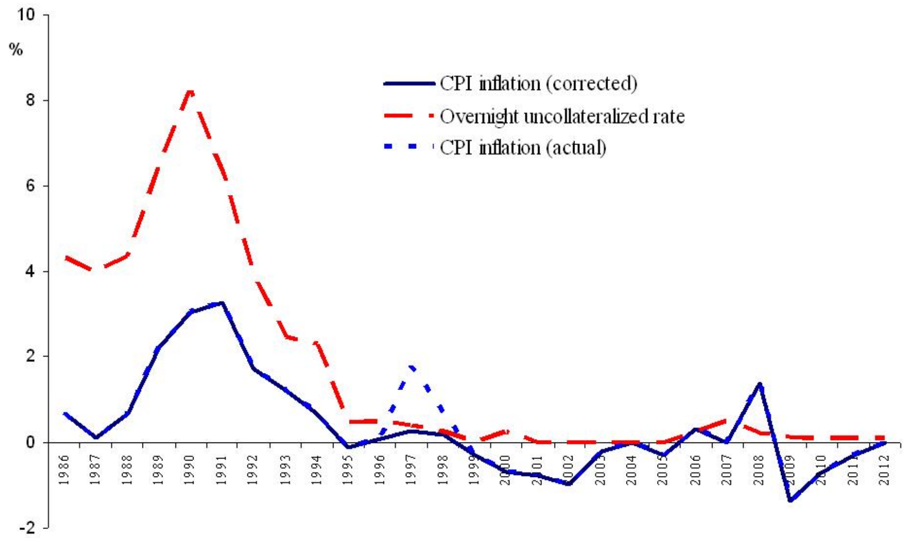

Figure 1: Inflation and the nominal short-term rate in Japan. In April 1997 the consumption tax in Japan was raised by 2 percent (from 3 to 5 percent). To obtain a measure of inflation net (corrected) of the change in the consumption tax, the graph assumes that it took 4 quarters for the tax hike to be completely incorporated into prices. Therefore, three fourths of the overall 2 percent VAT-induced inflation are attributed to the year 1997, and one fourth to 1998.

2 A primer on modeling liquidity traps

To put into context the contribution of the paper, in this section I provide a brief summary of the main current theoretical approaches to liquidity traps. From an empirical perspective I discuss the general stylized facts, some of which have not been previously highlighted in the literature, about the Japanese liquidity trap.

The most prominent example of what is usually intended with the expression "liquidity trap" is provided by Japan. Starting with 1992, and in correspondence with a bust in house prices and at the onset of a severe banking crisis, CPI inflation began to fall (Figure 1). In response to these developments, the Bank of Japan progressively cut the nominal interest rate which, in a time span of about 10 years, reached the zero lower bound. While the overall peak-to-trough reduction in inflation was quite sizeable (about 4 percentage points), it took a relatively long time to fully materialize (10 years, from 1992 to 2002). Once the CPI inflation is netted of the effects of changes in the consumption tax, it becomes also apparent that the deflationary path was roughly monotonic. After 2002 a mild deflation remained in place, bar short periods of price increases mostly due to commodity price spikes. Because of these stylized facts, the initial deflationary process in Japan gained the adjective of "creeping". On the fiscal side, the government reacted to the deflation through various measures of fiscal stimulus, that eventually caused the gross public debt to soar to over 200 percent of GDP. While it is sometimes argued that these measures have prevented an even worse outcome for Japan, it remains true that, from a purely observational perspective, they did not lift the country out of the deflation. Interestingly, over this rather long period of time, the Japanese authorities have reiterated their commitment to put in place a (arguably, Ricardian) long-run fiscal strategy to repay the public debt and avoid both an outright default and debt monetization[2].

Moving from the empirical to the theoretical description, the literature has mainly proposed two ways of explaining liquidity traps. The first sees liquidity traps as a consequence of *fundamental* shocks hitting the economy. The second, which is the one studied in this paper, focuses on liquidity traps that are generated by the presence of multiple *self-fulfilling* equilibria, whose existence is induced by specific government policy rules. Interestingly, both explanations arise from the same underlying model economy, i.e. from a standard dynamic general equilibrium model. This section provides a brief overview of these two alternative views on liquidity traps. As we shall see, depending on which of the two views is embraced, one may obtain quite different answers to the following questions: a) how do liquidity traps come about? b) are there any negative consequences for welfare when the economy ends up in a liquidity trap? c) what can and should policymakers do?

2.1 Shocks to fundamentals

The first type of explanation is based on Eggertsson and Woodford (2003, 2004), who in turn formalize ideas already present in Krugman (1998). Here, liquidity traps are generated by exogenous fundamental shocks that push the household's subjective rate of time preference to levels above one, with the consequence of driving to negative values the "natural rate" of interest (defined as the inverse of the subjective discount factor, minus one). If nominal prices are sticky, and since the nominal interest rate cannot fall below zero, reductions in the nominal rate may not be enough to accommodate negative natural real rates. As a consequence of the mismatch between a smaller (negative) "natural rate" and a larger equilibrium real rate, current aggregate demand falls, and the ensuing economic recession puts downward pressure on nominal prices, generating deflation and exacerbating further the recession. However, even though maneuvering the current nominal interest rate is not enough to re-establish the correct equilibrium real rate, policymakers could still resort to maneuvering future nominal rates: by promising to keep nominal rates at lower-than-usual levels even when natural rate has exogenously returned to normal levels, policy makers increase *future* inflation. Hence, working through the expectation channel of monetary policy, *current* inflation also rises, and this achieves the goal of reducing current real rates. This is in essence the role of the so called *forward guidance*. Indeed, any policy that is able to produce or to mimic the effects of a rise in current inflation has the potential of being beneficial in a liquidity trap of the Eggertsson-Woodford type. Examples of such policies are fiscal expansions (as in Christiano, Eichenbaum, Rebelo [2011]) or, alternatively, fine-tuned sequences of consumption tax hikes (as in Correira, Farhi, Nicolini, Teles [2013]). In conclusion, in Eggertsson-Woodford analysis: a) liquidity traps come about, and end, because of exogenous shocks to fundamentals; b) the presence of excessively high equilibrium real rates are associated with suboptimal equilibrium outcomes; c) policy makers can adopt "stimulative" measures (lowering to zero the current and future level of the nominal interest rate, temporarily increasing government spending, raising consumption taxes) that significantly lessen the negative consequences of the liquidity trap.

[2]For instance, recently Prime Minister Shinzo Abe has declared that the Bank of Japan's bond buying program is not tantamount to debt monetization (Thomson Reuters, May 24 2013, 02:45).

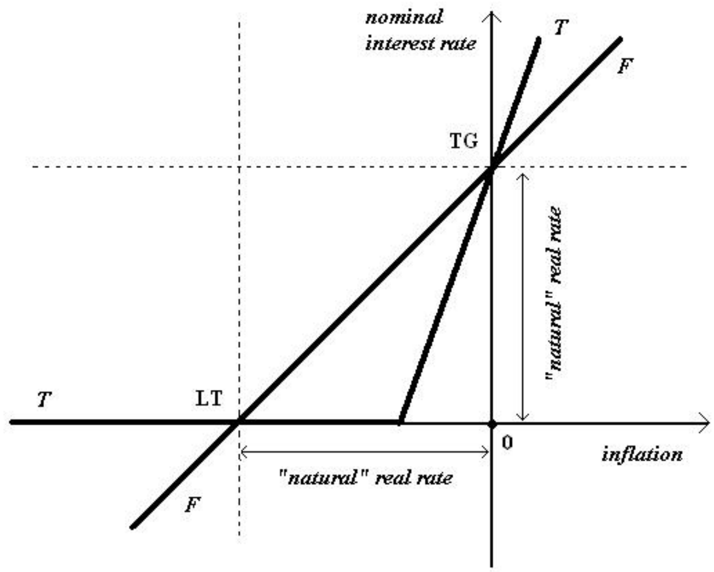

Figure 2: Taylor principle, zero lower bound and the liquidity trap steady state.

2.2 Self-fulfilling deflations

The second explanation of liquidity traps exploits the property that certain monetary policy rules give rise to self-fulfilling multiple equilibria, a fact that has been known at least since the work of Sargent and Wallace (1975). In particular, Benhabib, Schmitt-Grohé and Uribe (2001, 2002a, 2002b) show that liquidity traps having a self-fulfilling nature are generated when monetary policy follows a static Taylor rule satisfying the "Taylor principle" (Taylor [1993]), i.e. a rule that prescribes that the nominal interest rate reacts more than one-to-one to deviations of current inflation from the target inflation rate. Figure 2 provides a standard exemplification of why, in the presence of lower bounds on the nominal interest rate, this type of interest rate policies gives rise to self-fulfilling liquidity traps and multiple equilibria. The FF line is the (linearized) Fisher equation, i.e. the locus of points such that the difference between the nominal interest rate i and the inflation rate π equals a given natural real rate r^n. The line TT is the (linearized) Taylor rule, indicating how the nominal rate deviates from its target level i^{TG} as inflation deviates from the target $\pi^{TG} = i^{TG} - r^n$. Since the nominal rate reacts, at least around the target inflation, more than one-to-one to deviations of π from π^{TG}, the line TT has locally a larger slope than FF. Then, since by assumption the nominal rate is bounded below by zero and is a continuous function of inflation, there must necessarily exist an inflation π^{LT}, that we may identify with a liquidity trap equilibrium, where TT crosses FF a second time. Hence, the economy displays multiple steady states π^{TG} and π^{LT}. In conclusion in this class of models a) liquidity traps have a self-fulfilling nature, and are due to the presence of interest rate rules that follow the Taylor principle b) liquidity traps arise with or without sticky prices[3], but with sticky prices liquidity traps equilibria are inefficient, provided that the level of the target inflation corresponds to the Pareto optimal

[3]Intuitively, the qualitative properties of the lines FF and TT in Figure 2 don't depend on assumptions regarding price stickiness.

equilibrium. To understand the role of government policies, i.e. point c), in a self-fulfilling liquidity trap, it is important to remark that the liquidity trap equilibrium arises for the very reason that the interest rate policy is "stimulative", in the sense that it prescribes strong interest rate cuts whenever inflation falls below target. It is then not surprising that other forms of aggregate demand "stimulus", such as temporary increases in government spending or reductions in consumption taxes, are not effective policy tools in a self-fulfilling liquidity trap, as Mertens and Ravn (2012) very clearly explain.

Turning to the dynamic properties of self-fulfilling liquidity traps, the literature has emphasized that the convergence path of the economy to the liquidity trap steady state is in general quite complicated, displaying spiral patterns or even chaotic behavior, as shown in Benhabib, Schmitt-Grohé, Uribe M. (2001, 2002b). Moreover, since the degree of indeterminacy could be up to two (Benhabib, Schmitt-Grohé, Uribe [2001]), equilibrium selection issues typically arise in quantitative applications (Mertens and Ravn [2012]). Taken together, these problems have led some authors (Bullard 2010) to question the practical relevance of self-fulfilling liquidity traps. In particular, a large body of literature (Bullard and Mitra [2007], Eusepi [2007], Christiano and Eichenbaum [2012]) has investigated the learnability of multiple rational expectations equilibria and found that under various learning mechanisms self-fulfilling equilibria can be eliminated. This is an important strand of literature which is not directly addressed in this paper. On the other hand, Aruoba and Schorfheide F. (2012) estimate a Neo-Keynesian model by explicitly considering sunspot shocks to inflation expectations as a driving force of equilibrium dynamics. They find that a downward sunspot shift in inflation expectations was a policy relevant phenomenon for the US during the period 2008-2009. Their analysis is complementary to mine in that they focus on sunspot deflation that are different (less extreme, and thus more apt to explain the US experience) than the permanent deflationary path I consider in this paper.

2.3 Expected versus unexpected deflation

The choice between modeling liquidity traps as driven by fundamental shocks to preferences or by self-fulfilling multiple equilibria presents both theoretical and empirical challenges. In particular, for a successful model to replicate the main stylized facts about the Japanese deflation, the model must generate a *long lasting* liquidity trap and deflation, coupled with an *absence of above-trend output growth*. Eggertsson and Woodford (2003) show that such features can be obtained if, in every period and for a long time, the economy is unexpectedly hit by a negative shocks to the fundamental rate of time preference. One implication of the unexpected nature of the shock is such that, in every period and for a long time, actual inflation falls short of agents' ex-ante expected inflation[4]. On the contrary, in a model of self-fulfilling liquidity traps a protracted (actually, permanent) liquidity trap is generated as a perfect-foresight equilibrium, where ex-ante inflation expectations coincide with actual inflation.

We turn to the data to assess whether deflation in Japan was expected or unexpected. The left-hand side panel of Figure 3 plots the difference between ex-post and ex-ante inflation over a three-year forecasting period. Inflation expectations are taken from *Consensus Forecast.*

[4]In the Eggertsson and Woodford (2003) framework, a fully anticipated sequence of positive shocks to the rate of time preference can generate a liquidity trap with, in addition, no deviation between actual and expected inflation. However, in this case, the model also predicts that during the duration of the liquidity trap the growth rate of output should be particularly large. See Schmitt-Gore and Uribe (2013).

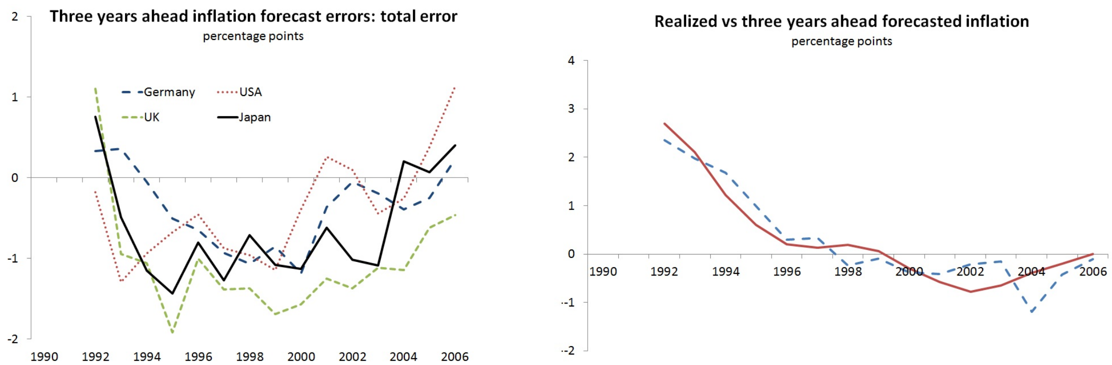

Figure 3: Left-hand side panel: average realized inflation minus ex-ante average forecasted inflation for the three years ending in each t. Right-hand side panel: three-year average realized inflation (solid line) and ex-ante average forecasted inflation (dashed line) once the global error is removed. Sources: Thomson Reuters Datastream, Statistics Bureau of Japan, Consensus Forecast (October survey).

