Essays on the Liquidity Trap, Oil Shocks, and the Great Moderation

A thesis presented

by

Anton Nakov

to

the Department of Economics and Business

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in Economics

Universitat Pompeu Fabra

Barcelona

November 2007

Dipòsit legal: B.9515-2008
ISBN: 978-84-691-4441-1

Universitat Pompeu Fabra
Department of Economics and Business

The undersigned hereby certify that they have read and recommend to the Faculty of Graduate Studies for acceptance a thesis entitled "Essays on the Liquidity Trap, Oil Shocks, and the Great Moderation" by Anton Nakov in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Date: November 2007

Supervisor: ______________________________
Prof. Jordi Galí

Readers: ______________________________

____

____

____

____

Universitat Pompeu Fabra
Department of Economics and Business

Author: Anton Nakov

Contact: Tellez 12 B Entreplanta "I"; 28007 Madrid; Spain
Tel: 647207752. Email: Anton.Nakov@Gmail.com

Title: Essays on the Liquidity Trap, Oil Shocks, and the Great Moderation

Degree: Doctor of Philosophy

Date: November 2007

To my family

Acknowledgments

More than anyone I would like to thank Prof. Jordi Galí for the invaluable guidance and inspiration that having him as a thesis advisor has meant to me. Without him this work would not have been possible.

I am indebted also to many people from the Barcelona economics community who, together with the graduate program at Universitat Pompeu Fabra, turned me from someone who was hardly aware about many of the issues into somebody who aspires to contribute to economic research at the highest level. In particular, I have benefited immensely from interaction with Albert Marcet, Ramon Marimon, Fabio Canova, Michael Reiter, Kosuke Aoki, and Jaume Ventura. Needless to say, I have learned a great deal from my co-authors Max Gillman and Andrea Pescatori, as well as from fellow students Krisztina Molnar, Lutz Weinke, Philip Saure, Gustavo Solorzano, Helmuts Azacis and Peter Vida. To all these people I am very, very grateful.

Part of the work on this thesis was undertaken while I was an intern at the ECB and I would like to thank the people there for their hospitality, especially Massimo Rostagno and Victor Lopez Perez. I enjoyed and learned much also during my internship at the IMF, for which I am grateful to Roberto Garcia-Saltos, Douglas Laxton and Pau Rabanal. Finally, I would like to thank my colleagues at Banco de España, especially Juan Francisco Jimeno and Fernando Restoy for useful comments and support during the final stage of my dissertation. Last but not least, I am grateful for financial support from Agencia Española de Cooperacion Internacional (Becas MAE).

Contents

Preface

The thesis studies three distinct issues in monetary economics using a common dynamic general equilibrium approach under the assumptions of rational expectations and nominal price rigidity.

The first chapter deals with the so-called “liquidity trap” – an issue which was raised originally by Keynes in the aftermath of the Great Depression. Since the nominal interest rate cannot fall below zero, this limits the scope for expansionary monetary policy when the interest rate is near its lower bound. The chapter studies the conduct of monetary policy in such an environment in isolation from other possible stabilization tools (such as fiscal or exchange rate policy). In particular, a standard New Keynesian model economy with Calvo staggered price setting is simulated under various alternative monetary policy regimes, including optimal policy. The challenge lies in solving the (otherwise linear) stochastic sticky price model with an explicit occasionally binding non-negativity constraint on the nominal interest rate. This is achieved by parametrizing expectations and applying a global solution method known as “collocation”. The results indicate that the dynamics and sometimes the unconditional means of the nominal rate, inflation and the output gap are strongly affected by uncertainty in the presence of the zero lower bound. Commitment to the optimal rule reduces unconditional welfare losses to around one-tenth of those achievable under discretionary policy, while constant price level targeting delivers losses which are only 60% larger than under the optimal rule. On the other hand, conditional on a strong deflationary shock, simple instrument rules perform substantially worse than the optimal policy even if the unconditional welfare loss from following such rules is not much affected by the zero lower bound per se.

The second thesis chapter (co-authored with Andrea Pescatori) studies the implications of imperfect competition in the oil market, and in particular the existence of a welfare-relevant trade-off between inflation and output gap volatility. In the standard New Keynesian model exogenous oil shocks do not generate any such tradeoff: under a strict inflation targeting policy, the output decline is exactly equal to the efficient output contraction in response to the shock. I propose an extension of the standard model in which the existence of a dominant oil supplier (such as OPEC) leads to inefficient fluctuations in the oil price markup, reflecting a dynamic distortion of the economy's production process. As a result, in the face of oil sector shocks, stabilizing inflation does not automatically stabilize the distance of output from first-best, and monetary policymakers face a tradeoff between the two goals. The model is also a step away from discussing the effects of exogenous oil price changes and towards analyzing the implications of the underlying shocks that cause the oil price to change in the first place. This is an advantage over the existing literature, which treats the macroeconomic effects and policy implications of oil price movements as if they were independent of the underlying source of disturbance. In contrast, the analysis in this chapter shows that conditional on the source of the shock, a central bank confronted with the same oil price change may find it desirable to either raise or lower the interest rate in order to improve welfare.

The third thesis chapter (co-authored with Andrea Pescatori) studies the extent to which the rise in US macroeconomic stability since the mid-1980s can be accounted for by changes in oil shocks and the oil share in GDP. This is done by estimating with Bayesian methods the model developed in the second chapter over two samples - before and after 1984 - and conducting counterfactual simulations. In doing so we nest two other popular explanations for the so-called “Great Moderation”: (1) smaller (non-oil) shocks; and (2) better monetary policy. We find that the reduced oil share can account for around one third of the inflation moderation, and about 13% of the GDP growth moderation. At the same time smaller oil shocks can explain approximately 7% of GDP growth moderation and 11% of the inflation moderation. Thus, the oil share and oil shocks have played a non-trivial role in the moderation, especially of inflation, even if the bulk of the volatility reduction of output growth and inflation is attributed to smaller non-oil shocks and better monetary policy, respectively.

Chapter 1
Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate

1.1 Introduction

An economy is said to be in a "liquidity trap" when the monetary authority cannot achieve a lower nominal interest rate in order to stimulate output. Such a situation can arise when the nominal interest rate has reached its zero lower bound (ZLB), below which nobody would be willing to lend, if money can be stored at no cost for a nominally riskless zero rate of return.

The possibility of a liquidity trap was first suggested by Keynes (1936) with reference to the Great Depression of the 1930s. At that time he compared the effectiveness of monetary policy in such a situation with trying to "push on a string". After WWII and especially during the high inflation period of the 1970s interest in the topic receded, and the liquidity trap was relegated to a hypothetical textbook example. As Krugman (1998) noticed, of the few modern papers that dealt with it most concluded that "the liquidity trap can't happen, it didn't happen, and it won't happen again".

With the benefit of hindsight, however, it did happen, and to no less than Japan. Figure 1.1 illustrates this, showing the evolution of output, inflation, and the short-term nominal interest rate following the collapse of the Japanese real estate bubble of the late 1980s. The figure exhibits a persistent downward trend in all three variables, and in particular the emergence of deflation since 1998 coupled with a zero nominal interest rate since 1999.

Motivated by the recent experience of Japan, the aim of the present study is to contribute a quantitative analysis of the ZLB issue in a standard sticky price model under alternative monetary policy regimes. One the one hand, the paper characterizes optimal monetary policy in the case of discretion and commitment.[1] And on the other hand, it studies the performance of several simple monetary policy rules, modified to comply with the zero floor, relative to the optimal policy. The analysis is carried out within a stochastic general equilibrium model with monopolistic competition and Calvo (1983) staggered price setting, under a standard calibration to the postwar US economy.

The main findings are as follows: the optimal discretionary policy with zero floor involves a deflationary bias, which may be significant for certain parameter values and which implies that any quantitative analyses of discretionary biases of monetary policy that ignore the zero lower bound may be misleading. In addition, optimal discretionary policy implies much more aggressive cutting of the interest rate when the risk of deflation is high, compared to the corresponding policy without zero floor. Such a policy helps mitigate the depressing effect of private sector expectations on current output and prices when the probability of falling into a liquidity trap is high.[2]

In contrast, optimal commitment policy involves less preemptive lowering of the interest rate in anticipation of a liquidity trap, but it entails a promise for sustained mon-

---

1 The part of the paper on optimal policy is similar to independent work by Adam and Billi (2006) and Adam and Billi (2007). The added value is to quantify and compare the performance of optimal policy to that of a number of suboptimal rules (including discretionary policy) in the same stochastic sticky price setup.

2 An early version of this paper comparing the performance of optimal discretionary policy with three simple Taylor rules was circulated in 2004; optimal commitment policy and more simple rules were added in a version circulated in 2005. Optimal discretionary policy was studied also independently by Adam and Billi (2004b), and optimal commitment policy by Adam and Billi (2004a).

etary policy easing following an exit from a trap. This type of commitment enables the central bank to achieve higher expected inflation and lower real rates in periods when the zero floor on nominal rates is binding.[3] As a result, under the baseline calibration, the expected welfare loss under commitment is only around one-tenth of the loss under optimal discretionary policy. This implies that the cost of discretion may be much higher than normally considered when abstracting from the zero lower bound issue.

The average welfare losses under simple instrument rules are 8 to 20 times bigger than under the optimal rule. However, the bulk of these losses stem from the intrinsic suboptimality of simple instrument rules, and not from the zero floor *per se*. This is related to the fact that under these rules the zero floor is hit very rarely - less than 1% of the time - compared to optimal policy, which visits the liquidity trap one-third of the time. On the other hand, conditional on a large deflationary shock, the relative performance of simple instrument rules deteriorates substantially vis-a-vis the optimal policy.

Issues of deflation and the liquidity trap have received considerable attention recently, especially after the experience of Japan.[4] In an influential article Krugman (1998) argued that the liquidity trap boils down to a credibility problem in which private agents expect any monetary expansion to be reverted once the economy has recovered. As a solution he suggested that the central bank should commit to a policy of high future inflation over an extended horizon.

---

3 This basic intuition was suggested already by Krugman (1998) based on a simpler model.

4 A partial list of relevant studies includes Krugman (1998), Wolman (1998), McCallum (2000), Reifschneider and Williams (2000), Klaeffling and Lopez-Perez (2003), Eggertsson and Woodford (2003), Coenen, Orphanides and Wieland (2004), Kato and Nishiyama (2005), Jung, Teranishi and Watanabe (2005), Adam and Billi (2006), Adam and Billi (2007).

More recently, Jung et al. (2005) have explored the effect of the zero lower bound in a standard sticky price model with Calvo price setting under the assumption of perfect foresight. Consistent with Krugman (1998), they conclude that optimal commitment policy entails a promise of a zero nominal interest for some time after the economy has recovered. Eggertsson and Woodford (2003) study optimal policy with zero lower bound in a similar model in which the natural rate of interest is allowed to take two different values. In particular, it is assumed to become negative initially and then to jump to its “normal” positive level with a fixed probability in each period. These authors also conclude that the central bank should create inflationary expectations for the future. Importantly, they derive a moving price level targeting rule that delivers the optimal policy in this model.

One shortcoming of much of the modern literature on monetary policy rules is that it largely ignores the ZLB issue or at best uses rough approximations to address the problem. For instance Rotemberg and Woodford (1997) introduce nominal rate targeting as an additional central bank objective, which ensures that the resulting path of the nominal rate does not violate the zero lower bound too often. In a similar vein, Schmitt-Grohe and Uribe (2004) exclude from their analysis instrument rules that result in a nominal rate the average of which is less than twice its standard deviation. In both cases therefore one might argue that for sufficiently large shocks that happen with a probability as high as 5%, the derived monetary policy rules are inconsistent with the zero lower bound.

On the other hand, of the few papers that do introduce an explicit non-negativity constraint on nominal interest rates, most simplify the stochastics of the model, for example by assuming perfect foresight (Jung et al. 2005), or a two-state low/high economy

(Eggertsson and Woodford (2003); Wolman (1998)). Even then, the zero lower bound is effectively imposed as an initial ("low") condition and not as an occasionally binding constraint.[5] While this assumption may provide a reasonable first-pass at a quantitative analysis, it may be misleading to the extent that it ignores the occasionally binding nature of the zero interest rate floor.

Other studies (e.g. Coenen et al. (2004)) lay out a stochastic model but knowingly apply inappropriate solution techniques which rely on the assumption of certainty equivalence. It is well known that this assumption is violated in the presence of a non-linear constraint such as the zero floor but nevertheless these researchers have imposed it for reasons of tractability (admittedly they work with a larger model than the one studied here). Yet forcing certainty equivalence in this case amounts to assuming that agents ignore the risk of the economy falling into a liquidity trap when making their optimal decisions.

The present study contributes to the above literature by solving numerically a stochastic general equilibrium model with monopolistic competition and sticky prices with an explicit occasionally binding zero lower bound, using an appropriate global solution technique that does not rely on certainty equivalence. It extends the analysis of Jung et al. (2005) to the stochastic case with an AR(1) process for the natural rate of interest.

After a brief outline of the basic framework adopted in the analysis (section 1.2), the paper characterizes and contrasts the optimal discretionary and optimal commitment policies (sections 1.3 and 1.4). It then analyzes the performance of a range of simple

[5] Namely, the zero floor binds for the first several periods but once the economy transits to the "high" state, the ZLB never binds thereafter.

instrument and targeting rules (sections 1.5 and 1.6) consistent with the zero floor.[6] Sections 1.4 to 1.6 include a comparison of the conditional performance of all rules in a simulated liquidity trap, while section 1.7 presents their average performance, including a ranking according to unconditional expected welfare. Section 1.8 studies the sensitivity of the findings to various parameters of the model, as well as the implications of endogenous inflation persistence for the ZLB issue and the last section concludes.

1.2 Model

While in principle the zero lower bound phenomenon can be studied in a model with flexible prices, it is with sticky prices that the liquidity trap becomes a real problem. The basic framework adopted in this study is a stochastic general equilibrium model with monopolistic competition and staggered price setting *a la* Calvo (1983) as in Galí (2003) and Woodford (2003). In its simplest log-linearized[7] version the model consists of three building blocks, describing the behavior of households, firms and the monetary authority.

The first block, known as the "IS curve", summarizes the household's optimal consumption decision,

$$x_t = E_t x_{t+1} - \sigma \left(i_t - E_t \pi_{t+1} - r_t^n \right). \tag{1.1}$$

[6] These rules include truncated Taylor-type rules reacting to contemporaneous, expected future, or past, inflation, output gap, or price level; with or without "interest rate smoothing"; truncated first-difference rules; price level targeting; and strict inflation targeting rules.

[7] It is important to note that, like in the studies cited in the introduction, the objective here is a modest one, in that the only source of non-linearity in the model stems from the ZLB. Solving the fully non-linear sticky price model with Calvo (1983) contracts can be a worthwile enterprise, however it increases the dimensionality of the computational problem by the number of states and co-states that one should keep track of (e.g. the measure of price dispersion and in the case of optimal policy the Lagrange multipliers associated with all expectational constraints).

It relates the "output gap" x_t (i.e. the deviation of output from its flexible price equilibrium) positively to the expected future output gap and negatively to the gap between the ex-ante real interest rate, $i_t - E_t\pi_{t+1}$, and the "natural" (i.e. flexible price equilibrium) real rate, r_t^n (which is known to all agents at time t). Consumption smoothing accounts for the positive dependence of current on expected future output demand, while intertemporal substitution implies the negative effect of the ex-ante real interest rate. The interest rate elasticity of output, σ, corresponds to the elasticity of intertemporal subsitution of the consumers' utility function.

The second building block of the model is a "Phillips curve"-type equation, which derives from the optimal price setting decision of monopolistically competitive firms under the assumption of staggered price setting *a la* Calvo (1983),

$$\pi_t = \beta E_t\pi_{t+1} + \kappa x_t, \tag{1.2}$$

where β is the time discount factor and κ, "the slope" of the Phillips curve, is related inversely to the degree of price stickiness[8]. Since firms are unable to adjust prices optimally every period, whenever they have the opportunity to do so, they choose to price goods as a markup over a weighted average of current and expected future marginal costs. Under appropriate assumptions on technology and preferences, marginal costs are proportional to the output gap, resulting in the above Phillips curve. Here this relation is assumed to hold exactly, ignoring the so-called "cost-push" shock, which sometimes is appended *ad-hoc* to generate a short-term trade-off between inflation and output gap stabilization.

---

[8] In the underlying sticky price model the slope κ is given by $[\theta(1+\varphi\varepsilon)]^{-1}(1-\theta)(1-\beta\theta)(\sigma^{-1}+\varphi)$,where θ is the fraction of firms that keep prices unchanged in each period, φ is the (inverse) wage elasticity of labor supply, and ε is the elasticity of substitution among differentiated goods.

The final building block models the behavior of the monetary authority. The model assumes a "cashless limit" economy in which the instrument controlled by the central bank is the nominal interest rate. One possibility is to assume a benevolent monetary policy maker seeking to maximize the welfare of households. In that case, as shown in Woodford (2003), the problem can be cast in terms of a central bank that aims to minimize (under discretion or commitment) the expected discounted sum of losses from output gaps and inflation, subject to the optimal behavior of households (1.1) and firms (1.2), and in addition, the zero nominal interest rate floor:

$$\underset{i_t,\pi_t,x_t}{\text{Min}}\ E_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2\right) \tag{1.3}$$

$$\text{s.t. } (1.1),\ (1.2)$$

$$i_t \geq 0 \tag{1.4}$$

where λ is the relative weight of the output gap in the central bank's loss function.

An alternative way of modeling monetary policy is to assume that the central bank follows some sort of simple decision rule that relates the policy instrument, implicitly or explicitly, to other variables of the model. An example of such a rule, consistent with the zero floor, is a truncated Taylor rule,

$$i_t = \max\left[0,\ r^* + \pi^* + \phi_\pi\left(\pi_t - \pi^*\right) + \phi_x x_t\right] \tag{1.5}$$

where r^* is an equilibrium real rate, π^* is an inflation target, and ϕ_π and ϕ_x are response coefficients for inflation and the output gap.

To close the model one needs to specify the behavior of the natural real rate. In the fuller model the latter is a composite of a variety of real disturbances, including shocks to technology, preferences, and government spending. Following Woodford (2003) here

I assume that the natural real rate follows an exogenous mean-reverting process,

$$\widehat{r}_t^n = \rho \widehat{r}_{t-1}^n + \epsilon_t, \tag{1.6}$$

where $\widehat{r}_t^n \equiv r_t^n - r^*$, is the deviation of the natural real rate from its unconditional mean, r^*; ϵ_t are i.i.d. $\mathrm{N}(0, \sigma_\epsilon^2)$ real shocks, and $0 \leq \rho < 1$ is a persistence parameter.

The equilibrium conditions of the model therefore include the constraints (1.1), (1.2), and either a set of first-order optimality conditions (in the case of optimal policy), or a simple rule like (1.5). In either case the resulting system of equations cannot be solved with standard solution methods relying on local approximation because of the non-negativity constraint on the nominal rate. Hence I solve them with a global solution technique known as "collocation". The rational expectations equilibrium with occasionally binding constraint is solved by way of parametrizing expectations (Christiano and Fischer 2000), and is implemented with the MATLAB routines developed by Miranda and Fackler (2001). Appendix 1.A outlines the simulation algorithm, while the following sections report the results.