It is clear that, between the forecasting period ending in 1992 and the one ending in 2006, actual inflation persistently fell short of expected inflation. Shall we then conclude that, as in Eggertsson and Woodford (2003), the Japanese liquidity trap was caused by a sequence of unexpected shocks to fundamentals and thus to inflation? While this is certainly a possibility, Figure 3 also makes clear that this conclusion is far from granted. Japan, in fact, was not the only country where actual inflation was persistently lower than expected inflation during the years 1992-2006, since all major advanced countries also experienced such phenomenon. Still, Japan was the only advanced country where negative inflation surprises turned into a liquidity trap. In other words, the presence of a *global* component in inflation forecasting errors does not account alone for the *idiosyncratic* deflationary experience of Japan (global inflation shocks may be due, for instance, to a global component of business cycles, to global shocks to commodity prices, or to a global and slow structural adjustment of inflation expectations from the high-inflation regime of the '80s). The right-hand side panel of Figure 3 depicts the result of removing the global forecasting error (proxied by the average forecasting errors across the US, the UK and Germany) from expected inflation in Japan. It is clear from the figure that this idiosyncratic-only measure of expected inflation tracked remarkably well actual inflation. We could summarize this interpretation of the evidence by saying that, during the period 1992-2006, the inflation process in Japan could be decomposed into two elements. One is a series of *global* unexpected negative shock to inflation. The other is an *underlying* and *fully anticipated* deflationary process, specific only to Japan, which slowly brought the country into the liquidity trap. Appendix C further elaborates on this conclusion, and shows that it is robust to the choice of a one-year forecast period or to the use of qualitative measures of Japanese households' inflation expectations, quantified through an appropriately modified Carlson-Parkin (1975) procedure.

Motivated also by the discussion in this section, I now move to present a formal analysis of liquidity traps in a standard sticky-prices general equilibrium model.

3 Self-fulfilling deflations in a three equation model

To lay out the main theoretical results of the paper in a simple and straightforward way it is worth proceeding backwards in the presentation of the model economy. Therefore, I start by introducing directly the three crucial equilibrium equations that characterize the aggregate behavior of the economy in Neo-Keynesian models. Then, only later in Section 4, I will give a more formal presentation of how these three equations are derived and generalized from a microfounded model. The starting point of our discussion is the case where monetary policy is the only policy tool in the hands of the government. Section 3.3 extends the analysis to include fiscal policy. To keep the exposition simple, I consider a non-stochastic environment, i.e. a setting where all the uncertainty (stemming from either fundamental or sunspot shocks) is resolved in the first period of time.

3.1 Properties of the monetary policy rule

This section presents the main conceptual building blocks for our study of self-fulfilling liquidity traps. First of all, I introduce the three equations that are commonly used to describe the aggregate behavior of the economy when monetary policy follows an interest rate feedback-rule. I intuitively explain why this system of equations can give rise, depending on the particular interest rate rule assumed, to multiple equilibrium that are uniquely pinned down by an initial *sunspot* state of the economy. I then move to provide a more concrete characterization of the set and dynamic properties of sunspot equilibria. Doing this requires postulating some minimal properties for the monetary policy rule. I focus on interest rate rules that have the common flavor of providing "stimulus" to the economy when inflation falls below a desired target. Given the assumptions on the rules, I characterize the set of possible equilibrium dynamics to the liquidity trap steady state, starting from an initial sunspot inflation at or below the target.

The three equations that define the aggregate behavior of the economy are the following:

$$\frac{\hat{Y}_{t+1}}{\hat{Y}_t} = \beta(1 + r_{t+1}) \tag{1}$$

$$\pi_t = -\chi + \kappa\hat{Y}_t + \beta\pi_{t+1} \tag{2}$$

$$r_{t+1} = r(\pi_t, \pi_{t+1}) \tag{3}$$

where $\hat{Y}_t = \frac{Y_t}{Y_t^{TG}}$ is the deviation of output Y_t from the Pareto optimal (target) output Y_t^{TG}. Moreover, π_{t+1} and r_{t+1} are respectively the inflation rate and the real interest rate between time t and time $t + 1$. The constants χ, κ and β are all strictly positive, with $\beta < 1$. The Wicksellian natural real interest rate r^n is defined as

$$r^n = \beta^{-1} - 1 > 0$$

Equation (1) can be interpreted as resulting from the Euler equation of the household's maximization problem. Equation (2), instead, defines a forward-looking Phillips curve, relating current inflation to current output and to future inflation. Finally, equation (3) is the monetary policy rule, that defines the real rate as a function of current and future inflation. As I show in Section 4, equation (2) is[5] the true forward-looking Phillips curve derived in the

[5] Up to a harmless change of variables.

microfounded model. This is important since, by avoiding linearizing the aggregate equations, I am thus able to legitimately use (1)-(3) to perform a global equilibrium analysis.

The interest rate rule in equation (3) was presented in the form of a *real* interest rate rule, rather than the more commonly used *nominal* interest rate rule $i_{t+1} = i(\pi_t, \pi_{t+1})$. This choice is made for the purpose of reducing the notation and simplifying the exposition, but is not restrictive in any way. The reason is that the Fisher condition always provides us with a one to one relation between a nominal and an real interest rate rule. In other words, for any given vector (π_t, π_{t+1}) we can always uniquely recover one rule from the other by using the equivalence[6]

$$1 + i(\pi_t, \pi_{t+1}) = [1 + r(\pi_t, \pi_{t+1})](1 + \pi_{t+1}) \tag{4}$$

If we substitute the rule (3) into (1) we obtain a non-linear system of two second order difference equations in the variables $\hat{Y}_t$ and π_t. Hence, given two arbitrary conditions (both for $\hat{Y}_t$, or both for π_t, or one condition for each of the two variables), the dynamic evolution of the system is uniquely determined. Even though we have not yet defined what an equilibrium is[7], it is nonetheless quite intuitive to realize that, since the system's dynamics are pinned down by two arbitrary conditions, then it is in principle possible that the aggregate economy could be subject to multiple equilibria. To be more concrete, let's define $S_0 = (\pi_0, \pi_1)$ as the initial sunspot *state* of the economy. Given S_0, the system (1)-(2) gives a unique evolution of the state $S_t(\pi_t, \pi_{t+1})$ at all $t \geq 0$. Of course, as a joint-product of the sequence $\{S_t\}_{t=0}^{\infty}$, we obtain also a unique sequence $\{\hat{Y}_t\}_{t=0}^{\infty}$[8]. To sharpen further our analysis, I postulate that any interesting candidate equilibrium for the economy must satisfy two sets of conditions:

$$\begin{aligned} S_t &\to S \in \mathbb{R}^2 \\ \hat{Y}(S_t) &> 0 \end{aligned} \tag{5}$$

The first of the two conditions restricts our attentions to dynamics where the variables converge to some finite steady state value. This is not necessarily an obvious requirement for a candidate

[6] It is worth emphasizing that casting monetary policy in terms of real interest rate rules is not equivalent to assuming that the monetary authority has the power to affect as it pleases the *equilibrium* real interest rate. To see this, consider the case of flexible prices where we know that, *in equilibrium*, the real interest rate must always equal the natural real rate, i.e. $r(\pi_t^*, \pi_{t+1}^*) = r^n$ for all t. Consider a nominal rule such that $1 + i(\pi_t, \pi_{t+1}) = (1 + r^n)(1 + \zeta\pi_{t+1})$ for some $\zeta > 1$. It is easy to verify that this nominal rule is equivalently expressed via a real rule $1 + r(\pi_t, \pi_{t+1}) = (1 + r^n)(1 + \zeta\pi_{t+1})/(1 + \pi_{t+1})$. The real rule prescribes a real rate strictly smaller (strictly greater) than the natural rate r^n whenever the future inflation rate π_{t+1} is strictly smaller (strictly greater) than the zero target. This is a perfectly legitimate choice for a rule, yet it does not imply that the monetary authority has the power to steer the real interest rate away from the natural rate. It only implies that any state where $\pi_{t+1} \neq 0$, and thus where $r(\pi_t, \pi_{t+1}) \neq r^n$, must be *off the equilibrium path*. Indeed, under such rule for any $t > 1$ there is a unique equilibrium for the economy, and this equilibrium coincides with the steady state $\pi_t^* = 0$, where indeed the real rule is consistent with $r(\pi_t^*, \pi_{t+1}^*) = r^n$. It is straightforward to see that, once a zero lower bound on the nominal rate is added, the (modified) rule above gives again rise to a second steady state, and thus to multiple equilibria, exactly as explained in Figure 2.

[7] In particular, I have not specified yet which of the infinite sequences $\left\{\hat{Y}_t, \pi_t\right\}_{t=0}^{\infty}$ that solve the nonlinear system are also legitimate equilibria for the economy.

[8] As mentioned above, the definition of an appropriate state of the system is arbitrary, and subject only to the requirement that it is a vector of two variables. For instance, instead of using as a state the vector (π_t, π_{t+1}), we could as well use the vector $(\pi_t, \hat{Y}_t)$. For any relevant purpose the particular choice we make is inconsequential. Still, I prefer to use the state $S_t = (\pi_t, \pi_{t+1})$ because notationally this choice lends itself better to collapsing, as I do below, the entire system (1)-(3) into one third order difference equation in just the variable π_t.

equilibrium (for instance, this requirement rules out not only cyclical behaviors in inflation, but also explosive inflationary paths), but it is routinely done in most of the Neo-Keynesian literature. The second condition naturally requires that ouput must always be strictly positive.

There is one specific path for the state S_t that is of particular importance for our analysis, namely the *target* steady state $S^{TG} = (\pi^{TG}, \pi^{TG})$, with the target inflation given by $\pi^{TG} = 0$. By setting $\kappa = \chi$, an assumption that I maintain throughout, then the target steady state becomes also the Pareto optimal level of output (i.e. $\hat{Y}_t = 1$). Clearly, for S^{TP} to be a steady state of the economy, we must require that

$$r(0, 0) = r^n \tag{6}$$

So far I have shown how, given an appropriate definition of the state, a set of initial sunspots uniquely characterizes, through the particular form of the feedback monetary policy rule, the dynamic evolution of the economy. In particular, if $S_0 = S^{TG}$ the economy is forever at the target steady state. The issue of the existence of multiple equilibria then boils down to finding whether, in addition an initial state $S_0 = S^{TG}$, we can find other initial states $S_0 \neq S^{TG}$ that also induce a dynamic evolution of the state satisfying (1)-(3) and (5). To answer this question, we need to provide some minimal structure to the real interest rate rule $r(S_t)$ beyond condition (6). A first basic requirement that I impose is that the nominal interest rate rule $i(S_t)$, that induces $r(S_t)$, must satisfy a zero lower bound constraint:

$$i(S_t) \geq 0 \qquad \forall S_t \in \mathbb{R}^2 \tag{7}$$

Another assumption is that the monetary policy rule is a continuous function of the state of the economy and it is monotonically increasing in π_t, i.e. it calls for a reduction in the nominal interest rate whenever current inflation goes down. As usual, by exploiting the Fischer condition (4), the assumption of monotonicity in π_t can be cast in terms of the real interest rate rule as follows,

Assumption 1. *At any state $S_t = (\pi_t, \pi_{t+1})$ with $S < S^{TP}$, the real rate policy $r(S_t)$ is continuous and increasing in π_t, and continuous in π_{t+1}. Moreover, $r(S)$ is differentiable at any S where $r(S) = r^n$.*

The additional requirement regarding differentiability is made for analytical convenience. Finally, I assume that

Assumption 2. *The real interest rate policy $r(S)$ satisfies*

$$r(\pi, \pi) < r^n \qquad \text{for some } \pi < \pi^{TG} \tag{8}$$

We can use Figure 2 to provide a graphical interpretation of Assumption 2. Recall that every point $S_t = (\pi_t, \pi_{t+1})$ of the two dimensional space $\mathbb{R}^2$ represents a possible state of the economy. Imagine that the x-axis in Figure 2 represents the subspace of states $\bar{S}_t = (\pi_t, \pi_t)$ with the property that $\pi_t = \pi_{t+1}$, and π_t is the inflation rate reported on the x-axis. Then Assumption 2 simply requires that there is at least one state $\bar{S}$ where the corresponding nominal interest rate $i(\bar{S})$ is below the 45-degree line (the FF curve), so that $r(\bar{S}) < r^n$.

Given Assumptions 1-2, we can now characterize the set of possible equilibrium paths. First of all, similarly to what depicted in Figure 2, we can show that, together with the target steady state S^{TG}, there always exists a liquidity trap steady state $S^{LT} = (\pi^{LT}, \pi^{LT})$,

Proposition 1. *There always exists one liquidity trap steady state.*

Proof. By Assumption 2, there is a state $\bar{S} = (\pi, \pi)$ with $\pi < \pi^{TG}$ and $r(\bar{S}) < r^n$. Then, by continuity of $r(\cdot)$ and the fact that the nominal rate is non-negative, we can always find a state $S^{LT} = (\pi^{LT}, \pi^{LT})$ such that, for π^{LT} sufficiently smaller than π, we have $r(S^{LT}) = r^n$. □

For expositional simplicity suppose that, as in Figure 2, there is only one liquidity trap steady state S^{LT}, and that $i(S^{LT}) = 0$, i.e. the zero lower bound on the nominal rate is binding at the liquidity trap steady state. The dynamic properties of any path starting from S_0 are obtained by studying the following third order difference equation, obtained by collapsing in one equation the system (1)-(3),

$$\frac{\chi + \pi_{t+1} - \beta\pi_{t+2}}{\chi + \pi_t - \beta\pi_{t+1}} = \frac{1 + r(\pi_t, \pi_{t+1})}{1 + r^n} \tag{9}$$

Let us restrict our attention to initial conditions S_0 with the property that $S_0 \in (S^{LT}, S^{TG})$, i.e. the initial sunspot values of π_0 and π_1 are in between the target π^{TG} and the liquidity trap steady state inflation π^{LT}. To characterize the possible equilibrium dynamic paths of the state S_t, it is useful to use the concept of *maximum reaction* of the monetary policy rule,

Definition 1. *The* maximum reaction $\bar{\phi}$ *of the monetary policy rule* $r(S_t)$ *is given by*

$$\bar{\phi} = \sup\left\{\phi \geq 1 | r(\pi, \phi\pi) - r^n < 0, \quad \text{for some } \pi \in (\pi^{LT}, \pi^{TG})\right\} \tag{10}$$

The maximum reaction of monetary policy gives the highest possible gross rate $\bar{\phi} = \pi_{t+1}/\pi_t$ at which inflation could fall while still allowing the real interest rate to be smaller than the natural rate. Notice that Assumption 2 guarantees that a $\bar{\phi} \geq 1$ actually exists. Monetary policy rules that are characterized by a higher value of $\bar{\phi}$ are referred to, in this paper, as policies that are more "stimulative".

It is instructive to show how $\bar{\phi}$ is characterized in the case of a monetary policy rule that has attracted much interest, i.e. a Taylor rule, defined by positive inflation and output gap parameters, respectively ϕ_π and ϕ_Y. Suppose, for the moment, that the static Taylor rule is unconstrained by the zero lower bound. Specifically, define the nominal interest rate policy $i_{t+1} = i(\pi_t, \pi_{t+1})$ in the following log-linear form

$$\log[1 + i(\pi_t, \pi_{t+1})] = \log(1 + r^n) + \phi_\pi[\log(1 + \pi_t) - \log(1 + \pi^{TG})] + \phi_Y \log(\hat{Y}_t - 1) \tag{11}$$

Since $\pi^{TG} = 0$, by subtracting $\log(1 + \pi_{t+1})$ from both sides of the above equation, we derive the real interest rate rule as,

$$\log[1 + r(\pi_t, \pi_{t+1})] = \log(1 + r^n) + \phi_\pi \log(1 + \pi_t) + \phi_Y \log(\hat{Y}_t - 1) - \log(1 + \pi_{t+1}) \tag{12}$$

Linearizing the above expression around the target steady state, and using (2) to substitute for $\hat{Y}_t$, we obtain the approximate expression[9]

$$\log \frac{1 + r(\pi, \phi\pi)}{1 + r^n} \approx -\pi \left[\left(1 + \frac{\beta\phi_Y}{\kappa}\right)\phi - \phi_\pi + \frac{\phi_Y}{\kappa}\right]$$

[9] Around the target steady state $\pi_t = \pi_{t+1} = 0$ and $\hat{Y}_t = 1$, so that the linearized equation takes the usual form $\log \frac{1+r(\pi_t,\pi_{t+1})}{1+r^n} = \phi_\pi\pi_t + \phi_Y(\hat{Y}_t - 1) - \pi_{t+1}$.

Hence, irrespective of which negative value for $\pi \in (\pi^{LT}, 0)$ we pick, the superior $\bar{\phi}$ in (10) is found by setting to zero the difference $r^n(\pi, \bar{\phi}\pi) - r^n$, giving

$$\bar{\phi} = \frac{\phi_\pi + \frac{\phi_Y}{\kappa}}{1 + \beta \frac{\phi_Y}{\kappa}} \tag{13}$$

In conclusion, up to a first order approximation, if the nominal interest rate follows a Taylor rule then the maximum reaction $\bar{\phi}$ can be analytically expressed as a function of the Taylor parameters. This result is unaffected by the inclusion of a zero lower bound on the nominal interest rate rule[10]. As mentioned above, recall that Assumption 2 guarantees that we are only considering monetary policy rules for which the maximum reaction parameter satisfies $\bar{\phi} > 1$. For the special case (13), the condition $\bar{\phi} > 1$ takes a well-known significance: it is exactly equivalent to the necessary and sufficient condition that must hold for a Taylor rule to satisfy the Taylor principle (see Woodford [2003], p. 254). This is a way to show how monetary rules of the Taylor form, together with the parametric restrictions usually imposed on them, turn out to be just a special case of the more general rules and restrictions that I consider here.

3.2 The equilibrium manifold

We are now ready to characterize the dynamic path of the economy starting from initial sunspots $S_0 \in (S^{LT}, S^{TG})$,

Proposition 2. *Denote with S^{LT} the liquidity trap steady state. Then,*

i) There exists a continuous function $\mathcal{M}(\pi)$ satisfying $\pi^{LT} < \mathcal{M}(\pi) < \pi$ and such that, if the initial sunspot state $S_0 = (\pi_0, \pi_1)$ has $\pi_1 = \mathcal{M}(\pi_0)$, the equilibrium inflation π_t^* monotonically decreases to π^{LT}.

Take $S_0^* \in (S^{LT}, S^{TG})$ with $\pi_1^* = \mathcal{M}(\pi_0^*)$.

ii) If $1 < \bar{\phi} \leq \frac{1}{\beta}$, then along the equilibrium path $\hat{Y}_t^* < 1$ for all t and $\mathcal{M}(\cdot)$ satisfies

$$\pi_0^* - \mathcal{M}(\pi_0^*) \to 0 \qquad \text{for } \pi_0^* \to \pi^{LT} \tag{14}$$

Proof. See Appendix A. □

Figure 4 provides a graphical representation of the results in Proposition 2. The figure depicts the two steady states S^{LT} and S^{TG} alongside with two possible versions of the equilibrium manifold $\pi_{t+1} = \mathcal{M}(\pi_t)$.

A *sufficient* condition to obtain a *Type I* saddle connection is that $1 < \bar{\phi} < 1/\beta$, i.e. the monetary policy rule is "not too stimulative". In this case, any arbitrarily small sunspot deviation π_1^* of inflation from the target is associated with a monotonic deflationary path leading to the (dis-anchored) deflationary steady state. In addition, output is always below the target

[10] If the static Taylor rule is constrained by the zero lower bound, then the nominal rate cannot be decreased indefinitely. Hence, for a level of inflation π_t low enough, over such states the constrained rule becomes *less* "stimulative" than the unconstrained rule. Still, the value of $\bar{\phi}$ is not affected when the zero lower bound constrained is added, since $\bar{\phi}$ is calculated using a *sup* operator, i.e. it is calculated over the states where the rule is *most* "stimulative".

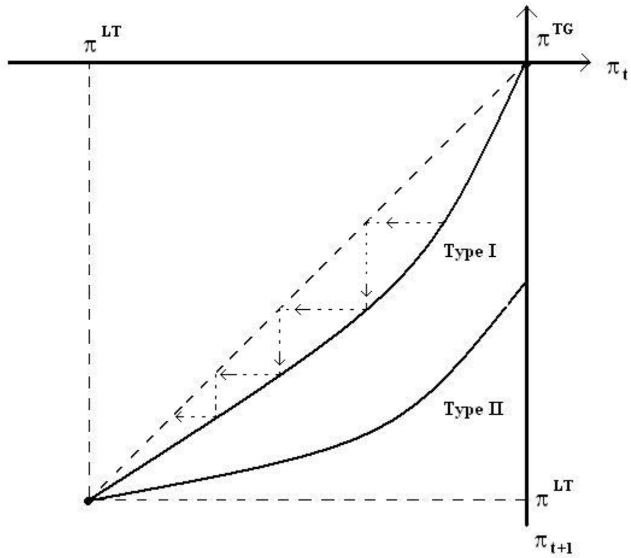

Figure 4: The two types of equilibrium manifold. The dashed line is the 45 degree line.

level. Notice that $1 < \bar{\phi} < 1/\beta$ is only a sufficient condition, meaning that *Type I* saddle connections *may* arise even for values $\bar{\phi} > 1/\beta$ (in this case, though, it is not guaranteed that $\hat{Y}_t^* < 1$ at all times). If the monetary rule is instead sufficiently stimulative, and thus $\bar{\phi}$ is large enough, then the equilibrium manifold is of *Type II.* In this case a discrete downward shift in inflation expectations π_1^* is required for a monotonic deflationary path to form.

Taylor rules: a special case. The assumption that the monetary authority follows a static Taylor rule underpins the work by Benhabib, Schmitt-Grohé, Uribe (2001), which was the first paper to prove the existence of a deflationary equilibrium path in a Neo-Keynesian sticky price model. Taylor rules are a special case of monetary policy rules that allows us to connect the general results of Proposition 2 to these seminal results on self-fulfilling deflations. As mentioned in the discussion about (13), any Taylor rule that also satisfies the Taylor principle is characterized by a $\bar{\phi} > 1$, and can be easily shown to satisfy Assumptions 1-2. For instance, to replicate the Taylor rule specification in Benhabib, Schmitt-Grohé, Uribe (2001), we would have to make the additional restriction that $\phi_Y = 0$, thus obtaining $\bar{\phi} = \phi_\pi$. In this case, Proposition 2 tells us that if ϕ_π is not only greater than 1, but is also large enough, then the equilibrium manifold is of *Type II.* In this special case, such manifold corresponds to the decreasing arm of the oscillatory path identified by Benhabib, Schmitt-Grohé, Uribe (2001)[11]. Instead, if ϕ_π is greater than 1 but small enough, then Proposition 2 guarantees that the equilibrium is of *Type I*, and thus no discrete jump in inflation expectations is needed for the monotonic deflationary path to form. The possible existence, in sticky price models, of a *Type I* saddle had not been previously identified in the literature, not even in the special setting

[11]If the stable manifold is of *Type II*, then the manifold clearly extends also to initial inflation levels $\pi > \pi^{TG}$. The characterization of this additional section of the manifold requires assumptions on the properties of $r(S)$ even at states that don't satisfy $S < S^{TG}$. Such analysis is beyond the scope of this paper, which is interested only in characterizing monotonic deflationary paths for initial sunspot expectations at or below the target.

of Benhabib et al. (2001). The existence of such paths was previously proved only when, in addition to the aforementioned Taylor assumptions, prices were assumed to be perfectly flexible (Benhabib et al. [2001], Cochrane [2011]). The quantitative analysis in Section 5 will further explore Taylor rule specifications and show that *Type I* saddle connections can arise for empirically relevant calibrations of the rule.

Proposition 2 proves that the existence of reasonably looking (i.e. monotonic) self-fulfilling deflationary paths is a general feature of sticky price model. The sufficient conditions in Assumptions 1-2 are not very restrictive and hold in a wide variety of cases, as for instance the following example indicates.

Example of an "unconventional" policy. Taylor rules certainly do no exhaust the set of monetary rules which, on the basis of optimality or of other arguments, one may be interested in considering. For instance, a central bank may be interested in knowing whether a new monetary rule can, hypothetically, rule out once and for all the possibility of self-fulfilling deflationary paths. All this raises the following question: once a rule has been proposed, how can we know whether it gives rise to monotonic deflationary paths? The usefulness of Proposition 2 is in providing a set of sufficient conditions that helps answering this question in an easy way. To see this, let us look at a simple example. Consider the following nominal rule

$$i(\pi_t, \pi_{t+1}) = r^n + \phi_\pi \pi_t + \phi_\Delta(\pi_t - \pi_{t+1})$$

for $\phi_\pi > 1$ and $\phi_\Delta > 0$. The proposed rule is like a standard Taylor rule but with the addition of a term $\phi_\Delta(\pi_t - \pi_{t+1})$ which responds to changes in inflation. The proposed rule is "unconventional", in the sense that there always are states where current and future inflation are strictly below the target $\pi^{TG} = 0$, and yet the rule calls for an *increase* in the nominal rate above r^n. In particular, this happens anytime the expected fall $\pi_t - \pi_{t+1}$ in inflation is large enough. Schmitt-Grohé and Uribe (2013) have recently advocated in favor of a (discontinuous) rule that *raises* the nominal interest rate once the economy is in a liquidity trap.

Is the rule above effective in preventing the economy from falling in monotonic deflationary paths? Answering this question in a traditional way would require trying to characterize the set of solutions to the nonlinear system of difference equations (1)-(3) under this new, previously unexplored, rule. This task may be hard (Fernandez-Villaverde [2014]), but Proposition 2 provides us with a simple shortcut. In fact, it is easy to verify that, under the usual ZLB constraint, the proposed rule satisfies Assumptions 1-2, and the only steady states are, as usual, π^{TG} and π^{LT}. We thus have the answer to our question, i.e. we now know that the proposed "unconventional" rule is again subject to self-fulfilling and monotonic deflationary paths for any initial sunspot $\pi^* \in (\pi^{LT}, \pi^{TG})$. We can also say something more once, by using the usual approximation technique, we compute the *maximum reaction* parameter as

$$\bar{\phi} = \frac{\phi_\pi + \phi_\Delta}{1 + \phi_\Delta} > 1$$

For any possible value of $\phi_\pi > 1$, there always exists a ϕ_Δ large enough for which $\bar{\phi} < \beta^{-1}$ which, by Proposition 2, implies that for large ϕ_Δ the saddle connection is of *Type I*. We can therefore conclude not only that a monetary authority who adopts the "unconventional" rule proposed above does not eliminate the possibility of self-fulfilling deflationary equilibria. But, in addition, we can also say that if its "unconventional" response ϕ_Δ is aggressive enough, then the insurgence of such paths may be facilitated.

The mechanics of self-fulfilling deflations. It is worth concluding this section with some intuition on why a monetary policy rule that satisfies the properties laid out in Assumptions 1-2 allow the emergence of deflationary equilibrium paths. Suppose that from time $t+1$ onwards inflation is *anchored* at the target level, which implies, in particular, that $\hat{Y}_{t+1} = 1$. Can a deflation at time t temporarily push the economy to a sub-optimal level (i.e. $\hat{Y}_t < 1$)? The answer is no, as long as monetary policy is sufficiently "stimulative of current demand". In fact, if $\hat{Y}_t < 1$ and simultaneously $r_{t+1} < r^n$, then the left hand side of (1) would be strictly bigger than one, while the right hand side would be smaller than one. Hence, as long as $\hat{Y}_{t+1} = 1$ and monetary policy stands ready to react in a sufficiently stimulative way, no equilibrium can form with $\hat{Y}_t < 1$. The target steady state is dynamically unstable and thus, locally, it is the unique equilibrium (Woodford [2003]).

However, as King (2000) points out, in the New Keynesian framework, which incorporates the rational expectations framework, «macroeconomic analysis can[not] be conducted by simple curve-shifting», since it is «necessary to solve simultaneously for current and expected future variables». This means that, in our example, the expected value of output $\hat{Y}_{t+1}$ has to be solved simultaneously with the rest of the equilibrium variables. In particular, consider again the possibility that $\hat{Y}_t < 1$. If $r_{t+1} < r^n$ then, again, monetary policy is stimulating the *relative* demand of goods at time t. However, differently from before, assume now that this intertemporal stimulus, instead of being associated with an increase in current demand $\hat{Y}_t$, is associated with a self-fulfilling pessimistic expectation of reduction in future demand, i.e. $\hat{Y}_{t+1} < \hat{Y}_t < 1$. An equilibrium of this type, where now $\hat{Y}_{t+1}$ is away from the target, can indeed exist. In conclusion, once we select an equilibrium with expectations *dis-anchored* from the target, then a sequence of low real (and nominal) interest rates is endogenously consistent with a series of pessimistic expectations about the evolution of future demand and thus, in turn, with an expected fall in inflation.

3.3 Fiscal policy

When the nominal interest rate has already reached the zero lower bound, the central bank cannot further reduce short-term rates in reaction to an ongoing deflation. In these cases, as mentioned in Section 2, it is often argued than fiscal policy could play an useful role in complementing the insufficient "stimulative" thrust of monetary policy. This section is devoted to showing instead that monotonic self-fulfilling deflationary paths are not eliminated when "stimulative" fiscal policies are brought into the picture. Rather, they can possibly be reinforced. In particular, when fiscal policy is assumed to be time-consistent, then the stable manifold is always of *Type I*, irrespective of how stimulative monetary policy is.

Fiscal policies, in the form of either distortionary taxation on consumption and production inputs, or in the form of wasteful government spending, have the effect of introducing "wedges" (Chari, Kehoe, McGrattan [2007], but see also Correia, Farhi, Nicolini, Teles [2013]) which distort the dynamics of the system (1)-(3). As I show in the next section, such distortions show up, in a microfounded model, in the form of time-varying parameters $\tilde{\beta}_{t+1} = \beta\omega^d_{t+1}$ and $\tilde{\kappa}_t = \kappa\omega^s_t$ which replace, respectively, the constants β and κ in the system (1)-(3). I call the variables ω^d_{t+1} and ω^s_t, respectively, the *dynamic* and *static wedges* at time t induced by fiscal policy in the following way

$$\begin{aligned}\omega^d_{t+1} &\equiv \frac{1+\sigma^G_{t+1}}{1+\sigma^G_t}\frac{1+\tau^C_t}{1+\tau^C_{t+1}} \\ \omega^s_t &\equiv \frac{1+\tau^C_t}{(1+\sigma^G_t)(1-\tau^N_t)}\end{aligned} \tag{15}$$

where τ^C_t is the tax rate on consumption, τ^N_t is the tax rate on labor (the only input in production), and σ^G_t indicates the amount of government consumption expressed as a share of private consumption.

When fiscal policy is conducted in the form of rules then fiscal wedges can be expressed as a function of the state S_t of the economy, i.e. we can write $\tilde{\beta}_{t+1} = \beta\omega^d_{t+1} = \beta\omega^d(S_t)$ and $\tilde{\kappa}_t = \kappa\omega^s_t = \kappa\omega^s(S_t)$. These two last sets of equalities are added to the original three equations[12] (1)-(3) to yield a five-equation model. After appropriately substituting the variables r_{t+1}, $\tilde{\beta}_{t+1}$ and $\tilde{\kappa}_{t+1}$ for the three corresponding rules, we again end up with a two-equation system

$$\frac{\hat{Y}_{t+1}}{\hat{Y}_t} = \beta\omega^s(S_t)[1+r(S_t)] \tag{16}$$

$$\pi_t = -\chi + \kappa\omega^s(S_t)\hat{Y}_t + \beta\omega^d(S_t)\pi_{t+1} \tag{17}$$

Again, given an initial state S_0 the evolution of the economy is uniquely characterized by (16)-(17). Once more, in order to be able to say something about the dynamic path of the economy, we need to impose some structure on the policies $\omega^d(\cdot)$ and $\omega^s(\cdot)$.