1.2.1 Baseline Calibration

The model's parameters are chosen to be consistent with the "standard" Woodford (2003) calibration to the US economy, which is based on Rotemberg and Woodford (1997) (Table 1.1). Thus, the slope of the Phillips curve (0.024), the weight of the output gap in the central bank loss function (0.003), the time discount factor (0.993), the mean (3% pa) and standard deviation (3.72%) of the natural real rate are all taken directly from Woodford (2003). The persistence (0.65) of the natural real rate is assumed to be between the one used by Woodford (2003) (0.35) and that estimated by

Structural parameters		
Discount factor	β	0.993
Real interest rate elasticity of output	σ	0.25
Slope of the Phillips curve	κ	0.024
Weight of the output gap in loss function	λ	0.003
Natural real rate parameters		
Mean (% per annum)	r^*	3%
Standard deviation (annual)	$\sigma(r^n)$	3.72%
Persistence (quarterly)	ρ	0.65
Simple instrument rule coefficients		
Inflation target (% per annum)	π^*	0%
Coefficient on inflation	ϕ_π	1.5
Coefficient on output gap	ϕ_x	0.5
Interest rate smoothing coefficient	ϕ_i	0

Table 1.1. Baseline calibration (quarterly unless otherwise stated)

Adam and Billi (2006) (0.8) using more recent data.[9] The real interest rate elasticity of aggregate demand (0.25)[10] is lower than the elasticity assumed by Eggertsson and Woodford (2003) (0.5), but as these authors point out, if anything, a lower degree of interest sensitivity of aggregate expenditure biases the results towards a more modest output contraction as a result of a binding zero floor.[11] In the simulations with simple rules, the baseline target inflation rate (0%) is consistent with the implicit zero target for inflation in the central bank's loss function. The baseline reaction coefficients on inflation (1.5), the output gap (0.5), and the lagged nominal interest rate (0) are standard in the literature on Taylor (1993)-type rules. Section 8 studies the sensitivity of the results to various parameter changes.

[9] These parameters for the shock process imply that the natural real interest rate is negative about 15% of the time on an annual basis. This is slightly more often than with the standard Woodford calibration (10%).

[10] This corresponds to a constant relative risk aversion of 4 in the underlying model.

[11] With the Woodford (2003) value of this parameter (6.25), the model predicts unrealistically large output shortfalls when the zero floor binds - e.g. an output gap around -30% for values of the natural real rate around -3%.

1.3 Optimal Discretionary Policy with Zero Floor

Abstracting from the zero floor, the solution to the discretionary optimization problem is well known (Clarida, Gali and Gertler 1999)[12]. Under full discretion, the central bank cannot manipulate the beliefs of the private sector and it takes expectations as given. The private sector is aware that the central bank is free to re-optimize its plan in each period and, therefore, in a rational expectations equilibrium, the central bank should have no incentives to change its plans in an unexpected way. In the baseline model with no endogenous state variables, the discretionary policy problem reduces to a sequence of static optimization problems in which the central bank minimizes current period losses by choosing the current inflation, output gap, and nominal interest rate as a function only of the exogenous natural real rate, r_t^n.

The solution without zero bound then is straightforward: inflation and the output gap are fully stabilized at their (zero) targets in every period and state of the world, while the nominal interest rate moves one-for-one with the natural real rate. This is depicted by the dashed lines in Figure 1.2. With this policy the central bank is able to achieve the globally minimal welfare loss of zero at all times.

With the zero floor, the basic problem of discretionary optimization (without endogenous state variables) can still be cast as a sequence of static problems. The Lagrangian can be written as,

$$L_t = \frac{1}{2}\left(\pi_t^2 + \lambda x_t^2\right) + \phi_{1t}\left[x_t - f_{1t} + \sigma\left(i_t - f_{2t}\right)\right] + \phi_{2t}\left[\pi_t - \kappa x_t - \beta f_{2t}\right] + \phi_{3t} i_t \quad (1.7)$$

where ϕ_{1t} is the Lagrange multiplier associated with the IS curve (1.1), ϕ_{2t} with the Phillips curve (1.2), and ϕ_{3t} with the zero constraint (1.4). The functions $f_{1t} = E_t\left(x_{t+1}\right)$,

[12] We restrict our attention to Markov-perfect equilibria here.

and $f_{2t} = E_t(\pi_{t+1})$ are private sector expectations which the central bank takes as given. Noticing that $\phi_{3t} = -\sigma\phi_{1t}$, the Kuhn-Tucker conditions for this problem can be written as:

$$
\begin{aligned}
\pi_t + \phi_{2t} &= 0 && (1.8)\\
\lambda x_t + \phi_{1t} - \kappa\phi_{2t} &= 0 && (1.9)\\
i_t\phi_{1t} &= 0 && (1.10)\\
i_t &\geq 0 && (1.11)\\
\phi_{1t} &\geq 0 && (1.12)
\end{aligned}
$$

Substituting (1.8) and (1.9) into (1.10), and combining the result with (1.1), (1.2), and (1.4), a Markov perfect rational expectations equilibrium should satisfy:

$$
\begin{aligned}
x_t - E_t x_{t+1} + \sigma\left(i_t - E_t\pi_{t+1} - r_t^n\right) &= 0 && (1.13)\\
\pi_t - \kappa x_t - \beta E_t\pi_{t+1} &= 0 && (1.14)\\
i_t\left(\lambda x_t + \kappa\pi_t\right) &= 0 && (1.15)\\
i_t &\geq 0 && (1.16)\\
\lambda x_t + \kappa\pi_t &\leq 0 && (1.17)
\end{aligned}
$$

Notice that (1.15) implies that the typical “targeting rule” involving inflation and the output gap is satisfied whenever the zero floor on the nominal interest rate is not binding,

$$
\begin{aligned}
\lambda x_t + \kappa\pi_t &= 0, && (1.18)\\
\text{if } i_t &> 0
\end{aligned}
$$

However, when the zero floor is binding, from (1.13) the dynamics are governed by

$$\begin{aligned} x_t + \sigma p_t - \sigma r_t^n &= E_t x_{t+1} + \sigma E_t p_{t+1} \\ \text{if } i_t &= 0 \end{aligned} \tag{1.19}$$

where p_t is the (log) price level. Note that it is no longer possible to set inflation and the output gap to zero at all times, for such a policy would require a negative nominal interest rate when the natural real interest rate falls below zero. Moreover, (1.19) implies that if the natural real rate falls so that the zero floor becomes binding, then since next period's output gap and price level are independent of today's actions, for expectations to be rational, the sum of the current output gap and price level must fall. The latter is true for any process for the natural real interest rate which allows it to take negative values.

An interesting special case, which replicates the findings of Jung et al. (2005), is the case of perfect foresight. By perfect foresight we mean that the natural real rate jumps initially to some (possibly negative) value, after which it follows a deterministic path (consistent with the expected path of an AR(1) process) back to its steady-state. In this case the policy functions are represented by the solid lines in Figure 1.2. As anticipated in the previous paragraph, at negative values of the natural real rate both the output gap and inflation are below target. On the other hand, at positive levels of the natural real rate, prices and output can be stabilized fully in the case of discretionary optimization with perfect foresight. The reason for this is simple: once the natural real rate is above zero, deterministic reversion to steady-state implies that it will never be negative in the future. This means that it can always be tracked one-for-one by the nominal rate (as in the case without zero floor), which is sufficient to fully stabilize prices and output.

One of the contributions of this paper is to extend the analysis in Jung et al. (2005) to the more general case in which the natural real rate follows a stochastic AR(1) process. Figure 1.3 plots the optimal discretionary policies in the stochastic environment. Clearly optimal discretionary policy differs in several important ways both from the optimal policy unconstrained by the zero floor and from the constrained perfect-foresight solution. First of all, given the zero floor, it is in general no longer optimal to set either inflation or the output gap to zero in any state of the world. In fact, in the solution with zero floor, inflation falls short of target at *any* level of the natural real rate. This gives rise to a "deflationary bias" of optimal discretionary policy, that is, an *average* rate of inflation below the target. Sensitivity analysis shows that for some plausible parameter values the deflationary bias becomes quantitatively significant[13]. This implies that any quantitative analysis of discretionary biases in monetary models that does not take into account the zero lower bound can be misleading.

Secondly, as in the case of perfect foresight, at negative levels of the natural real rate, both inflation and the output gap fall short of their respective targets. However, the deviations from target are larger in the stochastic case - up to 1.5 percentage points for the output gap, and up to 15 basis points for inflation at a natural real rate of -3% under the baseline calibration.

Third, above a positive threshold for the natural real rate, the optimal output gap becomes positive, peaking around +0.5%.

13 E.g. half a percentage point with $\rho = 0.8$ and $r^* = 2\%$.

Finally, at positive levels of the natural real rate the optimal nominal interest rate policy with zero floor is both *more expansionary* (i.e. prescribing a lower nominal rate), and *more aggressive* (i.e. steeper) compared to the optimal policy without zero floor[14]. As a result, the nominal rate hits the zero floor at levels of the natural real rate as high as 1.8% (and is constant at zero for lower levels of the natural real rate).

These results hinge on the combination of two factors: (1) the stochastic nature of the natural real rate; and (2) the non-linearity induced by the zero floor. The effect is an asymmetry in the ability of the central bank to respond to positive versus negative shocks when the natural real rate is close to zero. Namely, while the central bank can fully offset any positive shocks to the natural real rate because nothing prevents it from raising the nominal rate by as much as is necessary, it cannot fully offset large enough negative shocks. The most it can do in this case is to reduce the nominal rate down to zero, which is still higher than the interest rate consistent with zero output gap and inflation. Taking private sector expectations as given, the latter implies a higher than desired *real* interest rate, which depresses output and prices through the IS and Phillips curves.

The above asymmetry is reflected in the private sector's expectations: a positive shock in the following period is expected to be neutralized, while an equally probable negative one is expected to take the economy into a liquidity trap. This gives rise to a "deflationary bias" in the private sector's expectations, which in a forward-looking economy has an impact on the *current* evolution of output and prices. Absent an endogenous state, the current evolution of the economy is all that matters for welfare, and so it is optimal for the central bank to partially offset the depressing effect of expec-

[14] ... or compared to the optimal discretionary policy with zero floor and perfect foresight.

tations on today's outcome by more aggressive lowering of the nominal rate when the risk of deflation is high.

At sufficiently high levels of the natural real interest rate, the probability for the zero floor to become binding tends to zero. Hence optimal discretionary policy approaches the unconstrained case, namely zero output gap and inflation, and a nominal rate equal to the natural real rate. However, around the deterministic steady-state the difference between the two policies - with and without zero floor - remains significant.

We have seen that in the baseline model with no endogenous states, optimal discretionary policy is independent of history. This means that it is only the current risk of falling into a liquidity trap that matters for today's policy, regardless of whether the economy is approaching a liquidity trap or has just recovered from one. This is in sharp contrast with the optimal policy under commitment, which involves a particular type of history dependence as the following section shows.

1.4 Optimal Commitment Policy with Zero Floor

In the absence of the zero lower bound the equilibrium outcome under optimal discretion is globally optimal and therefore it is observationally equivalent to the outcome under optimal commitment policy. The central bank manages to stabilize fully inflation and the output gap while adjusting the nominal rate one-for-one with the natural real rate.

However, this observational equivalence no longer holds in the presence of a zero interest rate floor. While full stabilization under either regime is not possible, important gains can be obtained from the ability to commit to future policy. In particular, by committing to deliver inflation in the future, the central bank can affect private sector's expectations about inflation, and thus the real rate, even when the nominal interest rate is constrained by the zero floor. This channel of monetary policy in simply unavailable to a discretionary policy maker.

Using the Lagrange method as before, but this time taking into account the dependence of expectations on policy choices, it is straightforward to obtain the equilibrium conditions that govern the optimal commitment solution:

$$
\begin{aligned}
x_t - E_t x_{t+1} + \sigma \left(i_t - E_t \pi_{t+1} - r_t^n \right) &= 0 && (1.20) \\
\pi_t - \kappa x_t - \beta E_t \pi_{t+1} &= 0 && (1.21) \\
\pi_t - \phi_{1t-1} \sigma / \beta + \phi_{2t} - \phi_{2t-1} &= 0 && (1.22) \\
\lambda x_t + \phi_{1t} - \phi_{1t-1} / \beta - \kappa \phi_{2t} &= 0 && (1.23) \\
i_t \phi_{1t} &= 0 && (1.24) \\
i_t &\geq 0 && (1.25) \\
\phi_{1t} &\geq 0 && (1.26)
\end{aligned}
$$

From conditions (1.22) and (1.23) it is clear that the Lagrange multipliers inherited from the past period will have an effect on current policy. They in turn will depend on the history of endogenous variables and in particular on whether the zero floor was binding in the past. In this sense the Lagrange multipliers summarize the effect of optimal commitment, which in contrast to optimal discretionary policy, involves a particular type of history dependence.

Figures 1.4 to 1.6 plot the optimal policies in the case of commitment. The figures illustrate specifically the dependence of policy on ϕ_{1t-1}, the Lagrange multiplier associated with the zero floor, while holding ϕ_{2t-1} fixed. When the nominal interest rate is constrained by the zero floor ϕ_1 becomes positive, implying that the central bank commits to a lower nominal rate, higher inflation and higher output gap in the following period, conditional on the value of the natural real rate.

Since the commitment is assumed to be credible, it enables the central bank to achieve higher expected inflation and a lower real rate in periods when the nominal rate is constrained by the zero floor. The lower real rate reinforces expectations for higher future output and thus further stimulates current output demand through the IS curve. This, together with higher expected inflation stimulates current prices through the expectational Phillips curve. Commitment therefore provides an additional channel of monetary policy, which works through expectations and through the ex-ante real rate, and which is unavailable to a discretionary monetary policy maker.

A standard way to illustrate the differences between optimal discretionary and commitment policies is to compare the dynamic evolution of endogenous variables under each regime in response to a single shock to the exogenous natural real rate. Figures 1.7 and 1.8 plot the impulse-responses for a small and a large negative shock to the natural real rate respectively. Notice that in the case of a small shock to the natural real rate from its steady-state of 3% down to 2%, inflation and the output gap under optimal commitment policy (lines with circles) remain almost fully stabilized. In contrast, under discretionary optimization (lines with squares), inflation stays slightly below target and the output gap remains about half a percentage point above target, consistent with equation (1.18), as the economy converges back to its steady-state. The nominal interest rate under discretion is about 1% lower than the one under commitment throughout the simulation, yet it remains strictly positive at all times.

The picture changes substantially in the case of a large negative shock to the natural interest rate to -3%. Notably, under both commitment and discretion, the nominal interest rate hits the zero lower bound, and remains there until two quarters after the natural interest rate has returned to positive.[15] Under discretionary optimization, both inflation and the output gap fall on impact, consistent with equation (1.19), after which they converge towards their steady-state. The initial shortfall is significant, especially for the output gap, amounting to about 1.5%. In contrast, under the optimal commitment rule the initial output loss and deflation are much milder, owing to the ability of the central bank to commit to a positive output gap and inflation once the natural real rate has returned to positive.

An alternative way to compare optimal discretionary and commitment policies in the stochastic environment is to juxtapose the dynamic paths that they prescribe for endogenous variables under a particular path for the stochastic natural real rate[16]. The experiment is shown in figure 1.9, which plots a simulated "liquidity trap" under the two regimes. The line with triangles in the bottom panel is the assumed evolution of the natural real rate. It slips down from +3 percent (its deterministic steady-state) to -3 percent over a period of 15 quarters, then remains at -3 percent for 10 quarters, before recovering gradually (consistent with the assumed AR(1) process) to around +3 percent in another 15 quarters.

The first and the second panels of figure 1.9 show the responses of inflation and the output gap under each of the two regimes. Not surprisingly, under the optimal

---

[15] That the zero interest rate policy should terminate in the same quarter under commitment and under discretion is a coincidence in this experiment. The relative duration of a zero interest rate policy under commitment versus discretion depends on the parameters of the shock process as well as the particular realization of the shock.

[16] Notice that agents observe only the current state but form expectations about the future evolution of the natural real rate.

commitment regime both inflation and the output gap are closer to target than under the optimal discretionary policy. In particular, under optimal discretionary policy inflation is always below the target as it falls to -0.15% shadowing the drop in the natural real rate. Compared to that, under optimal commitment prices are almost fully stabilized, and in fact they even slightly increase while the natural real rate is negative. In turn, under discretionary optimization the output gap is initially around +0.4% but then it declines sharply to -1.6% with the decline in the natural real rate. In contrast, under optimal commitment, output is initially at its potential level and the largest negative output gap is only half the size of the one under optimal discretionary policy.

Supporting these paths of inflation and the output gap are corresponding paths for the nominal interest rate. Under discretionary optimization the nominal rate starts at around 2% and declines at an increasing rate until it hits zero two quarters before the natural real rate has turned negative. It is then kept at zero while the natural real rate is negative, and only two quarters after the latter has returned to positive territory does the nominal interest rate start rising again. Nominal rate increases following the liquidity trap mirror the decreases while approaching the trap, so that the tightening is more aggressive in the beginning and then gradually diminishes as the nominal rate approaches its steady-state.

In contrast, the nominal rate under optimal commitment begins closer to 3%, then declines to zero one quarter before the natural real rate turns negative. After that, it is kept at its zero floor until three quarters after the recovery of the natural real rate to positive levels, that is one quarter longer compared to optimal discretionary policy. Interestingly, once the central bank starts increasing the nominal rate, it raises it very quickly - the nominal rate climbs nearly 3 percentage points in just two quarters. This is equivalent to six consecutive monthly increases by 50 basis points each. The reason is that once the central bank has validated the inflationary expectations it has created (helping mitigate deflation during the liquidity trap), there is no more incentive (and it is costly) to keep inflation above the target.

Under discretion, the paths of inflation, output and the nominal rate are symmetric with respect to the midpoint of the simulation period because optimal discretionary policy is independent of history. Therefore, inflation and the output gap inherit the dynamics of the natural real rate, the only state variable on which they depend. This is in contrast with the asymmetric paths of the endogenous variables under commitment, reflecting the optimal history dependence of policy under this regime. In particular, the fact that under commitment the central bank can promise higher output gap and inflation in the wake of a liquidity trap is precisely what allows it to engage in less preemptive easing of policy in anticipation of the trap, and at the same time deliver a superior inflation and output gap performance compared to the optimal policy under discretion.

1.5 Targeting Rules with Zero Floor

In the absence of the zero floor, targeting rules take the form

$$\alpha_\pi E_t \pi_{t+j} + \alpha_x E_t x_{t+k} + \alpha_i E_t i_{t+l} = \tau, \tag{1.27}$$

where α_π, α_x, α_i are weights assigned to the different objectives, j, k and l are forecasting horizons, and τ is the target. These are sometimes called *flexible* inflation targeting rules to distinguish them from *strict* inflation targeting of the form $E_t \pi_{t+j} = \tau$. When

j, k or $l > 0$, the rules are called inflation *forecast* targeting to distinguish them from targeting contemporaneous variables.

As demonstrated by (1.19) in section 1.3, in general such rules are not consistent with the zero floor for they would require negative nominal interest rates at times. A natural way to modify targeting rules so that they comply with the zero floor is to write them as a complementarity condition,

$$\begin{aligned} i_t \left(\alpha_\pi E_t \pi_{t+j} + \alpha_x E_t x_{t+k} + \alpha_i E_t i_{t+l} - \tau \right) &= 0 \\ i_t &\geq 0 \end{aligned} \tag{1.28}$$

which requires that either the target τ is met, or else the nominal interest rate should be bounded below by zero. In this sense, a rule like (1.28) can be labelled "flexible inflation targeting with a zero interest rate floor".