Let us focus first on the dynamic wedge. First of all, notice that monetary policy and fiscal policy interact in equation (16) simply through the product

$$\omega^d(S)[1+r(S)] \equiv 1+\tilde{r}(S)$$

where we may call $\tilde{r}(S)$ the effective real interest rate. From the viewpoint of the "aggregate demand equation", as sometimes is called equation (16), there is no difference between a government intervention using the monetary policy tool $r(S)$ or using the fiscal policy tool $\omega^d(S)$. This point is made very clearly, for instance, in Correia, Farhi, Nicolini, Teles (2013). It makes then sense to require at the very least that the two policy rules $r(S)$ and $\omega^d(S)$ do not conflict with each other with respect to their effect on $\tilde{r}(S)$, which is all that matters for "aggregate demand" management. In particular, we may want to avoid setting up a rule for $\omega^s(S)$ such that if monetary policy tends to be "stimulative", in the sense that it tends to lower $\tilde{r}(S)$ at some S, fiscal policy goes the other direction by contributing to increasing $\tilde{r}(S)$. To guarantee that the fiscal and monetary policy are complementary to each-other in dealing with states $S \in (S^{LT}, S^{TG})$ where inflation is inefficiently below target, it is then natural to require that Assumptions 1-2 are inherited also by the rule $\omega^d(S)$.

Assumption 3. *At any state $S_t < S^{TG}$, the dynamic wedge $\omega^d(S_t)$ is continuous and increasing in π_t and is continuous in π_{t+1}. Moreover, $\omega^d(S)$ is differentiable at any S where $r(S) = r^n$. Finally, $\omega^d(S) \le 1$ for all $S \in (S^{LT}, S^{TG})$.*

[12] Where now $\tilde{\beta}_{t+1}$ and $\tilde{\kappa}_t$ are written in place of β and κ.

To guarantee that S^{TG} and S^{LT} are still steady states, we have of course to posit that $\omega^d(S) = 1$ for $S \in \{S^{LT}, S^{TG}\}$. Notice, finally, that the assumption that $\omega^d(S)$ is smaller than its steady state value of 1 for *all* $S \in \{S^{LT}, S^{TG}\}$ is stronger than the corresponding Assumption 2 for $r(S)$. This stronger restriction is rather inconsequential on the main results, and is done for analytical simplicity.

Fiscal policy rules that induce a dynamic wedge that satisfies Assumption 3 can obviously take a wide variety of forms. Two examples are particularly interesting, since they mirror some types of fiscal policy prescriptions that in the literature have been proposed as a means to "stimulate" the economy and fight a deflation. Assume that the economy is following the deflationary path, highlighted in the previous section, where π_t monotonically decreases towards π^{LT}. First, consider the case where the government tries to fight the ongoing deflation through an anticipated increase in spending σ^G. As time goes by, the fiscal impulse is re-absorbed, so that σ_t^G progressively decreases. Along a path with falling inflation, the decreasing path for σ_t^G is observationally equivalent to assuming that the government follows a rule $\sigma_t^G = \sigma^G(\pi_t)$ such that $\sigma^G(\cdot)$ is an increasing function. Such type of rule induces in (15) a rule for $\omega^s(S)$ which indeed satisfies Assumption 3. Similarly, take $\sigma_t^G = 0$, and consider the fiscal stimulus policy that calls for a temporary cut in the consumption tax τ_t^C, as for instance in Correia, Farhi, Nicolini, Teles (2013). Observationally, this is equivalent to assuming that, along the path with monotonically decreasing inflation, the consumption tax rule $\tau_t^C = \tau^C(\pi_t)$ is a decreasing function of π_t. Again, the induced dynamic wedge satisfies Assumption 3.

Having defined some basic properties for $\omega^d(\cdot)$, we are only left with characterizing the static wedge policy $\omega^s(\cdot)$. This can be done in a number of ways, and for various cases which are particularly intuitive and, possibly, of practical relevance, we can prove following result.

Proposition 3. *If the maximum reaction $\bar{\phi}$ calculated with respect to the appropriate effective real rate $\tilde{r}(S)$, then Proposition 2 holds under any of the following assumptions for the labor tax policy:*

a) the labor tax policy $\tau^n(\pi_t)$ is set so that the static wedge in (15) is constant.

b) the labor tax policy $\tau^n(\pi_t)$ is a weakly decreasing function of π_t.

Moreover, if the labor tax policy $\tau^n(\pi_t)$ is set in a time-consistent way, then the saddle connection is always of Type I *for any finite-valued $\bar{\phi}$.*

Proof. See appendix B. □

Adding fiscal policy to the model complicates the form of the equilibrium system of equations, which in general now feautures a time-varying dynamic wedge in both equation (16) and (17), and a time varying static wedge appearing only in (17). Notwithstanding these complications, Proposition 3 tells us that if fiscal policy satisfies some "minimally stimulative"[13] sufficient conditions, then the general and simple characterization of the saddle connections in Proposition 2 is still valid. This can be proven, in particular, under the assumption that

[13] One may argue that the assumption that the labor tax is not reduced as deflation deepens ($\tau^t(\pi_t)$ is a decreasing function) is an unreasonable way to conceive a stimulative fiscal policy. Yet, in standard the New-Keynesian framework that I am using, this is exactly the type of labor tax policy that, by fighting a "paradox of toil", has expansionary effects at the zero lower bound when inflation expectations are anchored (Eggertsson [2010]).

the labor tax is either not increased while inflation falls, or it is set to keep a constant static distortion, or it is set in a time-consistent way. In the latter case, we can actually prove a stronger result, i.e. that the saddle connection is always of *Type I*, regardless the value of $\bar{\phi}$.

4 The microfounded model

The equilibrium equations analyzed so far turn out to be a special case of those derive from the richer setting I now consider. Specifically, this section contains the description and equilibrium equations of a standard New-Keynesian model with Rotemberg sticky pricing. The proof in Appendix A shows that the results from the previous sections carry over to this richer environment.

4.1 The problem of the agents

There are three classes of economic agents: households, a government and firms.

Household. The household maximizes its lifetime utility

$$\sum_{t=0}^{\infty} \beta^t u(C_t, N_t) \tag{18}$$

$$u(C_t, N_t) = \ln C_t - \frac{N_t^{1+\psi}}{1+\psi}$$

$$C_t = \left[\int_{[0,1]} c_t^{\frac{\theta-1}{\theta}}(j) dj \right]^{\frac{\theta}{\theta-1}}$$

subject to the budget and cash constraints

$$(1+\tau_t^C)P_tC_t + B_{t+1} + M_{t+1} = M_t + (1+i_t)B_t + \frac{\theta}{\theta-1}(1-\tau_t^N)W_tN_t + \Pi_t - P_tT_t$$

$$\frac{M_{t+1}}{P_{t+1}} \geq m_{t+1}$$

with $m_t \geq 0$ an exogenous bounded sequence of real balances, and total spending P_tC_t on each of the consumption good variety j is defined by $P_tC_t = \int_{j\in[0,1]} p_t(j)c_t(j)dj$. Distortionary tax wedges on consumption and labor are, respectively, $1+\tau_t^C$ and $\frac{\theta}{\theta-1}(1-\tau_t^N)$. Lump sum taxation is T_t, Π_t are profits received by the household from firms, and B_t is then nominal government debt purchased by the household.

As we shall see, in equilibrium all firms charge the same price P_t, therefore the household's optimal demand of for variety j produced by a (deviant) firm charging the price $p_t(j)$ is given by

$$\frac{c_t(j)}{C_t} = \left[\frac{P_t}{p_t(j)} \right]^{\theta} \tag{19}$$

Since in equilibrium we will have $p_t(i) = P_t$ and then $c_t(i) = C_t$, the optimal intertemporal allocation of consumption requires that the following Euler equation holds

$$\frac{C_{t+1}}{C_t} = \beta \frac{1+\tau_t^C}{1+\tau_{t+1}^C}(1+r_{t+1}) \tag{20}$$

where $1 + r_{t+1} = \frac{P_t}{P_{t+1}(1+i_{t+1})}$ is the real interest rate between time t and time $t+1$.

The household's optimal choice of time allocation between leisure and labor gives the labor supply schedule

$$C_t N_t^{\psi} = \frac{\theta}{\theta - 1} \frac{(1-\tau_t^N)W_t}{(1+\tau_t^C)P_t} \tag{21}$$

Optimal monetary holdings are given by

$$\frac{M_{t+1}}{P_{t+1}} = m_t \qquad i_{t+1} > 0 \tag{22}$$

while the combination of M_t and B_t is undetermined if $i_{t+1} = 0$. As underlined below, the money quantity M_t will play no role in our analysis so that, following much of the current Neo-Keynesian literature, we may even consider a cashless economy where $M_t = 0$. Still, to stress the assumed Ricardian nature of the economy, it is instructive to explicitly maintain, at least for the moment, the presence in the equilibrium equations of the money quantity. In particular, to guarantee the optimality of the household's consumption plan, I impose the following transversality condition to the problem,

$$\lim_{t\to\infty} \beta^t u_{C,t} \frac{B_t + M_t}{P_t} = 0 \tag{23}$$

Equations (20), (21) and (22) provide, together with the individual demand (19) and the transversality condition (23), the necessary and sufficient conditions for the optimality of the household's plan.

Government. I consider a Ricardian environment where government policies are divided into two groups, *instrumental* and *non-instrumental*. Instrumental policies comprise the nominal interest rate i_{t+1}, the government spending on final goods, and the distortionary taxes (on consumption and labor) levied. Instrumental policies are those that the government uses to directly influence the equilibrium of the economy. In particular, government's aggregate purchases of final goods are given by

$$G_t = \left[\int_{[0,1]} g_t^{\frac{\theta-1}{\theta}}(j)dj \right]^{\frac{\theta}{\theta-1}}$$

while purchases of the individual variety j are

$$\frac{g_t(j)}{G_t} = \left[\frac{P_t}{p_t(j)} \right]^{\theta} \tag{24}$$

Conversely, non-instrumental policies are either set arbitrarily or are just an endogenous by-product of instrumental policies. Specifically, the set of non-instrumental policies comprises the supply M_t^s and B_t^s of money and of one period bonds, and the levying of lump sum taxation T_t. I assume that, given the evolution of prices P_t, the stock of money issued is set to

$$M_t^s = \bar{m} P_t \tag{25}$$

for given constant real balances $m_t = \bar{m} \geq 0$. Similarly, I assume that the nominal government debt B_t^s issued equals a constant $\bar{B} \geq 0$. Lump sum taxes T_t are raised so that the following budget constraint of the government is satisfied period by period

$$P_t G_t + i_t B_t^s + \left[\frac{\theta}{\theta - 1}(1 - \tau_t^N) - 1\right] W_t N_t = B_{t+1}^s + \tau_t^C C_t + P_t T_t + M_{t+1}^s - M_t^s$$

I will return on the distinction between instrumental and non-instrumental policies in Section 4.2. For now, it is enough to point out that this distinction qualifies the model as Ricardian. The model is Ricardian in the sense that the particular mix between debt issuance and lump-sum taxation is irrelevant for the equilibrium of the economy. This justifies fixing, as I did, an arbitrary evolution of nominal debt, while letting lump sum taxes T_t automatically adjust to cover the higher (lower) interest rate expenditure stemming from an arbitrary higher (lower) level of public debt $\bar{B}$. Second, the paper studies how different rules for the nominal interest rate instrument i_{t+1} influence the equilibrium evolution of the price level. Given the price level, the quantity of money adjusts automatically, according to (25).

A somewhat subtler issue is whether, under the assumed evolution of the government's nominal liabilities B_t^S and M_t^S, the transversality condition (23) represents or not a redundant constraint for the equilibrium of the economy. To settle this point, which is beyond the scope of this paper, I take $\bar{B} = 0$ which, together with (25), guarantees that the transversality (23) is always automatically satisfied. [14]. The restrictions on the evolution of nominal public liabilities B_t and M_t are assumed, here, only to guarantee that the government does not commit to follow *at the infinity* policy that violate (23) and thus "blow up the economy". For more details see Cochrane (2011), who clarifies why, if the economy is assumed to behave in a Ricardian way (rather than, for instance, behaving in a non-Ricardian way as in Cochrane [1998]), then the use of the transversality condition as an equilibrium selection device (as in Benhabib, Schmitt-Grohé, Uribe [2002a]) is problematic.

Firms. Each firm in the production sector is indexed by the variety j it produces. Firms sell their output o households and to the government in a monopolistically competitive final goods market. Firms also sell part of their output as intermediate goods to other firms.

The production technology of each monopolistically competitive firm transforms n_t units of labor input into $y_t = A_t n_t^\alpha$ of output, with $\alpha \in (0, 1]$. Firms are subject to Rotemberg (1982) price adjustment costs[15]. In particular, a firm j has to buy $\Xi_t(j) Y_t$ units of intermediate goods if it wants to change its selling price from p_{t-1} to p_t, where $Y_t = C_t + G_t$ is the GDP and $\Xi_t(j) = \frac{\xi}{2}\left[\frac{p_t(j)}{p_{t-1}(j)} - 1\right]^2$ is a quadratic cost, for a strictly positive constant ξ.

[14] As long, of course that, C_t is strictly bounded away from zero, which will always be the case in any equilibrium we consider.

[15] The model may also be formulated with price adjustment costs in the utility function (as in Benhabib, Schmitt-Grohé and Uribe M. [2001], rather than in the production set. The particular formulation of the Rotemberg adjustment costs is not crucial for the theoretical results obtained in the paper. Quantitatively, instead, the formulation with adjustment costs in the utility function allows for an equilibrium evolution of the labor input which is empirically more appealing, as emphasized in Section 5.

Given the price set by a firm, the quantity produced is demand-determined. As noted, total demand for a firm's output has two components. The first is the quantity $c_t(j) + g_t(j)$ demanded in the monopolistically competitive final goods market, and defined by (19) and (24). The second is the demand in the form of intermediate good by other firms, who use the goods purchased to pay for their adjustment costs. Since price adjustment costs are only a small share of GDP, I simplify the exposition and assume that firms disregard this part of their output demand when deciding the optimal price $p_t(j)$ at which to sell their output. This is equivalent to assuming that firms take as given the selling price P_t in the intermediate goods market, where equilibrium quantities are demand-determined, and demand $\Xi_t Y_t$ is assumed to be spread equally across firms. Hence, the problem of each firm is to choose sequences of prices $p_t(j)$ in order to maximize the firm's discounted flow of profits, which is equivalent to solving the following problem

$$\max_{\{p_t(j)\}} \sum_{t=0}^{\infty} \mathcal{D}_t \Pi_t(j) \tag{26}$$

where nominal period profits are given by

$$\Pi_t(j) = P_t \Xi_t Y_t + p_t(j) y_t(j) - W_t n_t(j) - \frac{\xi}{2} \left(\frac{p_t(j)}{p_{t-1}(j)} - 1 \right)^2 P_t Y_t \tag{27}$$

the discount factor is

$$\mathcal{D}_t = \prod_{s=1}^{t} \frac{1}{1 + i_t}$$

and firm's labor demand satisfies $n_t(j) = \left[\frac{y_t(j) + \Xi_t Y_t}{A_t} \right]^{\frac{1}{\alpha}}$. In a symmetric equilibrium, the optimal price $p_t^*(j)$ is the same for all firms, hence $p_t^*(j) = P_t$ and $y_t^*(j) = Y_t$. The first order condition for the optimal price then gives

$$P_t = \left[\frac{\theta}{\theta - 1} \right] \cdot \left[\frac{1}{1 - v_t} \right] \cdot \left[\left(\frac{Y_t}{A_t} \right)^{\frac{1}{\alpha} - 1} \frac{(1 + \Xi_t)^{\frac{1}{\alpha} - 1}}{\alpha A_t} \right] W_t \tag{28}$$

where

$$v_t = \frac{\xi}{\theta - 1} \left[-\pi_t + \frac{Y_{t+1}}{(1 + r_{t+1}) Y_t} \pi_{t+1} \right] \tag{29}$$

With a harmless abuse of terminology[16], I refer to $\pi_t = \frac{P_t}{P_{t-1}} \left(\frac{P_t}{P_{t-1}} - 1 \right)$ as the inflation rate at time t. Equation (28) indicates that the mark-up charged by firms is the product of three components: the elasticity of substitution between final goods (first term in brackets), deviations of current and future inflation from the zero inflation target (second term in brackets), and the nominal marginal cost of production, which equals the inverse of labor productivity

[16] The true price inflation between time $t-1$ and time t is $\tilde{\pi}_t = \frac{P_t}{P_{t-1}} - 1 > -1$. We can then write the relation between π_t and $\tilde{\pi}_t$ as $\pi_t = \tilde{\pi}_t(1 + \tilde{\pi}_t)$. As long as $\tilde{\pi}_t > -50\%$, which for any practical purpose we assume to be the case in our analysis, there is a monotonically increasing relation between π_t and $\tilde{\pi}_t$. Moreover, the sign of π_t is always identical to the sign of $\tilde{\pi}_t$ and, up to a first order approximation, there is no difference between π_t and $\tilde{\pi}_t$ in a neighborhood of $\tilde{\pi}_t = 0$. Overall, in the context of this paper, there is no relevant loss of information in referring to π_t as the inflation rate.

(third term in brackets) at the level $(1+\Xi_t)Y_t$ of gross production, times the nominal wage. Finally, equilibrium demand of labor and profits are, respectively,

$$N_t^d = \left[\frac{Y_t}{A_t}(1+\Xi_t)\right]^{\frac{1}{\alpha}} \tag{30}$$

$$\Pi_t = P_t Y_t - W_t N_t^d \tag{31}$$

4.2 Equilibrium

Definition 2. *An equilibrium is a sequence of exogenous productivities* $\{A_t\}_{t=0}^{\infty}$*, prices* $\{P_t, W_t, i_{t+1}\}_{t=0}^{\infty}$*, household quantities* $\{C_t, L_t, M_t, B_t,\}_{t=0}^{\infty}$*, taxes and government spending* $\{\tau_t^N, \tau_t^C, T_t, G_t\}_{t=0}^{\infty}$*, government assets* $\{M_t^s, B_t^s\}_{t=0}^{\infty}$*, and firms' labor demand and profits* $\{N_t^d, \Pi_t\}_{t=0}^{\infty}$ *such that i) given prices, taxes, and profits, the household quantities solve the household problem; ii) the government is solvent at every period; iii) given taxes and wages and final demand* $Y_t = C_t + G_t$*, the evolutions of aggregate prices* P_t *and of labor demand* N_t^d *solve the firms' problem, giving* Π_t *as the maximal profits; iv) markets clear, i.e.* $M_t^s = M_t$*,* $B_t^s = B_t$*,* $N_t = N_t^d$*,* $C_t + G_t = Y_t$*.*

There are two crucial equations that characterize any equilibrium path for the economy. The first is the Euler equation (20). The second equation is the forward-looking Phillips curve, which incorporates the optimal pricing condition for firms and is derived as follows. Define $\sigma_t^G = \frac{G_t}{C_t}$ the ratio of government spending over private spending. It is then just a matter of simple algebra[17] to derive the forward-looking Phillips curve as

$$\pi_t = -\chi + \tilde{\kappa}_t \hat{Y}_t^{\frac{1+\psi}{\alpha}} + \tilde{\beta}_t \pi_{t+1} \tag{32}$$

where $\hat{Y}_t \equiv \frac{Y_t}{Y_t^{TG}}$ with $Y_t^{TG} = \alpha^{\frac{\alpha}{1+\psi}} A_t$, and

$$\begin{aligned} \chi &\equiv \frac{\theta - 1}{\xi} > 0 \\ \tilde{\kappa}_t &\equiv \chi (1+\Xi_t)^{\frac{1+\psi}{\alpha} - 1} \frac{1+\tau_t^C}{(1+\sigma_t^G)(1-\tau_t^N)} \\ \tilde{\beta}_t &\equiv \beta \frac{1+\sigma_{t+1}^G}{1+\sigma_t^G} \frac{1+\tau_t^C}{1+\tau_{t+1}^C} \end{aligned} \tag{33}$$

Notice that, because of our modified definition of inflation as $\pi_t = \frac{P_t}{P_{t-1}}\left(\frac{P_t}{P_{t-1}} - 1\right)$, equation (32) is linear in both inflation and output deviations. This is important, because employing the actual forward-looking Phillips curve allows us to look at the global stability properties of the economy. Finally, substituting for the definition of $\hat{Y}_t$ and $\tilde{\beta}_t$, we can re-write the (20) as

$$\frac{\hat{Y}_{t+1}}{\hat{Y}_t} = \frac{\tilde{\beta}_t}{\gamma}(1 + r_{t+1}) \tag{34}$$

[17] To obtain the forward-looking Phillips curve start by eliminating the real wage W_t/P_t in the labor supply (21) using the value of the real wage derived from the labor demand (30). Then notice that consumption is written as $C_t = Y_t/(1+\sigma_t^G)$, and that by (20) we can write $\tilde{\beta}_t = \frac{Y_{t+1}}{Y_t(1+r_{t+1})}$.

Parameter	Description	Calibration
ξ	Adjustment cost coefficient	50
α	Labor share in production	0.7
$\frac{1}{\psi}$	Frisch elasticity	1
θ	Demand elasticity	10

Table 1: Calibration parameters. The model is calibrated in annual terms.

where $\gamma \equiv \frac{A_{t+1}}{A_t}$ is the productivity growth rate, which I take to be constant. It is straightforward to see that the aggregate equilibrium equations presented in Section 3 are obtained for $\psi = 0$ and $\alpha = \gamma = 1$. As the proofs to the propositions in Section 3 show, none of the theoretical results in Section 3 hinges on this particular calibration of the parameters.

5 Calibration

The calibrations in this section explore the quantitative properties of the self-fulfilling path to the liquidity trap steady state[18]. The model is calibrated in annual terms, with the parameters in Table 1 common across all exercises, while the values of subjective discout factor β and of productivity growth γ_t vary across calibrations. All calibrations assume that monetary policy is conducted according to a Taylor rule (12)[19]. As shown in the theoretical part of the paper, the qualitative characteristics of the deflationary path do not depend on this specific assumption. Yet the choice of focusig on Taylor rules is justified here since, in order to coherently compare the model to actual data, it is proper to assume rules that have, arguably, the ability to replicate actual past policy behaviors (Taylor [1993]).

The section presents four calibrations. The first explores the empirical relevance of *Type I* versus *Type II* saddle connections under the Taylor framework. The second looks at the full dynamic of the main macro variables under a *Type I* saddle connections. The third calibration is the most comprehensive one and assess the performance of the model in explaining the Japanese "lost decade" 1992-2002. The fourth calculates fiscal multipliers at the zero lower bound along a (dis-anchored) deflationary path.

Type I and Type II manifolds. Under the additional calibration $\beta = 0.97$ and $\gamma_t = 0$, Figure 5 depicts the equilibrium manifold for two different parametrizations of the Taylor parameters. By presenting the manifold in the $(\tilde{\pi}_t, \hat{Y}_t - 1)$ space, Figure 5 provides information on both inflation $\tilde{\pi}_t = P_t/P_{t-1} - 1$ and on the net output gap $\hat{Y}_t - 1$. By virtue of (12), the rule calibrated on the left-hand side of the figure yields a maximum reaction parameter of $\bar{\phi} = 1.5$ which, in the language of this paper, is much more stimulative than the calibration on the

[18]The algorithm used for the numerical calculation is based on finding the saddle path solution to equation (35) in Appendix A. The algorithm exploits the fact that the value of π_{t+2} can be expressed explicitly as a function of π_{t+1} and π_t. Recall that the initial sunspot π_0 is chosen exogenously. Solving (35) given the initial condition π_0 then means finding the unique value for π_1 such that the uniquely identified sequence π_0, π_1, π_2, ... converges monotonically to π^{LT}. Given π_0, the algorithm then searches for the value π_1 over the range (π^{LT}, π_0) such that the corresponding sequence of π_t is a) monotonically decreasing; b) as time grows large it lays between a radius of 0.005% from π^{LT} (by comparison, the absolute value of π^{LT} typically equals 3%).

[19]In all exercises and accompanying discussions actual inflation $\tilde{\pi}_t$ now replaces the convenient variable π_t.

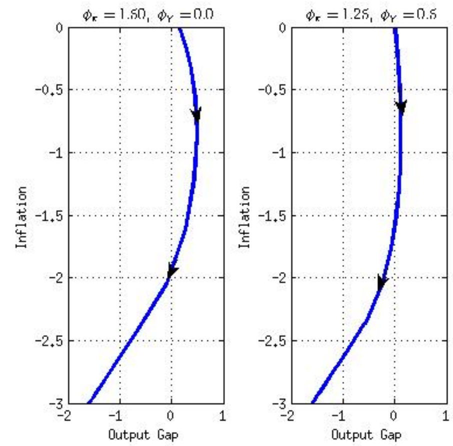

Figure 5: Equilibrium manifold in the $(\tilde{\pi}, \hat{Y} - 1)$ space. The manifold is of *Type II* for $\phi_\pi = 1.50$ and $\phi_Y = 0$, while it is of *Type I* for $\phi_\pi = 1.25$ and $\phi_Y = 0.5$. The subjective discount factor and productivity growth are, respectively, calibrated to $\beta = 0.97$ and $\gamma_t = 0$.

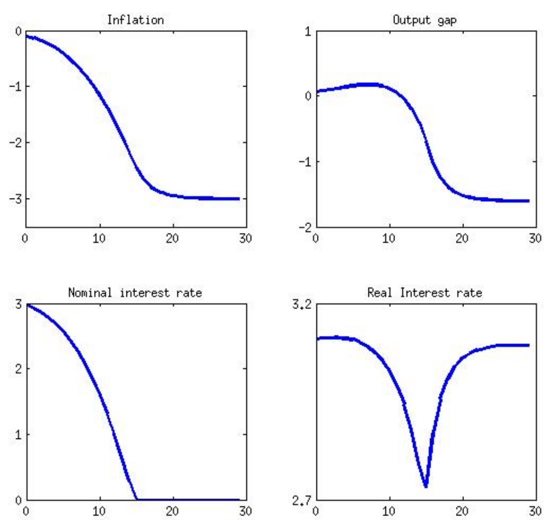

Figure 6: Dynamics of the *Type I* manifold ($\phi_\pi = 1.25$, $\phi_Y = 0.5$). The figure depicts the evolution of inflation $\tilde{\pi}_t$, the net output gap $\hat{Y}_t - 1$, the nominal rate i_t and the equilibrium real rate r_t starting from a sunspot $\tilde{\pi}_0 = -10^{-3}$. The subjective discount factor and productivity growth are, respectively, calibrated to $\beta = 0.97$ and $\gamma_t = 0$. Values in the figure are expressed in percentage points.

right-hand side, which instead yields $\bar{\phi} = 1.065$. Consistently with our theoretical results, the calibration with $\bar{\phi} = 1.5$ gives rise to a *Type II* manifold while the "not too stimulative" calibration with $\bar{\phi} = 1.065$ generates a *Type I* saddle connection. This conclusion can be drawn by observing that only in the right-hand side case the saddle connection reaches continuously, i.e. without a final jump, the target steady state $(\tilde{\pi}^{TG}, \hat{Y}^{TG} - 1) = (0, 0)$. Notice also that the *Type I* saddle connection arises even though $\bar{\phi} = 1.065$ is slightly above the *sufficient* condition threshold $\beta^{-1} = 1.031$. Since in this case $\bar{\phi} > \beta^{-1}$, it is not guaranteed that $\hat{Y}_t - 1 < 0$ for all t, and in fact output gaps are marginally positive when inflation is not to0 far away from the target.

Dynamics along a Type I saddle. Figure 6 presents the time evolution of the main macro variables along the *Type I* saddle connection identified in Figure 5, assuming an initial sunspot inflation $\tilde{\pi}_0^* = -0.1\%$, which is thus very close to the target. Starting from such value, it takes approximately 20 years for the inflation rate $\tilde{\pi}_t$ to be close to its liquidity trap steady state value. In about 15 years the nominal interest rate hits the zero lower bound. Overall, the deflationary path is slow, qualitatively replicating the features of a "creeping" deflation. With an output gap of roughly zero for the first 10 years, the monetary authority keeps lowering the nominal rate as inflation slowly falls. Since the natural real rate is calibrated to a constant value the equilibrium real interest rate, while moving quantitatively little, follows a V-shape path, with its initial and final values coinciding with an annual $r^n = \beta^{-1} - 1 = 3.1\%$.

Simulating the Japanese case: expectations dis-anchoring and a falling natural real rate. A casual observation of Figure 1, confirmed by a more thorough analysis, suggests that the equilibrium real interest rate in Japan has fallen in a roughly monotonic way since the early '90s. This is in sharp contrast with the simulation in Figure 6 where, as discussed above, the constancy of the natural real rate forces V-shape dynamics for the equilibrium real rate. A falling natural real rate seems to be a crucial element for any model that attempts to replicate the Japanese experience. Since the natural real rate depends on potential growth, I turn to Hayashi and Prescott (2002) for their analysis of growth in Japan. In particular, I take advantage of their approach to ask the following question: can a real shock to potential growth, jointly with a sunspot shock to inflation expectations, account for the evolution of both real and nominal variables during the lost decade 1992-2002?

Hayashi and Prescott (2002) report that TFP growth in Japan fell from 3.7% during the booming years of the late '80s, to just 0.3% during the '90s. To roughly replicate this observation I calibrate the TFP growth rate in the year $t = 1990$ to $\gamma_t = 3.7\%$ and I further assume that subsequently TFP growth falls monotonically and reaches a long-run value of 0.3%, according to the law

$$\gamma_t = 0.3\% + 3.4\% \cdot 2^{-\frac{t-1990}{2}}$$

The top-left panel in Figure 7 shows that the exogenous TFP process described reaches the the desired value of just above 0.3% by the year 2000.