In fact, section 1.3 showed that the optimal policy under discretion takes this form with $\alpha_i = 0$, $\alpha_\pi = \kappa$, $\alpha_x = \lambda$, $j = k = 0$, and $\tau = 0$, namely

$$\begin{aligned} i_t \left(x_t + \frac{\kappa}{\lambda} \pi_t \right) &= 0 \\ i_t &\geq 0 \end{aligned} \tag{1.29}$$

In the absence of the zero floor it is known that optimal commitment policy can be formulated as optimal "speed limit targeting",

$$\Delta x_t + \frac{\kappa}{\lambda} \pi_t = 0, \tag{1.30}$$

where $\Delta x_t = \Delta y_t - \Delta y_t^n$ is the growth rate of actual output relative to the growth rate of its natural (flexible price) counterpart (the "speed limit"). In contrast to discretionary optimization however, the optimal commitment rule with zero floor cannot be written in the form (1.28). This is so because with zero floor the optimal target involves a particular type of history dependence as shown by Eggertsson and Woodford (2003)[17]. In particular, manipulating the first order conditions of the optimal commitment problem, one can arrive at the following speed limit targeting rule with zero floor:

$$\begin{aligned} i_t \left[\Delta x_t + \frac{\kappa}{\lambda}\pi_t - \frac{1}{\lambda}\left(\frac{\kappa\sigma+\beta}{\beta}\phi_{1t-1} - \phi_{1t} + \frac{1}{\beta}\Delta\phi_{1t-1} \right) \right] &= 0 \\ i_t &\geq 0. \end{aligned} \tag{1.31}$$

Since the product $\kappa\sigma$ is small and β is close to one, and for plausible (small) values of ϕ_{1t} consistent with the assumed stochastic process for the natural real rate, the above rule can be approximated by

$$\begin{aligned} i_t \left[\Delta y_t + \frac{\kappa}{\lambda}\pi_t - \tau_t \right] &= 0 \\ i_t &\geq 0. \end{aligned} \tag{1.32}$$

where $\tau_t \approx \Delta y_t^n + \lambda^{-1}\Delta^2\phi_{1t}$ is a history-dependent target (speed limit). In normal circumstances when $\phi_{1t} = \phi_{1t-1} = \phi_{1t-2} = 0$, the target is equal to the growth rate of flexible price output, as in the problem without zero bound; however if the economy falls into a liquidity trap, the speed limit is adjusted in each period by the speed of change of the penalty (the Lagrange multiplier) associated with the non-negativity constraint. The faster the economy is plunging into the trap, therefore, the higher is the speed limit target which the central bank promises to deliver in the future, contingent on the natural interest rate's return to positive territory.

While the above rule is optimal in this framework, it is not likely to be very practical. Its dependence on the unobservable Lagrange multipliers makes it very hard to implement or communicate to the public. Moreover, as pointed out by Eggertsson

[17] These authors derive the optimal commitment policy in the form of a moving *price level* targeting rule. Alternatively, it can be formulated as a moving *speed limit* taregting rule as demonstrated here.

and Woodford (2003), credibility might suffer if all that the private sector observes is a central bank which persistently undershoots its target yet keeps raising it for the following period. To overcome some of these drawbacks Eggertsson and Woodford (2003) propose a simpler constant price level targeting rule, of the form

$$\begin{aligned} i_t \left[x_t + \frac{\kappa}{\lambda} p_t \right] &= 0 \\ i_t &\geq 0 \end{aligned} \tag{1.33}$$

where p_t is the log price level.[18]

The idea is that committing to a price level target implies that any undershooting of the target resulting from the zero floor is going to be undone in the future by positive inflation. This raises private sector expectations and eases deflationary pressures when the economy is in a liquidity trap. Figure 1.10 demonstrates the performance of this simpler rule in a simulated liquidity trap. Notice that while the evolution of the nominal rate and the output gap is similar to that under the optimal discretionary rule, the path of inflation is much closer to the target. Since the weight of inflation in the central bank's loss function is much larger than that of the output gap, the fact that inflation is better stabilized accounts for the superior performance of this rule in terms of welfare.

1.6 Simple Instrument Rules with Zero Floor

The practical difficulties with communicating and implementing optimal rules like (1.31) or even (1.33) have led many researchers to focus on simple instrument rules of the type proposed by Taylor (1993). These rules have the advantage of postulating a relatively

[18] Notice that the weight on the price level is optimal within the class of constant price level targeting rules. In particular, it is related to $\kappa/\lambda = \varepsilon$, the degree of monopolistic competition among intermediate goods producers.

straightforward relationship between the nominal interest rate and a limited set of variables in the economy. While the advantage of these rules lies in their simplicity, at the same time - absent the zero floor - some of them have been shown to perform close enough to the optimal rules in terms of the underlying policy objectives (Galí 2003). Hence, it has been argued that some of the better simple instrument rules may serve as a useful benchmark for policy, while facilitating communication and transparency.

In most of the existing literature, however, simple instrument rules are specified as linear functions of the endogenous variables. This is in general inconsistent with the zero floor because for large enough negative shocks (e.g. to prices), linear rules would imply a negative value for the nominal interest rate. For instance, a simple instrument rule reacting only to past period's inflation,

$$i_t = r^* + \pi^* + \phi_\pi \left(\pi_{t-1} - \pi^*\right), \tag{1.34}$$

where r^* is the equilibrium real rate, π^* is the target inflation rate, and ϕ_π is an inflation response coefficient, can clearly imply negative values for the nominal rate.

In the context of liquidity trap analysis a natural way to modify simple instrument rules is to truncate them at zero with the $\max(\cdot)$ operator. For example, the truncated counterpart of the above Taylor rule can be written as

$$i_t = \max\left[0,\ r^* + \pi^* + \phi_\pi \left(\pi_{t-1} - \pi^*\right)\right]. \tag{1.35}$$

In what follows I consider several types of truncated instrument rules, including:

1. Truncated Taylor Rules (TTR) that react to past, contemporaneous or expected future values of the output gap and inflation ($j = -1, 0, 1$),

$$i_t^{TTR} = \max\left[0,\ r^* + \pi^* + \phi_\pi \left(E_t\pi_{t+j} - \pi^*\right) + \phi_x \left(E_t x_{t+j}\right)\right], \tag{1.36}$$

2. TTR rules with partial adjustment or "interest rate smoothing" (TTRS),

$$i_t^{TTRS} = \max\left\{0,\ \phi_i i_{t-1} + (1-\phi_i)\, i_t^{TTR}\right\} \tag{1.37}$$

3. TTR rules that react to the price *level* instead of inflation (TTRP),

$$i_t^{TTRP} = \max\left[0,\ r^* + \phi_\pi (p_t - p^*) + \phi_x x_t\right] \tag{1.38}$$

   where p_t is the log price level and p^* is a constant price level target; and

4. Truncated "first-difference" rules (TFDR) that specify the *change* in the interest rate as a function of the output gap and inflation,

$$i_t^{TFDR} = \max\left[0,\ i_{t-1} + \phi_\pi (\pi_t - \pi^*) + \phi_x x_t\right]. \tag{1.39}$$

The last formulation ensures that if the nominal interest rate ever hits zero it will be held there as long as inflation and the output gap are negative, thus extending the potential duration of a zero interest rate policy relative to a truncated Taylor rule.

I illustrate the performance of each family of rules by simulating a liquidity trap and plotting the implied paths of endogenous variables under each regime. The evaluation of average performance and unconditional welfare is left for the following section.

Given the model's simplicity the focus here is not on finding the optimal values of the parameters within each class of rules but rather on evaluating the performance of alternative monetary policy regimes. To do that I use values of the parameters commonly estimated and widely used in simulations in the literature. I make sure that the parameters satisfy a sufficient condition for local uniqueness of equilibrium. Namely, the parameters are required to observe the so-called Taylor principle according to which the nominal interest rate must be adjusted more than one-to-one with changes in the rate of inflation, implying $\phi_\pi > 1$. I further restrict $\phi_x \geq 0$ and $0 \leq \phi_i \leq 0.8$.

Figure 1.11 plots the dynamic paths of inflation, the output gap and the nominal interest rate which result under regimes TTR and TTRP, conditional on the same path for the natural real rate as before. Both the truncated Taylor rule (TTR, lines with squares) and the truncated rule responding to the price level (TTRP, lines with circles) react contemporaneously with coefficients $\phi_\pi = 1.5$ and $\phi_x = 0.5$, and $\pi^* = 0$.

Several features of these plots are worth noticing. First of all, and not surprisingly, under the truncated Taylor rule, inflation, the output gap, and the nominal rate inherit the behavior of the natural real rate. Perhaps less expected though, while both inflation and especially the output gap deviate further from their targets compared to the optimal rules in figure 1.9, the nominal interest rate stays always above one percent, even when the natural real rate falls as low as -3 percent! This suggests that - contrary to popular belief - an equilibrium real rate of 3% may provide a sufficient buffer from the zero floor even with a truncated Taylor rule targeting zero inflation.

Secondly, figure 1.11 demonstrates that in principle the central bank can do even better than TTR by reacting to the price level rather than to the rate of inflation. The reason for this is clear - by committing to react to the price level the central bank promises to undo any past disinflation by higher inflation in the future. As a result when the economy is hit by a negative real rate shock current inflation falls by less because expected future inflation increases.

Figure 1.12 plots the dynamic paths of endogenous variables under regimes TTRS and TFDR, again with $\phi_\pi = 1.5$, $\phi_x = 0.5$ and $\pi^* = 0$. TTRS (lines with circles) is a partial adjustment version of TTR, with smoothing coefficient $\phi_i = 0.8$. TFDR (line with squares) is a truncated first-difference rule which implies more persistent deviations of the nominal interest rate from its steady-state level.

The figure suggests that interest rate smoothing (TTRS) may improve somewhat on the truncated Taylor rule (TTR), and may do a bit worse than the rule reacting to the price level (TTRP). However, it implies the least instrument volatility. On the other hand, the truncated first-difference rule (TFDR) seems to be doing the best job at stabilization in a liquidity trap among the examined four simple instrument rules. However, under this rule the nominal interest rate deviates most from its steady-state, hitting zero for five quarters. Interestingly, the paths for inflation and the output gap under this rule resemble, at least qualitatively, those under the optimal commitment policy. This suggests that introducing a substantial degree of interest rate inertia may be approximating the optimal history dependence of policy implemented by the optimal commitment rule.

It is important to keep in mind that the above simulations are conditional on one particular path for the natural real rate. It is of course possible that a rule which appears to perform well while the economy is in a liquidity trap, turns out to perform badly "on average". In the following section I undertake the ranking of alternative rules according to an unconditional expected welfare criterion, which takes into account the stochastic nature of the economy, time discounting, as well as the relative cost of inflation vis-a-vis output gap fluctuations.

1.7 Welfare Ranking of Alternative Rules

A natural criterion for the evaluation of alternative monetary policy regimes is the central bank's loss function. Woodford (2003) shows that under appropriate assumptions the latter can be derived as a second order approximation to the utility of the representative consumer in the underlying sticky price model.[19] Rather than normalizing the

[19] Arguably, Woodford's (2003) approximation to the utility of the representative consumer is accurate

weight of inflation to one, I normalize the loss function so that utility losses arising from deviations from the flexible price equilibrium can be interpreted as a fraction of steady-state consumption,

$$\begin{aligned} WL &= \frac{\overline{U}-U}{U_c C} = \frac{1}{2}E_0\sum_{t=0}^{\infty}\beta^t\left[\varepsilon\left(1+\varphi\varepsilon\right)\zeta^{-1}\pi_t^2+\left(\sigma^{-1}+\varphi\right)x_t^2\right] \\ &= \frac{1}{2}\varepsilon\left(1+\varphi\varepsilon\right)\zeta^{-1}E_0\sum_{t=0}^{\infty}\beta^t L_t \end{aligned} \tag{1.40}$$

where $\zeta=\theta^{-1}\left(1-\theta\right)\left(1-\beta\theta\right)$; θ is the fraction of firms that keep prices unchanged in each period, φ is the (inverse) elasticity of labor supply, and ε is the elasticity of substitution among differentiated goods. Notice that $\left(\sigma^{-1}+\varphi\right)\left[\varepsilon\left(1+\varphi\varepsilon\right)\right]^{-1}\zeta=\kappa/\varepsilon=\lambda$ implies the last equality in the above expression, where L_t is the central bank's period loss function, which is being minimized in (1.3).

I rank alternative rules on the basis of the *unconditional* expected welfare. To compute it, I simulate 2000 paths for the endogenous variables over 1000 quarters and then compute the average loss per period across all simulations. For the initial distribution of the state variables I run the simulation for 200 quarters prior to the evaluation of welfare. Table 1.2 ranks all rules according to their welfare score. It also reports the volatility of inflation, the output gap and the nominal interest rate under each rule, as well as the frequency of hitting the zero floor.

A thing to keep in mind in evaluating the welfare losses is that in the benchmark model with nominal price rigidity as the only distortion and a shock to the natural real rate as the only source of fluctuations, absolute welfare losses are quite small - typically less than one hundredth of a percent of steady-state consumption for any sensible mon-

to second order only in the vicinity of the steady-state with zero inflation. To the extent that the shock inducing a zero interest rate pushes the economy far away from the steady-state, the approximation error could in principle be large. In that case, the welfare evaluation provided here can be interpreted as a relative ranking of alternative policies based on an *ad hoc* central bank loss criterion. Studying the welfare implication of different rules in the fully non-linear model lies outside the scope of this paper.

etary policy regime.[20] Therefore, the focus here is on evaluating rules on the basis of their welfare performance *relative* to that under the optimal commitment rule.

In particular, in terms of unconditional expected welfare, the optimal discretionary policy (ODP) delivers losses which are nearly eight times larger than the ones achievable under the optimal commitment policy (OCP). Recall that abstracting from the zero floor and in the absence of shocks other than to the natural real rate, the outcome under discretionary optimization is the same as under the optimal commitment rule. Hence, the cost of discretion is substantially understated in analyses which ignore the existence of the zero lower bound on nominal interest rates. Moreover, *conditional* on the economy's fall into a liquidity trap the cost of discretion is even higher.

Interestingly, the frequency of hitting the zero floor is quite high - around one third of the time under both optimal discretionary and optimal commitment policy. This however depends crucially on the central bank's targeting zero inflation in the baseline model without indexation. If instead the central bank targets a rate of inflation of, say 2%, the frequency of hitting the zero floor would decrease to around 12% of the time (which is still higher than what is observed in the US).

Table 1.2 further confirms Eggertsson and Woodford (2003) intuition about the desirable properties of an (optimal) constant price level targeting rule (PLT) - here losses are only 56% greater than those under the optimal commitment rule. It also involves hitting the zero floor around one third of the time.

In comparison, losses under the truncated first-difference rule (TFDR) are 7.5 times as large as those under the optimal commitment rule (even though its performance in a liquidity trap seemed comparable to that of PLT). Interestingly, however, TFDR

[20] To be sure, output gaps in a liquidity trap are considerable; however the output gap is attributed negligible weight in the central bank loss function of the benchmark model.

	OCP	PLT	TFDR	ODP	TTRP	TTRS	TTR
std(π) x 10^2	1.04	3.47	4.59	3.85	7.23	9.12	12.9
std(x)	0.45	0.69	1.04	0.71	1.61	1.91	1.90
std(i)	3.21	3.20	1.36	3.27	1.06	0.56	1.14
Loss x 10^5	6.97	10.9	52.3	54.2	62.9	103	147
Loss/OCP	1	1.56	7.50	7.77	9.01	14.8	21.1
$\Pr(i=0)$ %	32.6	32.0	1.29	36.8	0.24	0.00	0.44

Table 1.2. Properties of Optimal and Simple Rules with Zero Floor

narrowly outperforms optimal discretionary policy. Even though the implied volatility of inflation and the output gap is slightly higher under this rule, it does a better job than ODP at keeping inflation and the output gap closer to target on average. An additional advantage - albeit one that is not reflected in the benchmark welfare criterion - is that instrument volatility is less than half of that under any of the policies ODP, OCP or PLT. This implies that the zero floor is hit only around 1.3% of the time under this rule.

Similarly, losses under the truncated Taylor rule reacting to the price level (TTRP) are nine times larger than under OCP, and only slightly worse than optimal discretionary policy. Moreover, instrument volatility under this rule is smaller than under TFDR, which implies hitting the zero floor even more rarely - only one quarter every 100 years on average.

The rule with the least instrument volatility among the studied simple rules - less than one-fifth of that under OCP - is the truncated Taylor rule with smoothing (TTRS). As a consequence, under this rule the nominal interest rate virtually never hits the zero lower bound. However, welfare losses are almost fifteen times larger than under OCP.

Finally, under the simplest truncated Taylor rule (TTR) without smoothing the zero lower bound is hit only two quarters every 100 years, while welfare losses are around 20 times larger than those under OCP. Nevertheless, even under this simplest rule losses are very small in absolute terms.

The fact that the zero lower bound is hit so rarely under the four simple instrument rules suggests that the zero constraint may not be playing a big role for unconditional expected welfare under these regimes. Indeed, computing their welfare score without the zero floor (by removing the max operator), reveals that close to 99% of the welfare losses associated with the four simple instrument rules stem from their intrinsic sub-optimality rather than from the zero floor per se. Put differently, if one reckons that the stabilization properties of a standard Taylor rule are satisfactory in an environment in which nominal rates can be negative, then adding the zero lower bound to it leaves unconditionally expected welfare virtually unaffected. Nevertheless, as was illustrated in the previous section, *conditional* on a negative evolution of the natural real rate, the losses associated with most of the studied simple instrument rules are substantially higher relative to the optimal policy.

1.8 Sensitivity Analysis

In this section I analyze the sensitivity of the main findings with respect to the parameters of the shock process, the strength of reaction and the timing of variables in truncated Taylor-type rules, and an extension of the model with endogenous inflation persistence.

1.8.1 Parameters of the Natural Real Rate Process

Larger variance

Table 1.3 reports the effects of an increase of the standard deviation of r^n to 4.5% (a 20% increase), while keeping the persistence constant, under three alternative

	OCP	ODP	TTR
std(r^n)	4.46	4.46	4.46
std(π)	x 1.52	x 1.50	x 1.20
std(x)	x 1.46	x 1.45	x 1.20
std(i)	x 1.14	x 1.14	x 1.19
Loss	x 2.12	x 2.60	x 1.42
$\Pr(i = 0)$ %	x 1.23	x 1.19	x 3.24

Table 1.3. Properties of Selected Rules with Higher std(r^n)

regimes - optimal commitment policy, discretionary optimization and a truncated Taylor rule.

Under OCP, the zero floor is hit around 23% more often, while welfare losses more than double. Figure 1.13 shows that the higher volatility implies that both the preemptive easing of policy and the commitment to future loosening are slightly stronger. In contrast, figure 1.14 shows that under ODP preemptive easing is much stronger, the deflation bias bigger, and table 1.3 shows that welfare losses increase by a factor of 2.6. Finally, under TTR, the zero floor is hit three times more often, while welfare losses are up by 40%.