In addition, Hayashi and Prescott (2002) also calculate that in 1990 the post-tax real return on capital was 5%. For the same year, the difference in Figure 1 between the uncollateralized overnight rate and the inflation rate gives a real return of 4.3%. I take the average of these two numbers, and I set to 4.65% the value of the natural real rate in 1990. Using the condition $1 + r^n_{1990} = (1 + \gamma_{1990})/\beta$ I can then back out a calibration of $\beta = 0.991$ for the subjective

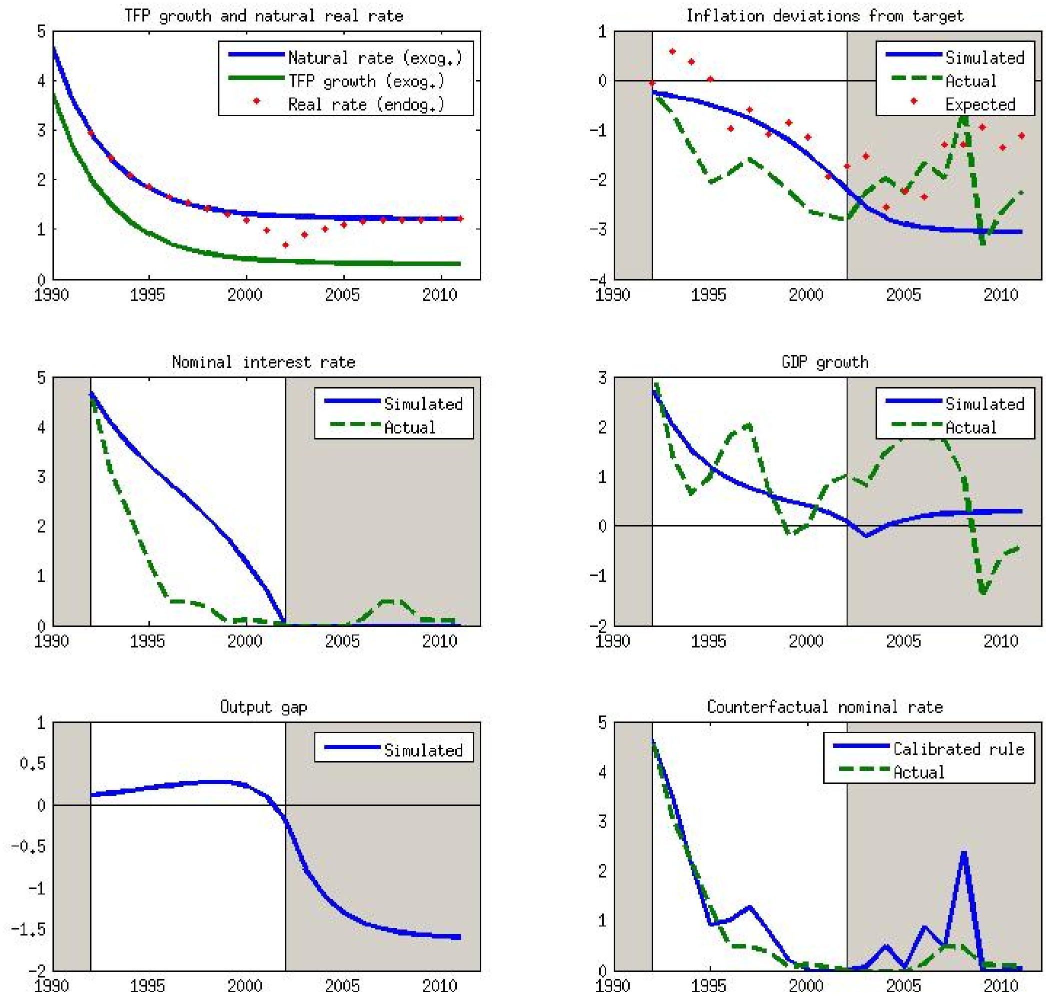

Figure 7: A calibrated model for the Japanese lost decade 1992-2002. The "potential" growth rate γ_t falls from 3.7% in 1990 towards 0.3%, following a process with constant half-life equal to 2 years. The subjective discount factor is calibrated to $\beta = 0.991$, the initial (normalized) inflation sunspot is $\pi_0^* = -0.24\%$, and the Taylor rule has parameters $\phi_\pi = 1.38$, $\phi_Y = 0.5$. Values in the figure are expressed in percentage points.

discount factor. Given β and the evolution of γ_t, the top-left panel of Figure 7 depicts the exogenous evolution of the calibrated natural real rate. To check whether the entire path for r_t^n, and not just its initial value, is reasonable I compare the calibrated value of r_t^n against the data in the year $t = 2000$, i.e. ten years after the initial calibration date. It turns out that the calibrated value of $r_{2000}^n = 1.28\%$ falls exactly between the 2% value calculated by Hayashi and Prescott (2002) for the year 2000 and the 0.8% computed, as above, using the overnight uncollateralized rate. Similar conclusions are reached if the data are compared to the equilibrium real interest rate generated by the model (dotted line in the top left-hand side panel of Figure 7).

Having calibrated the productivity shock, we need to calibrate the sunspot shock to inflation expectations. To do this, we first have to define a target inflation rate in the early '90s in Japan. This is a problematic task since the Bank of Japan was not operating under an explicit target at that time. I deal with this issue by assuming that the inflation target was 2%, a value consistent with both the inflation targets adopted at the time by various advanced countries' central banks, and with the current Bank of Japan's target. With this target in mind, I set the sunspot inflation at the beginning of the lost decade to a small deviation $\tilde{\pi}_{1992} = -0.24\%$ from the target. I take 1992 as the first year of the deflationary path because 1992 was in fact the first when actual inflation (1.76%) fell below the 2% target. We also need to define a monetary policy rule and calibrate its parameters. Once more, we face the issue that the Bank of Japan did not explicitly follow a monetary policy rule. I resort to the assumption, common in the literature, that the BoJ was following a standard Taylor rule[20]. I calibrate the reaction parameter to the output gap, which is an unobservable variable, to the standard value of $\phi_Y = 0.5$. Given this, the reaction parameter ϕ_π is calibrated using observables. Specifically I set ϕ_π so that, given the initial inflation deviation $\pi_{1992} = -0.24\%$, the endogenous nominal rate i_{1992} generated by the model matches exactly the observed value of 4.66% for the uncollateralized overnight rate in 1992. This calibration strategy for the monetary policy rule yields a parameter value of $\phi_\pi = 1.38$, which is "not too stimulative" relative to a standard calibration of 1.5 as in Taylor (1993).

Figure 7 compares to the data the deflationary path obtained from the calibrated model. The period of interest, for which the calibration is relevant, is the *lost decade* 1992-2002. As mentioned in the introduction, this decade is defined by the peak-to-trough monotonic fall in the inflation rate, and by a substantial reduction in GDP growth[21]. The model matches almost perfectly the timing of the trough reached by actual inflation and it matches exactly the year when the nominal rate reached the zero lower bound. Notice also that simulated inflation is, for most of the lost decade, very close to measured expected inflation but is higher than actual inflation. This makes sense in light of my interpretation of Figure 3. Recall, in fact, that the model is simulated under a perfect foresight assumption, i.e. in the absence of *unexpected*

[20]The model can be easily modified to include a non zero target inflation rate. This simply amounts to a normalization of the inflation variables. In particular, call $\bar{\pi}^{TG}$ the target (net) growth rate of prices. Renormalize the actual inflation rate so that now $1+\tilde{\pi}_t = \frac{P_t}{P_{t-1}(1+\bar{p}i^{TG})}$. With this normalization, price adjustment costs continue to be expressed as $\Xi(\tilde{\pi}_t) = \xi/2\tilde{\pi}_t^2$. Similarly, the convenient variable π_t that we have used in Sections 3-4 continue to be defined as $\pi_t = \tilde{\pi}_t(1+\tilde{\pi}_t)$. Finally, under the new definition of the actual inflation $\tilde{\pi}_t$, the real policy rule derived from the Taylor rule is still expressed by (12), with π_t replacing, as usual, the convenient variable π_t.

[21]The decade 2003-2012 is also reported in the graphs. This decade is marked by important exogenous shocks (e.g. the rise in commodity prices prior to 2008, the global financial crisis, the Japan earthquake and tsunami of 2011) that are not incorporated in the model simulation.

shocks to inflation. A fair comparison between the model and the empirical values would then require us to filter out any unexpected inflation shock from the data or, equivalently, to compare the simulated inflation path to that of measured *expected* inflation rather than to actual inflation. When this latter comparison is done then, as Figure 3 shows, the model not only captures the overall and monotonic reduction in inflation, but it also matches the slow speed of such process. Another way to phrase this conclusion is to say that the model simulates the counterfactual path for actual inflation and the nominal interest rate in the absence of the negative and *global* sequence of inflation shocks that, among other countries, hit Japan during its lost decade (see the discussion in Section 2.3). Indeed, these global shocks can also fully account for the somewhat slower fall, relative to the data, of the simulated nominal interest rate. To see this, assume that the (unobservable) output gap was zero during the lost decade, as roughly suggested by the model's simulation (bottom-left panel of Figure 7). Under this assumption, we can plug the sequence of actual inflation into our calibrated Taylor rule, thus obtaining an implied sequence of nominal rates that now incorporate also the reaction of the monetary authority to the unexpected global inflation shocks. As the figure makes clear (bottom right-hand side panel), once these unexpected global shocks are factored in, the rule-implied nominal rate practically coincide with the data.

Overall the double-shock approach – a shock to potential growth plus a simultaneous sunspot dis-anchoring of inflation expectations – simulated by the model accounts quite well for the evolution of key variables, both real and nominal, during the Japanese lost decade. In particular, once the sequence of unexpected global shocks to inflation are removed from the data, the model allows to interpret the resulting underlying deflationary process in Japan as rooted in a self-fulfilling dis-anchoring of inflation expectations. Such dis-anchoring may have been facilitated by a monetary policy rule that was "not too stimulative".

Government spending at the ZLB. Consider again the calibration where the natural real rate is constant, with $\beta = 0.97$ and $\gamma_t = 0$, and set the initial arbitrary sunspot inflation deviation to $\tilde{\pi}_0 = -1\%$. Consider now the effect of carrying out an anticipated fiscal expansion while the nominal interest rate is kept fixed at the ZLB. Specifically, assume that government spending σ^G initially increases to 3% of private spending and then, over the following four years, progressively returns to zero.

The solid lines in Figure 8 depict the evolution of the equilibrium variables under the announced fiscal stimulus plan, while the dashed line represents the corresponding baseline path in the absence of stimulus, i.e. for $\sigma_t^G = 0$. The simulation indicates that, while on impact the fiscal expansions allows for a sharp improvement in the output gap, inflation is nonetheless always lower than without fiscal stimulus. This can be interpreted in light of the result of the theoretical part of the paper. As seen, an expansion in government spending creates a dynamic wedge that operates in a way very similar to a lowering of the nominal interest rate. Hence, we can apply to a spending expansion the same intuition as the one developed in the last two paragraphs of Section 3.1: a more "stimulative" policy today is associated with an increase in current *relative* demand, which translates into a faster fall in future demand, and thus into a faster fall in inflation.

The association of an anticipated fiscal expansion with a fall in the inflation rate is only apparently in stark contrast with other results in the literature (Eggertsson [2010], Christiano, Eichenbaum, Rebelo [2011]). There is no contrast because the two types of results are obtained from essentially the same model but under two different assumptions for the equilibrium

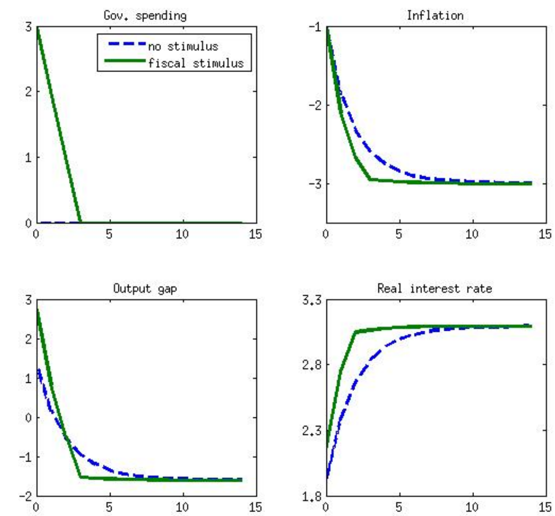

Figure 8: Dynamics with a fiscal expansion at the ZLB. The calibrated parameters are as in Table 1. The baseline scenario without stimulus is thus the same as the one in Figure 6, but holding fixed to zero the nominal interest rate. Values in the figure are expressed in percentage points.

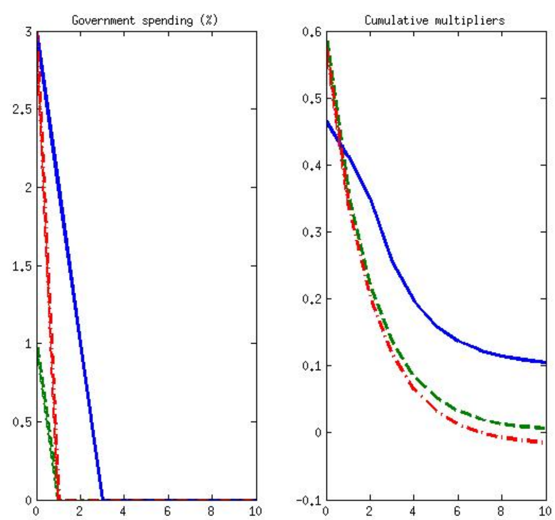

Figure 9: The left-hand side panel shows the path of the share of government spending σ^G starting from an inflation $\tilde{\pi}_0^* = -1\%$, while the right-hand side gives the corresponding cumulative spending multipliers at different horizons.

selection issue (Christiano, Eichenbaum [2012]). In particular, under the assumption that long-term inflation expectations are *anchored*, a fiscal expansion at the ZLB generates an increase in the inflation rate, a reduction in the real interest rate, and thus a stimulus to *private* demand: government spending crowds-in private spending and thus government spending multipliers are greater than one. The opposite holds, instead, when an equilibrium with *dis-anchored* inflation expectations is selected: fiscal expansions are anticipated to occur in conjunction with a faster fall in inflation, an increase in the real rate and thus with a crowding-out of private spending, leading to fiscal multipliers that are smaller than one.

To confirm this intuitive conclusion, Figure 9 calculates government spending multipliers at various horizons for three different sizes of the initial spending expansion. The solid line corresponds to the large and prolonged fiscal expansion in Figure 8, the dotted-dashed line corresponds to a large and short-lived expansion, while the dashed line gives a small and short-lived stimulus. The comparison among fiscal multipliers for different sizes of the stimulus is interesting because, contrary to what is often done in the Neo-Keynesian literature, my solution to the model is fully non-linear, leaving scope for non-linearity in the multipliers as well. For a given time horizon of t years, the *cumulative* multipliers are calculated as the ratio of two quantities. The numerator is set equal to the cumulative difference from time 0 up to time t between output under stimulus and output under the baseline of no stimulus (the numerator is thus the integral between the solid and the dashed output gap lines in Figure 8). Similarly, the denominator is equal to the cumulative difference between the absolute level of government spending under stimulus and the spending under no stimulus (zero spending). Figure 9 confirms that, under the assumption of dis-anchored inflation expectations, spending multipliers are always well below 1. In the short-run there is no evidence of strong non-linear effects, and on impact multipliers are clustered around 0.5 in all three cases. Over the longer run, sharper (i.e. larger and short-lived) fiscal expansions are associated with significantly lower cumulative multipliers.

The concept of fiscal "multipliers" is arguably quite problematic, and maybe of little use, in our context with multiple equilibria. Still, the calculations in Figure 9 are useful to confirm that, as emphasized for instance in Mertens and Ravn (2012) and in Christiano and Eichenbaum (2012), fiscal multipliers are rather small in self-fulfilling liquidity traps. The theoretical analysis of Sections 3.1-3.3 helps, in turn, to account for this result. Equilibrium selection choices between anchored versus dis-anchored equilibria are far from harmless for policy analysis in standard sticky price models (Cochrane [2011], Cochrane [2013]).