Stronger persistence of shocks

Table 1.4 and figures 1.15-1.17 show the effect of an increase in the persistence of shocks to the natural real rate to 0.8, while keeping the variance of r^n unchanged.

Under OCP (figure 1.15), preemptive easing is a bit stronger while future monetary loosening is much more prolonged. As a result of the stronger persistence, welfare losses under OCP more than double. Under ODP (figure 1.16), preemptive easing is much stronger, the deflation bias is substantially larger, and welfare losses increase by a factor of 5.5. And under TTR (figure 1.17), deviations of inflation and the output gap

	OCP	ODP	TTR
$\rho(r^n)$	0.80	0.80	0.80
std(π)	x 2.14	x 2.94	x 2.44
std(x)	x 1.55	x 1.76	x 1.42
std(i)	x 1.02	x 1.04	x 1.52
Loss	x 2.69	x 5.47	x 4.39
$\Pr(i = 0)$ %	x 1.09	x 1.21	x 11.2

Table 1.4. Properties of Selected Rules with More Persistent $\widehat{r}^n$

	OCP	ODP	TTR
r^*	2%	2%	2%
std(π)	x 1.74	x 1.79	x 1.00
std(x)	x 1.55	x 1.62	x 1.00
std(i)	x 0.89	x 0.87	x 0.98
Loss	x 2.55	x 4.47	x 1.48
$\Pr(i = 0)$ %	x 1.51	x 1.59	x 9.27

Table 1.5. Properties of Selected Rules with Lower r^*

from target become larger and more persistent, the frequency of hitting the zero floor increases by a factor of 11, and welfare losses more than quadruple.

Lower mean

The effects of a lower steady state of the natural real rate at 2% - keeping the variance and persistence of r^n constant - are illustrated in figures 1.18 and 1.19 and summarized in table 1.5.[21]

Under OCP, preemptive easing is a bit stronger while future monetary policy loosening is much more prolonged; losses more than double. Interestingly, under ODP preemptive easing is so strong that the nominal rate is zero more than half of the time. The deflation bias is larger, and losses increase by a factor of 4.5. And under TTR, the zero floor is hit nine times more often while losses increase by 50%.

[21] Notice that for simple rules such as TTR, it is the sum $r^* + \pi^*$ which provides a "buffer" against the zero lower bound. Therefore, up to a shift in the rate of inflation, varying r^* is equivalent to testing for sensitivity with respect to π^*.

ϕ_π	1.01	1.5	2	3	10	100
ϕ_x						
0	x 1.89	x 1.80	x 1.71	x 1.57	x 0.95	x 0.18
0.5	x 1.03	x 1	x 0.97	x 0.90	x 0.59	x 0.18
1	x 0.65	x 0.63	x 0.62	x 0.58	x 0.50	x 0.17
2	n.a.	x 2.05	x 1.13	x 0.64	x 0.41	x 0.23
3	n.a.	n.a.	x 4.44	x 1.92	x 0.45	x 0.27

Table 1.6. Relative Losses under TTRs with Different Response Coefficients

1.8.2 Instrument Rule Specification

The strength of response

Table 1.6 reports the dependence of welfare losses on the size of response coefficients in a truncated Taylor rule. It turns out that losses can be reduced substantially by having the interest rate react more aggressively to inflation and output gap deviations from target. For instance, losses are almost halved with $\phi_\pi = 3$ and $\phi_x = 1$, relative to the benchmark case with $\phi_\pi = 1.5$ and $\phi_x = 0.5$. And they are reduced further to one fifth, with $\phi_\pi = 100$.[22]

Forward, contemporaneous or backward-looking reaction

For given response coefficients of a truncated Taylor rule, welfare losses turn out to be smallest under a backward-looking rule and highest under a forward-looking specification. While losses are still small in absolute value, with $\phi_\pi = 1.5$ and $\phi_x = 0.5$ they are up by 25% under the forward-looking rule and are around 15% lower under the backward-looking rule, relative to the contemporaneous one. The frequency of hitting the zero floor is similarly higher under a forward-looking specification and lower under a backward-looking one. The reason for the dominance of the backward looking rule

[22] It is interesting to explore the behavior of truncated Taylor rules as $\phi_\pi \to \infty$ and $\phi_x \to \infty$. However, the algorithm used here runs into convergence problems with $\phi_\pi > 100$ and $\phi_x > 3$.

can be that under it the interest rate tends to be kept lower following periods of deflation, in a way that resembles the optimal history dependence under commitment. On the other hand, under forward-looking rules, the effective response to a given shock to the natural real rate is lower, given the assumed autoregressive nature of the natural real rate.

1.8.3 Endogenous Inflation Persistence and the Zero Floor

Wolman (1998) among others has argued that stickiness of inflation is crucial in generating costs of deflation associated with the zero floor. To follow up on this hypothesis, I extend the present framework by incorporating endogenous inflation persistence.[23] One way lagged dependence of inflation may result is if firms that do not reoptimize prices index them to past inflation. In this case the (log-linearized) inflation dynamics can be represented with the following modified Phillips curve (Christiano, Eichenbaum and Evans (2001), Woodford (2003)):

$$\widehat{\pi}_t = \beta E_t \widehat{\pi}_{t+1} + \kappa x_t \tag{1.41}$$

where $\widehat{\pi}_t = \pi_t - \gamma \pi_{t-1}$ is a quasi-difference of inflation and γ measures the degree of price indexation.

A thing to keep in mind is that in principle the welfare-relevant loss function is endogenous to the structure of the model. Hence strictly speaking one cannot compare welfare in the two environments - with and without inflation persistence - using the same loss criterion. On the other hand, Woodford (2003) shows that in the case of indexation to past inflation, the welfare-relevant loss function takes the same form as

---

23 In this case the discretionary optimization problem becomes dynamic with lagged inflation as a relevant state variable.

	PLT	PLT	ODP	ODP	TTR	TTR
γ	0.5	0.8	0.5	0.8	0.5	0.8
std(π)	x 1.46	x 2.21	x 1.80	x 1.02	x 1.54	x 2.35
std(x)	x 1.01	x 1.02	x 1.12	x 0.81	x 0.98	x 0.95
std(i)	x 1.00	x 1.00	x 1.01	x 1.02	x 1.06	x 1.10
Loss	x 1.50	x 2.77	x 3.74	x 2.66	x 1.83	x 3.70
Pr($i = 0$) %	x 1.00	x 1.00	x 1.04	x 1.15	x 1.46	x 2.00

Table 1.7. Performance of Selected Rules with Endogenous Inflation Persistence

(1.3), except that inflation is replaced by its quasi-difference $\widehat{\pi}_t$. This implies that inflation persistence (as measured by γ) does not affect welfare under an optimal targeting rule which takes into account the existing degree of economy-wide indexation. Nevertheless, for the sake of completeness, I compute and report *ad-hoc* the same criteria as in the baseline model, in addition to other reported statistics in the case of intrinsic inflation persistence.

Table 1.7 reports the properties of selected regimes relative to an environment without endogenous inflation persistence. Under the optimal constant price level targeting rule, an increase in the persistence of inflation to 0.8 results in doubling of inflation volatility and almost tripling of the baseline loss measure. Similarly, inflation volatility more than doubles and losses nearly quadruple under the baseline truncated Taylor rule when the stickiness of inflation increases to 0.8. Interestingly, the properties of optimal discretionary policy are found to depend in a non-linear way on the degree of inflation persistence. Namely, while an increase in inflation persistence to 0.5 raises inflation volatility by 80% and nearly quadruples losses, a further increase of inflation persistence to 0.8 leads to relatively smaller inflation volatility and losses. The reason is that high inflation persistence in this case serves as an additional channel of policy, making it possible for the central bank to “steer away” from an approaching liquidity trap by choosing higher current inflation.

1.9 Conclusions

Recent treatments of the zero lower bound issue have suffered from some important limitations. These include assuming perfect foresight or forcing certainty equivalence, or treating the zero floor as an initial condition rather than an occasionally binding non-negativity constraint. This paper addresses these issues, providing a global solution to a standard stochastic sticky price model with an explicit occasionally binding ZLB on the nominal interest rate. As it turns out, the dynamics, and in some cases the unconditional means, of the nominal rate, inflation and the output gap are strongly affected by uncertainty in the presence of the zero interest rate floor.

In particular, optimal discretionary policy involves a deflationary bias and interest rates are cut more aggressively when the risk of deflation is high, implying that they are kept lower both before and after a liquidity trap. The extent of such lowering of rates is found to increase in the variance and persistence of shocks to the natural real rate, and to decrease in its unconditional mean. Moreover, the preemptive lowering of rates is even more important under discretionary policy in the presence of endogenous inflation persistence. Compared to that, under optimal commitment policy the need for preemptive lowering of interest rates is limited since, conditional on a bad state, the central bank can commit to a period of looser monetary policy in the future, once the economy's has recovered from a possible liquidity trap.

Imposing the zero lower bound correctly in the stochastic model allows us to evaluate quantitatively the performance of a variety of monetary policy regimes. Thus, commitment to the optimal rule reduces welfare losses to one-tenth of those achievable under discretionary policy. Constant price level targeting delivers losses which are only

60% bigger than under the optimal policy. In contrast, under a truncated Taylor rule losses are 20 times bigger than under the optimal policy. Another interesting finding is that the unconditional welfare losses associated with simple instrument rules are almost unaffected by the zero lower bound *per se* and instead derive from the suboptimal responses to shocks characteristic of simple rules. This is related to the fact that under simple instrument rules the zero lower bound is hit very rarely, while optimal policy involves a zero nominal interest rate around one third of the time.

In fact in an extension of the model with money, optimal policy might be expected to visit the liquidity trap even more often. Hitting the zero lower bound in that case would be good because it eliminates the opportunity cost of holding cash balances. An interesting question to address in that setup would be how the optimal mean of the nominal interest rate is affected by the existence of the zero lower bound. Solving the fully non-linear problem would be another useful extension, which however increases the dimensionality of the computational problem. A limitation of the solution technique employed here is that it is practical only for models with a limited number of states.

1.A Numerical Algorithm

This section illustrates the algorithm used to solve the problem in the case of discretionary optimization. The cases with commitment and with simple rules are solved in a similar way. I apply the routines for rational expectations models included in the COMPECON toolkit of Miranda and Fackler (2001). These solve for the optimal response x as a function of the state s, when equilibrium responses are governed by an arbitrage-complementarity condition of the form

$$f\left[s_t, x_t, E_t h\left(s_{t+1}, x_{t+1}\right)\right] = \phi_t \tag{1.42}$$

where s follows the state transition function

$$s_{t+1} = g\left(s_t, x_t, \varepsilon_{t+1}\right) \tag{1.43}$$

and x_t and ϕ_t satisfy the complementarity conditions

$$a(s_t) \leqslant x_t \leqslant b(s_t), \quad x_{jt} > a_j(s_t) \Rightarrow \phi_{jt} \leqslant 0, \quad x_{jt} < b_j(s_t) \Rightarrow \phi_{jt} \geqslant 0, \tag{1.44}$$

where ϕ_t is a vector whose j^{th} element, ϕ_{jt}, measures the marginal loss from activity j. In equilibrium, ϕ_{jt} must be non-positive (nonnegative) if x_{jt} is greater (less) than its lower (upper) bound, otherwise agents can gain by reducing (increasing) activity j. If x_{jt} is neither at its upper nor at its lower bound, ϕ_{jt} must be zero to preclude arbitrage possibilities.

In the context of the monetary policy model under discretion, f_{jt} is the derivative of the complementarity condition (1.15) with respect to the nominal interest rate, and ϕ_{jt} is the Lagrange multiplier ϕ_{1t} associated to the non-negativity constraint on the nominal interest rate:

$$-\left(\lambda x_t + \kappa \pi_t\right) = \phi_{1t} \tag{1.45}$$

Since there is no upper bound on the interest rate, $b(s_t) = +\infty$, and $x_t < b(s_t)$ always holds so that ϕ_{1t} is non-negative. This, together with $a(s_t) = 0$, implies that in the case of discretionary optimization, the above complementarity problem reduces to

$$i_t \geqslant 0, \quad \phi_{1t} \geqslant 0, \quad i_t > 0 \Rightarrow \phi_{1t} = 0, \tag{1.46}$$

which can be written also as

$$i_t \geqslant 0, \quad \phi_{1t} \geqslant 0, \quad i_t\phi_{1t} = 0 \tag{1.47}$$

An approximate solution is obtained with the method of collocation, which in this case consists of approximating the expectation functions $E_t x_{t+1}$ and $E_t \pi_{t+1}$ by linear combinations of known basis functions, θ_j, whose coefficients, c_j, are determined by requiring the approximants to satisfy the equilibrium equations exactly at n collocation nodes:

$$h[s, x(s)] \approx \sum_{j=1}^{n} c_j \theta_j(s) \tag{1.48}$$

The coefficients are determined by the following algorithm. For a given value of the coefficient vector c, the equilibrium responses x_i are computed at the n collocation nodes s_i by solving the complementarity problem (which is transformed into a standard root-finding problem). Then, given the equilibrium responses x_i at the collocation nodes s_i, the coefficient vector c is updated solving the n-dimensional linear system

$$\sum_{j=1}^{n} c_j \theta_j(s_i) = h(s_i, x_i) \tag{1.49}$$

This iterative procedure is repeated until the distance between successive values of c becomes sufficiently small (Miranda and Fackler 2001).

To approximate the expectation functions, $E_t x_{t+1}$ and $E_t \pi_{t+1}$, one needs to discretize the shock to r^n. Here the normal shock to the natural rate of interest is discretized using a K-node Gaussian quadrature scheme:

$$Eh\left[s, x\left(s\right)\right] \approx \sum_{k=1}^{K} \sum_{j=1}^{n} \omega_k c_j \theta_j \left[g\left(s_i, x, \varepsilon_k\right)\right] \tag{1.50}$$

where ε_k and ω_k are Gaussian quadrature nodes and weights chosen so that the discrete distribution approximates the continuous univariate normal distribution $N\left(0, \sigma_\epsilon^2\right)$.

In the discretionary optimization problem I use linear splines on a uniform grid of 2000 points for values of the natural rate of interest between –10% and +10%, so that each point on the grid corresponds to 1 basis point. In this problem, linear splines work better than Chebychev polynomials or cubic splines because the response function has a kink in the place where the zero bound becomes binding.

There are two types of approximation errors. On the one hand are the deviations from the equilibrium first-order conditions. In this case the "arbitrage benefits" are negligible for each of the three equilibrium equations (1.13), (1.14), and (1.15). Specifically, they are of the order of 10^{-16} for the IS and the Phillips curves, and 10^{-19} for the complementarity condition. On the other hand are the residuals from the approximation of the expectation functions. Except for a few residuals of the order of 10^{-4}, concentrated mostly in the place where the zero constraint becomes binding, the rest of the residuals are of the order of 10^{-8}. Given the measurement units, a residual of 10^{-4} corresponds to 0.001% of annual inflation or output gap error, which is a satisfactory level of accuracy for the problem at hand. In principle, the expectations residuals can be reduced further by concentrating more evaluation points in the neighborhood of the kink, and by using more quadrature nodes, albeit at the cost of computing time.

In the case of commitment, the problem is first cast in the form specified by (1.42), (1.43), (1.44) by substituting out ϕ_{2t} from (1.22) into (1.23) and the resulting expression for ϕ_{1t} into (1.24). In addition, the state transition vector is augmented by the two "costate" variables ϕ_{1t} and ϕ_{2t}, which are expressed in recursive form using (1.22) and (1.23):

$$\phi_{1t} = \phi_{1t-1}(1+\kappa\sigma)/\beta + \kappa\phi_{2t-1} - \lambda x_t - \kappa\pi_t \tag{1.51}$$

$$\phi_{2t} = \phi_{1t-1}\sigma/\beta + \phi_{2t-1} - \pi_t \tag{1.52}$$

With simple rules, the system is in the required form, and the only necessary adjustments are to the state transition vector in those cases where past endogenous variables enter the rule.

1.B Figures

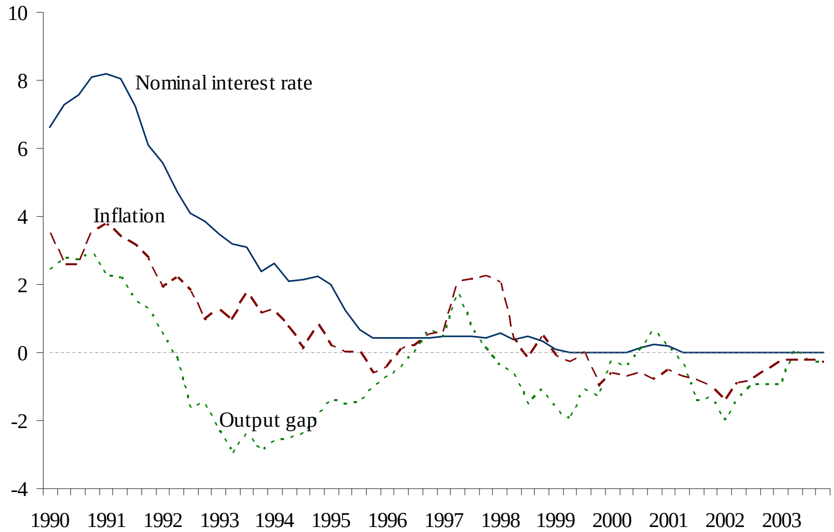

Fig. 1.1. Japan's fall into a liquidity trap

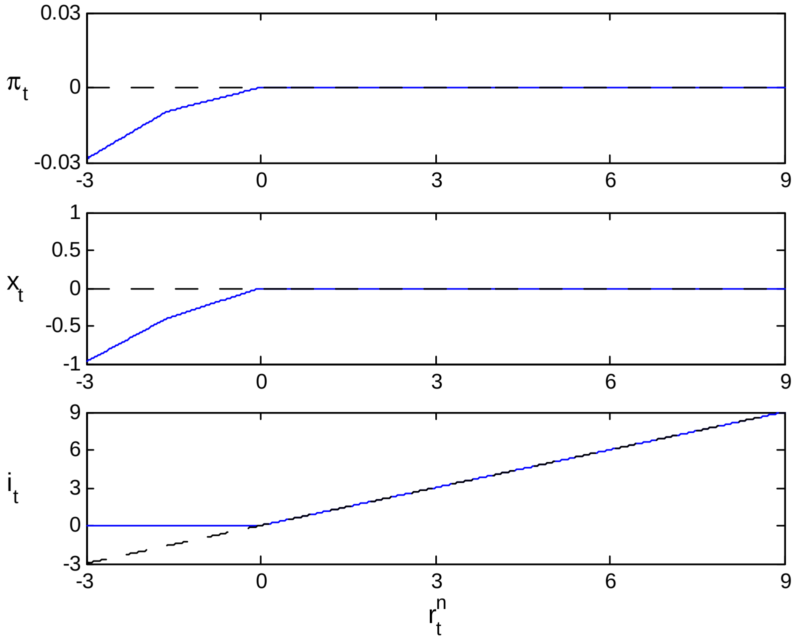

Fig. 1.2. Optimal discretionary policy with perfect foresight

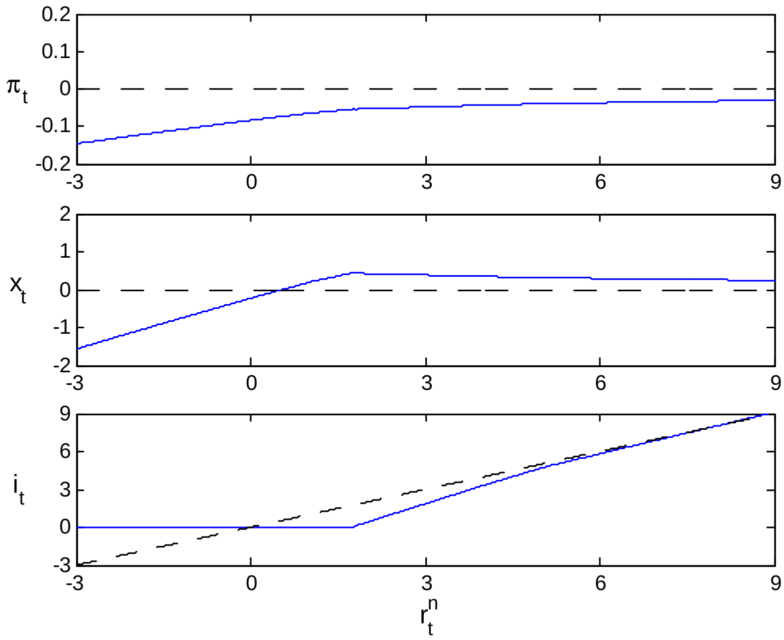

Fig. 1.3. Optimal discretionary policy in the stochastic case

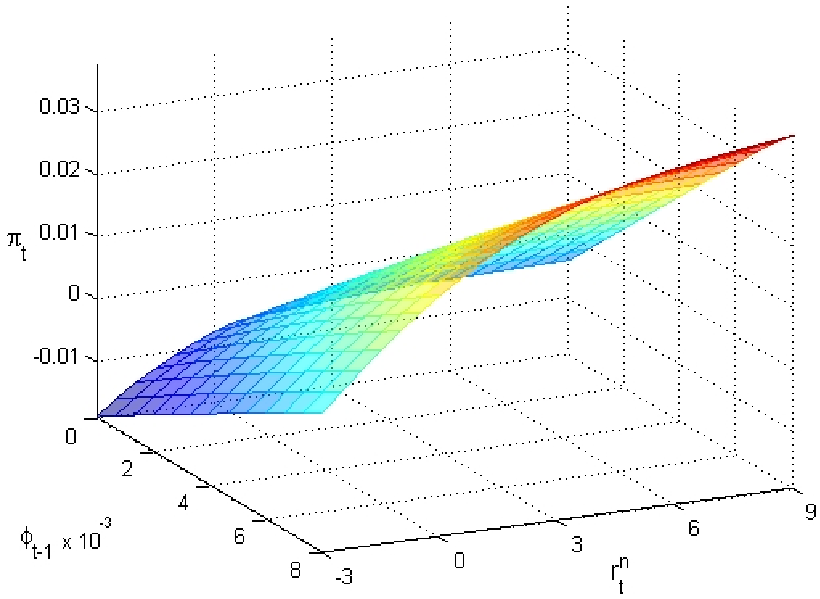

Fig. 1.4. Optimal commitment policy (inflation)

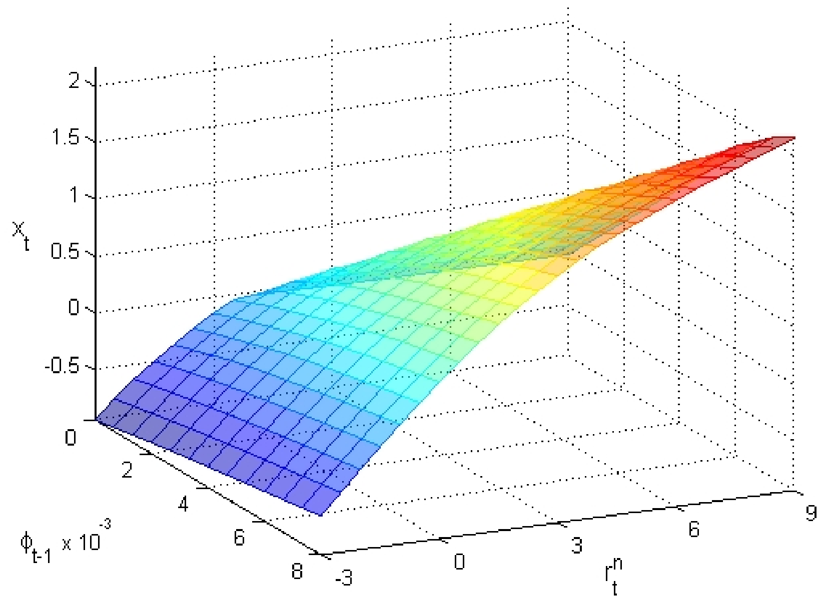

Fig. 1.5. Optimal commitment policy (output gap)

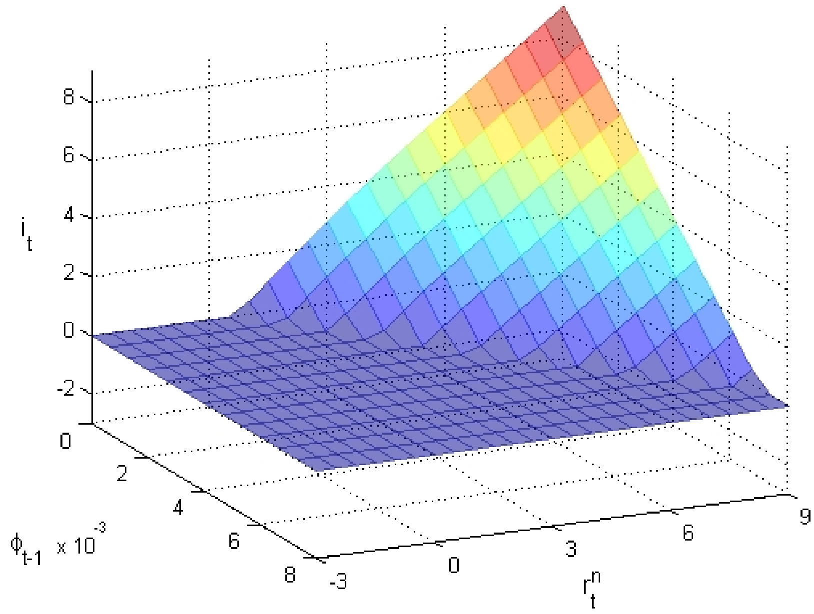

Fig. 1.6. Optimal commitment policy (nominal interest rate)