6 Conclusions

This paper provides a general analysis of deflationary paths in a standard sticky price model. I find sufficient conditions on the form of the monetary and fiscal policy rules that guarantee the existence of monotonic deflationary paths triggered by a sunspot dis-anchoring of inflation expectations. Quantitatively, the model is able to replicate the evolution of the main macro variables, both nominal and real, during the Japanese lost decade. This result is obtained under three main assumptions. First, the model is calibrated to include a sunspot that dis-anchors inflation expectations. Second, following Hayashi and Prescott (2002), the reduction in potential growth during the lost decade causes a permanent fall in the natural real interest rate. Third, monetary policy is assumed to follow a "not too stimulative" Taylor rule. Overall, the simulation aligns well with evidence suggesting the self-fulfilling nature of the Japanese deflationary experience. The quantitative analysis also highlights the central role, in policy exercises such as simulations of fiscal expansions, of the equilibrium selection choice operated on inflation expectations.

Monotonic self-fulfilling deflationary paths are a general feature of sticky price models. Moreover, such equilibria are not only pervasive in theory, but could also be interesting empirical phenomena which standard models are able to replicate. These findings open the door to further issues. One is how to model in a deeper way the triggers of inflation expectation shocks. Sunspots may be linked, for instance, to the occurrence of other, real or nominal, shocks. Another issue is how the exit from the self-fulfilling deflationary steady state. For example, in a world of multiple equilibria, can "shock and awe" policy announcements act as a coordination device that re-anchors inflation expectations? Answering these questions is ground for future research.

7 References

Aruoba B., Schorfheide F. (2012), *Macroeconomic Dynamics Near The ZLB: A Tale Of Two Equilibria*, manuscript, University of Pennsylvania.

Ayashi F., Prescott E. C. (2002), *The 1990s in Japan: A Lost Decade*, Review of Economic Dynamics 5, 206-235.

Benhabib J., Schmitt-Grohé S., Uribe M. (2001), *The Perils of Taylor Rules*, Journal of Economic Theory, 96, 40-69.

Benhabib J., Schmitt-Grohé S., Uribe M. (2002a), *Avoiding Liquidity Traps*, Journal of Political Economy, University of Chicago Press, 110, 535-563.

Benhabib J., Schmitt-Grohé S., Uribe M. (2002b), *Chaotic Interest-Rate Rules*, American Economic Review, 92, 72-78.

Bernanke B., Gertler M. (2000), *Monetary policy and asset price volatility*, NBER Working Paper 7559.

Bullard J., Mitra K. (2007), *Determinacy, Learnability, and Monetary Policy Inertia*, Journal of Money, Credit and Banking, 39, 1177-1212.

Bullard J. (2010), *Seven Faces of The Peril*, Federal Reserve Bank of St. Louis Review, 92, 339-352.

Carlson J.A., Parkin M. (1975), *Inflation expectations*, Economica 42, 123-138.

Chari V.V., Kehoe P., McGrattan E. (2007), *Business Cycle Accounting*, Econometrica, 75, 731-836.

Christiano L., Eichenbaum M. , Rebelo S. (2011), *When Is the Government Spending Multiplier Large?*, Journal of Political Economy, 119, 78-121.

Christiano L., Eichenbaum M. (2012), *Notes on Linear Approximations, Equilibrium Multiplicity and E-learnability in the Analysis of the Zero Lower Bound*, manuscript, Northwestern University.

Cochrane J. (2001), *Long-term Debt and Optimal Policy in the Fiscal Theory of the Price Level*, Econometrica, 69, 69-116.

Cochrane J. (2011), *Determinacy and Identification with Taylor Rules*, Journal of Political Economy, 119, 565-615.

Cochrane J. (2013), *The New-Keynesian Liquidity Trap*, mimeo, October 24.

Correia I., Farhi E., Nicolini J.P., Teles P. (2013), *Unconventional Fiscal Policy at the Zero Bound*, American Economic Review, 103, 1172-1211.

Eggertsson, G. B. (2010), *What Fiscal Policy Is Effective at Zero Interest Rates?*, NBER Macroeconomics Annual 2010, 25, 59-112.

Eggertsson G. B. (2010), *The Paradox of Toil*, Federal Reserve Bank of New York Staff Report 433.

Eggertsson, G. B., Woodford M. (2003), *The Zero Bound on Interest Rates and Optimal Monetary Policy*, Brookings Papers on Economic Activity, 139-211.

Eggertsson G. B., Woodford M. (2004), *Policy Options in a Liquidity Trap*, American Economic Review, 94, 76-79.

Eusepi S. (2007), *Learnability and Monetary Policy: a Global Perspective*, Journal of Monetary Economics, 54, 1115-1131.

Fernandez-Villaverde J. (2014), *Discussion of"Can Structural Reforms Help Europe?" by Gauti Eggertsson, Andrea Ferrero, and Andrea Raffo*, Journal of Monetary Economics 61, 23-31.

Gordon, R. (2012), *Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*, NBER Working Paper No. 18315.

King R.G. (2000), *Te New IS-LM Model: Language, Logic, and Limits*, Federal Reserve Bank of Richmond, Economic Quarterly 86, 45-103.

Krugman P. (1998), *Its Baaack! Japan's Slump and the Return of the Liquidity Trap,*Brookings Papers On Economic Activity, 137-87.

Mertens K., and Ravn M.O. (2012), *Fiscal Policy In An Expectations Driven Liquidity Trap*, manuscript, University College London.

Rotemberg J. J. (1982), emphSticky prices in the United States, Journal of Political Economy, 90, 1187-1211.

Sargent T. J., Wallace N. (1975), *Rational Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule*, Journal of Political Economy, 83, 241-254

Schmitt-Grohé S., Uribe M. (2013), *The Making Of A Great Contraction With A Liquidity Trap and A Jobless Recovery*, manuscript, Columbia University.

Taylor J. B. (1993), *Discretion Versus Policy Rules in Practice*, Carnegie-Rochester Conference Series On Public Policy 39, 195-214.

Ueda K. (2010), *Determinants of Households' Inflation Expectations in Japan and the United*

States, Journal of the Japanese and International Economies 24, 503-518, 2010.

Woodford M. (2003), *Interest & Prices*, Princeton University Press.

8 Appendix A

Proof. Since we are going to focus only on equations (32) and (34), to reduce notation we can normalize the constants so that $\psi/\alpha = \psi$ or $\alpha = 1$. Combining (32) and (34) we obtain

$$\frac{\chi + \pi_{t+1} - \beta\omega^d(\pi_{t+1}, \pi_{t+2})\pi_{t+2}}{\chi + \pi_t - \beta\omega^d(\pi_t, \pi_{t+1})\pi_{t+1}} = \left[\frac{1 + \tilde{f}(\pi_t, \pi_{t+1})}{1 + r^n}\right]^{1+\psi} \tag{35}$$

where,

$$1 + \tilde{f}(\pi_t, \pi_{t+1}) = [1 + r(\pi_t, \pi_{t+1})] \left[\frac{1 + \Xi(\pi_{t+1})}{1 + \Xi(\pi_t)}\omega^d(\pi_t, \pi_{t+1})\right]^{\frac{\psi}{1+\psi}} \tag{36}$$

For brevity, let us make the auxiliary assumption that $\omega^d(S) = 1$ not just at S^{LT} but also in a small neighborhood of S^{LT}. In such small neighborhood (35) becomes

$$\frac{1 - v(\pi_{t+1}, \pi_{t+2})}{1 - v(\pi_t, \pi_{t+1})} = \left[\frac{1 + r(\pi_t, \pi_{t+1})}{1 + r^n}\right]^{1+\psi} \left[\frac{1 + \Xi(\pi_{t+1})}{1 + \Xi(\pi_t)}\right]^{\psi}$$

where as usual $v(\pi_t, \pi_{t+1}) = -\frac{1}{\chi}(\pi_t - \beta\pi_{t+1})$. Linearizing around π^{LT}, we obtain the second order difference equation,

$$\hat{\pi}_{t+2} - \chi_1\hat{\pi}_{t+1} + \chi_0\hat{\pi}_t = 0$$

where

$$\chi_0 = \frac{1}{\beta}\left[1 - \chi\psi\Xi'(\pi^{LT})\frac{1 - v(\pi^{LT}, \pi^{LT})}{1 + \Xi(\pi^{LT})}\right] > 1$$
$$\chi_1 = 1 + \chi_0 - r_2(\pi^{LT}, \pi^{LT})(1 + \psi)\chi[1 - v(\pi^{LT}, \pi^{LT})] > 1 + \chi_0$$

and $\hat{\pi}_t = \pi_t - \pi^{LT}$. To establish the above inequalities, notice that $\Xi'(\pi^{LT}) < 0$. Morever, since $i(S_t) = 0$ for $S_t \le S^{LT}$ and by assumption $i(S)$ is continuously differentiable at $S = S^{LT}$, then we have two results. The first is that $r_2(\pi^{LT}, \pi^{LT}) = \lim_{\pi\uparrow\pi^{LT}} \frac{\partial r(\pi^{LT}, \pi)}{\partial\pi} > 0$. The second is that $r_1(\pi^{LT}, \pi^{LT}) = \lim_{\pi\uparrow\pi^{LT}} \frac{\partial r(\pi, \pi^{LT})}{\partial\pi} = 0$ (since $r_1(S^{LT}) = 0$ it does not appear in the expressions for the coefficients χ_0 and χ_1). The difference equation has two roots

$$\lambda_1 = \frac{\chi_1 + \sqrt{\chi_1^2 - 4\chi_0}}{2}$$

$$\lambda_2 = \frac{\chi_1 - \sqrt{\chi_1^2 - 4\chi_0}}{2}$$

Since $\chi_1 > 1 + \chi_0$ then $\chi_1^2 - 4\chi_0 > (\chi_0 - 1)^2 > 0$, so that both roots are positive and real. Moreover, since λ_1 is increasing in χ_1, the condition $\chi_1 > 1 + \chi_0$ implies that $\lambda_1 > \chi_0 > 1$. Similarly, since λ_2 is stricly decreasing in χ_1, the condition $\chi_1 > 1 + \chi_0$ implies $\lambda_2 < 1$. We have proven that around the liquidity trap equilibrium, there is a saddle path (approximately) described by the stable manifold $\pi_{t+1} = \mathcal{M}(\pi_t) = \pi^{LT} + \lambda_2(\pi_t - \pi^{LT})$. Notice that, locally, $S(\pi)$ is continuous, monotonically increasing and such that $\pi^{LT} < \mathcal{M}(\pi) < \pi$ for $\pi > \pi^{LT}$.

I now show how to extend the function $\mathcal{M}(\pi)$ from a neighbourhood of π^{LT} to the entire interval (π^{LT}, π^{TG}). I consider the general case where Assumption 3 holds (a special case is

$\tilde{\beta}_t = \beta$ for all t). At any point in time, inflation rates π_{t+2}, π_{t+1} and π_t must solve (35). For π_{t+1} close to π^{LT} we can write $\pi_{t+2} = \mathcal{M}(\pi_{t+1})$. Moreover, given π_{t+1}, we can show that there is one and only one value of π_t that solves (35). In particular, uniqueness is guaranteed by the fact that the left hand side of the equation is decreasing in π_t while, by assumption, the right hand side is increasing in π_t. Existence is established as follows. For $\pi_t = \pi_{t+1} \in (\pi^{LT}, \pi^{TG})$, the left hand side is greater than one, and thus is strictly bigger than the right hand side, which is strictly smaller than one. Moreover, for π_t large enough (possibly positive), the left hand side is smaller than the right hand side since, by non-negativity of nominal rates, the right hand side is bounded away from zero by some strictly positive constant. The continuity of the interest rate policy then implies existence of one and only one value $\pi_t = \mathcal{M}^{-1}(\pi_{t+1})$ that satisfies (35), with $\pi_t > \pi_{t+1}$ and $\mathcal{M}^{-1}(\cdot)$ a continuous function. Notice that for $\pi_{t+1} < \pi^{LT}$ close enough to π^{LT} we are certain that $\pi_t = \mathcal{M}^{-1}(\pi_{t+1}) \in (\pi_{t+1}, \pi^{TG})$. Therefore, by continuity of $\mathcal{M}^{-1}(\cdot)$, we can increase π_{t+1} towards π^{TG} until $\mathcal{M}^{-1}(\pi_{t+1})$ is arbitrarily close to π^{TG} (there is no fixed point for $\mathcal{M}^{-1}$ between π^{LT} and π^{TG}). We have thus extended $\mathcal{M}(\pi)$ to the entire set (π^{LT}, π^{TG}).

Next, assume that $1 \leq \bar{\phi} \leq \frac{1}{\beta}$, where $\bar{\phi}$ is calculated with respect to $\tilde{f}(\cdot,\cdot)$. For π^*_{t+1} sufficiently close to π^{LT} we have $\hat{Y}^*_{t+1} < 1$. By (32),

$$\hat{Y}_t^{*1+\psi} = 1 + \frac{1}{\chi}(\pi^*_t - \tilde{\beta}_t \pi^*_{t+1}) \tag{37}$$

Since $\tilde{\beta}(\pi_t, \pi_{t+1}) = \beta\omega^d(\pi_t, \pi_{t+1})$ is increasing in π_t, there exists a unique $\pi_t \in (\pi^*_{t+1}, 0)$ such that $\pi_t - \tilde{\beta}(\pi_t, \pi^*_{t+1})\pi^*_{t+1} = 0$. Using expression (37) to define $\hat{Y}_t$ in correspondence of such π_t, we obtain $\hat{Y}_t = 1$, which implies

$$\frac{Y^*_{t+1}}{\hat{Y}_t} < 1$$

Given that $\tilde{\beta}_t \leq \beta$ and that $\bar{\phi} \leq \frac{1}{\beta}$, we have

$$\tilde{f}(\pi_t, \pi^*_{t+1}) = \tilde{f}(\tilde{\beta}_t \pi^*_{t+1}, \pi^*_{t+1}) \geq \tilde{f}(\beta \pi^*_{t+1}, \pi^*_{t+1}) \geq \tilde{f}(\pi^*_{t+1}/\bar{\phi}, \pi^*_{t+1}) \geq r^n$$

Hence, we cannot have $\pi^*_t = \pi_t$, otherwise the last two inequalities would imply that the left-hand-side of (34) is strictly smaller than one, while the right-hand-side is larger than one. Hence, necessarily, the equilibrium value of π^*_t has to be smaller than π_t, giving $\pi^*_t - \tilde{\beta}^*_t \pi^*_{t+1} < \pi_t - \tilde{\beta}_t \pi^*_{t+1} = 0$ and consequently $\hat{Y}^*_t < 1$. We can then continue the process by induction at time $t-1, t-2, ...0$ and thus prove $\hat{Y}^*_t < 1$ for all t. In particular, as $\pi^*_0 \uparrow 0$, the condition $\pi^*_0 - \tilde{\beta}^*_1 \pi^*_1 < 0$ implies $\pi^*_1 - \mathcal{M}(\pi^*_0) \uparrow 0$. $\square$

9 Appendix B

Constant static wedge. Assume that $\omega^s(S) = 1$. By equation (15), the labor tax is set according to,

$$1 - \tau_t^N = \frac{1 + \sigma_t^G}{1 - \tau_t^C}$$

The system of equations (16)-(17) reduces to a third order difference equation,

$$\frac{\chi + \pi_{t+1} - \beta\omega^d(\pi_{t+1}, \pi_{t+2})\pi_{t+2}}{\chi + \pi_t - \beta\omega^d(\pi_t, \pi_{t+1})\pi_{t+1}} = \frac{1 + \tilde{r}(\pi_t, \pi_{t+1})}{1 + r^n} \tag{38}$$

Because of Assumption 3, the effective real rate function $\tilde{r}(S)$ satisfies all the properties in Assumptions 1-2.