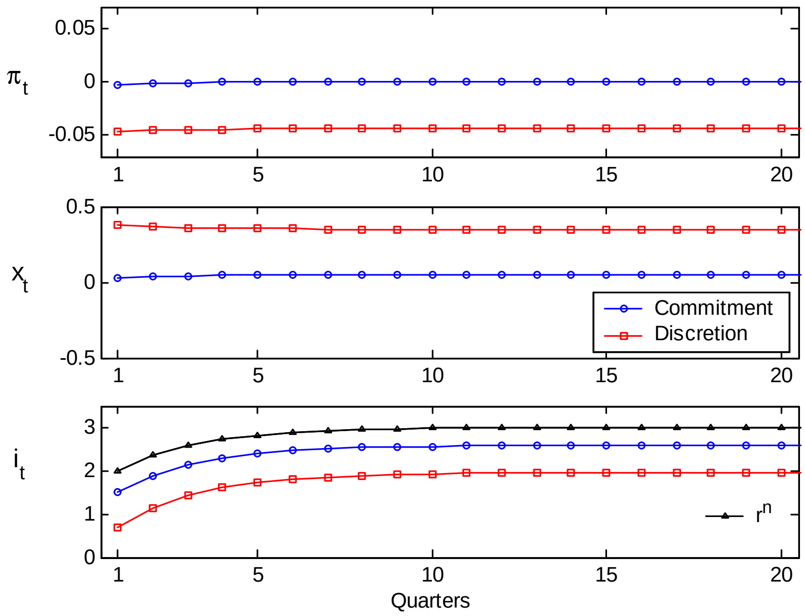

Fig. 1.7. Impulse-responses to a small shock: commitment vs discretion

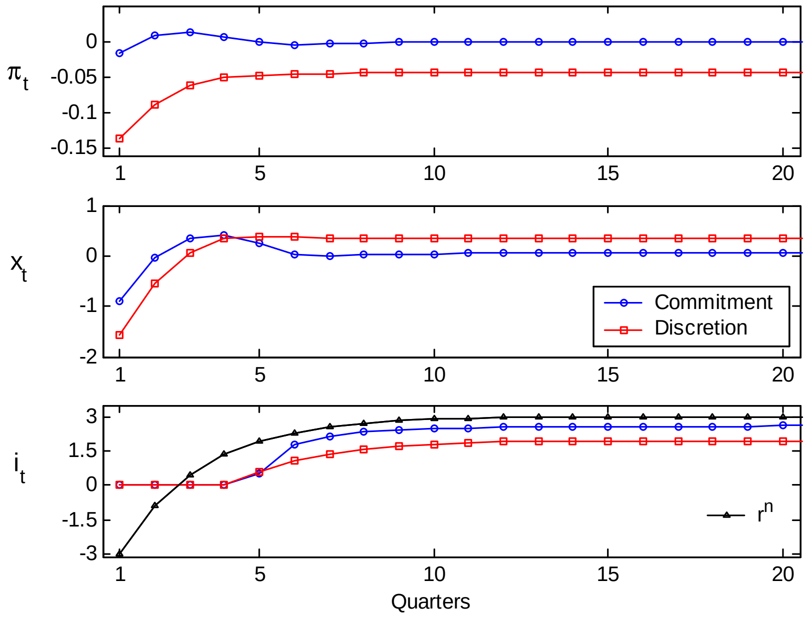

Fig. 1.8. Impulse-responses to a large shock: commitment vs. discretion

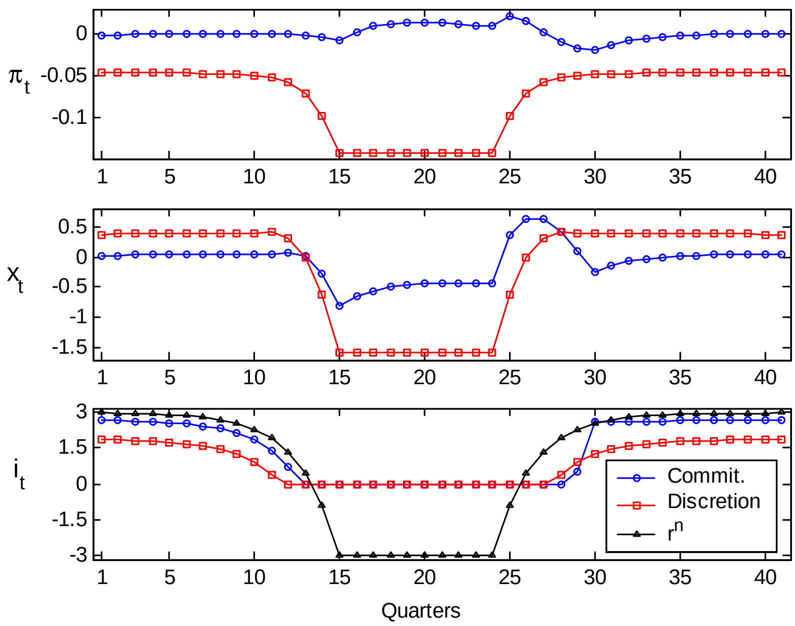

Fig. 1.9. Optimal paths in a liquidity trap - commitment vs discretion

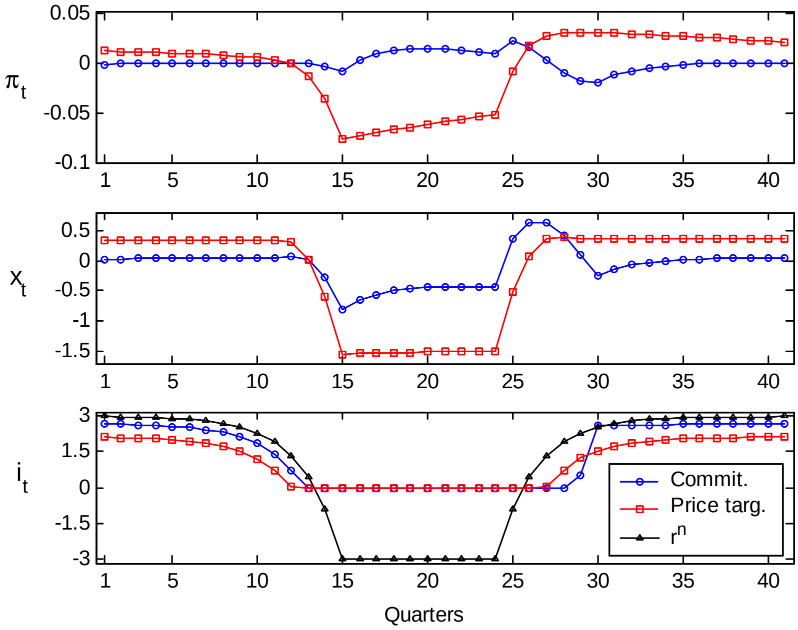

Fig. 1.10. Dynamic paths under constant price level targeting

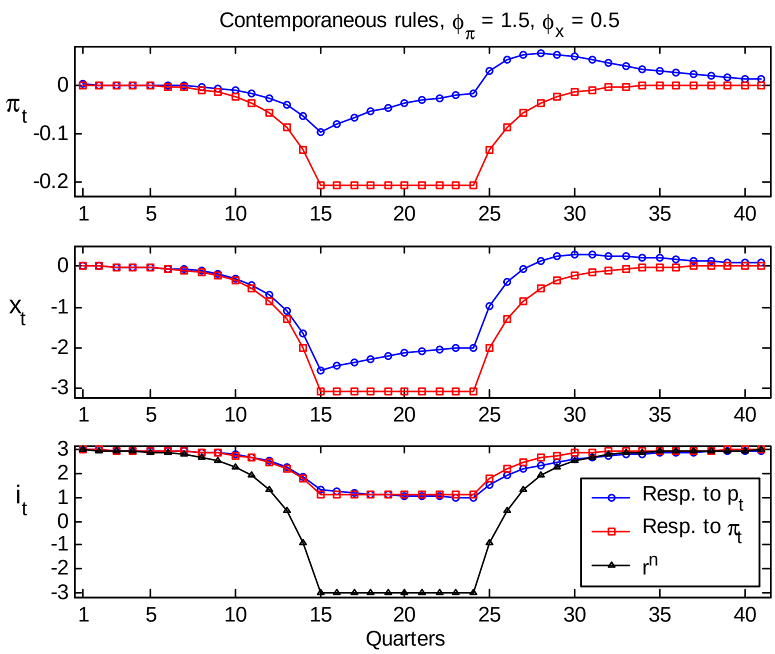

Fig. 1.11. Tuncated Taylor rules responding to the price level or to inflation

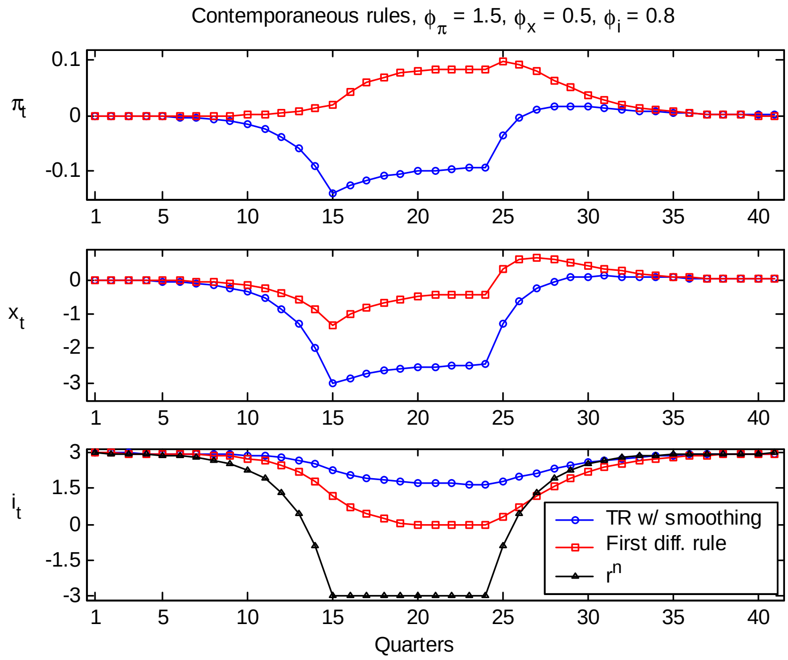

Fig. 1.12. Truncated Taylor rule with smoothing vs first-difference rule

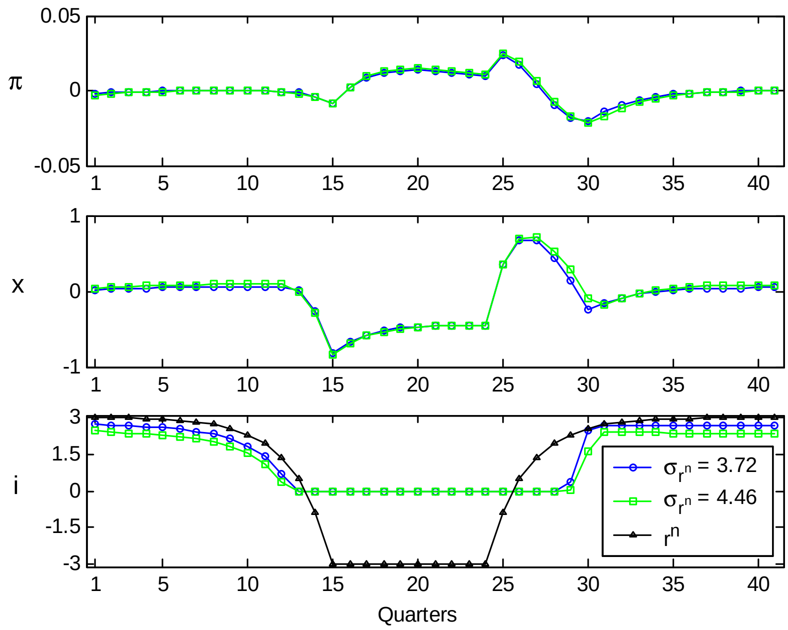

Fig. 1.13. Sensitivity of OCP to $\sigma(r^n)$

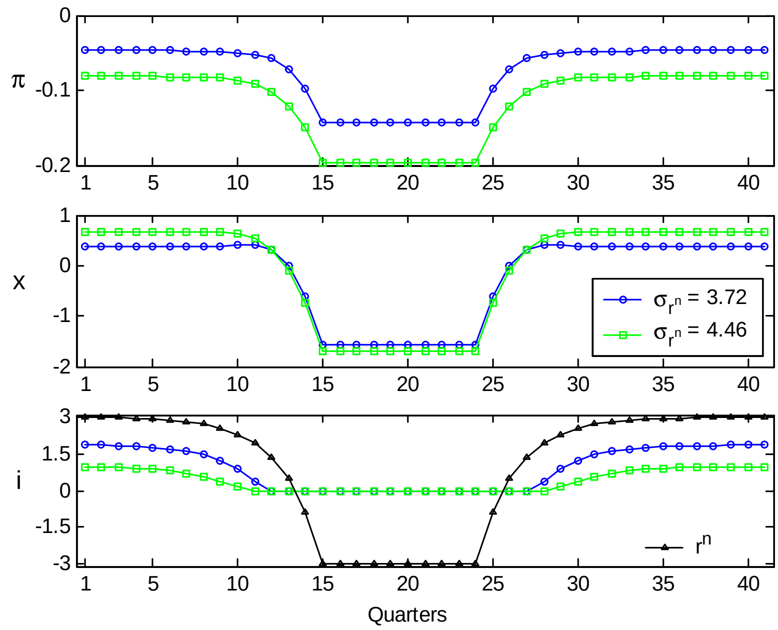

Fig. 1.14. Sensitivity of ODP to $\sigma(r^n)$

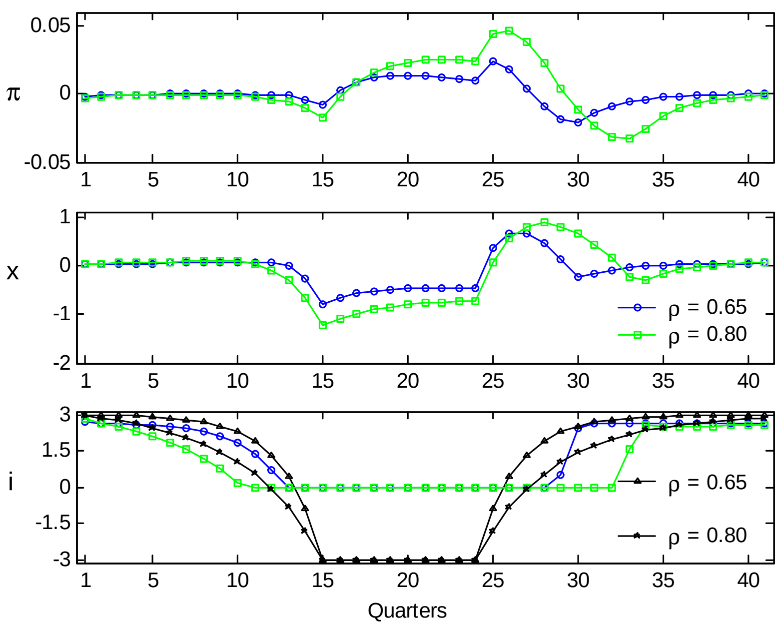

Fig. 1.15. Sensitivity of OCP to ρ

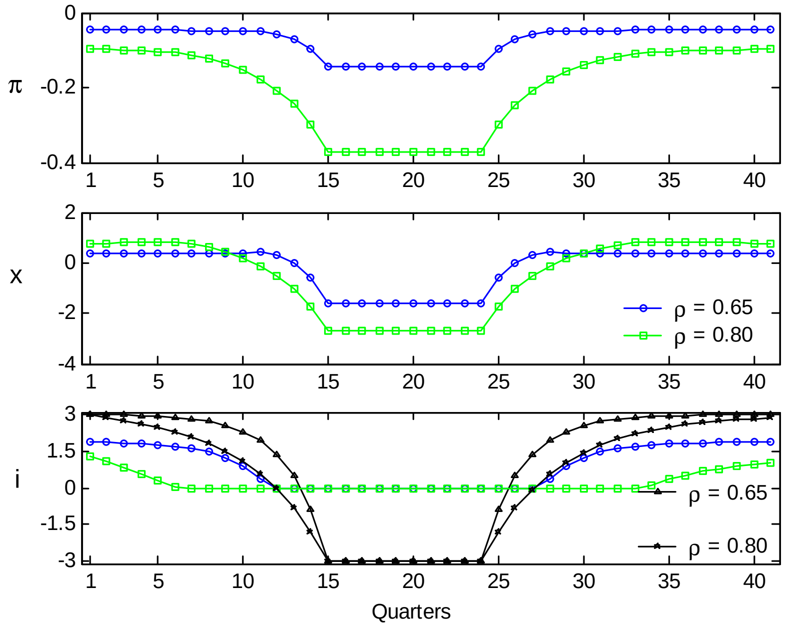

Fig. 1.16. Sensitivity of ODP to ρ

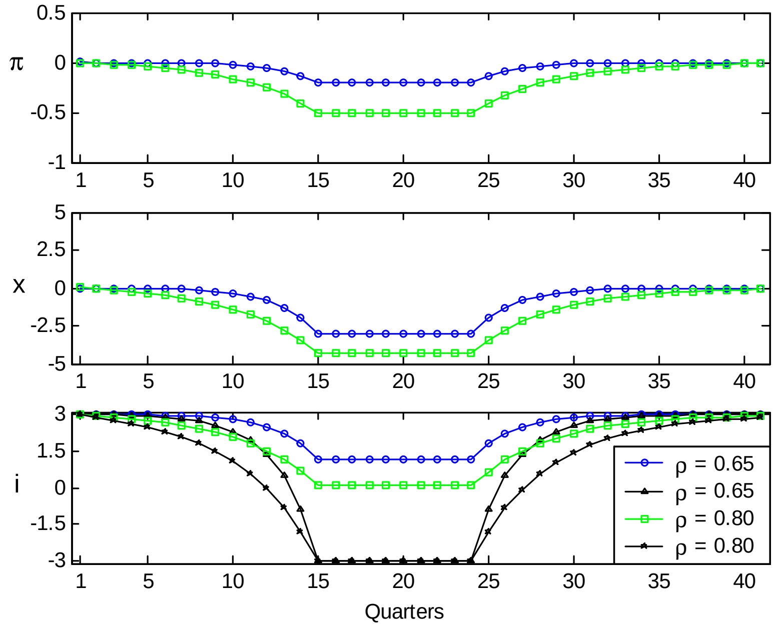

Fig. 1.17. Sensitivity of TTR to ρ

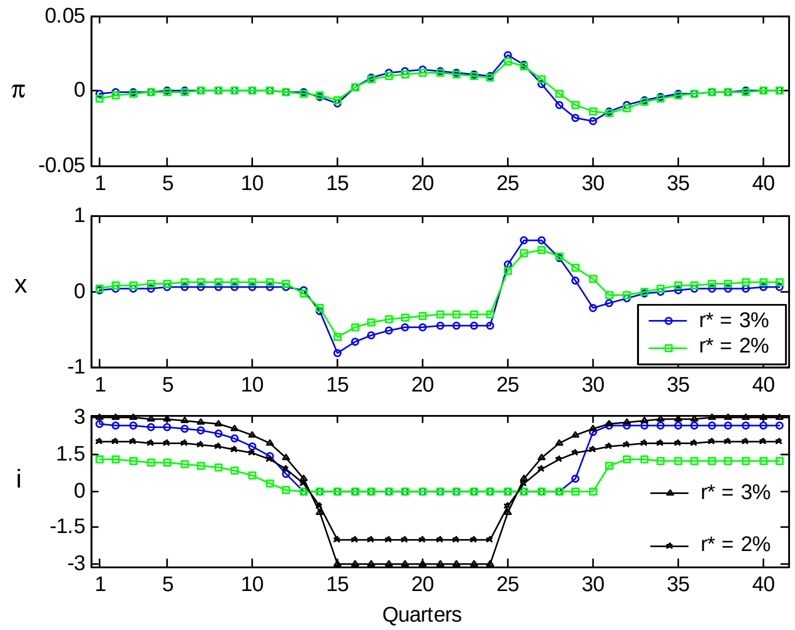

Fig. 1.18. Sensitivity of OCP to r^*

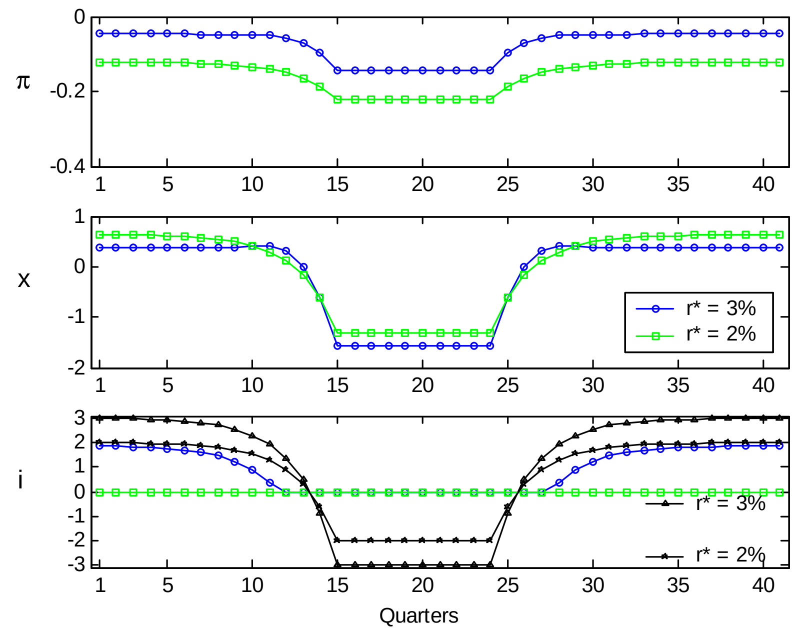

Fig. 1.19. Sensitivity of ODP to r^*

Chapter 2
Inflation–Output Gap Trade-off with a Dominant Oil Supplier

2.1 Introduction

Over the past five years the price of oil has tripled in real terms, from \$20 per barrel in 2002 to \$60 per barrel in 2006 (at constant prices of year 2000). This has rekindled memories of the sharp oil price rises in the 1970s when the real oil price tripled in 1973 and then again more than doubled in 1979 (see Figure 2.1). The former oil price hikes coincided with dramatic declines in US GDP growth and double-digit inflation.[1] And while so far the recent oil price build-up has been accompanied with only a modest pick up in inflation and more or less stable GDP growth, it has reignited discussions about the causes and effects of oil price fluctuations, as well as the appropriate policy responses to oil sector shocks (e.g. Bernanke, 2006).

Most of the existing academic and policy-oriented literature treats oil price movements as unexpected exogenous shifts in the price of oil, unrelated to any economic fundamentals. Seen in this way, oil price shocks are the typical textbook example of a supply-side disturbance which raises inflation and contracts output (e.g. Mankiw, 2006). Thus, for a central bank that cares about inflation and output stability, oil price shocks create a difficult policy trade-off: if the central bank raises the interest rate in order to fight off inflation, the resulting output loss will be larger. And if instead it lowers

---

[1] In fact, Hamilton (1983) observed that all but one US recessions since World War II (until the time of his publication) were preceded by increases in the price of crude oil.

the rate to prevent output from falling, the ensuing inflation rise will be higher. In any case, the central bank simply cannot stabilize both prices *and* output at their respective levels before the shock.

Modern theories of the business cycle have questioned the appropriateness of stabilizing output at its level before the shock. In particular, RBC theory points that in response to a negative productivity shock – which in that framework is equivalent to an exogenous oil price increase – the efficient (first-best) level of output declines, as firms find it optimal to scale down production (and households to give up some consumption for additional leisure). An implication of this for a world with nominal price rigidities, is that in the face of an oil price shock, the central bank should not attempt to stabilize output, but instead should seek to align the output response with the first-best reaction to the oil price change. That is, it should try to stabilize the *output gap*, defined as the distance between actual output and its efficient level given the shock.

Our first result is to show that in the standard New Keynesian model extended with oil as an additional productive input, if the oil price is taken to be exogenous (or perfectly competitive), then there is no tradeoff between inflation and output gap volatility. In other words, even in the face of oil price shocks, there is a "divine coincidence" in the sense of Blanchard and Galí (2006): a policy of price stability automatically stabilizes the distance of output from first-best. This result is important because, if it is true in general and is not just an artefact of some simplifying assumptions, it implies that the task of central banks is much easier and that monetary policy can focus exclusively on price stability.

Our second contribution is to demonstrate that the above "coincidence" breaks down when one relaxes the assumption of exogenous oil price and models explicitly the oil sector's supply behavior. To show this, we model in general equilibrium the behavior of OPEC as a dominant producer which seeks to maximize the welfare of its owner, internalizing the effect of its supply decision on the oil price. Operating alongside a competitive fringe of price-taking oil suppliers, the dominant oil exporter sells its output to an oil importing country (the US), which uses it to produce final goods.

The steady-state of this environment is characterized by an inefficiently low level of oil supply by OPEC, a positive oil price markup, and a suboptimal level of output in the oil importing country. Importantly, shocks in this setup induce inefficient fluctuations in the oil price markup, reflecting a dynamic distortion of the economy's production process. As a result, stabilizing inflation does not fully stabilize the distance of output from first-best, and monetary policy-makers face a tradeoff between the two goals.[2]

Our model allows us to move away from discussing the effects of exogenous oil price changes and towards analyzing the implications of the underlying shocks that cause the oil price to change in the first place. This is a clear advantage over the existing literature, which treats the macroeconomic effects and policy implications of oil price movements as if they were independent of the underlying source of disturbance.[3] In our case there are four structural shocks – to US total factor productivity, to monetary policy, to oil production technology, and to the total capacity of the competitive fringe, each of which affects the oil price through a different channel. Notably, the effects of each of these shocks on macroeconomic variables, and their policy implications, are

---

[2] Rotemberg and Woodford (1996) allow for *exogenous* variation in the oil price markup in a model very different from ours.

[3] See for example Kim and Loungani (1992), Leduc and Sill (2004), and Carlstrom and Fuerst (2005); see Killian (2006) for an exception.

quite different. In particular, conditional on the source of the shock, a central bank confronted with the same oil price increase may find it desirable to either raise or lower the interest rate in order to improve the real allocation.

Finally, we touch on the debate of the relevant inflation target, that is, "core" versus "headline" inflation. If the central bank targets headline inflation, then it implicitly reacts to movements in energy prices roughly in proportion to the share of energy in CPI. Yet our analysis suggests that oil sector developments affect stabilization performance through a different channel, and as such should be treated separately from the CPI index. In particular, we find that a relevant variable to target is the oil price *markup* (which under the assumptions of our model is related to OPEC's market share). This is quite different from advocating a uniform Taylor-type reaction to changes in the oil *price* (and indeed we show that, in general, the latter policy may not improve much on a rule which targets inflation only).

The following section presents the model and the baseline calibration; section 2.3 discusses the steady-state and comparative statics; section 2.4 analyzes the dynamic properties of the model, including impulse-responses and policy implications; section 2.5 reports the dependence of the effects of oil sector shocks on the oil share in production as well as on the monetary regime in place; and the last section concludes.

2.2 The Model

There are two large countries (or regions) — an oil importing and an oil exporting one, and a fringe of small oil exporting countries in the rest of the world. The oil importing country (the US) produces no oil itself but needs it to produce final goods of which it is the only exporter.[4] Oil is a homogenous commodity supplied to the US by two different types of producers: a dominant oil exporter (OPEC) who fully internalizes his effect on the global economy, and a competitive fringe of atomistic exporters, who choose their supply taking prices as given. Oil exporters produce oil only, using as inputs a fraction of the final goods sold to them by the US. In addition, they buy from the US a fraction of final goods which they use for consumption, with the rest of final goods output consumed by the US itself. There is no borrowing across regions (regional current accounts are balanced in each period) and trade is carried out in a common world currency (the dollar).

Two main features distinguish our model from the rest of the literature: the endogeneity of the oil price and the existence of a dominant oil supplier. These assumptions are consistent with a number of observations in the literature regarding the nature of the oil market. In particular, Mabro (1998) argued convincingly that oil demand and the oil price are affected significantly by global macroeconomic conditions.[5] At the same time, Adelman and Shahi (1989) estimated the marginal cost of oil production well below the actual oil price. Indeed, it is obvious that the world's oil industry is not characterized by a continuum of measureless "Mom and Pop" oil extractors. Instead, there is one cartel (OPEC) with more power than any other producer, yet other producers exist and collectively can restrain the exercise of monopoly power by the cartel (Salant, 1976).[6] Empirical evidence by Griffin (1985), Jones (1990), and Dahl and Yucel (1991)

---

4 The US accounts for around 30% of global output, and 30% of OPEC's oil exports (IMF, 2007).

5 Moreover, when testing the null hypothesis that the oil price is not Granger-caused collectively by US output, unemployment, inflation, wages, money and import prices, Hamilton (1983) obtained a rejection at the 6% significance level. In the same article he explicitly referred to the possibility that the oil price was affected by US inflation.

6 Currently OPEC accounts for around 40% of the world's oil production (EIA, 2007).

also suggests that OPEC behavior is closer to that of a cartel than a confederation of competitive suppliers.

2.2.1 Oil Importing Country

The oil importing country is a canonical sticky price economy with oil included as an additional input in production, monopolistic competition, and Calvo (1983) contracts. We call this country "the US" for short.

Households

The country is populated by a representative household, which seeks to maximize the expected present discounted flow of utility streams,

$$\max E_o \sum_{t=0}^{\infty} \beta^t U(C_t, L_t), \tag{2.1}$$

subject to a budget constraint. The period utility function depends on consumption, C_t, and labor L_t, and we assume that it takes the form

$$U(C_t, L_t) = \log(C_t) - \frac{L_t^{1+\psi}}{1+\psi}. \tag{2.2}$$

The period t budget constraint,

$$P_t C_t + B_t R_t^{-1} = B_{t-1} + w_t P_t L_t + r_t P_t \bar{K} + \Pi_t^f, \tag{2.3}$$

equates nominal income from labor, $w_t P_t L_t$, capital $r_t P_t \bar{K}$, dividends from the final goods firms owned by the household, Π_t^f, and nominally riskless bonds, B_{t-1}, to outlays on consumption, $P_t C_t$, and bonds, $B_t R_t^{-1}$. The aggregate stock of capital which the household rents out to firms is assumed to be constant, $\bar{K}$, normalized to one.

The consumption good C_t is a Dixit-Stiglitz aggregate of a continuum of differentiated goods $C_t(i)$,

$$C_t = \left[\int_0^1 C_t(i)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}} \tag{2.4}$$

with associated price index,

$$P_t^{1-\epsilon} = \int_0^1 P_t(i)^{1-\epsilon} di, \tag{2.5}$$

where $P_t(i)$ is the price of good i.

The household chooses the sequence $\{C_t, L_t, B_t\}_{t=0}^{\infty}$ in order to maximize the expected present discounted utility (2.1) subject to the budget constraint (2.3). In addition, it allocates expenditure among the different goods $C_t(i)$ so as to minimize the cost of buying the aggregate bundle C_t.

Final Goods Sector

Final goods are produced under monopolistic competition with labor, capital, and oil according to

$$Q_t(i) = A_t L_t(i)^{\alpha_1} K_t(i)^{\alpha_2} O_t(i)^{1-\alpha_1-\alpha_2} \tag{2.6}$$

where A_t denotes aggregate total factor productivity. The latter evolves exogenously according to

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a \tag{2.7}$$

where $a_t \equiv \log(A_t)$ and $\varepsilon_t^a \sim i.i.d.N.\left(0, \sigma_a^2\right)$.

Individual firms are small and take all aggregate variables as given. In particular, firms take factor prices as given as they compete for inputs on economy-wide factor markets in order to minimize the total cost of production. In addition, firms reset their prices infrequently *a la* Calvo (1983). In each period a constant random fraction θ of all firms is unable to change their price and must satisfy demand at whatever price they posted in the previous period. Whenever they get a chance to change their price $P_t(i)$, firms seek to maximize the expected present discounted stream of profits,

$$\max E_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k}[P_t(i)Q_{t+k}(i) - P_{t+k}C\left(Q_{t+k}(i)\right)] \tag{2.8}$$

subject to a downward sloping demand schedule,

$$Q_{t+k}(i) = \left(\frac{P_t(i)}{P_{t+k}}\right)^{-\epsilon} Q_{t+k}, \tag{2.9}$$

where $Q_{t+k}(i)$ is demand for the output of firm i, $C\left(Q_{t+k}(i)\right)$ is the real cost of producing that output, and $\Lambda_{t,t+k}$ is the discount factor for nominal payoffs.

Monetary Policy

The central bank in the oil importing country is committed to set the nominal interest rate according to the rule

$$\frac{R_t}{\bar{R}} = e^{r_t}\left(\frac{R_{t-1}}{\bar{R}}\right)^{\phi_R}\left(\frac{\Pi_t}{\bar{\Pi}}\right)^{\phi_\pi}\left(\frac{p_{ot}}{p_{ot-1}}\right)^{\phi_o}, \tag{2.10}$$

where $\bar{R} \equiv \bar{\Pi}/\beta$ and $\bar{\Pi}$ is the target rate of inflation; r_t is an i.i.d. "interest rate shock", distributed normally with mean zero and variance σ_r^2. ϕ_R is an "interest rate smoothing" parameter, and ϕ_π and ϕ_o are policy reaction coefficients.

We allow for a possible non-zero reaction of the central bank to the change in the real price of oil. While our analysis in section 2.6 shows that the welfare-relevant target variable is not this but the oil price *markup*, the latter depends on the current marginal cost of oil production, which we assume to be unobservable by the monetary authority.

2.2.2 Oil Exporting Countries

Modelling the oil industry as a dominant firm with competitive fringe dates back to Salant (1976). He argued that neither perfect competition, nor a single monopolist owning all the oil, bear much resemblance to the actual structure of the world oil industry. While his focus was on the Cournot-Nash equilibrium of the game between the competitive fringe and the dominant extractor of exhaustible oil, our interest lies in the links between the dominant oil supplier and the oil importer. As we shall see, the existence of competitive oil producers affects in important ways the equilibrium behavior of the dominant oil supplier.

Dominant Oil Exporter

The large oil exporting country, called "OPEC", is populated by a representative household that seeks to maximize its expected present discounted flow of utility streams,

$$\max E_o \sum_{t=0}^{\infty} \beta^t U(\tilde{C}_t), \tag{2.11}$$

where the period utility function is logarithmic in consumption,

$$U\left(\tilde{C}_t\right) = \log(\tilde{C}_t). \tag{2.12}$$

The household faces a period budget constraint,

$$P_t \tilde{C}_t = \Pi_t^o, \tag{2.13}$$

which equates consumption expenditure to dividends from OPEC, Π_t^o, which is wholly owned by the household. As such, the representative household's objective of expected utility maximization is consistent with maximizing the expected present discounted value of the logarithm of real profits from oil production, where period profits are given by

$$\frac{\Pi_t^o}{P_t} = p_{ot} O_t - \tilde{I}_t. \tag{2.14}$$

OPEC produces oil according to

$$O_t = Z_t \tilde{I}_t, \tag{2.15}$$

where Z_t is an exogenous productivity shifter, and $\tilde{I}_t$ is an intermediate good used in oil production and bought from the oil importing country. The productivity of OPEC evolves exogenously according to

$$z_t = \rho_z z_{t-1} + \varepsilon_t^z, \tag{2.16}$$

where $z_t \equiv \log(Z_t)$ and $\varepsilon_t^z \sim i.i.d.N\left(0, \sigma_z^2\right)$.

The consumption good $\tilde{C}_t$ and the intermediate good $\tilde{I}_t$ are Dixit-Stiglitz aggregates of a continuum of differentiated goods of the same form (2.4) and with the same price index (2.5) as before. OPEC allocates expenditure among the different intermediate and final goods so as to minimize the cost of buying the aggregate bundles $\tilde{I}_t$ and $\tilde{C}_t$. It chooses a level of oil output, so as to maximize the expected present discounted utility of the representative household, subject to the behavior of competitive oil exporters, and households, firms and monetary authority in the US.

Competitive Fringe of Small Oil Exporters

Apart from the dominant oil exporter, in the rest of the world there is a continuum of atomistic oil firms, indexed by $i \in [0, \Omega_t]$. Each firm produces a quantity $X_t(i)$ of oil according to the technology

$$X_t(i) = \xi(i) Z_t \hat{I}_t(i), \tag{2.17}$$

subject to the capacity constraint,

$$X_t(i) \in [0, \bar{X}], \tag{2.18}$$

where $[\xi(i)Z_t]^{-1}$ is the marginal cost of oil production of firm i; $1/Z_t$ is a component of marginal cost common to all oil firms, while $1/\xi(i)$ is a constant firm-specific component distributed according to some probability distribution function $F(1/\xi(i))$. The input $\hat{I}_t(i)$ is purchased from the oil importer as is consumption of the representative household owning each oil firm, $\hat{C}_t(i)$, which is equal to the real profit from oil production.[7] Both $\hat{I}_t(i)$ and $\hat{C}_t(i)$ are Dixit-Stiglitz aggregates of differentiated goods analogous to those of the dominant oil firm.

The total mass (or total capacity) of competitive fringe producers Ω_t is allowed to vary according to a stationary stochastic process,

$$\hat{\omega}_t = \rho_\omega \hat{\omega}_{t-1} + \varepsilon_t^\omega \tag{2.19}$$

where $\hat{\omega}_t \equiv \log\left(\Omega_t/\bar{\Omega}\right)$ and $\varepsilon_t^\omega \sim i.i.d.N\left(0, \sigma_\omega^2\right)$. We make this allowance to capture the fact that some oil fields of the fringe are used up, while new ones are discovered and so the total amount of oil recoverable by the competitive fringe is not constant over time. In section 4 we evaluate the effects of a transitory change in the availability of oil outside OPEC's control on the equilibrium oil price and macroeconomic aggregates. As we will see, it is the only shock in our model which induces a negative correlation between the supply of OPEC and the output of the competitive fringe, a feature of the data which is prominent in the 1980-s and early 1990-s (see figure 2.10).

The produced oil can either be sold at the international price p_{ot}, which the atomistic exporters take as given, or it is lost. Each small supplier chooses the amount of oil

7 We assume perfect risk-sharing among competitive fringe producers.

to produce in each period so as to maximize profits,

$$\max \left\{ p_{ot} X_t(i) - X_t(i) / \left[\xi(i) Z_t \right] \right\} \tag{2.20}$$
$$\text{s.t.} \quad X_t(i) \in [0, \bar{X}]$$

The existence of competitive producers restrains significantly the exercise of market power by the dominant oil firm. In our case, the measure of non-OPEC competitors (calibrated to match their average market share) reduces the average oil price markup from 20 (in the case of full oil monopoly) to 1.36 times marginal cost (in the case of a "dominant firm"). Moreover, the introduction of a competitive fringe allows us to model transitory shifts in the market share of OPEC. Figure 2.2 shows that this share has not been constant over the last four decades: it was around 50% in the 1970s, then dropped down to 30% in the 1980s, before recovering to around 40% in the last two decades. Since around 70% of the world's "proven reserves" are under OPEC control (EIA, 2007), some observers suggest that in the absence of any new major oil discoveries or technological advances in non-OPEC countries, the cartel's market share would rise steadily in the future (however, see Adelman (2004) for a forceful refutation of the idea that oil is running out and on the meaninglessness of the concept of "proven reserves").

Most importantly for the oil importing country, the asymmetric distribution of market power between the two types of oil suppliers induces a dynamic production distortion reflected in variation of the oil price markup in response to all shocks. This is what ultimately breaks the "divine coincidence" between stabilizing inflation and stabilizing the welfare-relevant output gap, creating a tension between the two stabilization objectives.

2.2.3 Equilibrium Conditions for a Given Oil Supply

Optimality conditions

The first-order optimality conditions of the representative US household are:

$$
\begin{aligned}
C_t(i) &= \left(\frac{P_t(i)}{P_t}\right)^{-\epsilon} C_t && (2.21)\\
C_t L_t^{\psi} &= w_t && (2.22)\\
1 &= \beta R_t E_t \left[\frac{C_t}{C_{t+1}} \frac{P_t}{P_{t+1}}\right]. && (2.23)
\end{aligned}
$$

Condition (2.21) states that the relative demand for good i is inversely related to its relative price. Equation (2.22) is a standard labor supply curve equating the marginal rate of substitution between consumption and leisure to the real wage; and (2.23) is a standard consumption Euler equation.

Cost minimization by final goods firms implies

$$
\begin{aligned}
w_t L_t(i) &= \alpha_1 mc_t Q_t(i) && (2.24)\\
r_t K_t(i) &= \alpha_2 mc_t Q_t(i) && (2.25)\\
p_{ot} O_t(i) &= (1 - \alpha_1 - \alpha_2) mc_t Q_t(i) && (2.26)
\end{aligned}
$$

where w_t is the real wage, p_{ot} is the real price of oil, r_t is the real rental price of capital, and mc_t are real marginal costs, which are common across all firms. The above conditions equate marginal costs of production to the factor price divided by the marginal factor product for each input of the production function for final goods. At the same time, with Cobb-Douglas technology, marginal costs are given by

$$
mc_t = \frac{w_t^{\alpha_1} r_t^{\alpha_2} p_{ot}^{1-\alpha_1-\alpha_2}}{A_t \alpha_1^{\alpha_1} \alpha_2^{\alpha_2} (1 - \alpha_1 - \alpha_2)^{1-\alpha_1-\alpha_2}}. \qquad (2.27)
$$

The optimal price-setting decision of firm i implies that the optimal reset price $P_t^*(i)$ satisfies

$$p_t^* \equiv \frac{P_t^*(i)}{P_t} = \frac{N_t}{D_t}, \tag{2.28}$$

where N_t and D_t are governed by

$$\begin{aligned} D_t &= \frac{Q_t}{C_t} + \beta\theta E_t \left[\Pi_{t+1}^{\epsilon-1} D_{t+1}\right] && (2.29) \\ N_t &= \mu mc_t \frac{Q_t}{C_t} + \beta\theta E_t \left[\Pi_{t+1}^{\epsilon} N_{t+1}\right] && (2.30) \end{aligned}$$

with $\mu \equiv \frac{\epsilon}{\epsilon-1}$. These conditions imply that whenever a firm is able to change its price, it sets it at a constant markup μ over a weighted average of current and expected future marginal costs, where the weights associated with each horizon k are related to the probability that the chosen price is still effective in period k.

All resetting firms face an identical problem and hence choose the same price. Given that the fraction of firms resetting their price is drawn randomly from the set of all firms, and using the definition of the aggregate price index, we have

$$P_t^{1-\epsilon} = \theta P_{t-1}^{1-\epsilon} + (1-\theta) P_t^{\star 1-\epsilon} \tag{2.31}$$

which implies

$$1 = \theta \Pi_t^{\epsilon-1} + (1-\theta) p_t^{\star 1-\epsilon}. \tag{2.32}$$

Denoting the relative price dispersion by

$$\Delta_t \equiv \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\epsilon} di, \tag{2.33}$$

we can derive a law of motion for this measure as

$$\Delta_t = \theta \Pi_t^{\epsilon} \Delta_{t-1} + (1-\theta) p_t^{\star -\epsilon}. \tag{2.34}$$

Finally, each competitive fringe exporter finds it profitable to produce oil if and only if the current market price of oil p_{ot} is greater than his marginal cost. Thus, competitive oil firm i produces $\bar{X}$ if $[\xi(i)Z_t]^{-1} \leq p_{ot}$ and zero otherwise.

Aggregation

Aggregating the demand for labor, capital and oil by final goods firms yields,

$$L_t = \int_0^1 L_t(i)di \tag{2.35}$$

$$K_{dt} = \int_0^1 K_t(i)di \tag{2.36}$$

$$O_{dt} = \int_0^1 O_t(i)di \tag{2.37}$$

In turn, aggregate demand for final goods output is given by,

$$Q_t = \left[\int_0^1 Q_t(i)^{\frac{\epsilon-1}{\epsilon}} di\right]^{\frac{\epsilon}{\epsilon-1}}. \tag{2.38}$$

Analogous expressions describe the aggregate consumption and intermediate goods import components of aggregate demand for each country.

The above, together with (2.9), imply that the following aggregate demand relationships hold,

$$p_{ot}O_{dt} = (1-\alpha_1-\alpha_2)mc_tQ_t\Delta_t \tag{2.39}$$

$$w_tL_t = \alpha_1 mc_tQ_t\Delta_t \tag{2.40}$$

$$r_tK_{dt} = \alpha_2 mc_tQ_t\Delta_t, \tag{2.41}$$

where aggregate output satisfies

$$Q_t = \frac{A_t}{\Delta_t} L_t^{\alpha_1} K_{dt}^{\alpha_2} O_{dt}^{1-\alpha_1-\alpha_2}. \tag{2.42}$$

Notice in particular the distortionary effect of aggregate price dispersion in (2.42), which acts like a tax on aggregate output, in a way similar to a negative productivity shock.

Aggregate real profits of final goods firms in the oil importing country are given by,

$$\frac{\Pi_t^f}{P_t} = Q_t - p_{ot}O_{dt} - w_t L_t - r_t \bar{K}. \tag{2.43}$$

Finally, the amount of oil produced by the competitive fringe as a whole is given by

$$X_t \equiv \int_0^{\Omega_t} X_t(i)di = \Omega_t F(p_{ot}Z_t) \tag{2.44}$$

To simplify, we assume that the idiosyncratic component of marginal costs $1/\xi(i)$ is distributed uniformly in the interval $[a, b]$. In that case

$$X_t = \begin{cases} \Omega_t \bar{X}, & p_{ot}Z_t > b \\ \Omega_t \bar{X}\frac{p_{ot}Z_t - a}{b-a}, & a < p_{ot}Z_t \leq b \\ 0, & p_{ot}Z_t \leq a \end{cases} \tag{2.45}$$

We further assume without loss of generality[8] that $a = 0$ and normalize $b = \bar{X} > 1$ which we choose sufficiently large that at least some competitive fringe producers (potential entrants) are always priced out of the market by the dominant oil firm. With these assumptions the output of the competitive fringe is a product of the price of oil (p_{ot}), productivity of the oil sector (Z_t), and a component related to the depletion and discovery of new oil deposits by the competitive fringe (Ω_t):

$$X_t = \Omega_t p_{ot} Z_t. \tag{2.46}$$

---

[8] Our main results are unaffected if we assume instead that OPEC is the most efficient oil supplier by setting $a = 1$.

Market clearing

Bonds are in zero net supply and the supply of capital is fixed at the aggregate level. Hence, in equilibrium, we have

$$
\begin{aligned}
B_t &= 0 && (2.47) \\
K_{dt} &= \bar{K} = 1 && (2.48)
\end{aligned}
$$

which, substituting into the budget constraint of the oil importing country's household, implies

$$C_t = w_t L_t + r_t \bar{K} + \frac{\Pi_t^f}{P_t}. \tag{2.49}$$

Substituting aggregate real profits from (2.43) in the above equation yields,

$$C_t = Q_t - p_{ot} O_{dt}. \tag{2.50}$$

Further, aggregate oil demand is equal to the supply of the dominant oil firm plus the aggregate output of the competitive fringe of oil exporters:

$$O_{dt} = O_t + X_t. \tag{2.51}$$

Finally, the aggregate consumption of small oil exporters equals their aggregate real profits,

$$\hat{C}_t = p_{ot} X_t - \hat{I}_t \tag{2.52}$$

With these conditions we can verify that the aggregate resource constraint holds,

$$Q_t = C_t + \tilde{C}_t + \tilde{I}_t + \hat{C}_t + \hat{I}_t, \tag{2.53}$$

whereby global final goods output is equal to global final goods consumption plus global intermediate input purchases.

2.2.4 The Dominant Oil Exporter's Problem

We assume that OPEC solve a Ramsey-type problem. Namely, they seek to maximize the expected welfare of the representative household-owner of OPEC, subject to the behavior of all other agents and the global resource constraint. Formally, in our setup this is equivalent to maximizing the expected present discounted value of the logarithm of oil profits,

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \log \left[p_{ot} O_t - O_t / Z_t\right] \tag{2.54}$$

subject to the constraints imposed by the optimal behavior of the competitive fringe,

$$X_t = \Omega_t p_{ot} Z_t, \tag{2.55}$$

of households,

$$\begin{aligned} w_t &= C_t L_t^{\psi} && (2.56) \\ 1 &= \beta R_t E_t \left[\frac{C_t}{C_{t+1}} \frac{P_t}{P_{t+1}}\right], && (2.57) \end{aligned}$$

and final goods firms in the oil importing country,

$$\begin{aligned} D_t &= \frac{Q_t}{C_t} + \beta \theta E_t \left[\Pi_{t+1}^{\epsilon-1} D_{t+1}\right] && (2.58) \\ N_t &= \mu m c_t \frac{Q_t}{C_t} + \beta \theta E_t \left[\Pi_{t+1}^{\epsilon} N_{t+1}\right] && (2.59) \\ 1 &= \theta \Pi_t^{\epsilon-1} + (1-\theta) \left(\frac{N_t}{D_t}\right)^{1-\epsilon} && (2.60) \\ \Delta_t &= \theta \Pi_t^{\epsilon} \Delta_{t-1} + (1-\theta) \left(\frac{N_t}{D_t}\right)^{-\epsilon} && (2.61) \\ p_{ot} &= (1-\alpha_1-\alpha_2) m c_t Q_t \Delta_t / (O_t + X_t) && (2.62) \\ L_t &= \alpha_1 m c_t Q_t \Delta_t / w_t && (2.63) \\ Q_t &= \frac{A_t}{\Delta_t} L_t^{\alpha_1} \bar{K}_t^{\alpha_2} (O_t + X_t)^{1-\alpha_1-\alpha_2}, && (2.64) \end{aligned}$$

the rule followed by the monetary authority,

$$\frac{R_t}{\bar{R}} = e^{r_t} \left(\frac{R_{t-1}}{\bar{R}}\right)^{\phi_R} \left(\frac{\Pi_t}{\bar{\Pi}}\right)^{\phi_\pi} \left(\frac{p_{ot}}{p_{ot-1}}\right)^{\phi_o}, \tag{2.65}$$

and the global resource constraint,

$$C_t = Q_t - p_{ot}\left(O_t + X_t\right). \tag{2.66}$$

We assume throughout that OPEC can commit to the optimal policy rule that brings about the equilibrium which maximizes expression (2.54) above. Furthermore, we restrict our attention to Markovian stochastic processes for all exogenous variables, and to optimal decision rules which are time-invariant functions of the state of the economy.

2.2.5 Flexible Price Benchmarks

We begin by characterizing the equilibrium allocation in two benchmark scenarios which we will use later to evaluate alternative monetary strategies. One is the *natural* allocation, which corresponds to the equilibrium that would obtain if all prices were fully flexible. And the other is the *efficient* allocation, which we define as the allocation that would obtain if prices were fully flexible and there was perfect competition in oil production.