Labor tax as a decreasing function. For a general labor tax rule $\tau_t^N = \tau^N(\pi)$, simple calculations show that the effective real rate $\tilde{r}(S)$ in (38) is replaced by a function $\tilde{f}(S)$ with

$$\tilde{f}(S_t) = r(S_t)\frac{1 - \tau^N(\pi_t)}{1 - \tau^N(\pi_{t+1})}$$

Hence, the condition that $\tau^N(\pi)$ is a weakly decreasing function is sufficient to guarantee that $\tilde{f}(S)$ satisfies the usual Assumptions 1-2.

Time-consistent labor tax. Assume that each period the labor tax rate τ_t^N is set only after nominal wages, prices, interest rates, consumption tax rates and government spending have been fixed. Given the values of this variables, the benevolent government sets at every time t the labor tax τ_t^N, so that, given the tax, the allocation chosen by the agents maximizes time t period utility. As shown in Section 4 devoted to microfoundations, a time-consistent labor tax elicits a ratio $\hat{Y}_t$ of actual output relative to the level of output in the Pareto optimal equilibrium equal to

$$\hat{Y}_t = \frac{1}{1 + \Xi_t} \tag{39}$$

where, in a sticky-prices model of the Rotemberg (1982) type, the variable $\Xi_t = \Xi(\pi_t)$ represents the time t price adjustment cost (as a share of final output) associated with a deviation π_t of inflation from the optimal price stability value $\pi^{TG} = 0$. The function $\Xi(\pi)$ is U-shaped, with the minimum value $\Xi(\pi^{TG}) = 0$ reached at the target inflation. From (17) and the definition (15) of the static wedge, we have that the time-consistent labor tax satisfies

$$\frac{1 + \tau_t^C}{(1 + \sigma_t^G)(1 - \tau_t^N)} = (1 - v_t)(1 + \Xi_t) \tag{40}$$

where,

$$1 - v_t = 1 + \frac{1}{\chi}(\pi_t - \beta\omega_{t+1}^d\pi_{t+1})$$

Equation (40) is the generalization to the case where $\pi_t \neq 0$ and $\sigma_t^G \neq 0$ of the condition for optimal distortionary taxation presented in Correia, Farhi, Nicolini, Teles (2013). Finally, substituting (39) into (17), we now obtain a second order difference equation,

$$\frac{1 + \Xi(\pi_t)}{1 + \Xi(\pi_{t+1})} = \frac{1 + \tilde{r}(\pi_t, \pi_{t+1})}{1 + r^n} \tag{41}$$

Fix a value for $\pi_{t+1} \in (\pi^{LT}, \pi^{TG})$. For $\pi_t = \pi_{t+1}$ the left-hand-side of (41) is strictly greater than the-right-hand side, which is strictly smaller than one by assumption. The opposite relation holds for $\pi_t = \pi^{TG} = 0$, since $\bar{\phi} < \infty$, so that the right hand side is bigger than one. Moreover, since the left-hand-side of (41) is stricly decreasing in $\pi_t < \pi^{LT}$, while the right-hand-side is increasing in π_t then we can always find one and only one value $\pi_t \in (\pi_{t+1}, \pi^{TG})$ such that (41) holds. In this way we have implicitly defined and characterized the function $\mathcal{M}(\pi_t) = \pi_{t+1}$ as a *Type I* manifold.

10 Appendix C. Global and idiosyncratic inflation expectations

10.1 *Consensus Forecast* expectations

The analysis of inflation expectations in Japan presented in Figure 3 is based on the following identity,

$$\frac{1}{n}\sum_{j=0}^{n-1}\pi_{t-j} = \frac{1}{n}\sum_{j=0}^{n-1}\pi^{e}_{t-j|t-n} + \bar{\varepsilon}^{g}_{t} + \bar{\varepsilon}^{i}_{t} \tag{42}$$

where $\pi^{e}_{t-j|t-n}$ indicates the expectation form at time $t-n$ for the value of the inflation rate π_{t-j} at time $t-j$. The constant n indicates the forecasting horizon. The *total* average forecasting error equals the average deviation of realized inflation (the left-hand side of (42)) from the average expected inflation (the first term on the right-hand side of (42)). The total error is decomposed into an average *global* forecasting error $\bar{\varepsilon}^{g}_{t}$ and into an average *idiosyncratic* forecasting error $\bar{\varepsilon}^{i}_{t}$.

For any given forecasting horizon, the corresponding time t global forecasting error $\bar{\varepsilon}^{g}_{t}$ is estimated to be the average of the total forecasting errors calculated for Germany, the UK, and the US. This procedure implicitly assumes that the idiosyncratic forecasting errors specific to Germany, the UK and the US are washed away when taking cross-country averages, and thus only the global shock remains. The sample period runs from $t = 1990$ to $t = 2006$. During the years 1990-2006 there was no permanent deflationary process taking place in either Germany, the UK, or the US, which implies that the global shocks $\bar{\varepsilon}^{g}_{t}$ should not contain any permanent, unexpected, global deflationary process. Consequently, if Japan has indeed experienced a permanent deflationary process that was largely unexpected, such deflationary shock cannot be embedded in the $\bar{\varepsilon}^{g}_{t}$ shock, but should rather show up as a sequence of consistently negative values for the estimated *idiosyncratic* shock $\bar{\varepsilon}^{i}_{t}$.

At any time t, the solid line in the right-hand side panel of Figure 3 depicts, for a forecast horizons $n = 3$, the average realized inflation rates between time $t-n$ and time t, i. e. the left hand side of equation (42). The difference at t between the solid and the dashed line is by construction equal to the estimated *idiosyncratic* forecasting error $\bar{\varepsilon}^{i}_{t}$, obtained by plugging the estimated global forecasting error ε^{g}_{t} into (42).

The main takeaway from Figure 3 is that the idiosyncratic errors of inflation expectations in Japan were small and on average equal to zero, so that the ex-post underperformance of actual inflation relative to ex-ante expectations can be almost entirely attributed to global factors. This conclusion is robust to using the cross-country median, instead of the mean, in the calculation of $\bar{\varepsilon}^{g}_{t}$. The result also holds, albeit it is slightly weakened, when we move from medium-term ($n = 3$ years) inflation expectations to relatively short-term expectations

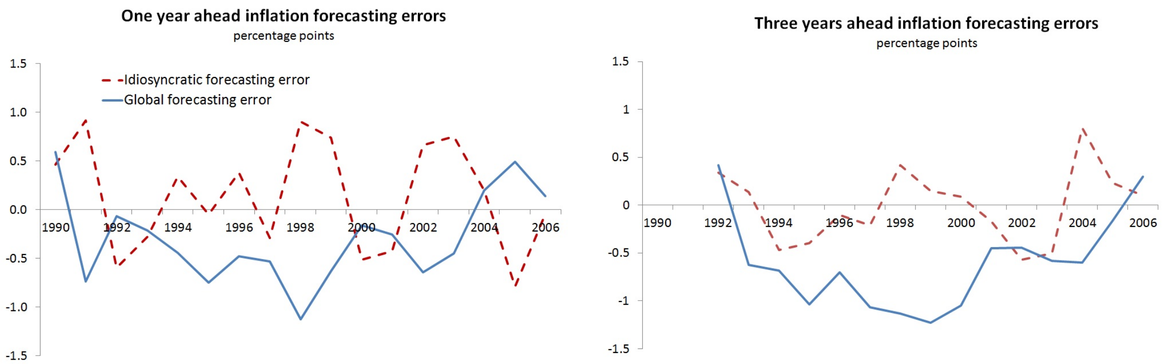

Figure 10: Global (solid line) and idiosyncratic (dashed line) errors. Sources: Thomson Reuters Datastream, Statistics Bureau of Japan, Consensus Forecast (October survery).

($n = 1$ year). This can be more clearly seen in Figure 10, which plots the idiosyncratic and global components of the forecasting errors. At the one year horizon (left-hand side panel), the idiosyncratic error has mean equal to -0.05% and standard deviation of 0.55%, while at the three-years horizon[22] (right-hand side panel) the mean is -0.01% and the standard deviation is 0.35%.

10.2 Robustness: *Consumer Confidence* expectations

The scarcity of long time series on inflation expectations makes it difficult to ascertain whether the results of the previous section are robust to the use of alternative measures of expectations. To partially address this issue, I compare one-year inflation expectations taken from *Consensus Forecast* with those that can be obtained from the *Consumer Confidence* survey of the Bank of Japan. I find that the two measures are highly correlated, a conclusion which provides at least a partial proof of robustness to my results.

The Consumer Confidence survey provides a long time series of answers to a question asking consumers whether in the next year prices would go up, remain unchanged, or go down. I transform these qualitative responses into quantitative ones by introducing a variation of the methodology developed by Carlson and Parkin (1975) and applied by Ueda (2010), whose original dataset I employ. In a nutshell, the original procedure of Carlson and Parkin (1975) is the following. Assume that agents believe that inflation between time $t-1$ and time t is a random variable $\tilde{\pi}_t$ which follows a normal distribution with mean $\pi^e_{t|t-1}$ and standard deviation $\sigma_{t|t-1}$. The goal of the procedure is to use the qualitative information from the survey in order to “pin down" the shape of the normal distribution for $\tilde{\pi}_t$, thus obtaining in particular the expected value of inflation $\pi^e_{t|t-1}$. The shape of the distribution is “pinned down" by making two assumptions. The first is that there exists a time-invariant parameter $\delta > 0$Ãň such that the probability that $\tilde{\pi}_t$ falls below the threshold $-\delta$ equals the share of individuals that indicated that they expect prices to fall, and the probability that $\tilde{\pi}_t$ rises above the threshold δ equals the share of individuals that indicated that they expect prices to increase. Clearly, the probability that $\tilde{\pi}_t$ falls between $-\delta$ and δ must be equal to the

[22] At the three-years horizon, the available observations are reduced from 17 to 15.

share of respondents that indicate that prices will not change. Thus, a natural interpretation of δ is that it provides the *precision* with which respondents make explicit that they expect inflation to be different from zero. Having used the shares of respondents to pin down two percentiles of the normal distribution, it is possible to show that the time varying mean can now be written as a function $\pi^e_{t|t-1}(\delta)$ of the unknown parameter δ. In turn, δ is obtained through an estimation procedure based on the assumption that expectations are unbiased on average over the sample period. Using our usual notation, this assumption implies setting

$$\sum_{t=1}^{T}[\pi_t - \pi^e_{t|t-1}(\delta)] = 0 \tag{43}$$

where $t = 1, ..., T$ indicates the *entire* sample over which the transformation from qualitative to quantitative expectations is performed. Having obtained an estimate $\hat{\delta}$ by plugging into (43) the actual inflation rates π_t, one can then proceed to contruct for every t the quantitative value $\pi^e_{t|t-1} = \pi^e_{t|t-1}(\hat{\delta})$ of the one year ahead inflation expectation.

The procedure outlined above, while useful for the goal of obtaining quantitative inflation expectations, has an important drawback for the purpose of our empirical investigation. It requires to assume, through (43), that inflation expectations are unbiased, which is exactly the properties that in the previous section I tried to establish. To make this point clear, assume that we apply the procedure above to the sample period $t = 1990$ to $t = 2006$. Then, by construction, we are bound to obtain a series of inflation expectations whose average forecasting error during the Japanese deflationary path is zero. To overcome this problem, I propose to follow an alternative route, which I lay out in two steps.

For the first step I keep assuming that δ is constant over time, but I modify the estimation procedure in (43) by using inflation rates and responses to the Consumer Confidence survey for the quarterly five-year window 1980Q1-1989Q4 [23]. The estimated $\hat{\delta}$ is then used, as in an out-of-sample procedure, to calculate inflation expectations $\pi^e_{t|t-1}(\hat{\delta})$ for the period 1990-2006. The left hand side panel in Figure 11 compares the result from this modified Carlson-Parkin procedure to the *Consensus Forecast* expectations. Two aspects deserve attention. First, the correlation between the two series is very high (87%), providing an indication that variations in *Consensus* expectations are coherent with variations in expectations obtained from other sources. The second is that in the second part of the sample, which corresponds to the period where annual inflation had already stabilized around very low values there is a persistent level gap between *Consensus* and the modified Carlson-Parkin expectations.

One reasonable cause that can justify the opening of this gap is that as Japan progressively moved from a largely positive inflation regime to a regime with zero or slightly negative inflation, consumers started to communicate with greater *precision* whether they indeed expected inflation to be exactly at zero or if they expected it to have even small positive or negative variations. If this were the case, then the precision threshold δ would not be constant, but instead would vary over time. To check for this possibility, I introduce a second modification to the Carlson-Parkin method, which allows for the construction of inflation expectations $\pi^e_{t|t-1}(\delta_{t-1})$ based on a time varying parameter δ_{t-1} estimated using only past information. Specifically, I estimate δ_{t-1} by applying the condition (43) to the five years rolling window period ending in

[23] While δ is estimated using quarterly data, Figure 11 reports only the inflation expectations for the fourth quarter of each year. Selecting the Consumer Confidence inflation expectation formed in the fourth quarter of each year is consistent with the choice of using the October survey from Consensus Forecast.

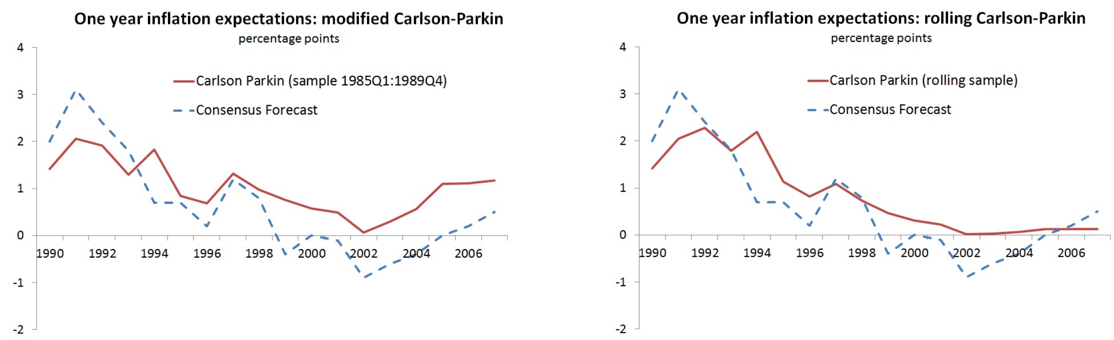

Figure 11: Inflation expectations from the modified Carlson-Parkin methods (solid line) and Consensus Forecast (dashed line) errors. These series are not corrected for the VAT tax increase. Sources: data for the calculations are from Ueda (2010).

with the year $t-1$[24]. The estimated $\hat{\delta}_{t-1}$ is then used to derive the expectation $\pi^e_{t|t-1}(\delta_{t-1})$ at time $t-1$. The results of this procedure are displayed in the right hand side panel of Figure 11. The main conclusion to draw is that the persistent gap between the two series in the later part of the sample is now reduced. The average distance between Consensus and Consumer Confidence inflation expectations is now 0.25% (from 0.38% when the fixed estimation sample is used), while the correlation remains high at 86%. Hence, since the two series are highly correlated and very close to each other on average, using either of them makes no significant difference when calculating the idiosyncratic forecast.

[24] I also impose that δ_{t-1} must be equal to the maximum between the value obtained from the estimation using (43) and an exogenous lower bound which I set equal to 0.1%. The reason for this choice is that for the period 2002-2006 the value $\hat{\delta}$ estimated through (43) would be slightly negative. This modified estimated procedure gives values (not reported) for $\hat{\delta}_{t-1}$ that confirm the conjecture that over time survey's respondents have progressively communicated with higher precision (i.e. $\hat{\delta}_{t-1}$ falls over time) whether they expect inflation to deviate form zero.