We make use of the following relation for equilibrium labor which holds regardless of the behavior of the oil sector. Substituting (2.22), (2.39), and (2.40) into (2.50), we can solve for equilibrium labor as a function of marginal cost and relative price dispersion in the US:

$$L_t = \left[\frac{\alpha_1 mc_t \Delta_t}{1-(1-\alpha_1-\alpha_2)mc_t\Delta_t}\right]^{\frac{1}{1+\psi}}. \tag{2.67}$$

Efficiency: perfect competition in oil and flexible prices

The efficient allocation (denoted by the superscript "e") is the one which would obtain under perfect competition in oil production and fully flexible prices.[9]

Will full price flexibility (attained by setting $\theta = 0$) all firms charge the same price and hence in the symmetric equilibrium there is no price dispersion,

$$\Delta_t^e = 1. \tag{2.68}$$

Moreover, in this case marginal costs are constant and equal to the inverse of the optimal markup of final goods firms (related to the elasticity of substitution among final goods)

$$mc_t^e = \mu^{-1} = \frac{\epsilon - 1}{\epsilon}. \tag{2.69}$$

With these substitutions, equation (2.67) reduces to

$$L_t^e = \left[\frac{\alpha_1}{\mu - (1 - \alpha_1 - \alpha_2)}\right]^{\frac{1}{1+\psi}} \equiv \bar{L}, \tag{2.70}$$

which implies that equilibrium labor is constant, unaffected by shocks. At the same time, equation (2.39) becomes

$$p_{ot}^e O_{dt}^e = (1 - \alpha_1 - \alpha_2)\mu^{-1} Q_t^e. \tag{2.71}$$

If, in addition, the dominant oil exporter operated as a perfect competitor, the real price of oil would be equal to its marginal cost,[10]

$$p_{ot}^e = mc_{ot} = Z_t^{-1}, \tag{2.72}$$

which is exogenously given. We can establish the following

---

9 Without loss of generality, we keep in the definition the static distorion due to monopolistic competition in the oil importing country.

10 Since our focus is on OPEC, we rule out the corner solution in which the collective supply of the more efficient fraction of the competitive fringe is sufficient to meet all demand and price OPEC out of the market.

Proposition 1 *With exogenous or competitive oil prices and full price flexibility, a shock to the oil price (or to the marginal cost of oil production) is equivalent to a total factor productivity shock.*

Proof. Equations (2.72) and (2.71) combined with (2.42) imply

$$Q_t^e = \left[A_t Z_t^{1-\alpha_1-\alpha_2}\right]^{\frac{1}{\alpha_1+\alpha_2}} \bar{L}^{\frac{\alpha_1}{\alpha_1+\alpha_2}} \bar{K}^{\frac{\alpha_2}{\alpha_1+\alpha_2}} \left[(1-\alpha_1-\alpha_2)\mu^{-1}\right]^{\frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}} \tag{2.73}$$

Labor and real marginal costs are constant, and all other real endogenous variables of the oil importer (w_t^e, r_t^e, C_t^e, and O_{dt}^e) can be expressed in terms of Q_t^e. For example, the efficient level of consumption (or value added) is given by,

$$C_t^e = (1-(1-\alpha_1-\alpha_2)\mu^{-1})Q_t^e. \tag{2.74}$$

■

In other words, apart from a possible scaling down by the share of oil in output, an oil price shock (a change in Z_t) affects the efficient level of output and all real variables in the same way as a TFP shock (a change in A_t).

Corollary 2 *With an exogenous or competitive oil sector any movements in the oil price caused by real shocks reflect opposite shifts in the efficient level of output.*

Replicating the efficient allocation under sticky prices

The above corollary suggests that one thing that monetary policy should not attempt is to "neutralize" shifts in competitively set (or exogenous) oil prices. We can show that in a scenario with sticky goods prices and an exogenous or competitive oil price, monetary policy can replicate the efficient equilibrium by targeting inflation alone, as stated in the following

Proposition 3 *If the oil price is exogenous or competitive and there is no price dispersion initially, then the optimal monetary policy is full price stability.*

Proof. See Appendix 2.B ■

In other words, with an exogenous or competitive oil price, there is a "divine coincidence" of monetary policy objectives in the sense of Blanchard and Galí (2006): stabilizing inflation will automatically stabilize the distance between output and its efficient level.

The intuition for this result is straightforward: with a competitive or exogenous oil price, there is only one source of distortion in the economy – the one associated with nominal price rigidity. A policy of full price stability eliminates this distortion and replicates the efficient allocation.

The following sections show how this result can be overturned with a dominant oil supplier.

Natural allocation: market power in oil and flexible prices

The natural allocation (denoted by the superscript "n") is defined as the one which would obtain if all prices were fully flexible. In this case, it is straightforward to show that equilibrium labor supply is constant and given by equation (2.70). We can use this fact to derive a relationship between the oil price and the demand for oil that obtains under flexible prices,

$$p_{ot}^{n} = (1-\alpha_1-\alpha_2)\mu^{-1} A_t \bar{L}^{\alpha_1} \bar{K}^{\alpha_2} \left(O_{dt}^{n}\right)^{-\alpha_1-\alpha_2} . \tag{2.75}$$

Consecutive substitution of (2.55) into (2.51) and the resulting expression into the equation above yields an oil demand curve which relates directly the natural price of oil to the demand for OPEC's output independently of any other endogenous variables.

This greatly simplifies the problem of OPEC (2.54) since now the only relevant constraint for the maximization of profits is a single demand curve (2.77). Hence, OPEC solves

$$\max_{O_t^n} \; E_0 \sum_{t=0}^{\infty} \beta^t \log[p_{ot}^n O_t^n - O_t^n / Z_t] \tag{2.76}$$

$$\text{s.t. } \; p_{ot}^n = (1 - \alpha_1 - \alpha_2)\mu^{-1} A_t \bar{L}^{\alpha_1} \bar{K}^{\alpha_2} \left(O_t^n + \Omega_t p_{ot}^n Z_t\right)^{-\alpha_1 - \alpha_2} \tag{2.77}$$

The solution to this problem implies that the price of oil is a time-varying markup ν_t^n over marginal cost mc_{ot},

$$p_{ot}^n = \nu_t^n mc_{ot}, \tag{2.78}$$

where marginal cost is given by

$$mc_{ot} = Z_t^{-1} = p_{ot}^e \tag{2.79}$$

while the optimal markup is inversely related to the (absolute) price elasticity of demand for OPEC's oil:

$$\nu_t^n = \frac{\left|\varepsilon_t^{O^n, p_o^n}\right|}{\left|\varepsilon_t^{O^n, p_o^n}\right| - 1}. \tag{2.80}$$

The latter can be derived from constraint (2.77) as

$$\left|\varepsilon_t^{O^n, p_o^n}\right| \equiv \left|\frac{\partial O_t^n}{\partial p_{ot}^n} \frac{p_{ot}^n}{O_t^n}\right| = \frac{1}{\eta s_t^n} - 1, \tag{2.81}$$

where $\eta \equiv \frac{\alpha_1 + \alpha_2}{1 + \alpha_1 + \alpha_2}$, and $s_t^n = \frac{O_t^n}{O_t^n + X_t^n}$ is the natural market share of OPEC.

Since $(\alpha_1 + \alpha_2) \in (0, \; 1)$ implies $\eta \in (0, \; 0.5)$, and given that $s_t^n \in [0, 1]$, we have $\eta s_t^n \in (0, 0.5)$ and therefore $\left|\varepsilon_t^{O, po}\right| \in (1, +\infty)$. This implies that the profit-maximizing dominant firm produces always on the elastic segment of its effective demand curve and that the oil price markup is positive ($\nu_t^n > 1$).

Moreover, from (2.81) we see that the (absolute) price elasticity of demand for OPEC's oil is a decreasing function of OPEC's market share. Hence, a negative shock to the supply of the competitive fringe which increases OPEC's market share, makes the demand for OPEC's oil less price-elastic, raising the optimal markup charged by OPEC.

Substituting (2.81) into (2.80) we can obtain a direct relationship between the optimal oil price markup and the market share of the dominant oil exporter,

$$\nu_t^n = \frac{\eta s_t^n - 1}{2\eta s_t^n - 1}, \tag{2.82}$$

which in a first-order approximation around the steady state becomes

$$\hat{\nu}_t^n = \frac{\eta}{(2\eta\bar{s} - 1)^2}\hat{s}_t^n. \tag{2.83}$$

This implies that, up to a first-order approximation, the oil price markup co-moves with OPEC's market share,

$$corr(\nu_t^n,\ s_t^n) \approx 1. \tag{2.84}$$

Full Monopoly in Oil Production

It is informative to consider the special case of a single oil supplier with full monopoly power (corresponding to $\Omega_t = 0$ and $s_t^n = 1$). The solution (denoted by the superscript "m") implies:

$$O_t^m = \left[(1 - \alpha_1 - \alpha_2)^2 \mu^{-1} A_t Z_t \bar{L}^{\alpha_1} \bar{K}^{\alpha_2}\right]^{\frac{1}{\alpha_1 + \alpha_2}}, \tag{2.85}$$

$$p_{ot}^m = \frac{1}{Z_t\left[1 - \alpha_1 - \alpha_2\right]} = \nu^m p_{ot}^e \tag{2.86}$$

The price of oil is a constant markup over marginal cost, where the optimal markup $\nu^m = [1 - \alpha_1 - \alpha_2]^{-1}$ is the inverse of the elasticity of oil in final goods production. For instance, if $1 - \alpha_1 - \alpha_2 = 0.05$, the optimal markup ν^m would be 20!

The intuition for this result is straightforward: with $s_t^n = 1$ the price elasticity of demand for the monopolist's oil (2.81) reduces to

$$\left|\varepsilon_t^{O^m, p_o^m}\right| = \left|\frac{\partial O_t^m}{\partial p_{ot}^m}\frac{p_{ot}^m}{O_t^m}\right| = \frac{1}{\alpha_1 + \alpha_2} = \frac{1}{1 - (1 - \alpha_1 - \alpha_2)}. \tag{2.87}$$

In words, with a single oil monopolist the (absolute) price elasticity of oil demand is positively related to the elasticity of oil in production. Therefore, a small share of oil in output implies that oil demand is quite insensitive to the price, which allows the monopolist to charge a high markup.

Finally, notice that the existence of a competitive fringe greatly reduces OPEC's optimal markup. For example, if in steady-state the supply of the competitive fringe is roughly equal to that of OPEC ($O_t^n = X_t^n$), OPEC's optimal markup reduces to a level which is an order of magnitude lower than the full monopoly markup,

$$s_t^n = 0.5 \Longrightarrow \nu^n = 1 + \frac{\alpha_1 + \alpha_2}{2} = 1.475 << \nu^m = 20. \tag{2.88}$$

The natural GDP gap

Since part of gross output is spent on intermediate (oil) imports, the output variable relevant for US welfare is value added (or GDP), which from the resource constraint (2.50) equals simply consumption,

$$Y_t = Q_t - p_{ot} O_{dt} = C_t. \tag{2.89}$$

We call "natural GDP gap", the distance between the natural level of value added, Y_t^n, and its efficient counterpart, Y_t^e, and denote this distance by $\hat{Y}_t^n$. It is straightforward to show that this distance is a function only of the natural oil price gap (p_{ot}^n / p_{ot}^e), which from (2.78) and (2.79) is equal to the oil price markup in the natural allocation,

$$\hat{Y}_t^n \equiv Y_t^n / Y_t^e = (p_{ot}^n / p_{ot}^e)^{-\frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}} = (\nu_t^n)^{-\frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}}. \tag{2.90}$$

Since we have seen in (2.80) that with a dominant oil supplier the oil price markup is always greater than one, the natural equilibrium is characterized by underproduction in the US, related to an inefficiently low oil supply by OPEC. Moreover, contrary to the polar cases of perfect competition or full monopoly power in oil, in the intermediate case with a dominant firm, the oil price markup fluctuates in response to all real shocks. And while these fluctuations are optimal responses from the point of view of OPEC, they are distortionary from the point of view of the US economy. Therefore, if US monetary policy can affect the actual evolution of output (and hence the average markup of final goods producers), it would make sense to counter, at least to some extent, fluctuations in the oil price markup, in addition to targeting inflation.

2.2.6 Equilibrium with Sticky Prices

Given a certain degree of price stickiness, monetary policy can affect the real economy in the short run. In particular, it can affect US output (and final goods producers' markup), and indirectly the demand for oil and its price.

The equilibrium with sticky prices and a dominant oil supplier is defined by a set of time-invariant decision rules for the endogenous variables as functions of the state and the shocks observed in the beginning of each period, which satisfy constraints (2.55) - (2.66) and which solve the dominant oil supplier's problem in (2.54). Appendix 2.A lists the first-order conditions of the optimal oil pricing problem.

We derive an expression for the *welfare-relevant* GDP gap, $\tilde{y}_t$ (to which we refer sometimes simply as "the output gap"), defined as the (log) distance between actual value added and its efficient level given by (2.73),

$$\tilde{y}_t \equiv y_t - y_t^e. \tag{2.91}$$

As shown in Appendix 2.B, this gap is related to real marginal costs mc_t – a standard result in the New Keynesian literature – but in our model also to the oil price markup ν_t. Thus, up to a first-order approximation, fluctuations in the output gap are related to shifts in these two variables,

$$\tilde{y}_t = \kappa_{mc}\hat{mc}_t - \kappa_\nu \hat{\nu}_t, \tag{2.92}$$

where κ_{mc} and κ_ν are related to the structural parameters of the model and are defined in Appendix 2.B; $\hat{mc}_t$ are real marginal costs in the final goods sector, and $\hat{\nu}_t = \hat{p}_{ot} - \hat{p}_{ot}^e$ is the oil price markup, expressed in log-deviations from the efficient equilibrium.[11]

Proposition 4 *In the presence of a dominant oil supplier, optimal monetary policy seeks to strike a balance between stabilizing inflation and stabilizing the output gap.*

From equation (2.92) we see that a policy aimed at full price stability would set $\hat{mc}_t$ equal to zero and would thus stabilize the gap between actual value added and its natural level. Yet this would not stabilize fully the welfare-relevant output gap,

$$\tilde{y}_t = (y_t - y_t^n) + (y_t^n - y_t^e), \tag{2.93}$$

since in response to real shocks OPEC's profit maximizing behavior induces inefficient fluctuations in the oil price markup $\hat{\nu}_t$, reflected in a time-varying wedge between the natural and the efficient level of output. The above result breaks the "divine coincidence" of monetary policy objectives and provides a rationale for the central bank to mitigate to a certain extent inefficient output gap fluctuations by tolerating some deviation from full price stability.

[11] Appendix 2.B shows that defining the output gap in terms of gross output instead of value added yields an identical expression with a minor reparametrization of κ_{mc} (provided the share of oil in GDP is not large).

2.2.7 Calibration

We calibrate our model so that it replicates some basic facts about the US economy and OPEC. Table 2.1 shows the parameters used in the baseline calibration. The quarterly discount factor corresponds to an average real interest rate of 3% per annum. Utility is logarithmic in consumption and we assume a unit Frisch elasticity of labor supply. We set the elasticity of labor in production equal to 0.63 and the elasticity of capital to 0.32, consistent with measures of the average labor and capital shares in output. This implies an elasticity of oil of 0.05 and an oil share of $0.05/\mu \approx 0.04$, which roughly corresponds to the value share of oil consumption in US GDP. The Calvo price adjustment parameter is set equal to 0.75, implying an average price duration of one year. The elasticity of substitution among final goods is assumed to be 7.66 corresponding to a steady-state price markup of 15%. And the mean of the total capacity of non-OPEC producers is set to match the average market share of OPEC of around 42%.

We choose the baseline parameters of the monetary policy rule as follows. We set the target inflation rate equal to zero, consistent with the optimal long-run inflation in our model.[12] The short-run reaction coefficient on inflation is set to 0.4, while the interest rate smoothing parameter is set to 0.8, implying a long-run inflation coefficient of 2. These values are similar to the estimates by Clarida, Gali and Gertler (2000) for the Volcker-Greenspan period. The baseline short-run coefficient on oil price inflation is set equal to zero.

There are three real and one nominal exogenous variables in our model. For US total factor productivity we assume an AR(1) process with standard deviation of the innovation of 0.007 and an autoregressive parameter of 0.95, similar to those calibrated

---

[12] More on this in the following section.

by Prescott (1986) and Cooley (1997). With these values we are able to match the standard deviation and persistence of US GDP growth from 1973:I to 2007:I. Similarly, the processes for oil technology and the capacity of non-OPEC producers are parametrized to match the volatility of the oil price (about 20 times more volatile than US GDP), its autoregressive coefficient (0.97), as well as the relative volatility of OPEC versus non-OPEC output (the former is five times more volatile) over the same period.[13] Finally, the interest rate shock is assumed to be $i.i.d.$ with standard deviation corresponding to a 25 basis points disturbance of the interest rate rule.

In the following section we study the steady-state properties of the model and perform comparative statics exercises varying some of the above parameters.

[13] Quarterly data on OPEC and non-OPEC oil output are taken from EIA (2007), and on US GDP from FRED II. Actual and model-generated data are made comparable by taking growth rates and then subtracting the mean growth rate for each variable. Volatility is measured as the standard deviation of the demeaned growth rate series.

Structural parameters		
Quarterly discount factor	β	0.9926
Elasticity of output wrt labor	α_1	0.63
Elasticity of output wrt capital	α_2	0.32
Elasticity of output wrt. oil		0.05
Price adjustment probability	θ	0.75
Price elasticity of substitution	ϵ	7.66
Mean of non-OPEC capacity	$\bar{\Omega}$	0.0093
Inv. Frisch labor supply elast.	ψ	1
Monetary policy		
Long run inflation target	$\bar{\Pi}$	1
Interest rate smoothing coeff.	ϕ_R	0.8
Inflation reaction coefficient	ϕ_π	0.4
Oil price reaction coefficient	ϕ_{po}	0
Shock processes		
Std of US TFP shock	σ_a	0.007
Persistence of US TFP shock	ρ_a	0.95
Std of oil tech. shock	σ_z	0.12
Persistence of oil tech. shock	ρ_z	0.95
Std of non-OPEC capacity	σ_ω	0.10
Persist. of non-OPEC capacity	ρ_ω	0.975
Std dev of int. rate innovation	σ_r	0.001

Table 2.1. Baseline calibration

2.3 Steady State and Comparative Statics

We focus our attention on the steady-state with zero inflation. The reason is that for an empirically plausible range of values for the reaction coefficients of the monetary policy rule, the optimal long-run rate of inflation in our model (from the point of view of the US consumer) is essentially zero.

The zero inflation steady-state is characterized by an inefficiently low oil supply by OPEC[14], a positive oil price markup, and underproduction of final goods in the US. In particular, under our baseline calibration OPEC produces only 45% of the amount of

[14] This result ignores any longer term costs of oil associated with environmental pollution and global warming.

oil that it would produce if it operated as a competitive firm. This allows it to charge a markup of around 36% over marginal cost, and make a positive profit of around 0.5% of US output (or around $65 billion per annum based on nominal US GDP in 2006). At the same time, imperfect competition in the oil market opens a steady-state output gap in the US of 1.6% ($208 billion per annum).

Figures 2.3 and 2.4 show two comparative statics exercises. Figure 2.3 illustrates the sensitivity of the steady-state to the availability of oil outside OPEC. In the face of a 50% reduction of the capacity of competitive oil producers with respect to the baseline, OPEC's output increases only by 10%. The market share of OPEC increases, and by (2.82) the oil price markup jumps from 35% to 75% over OPEC's marginal cost. This widens the US output gap to 3%, while doubling OPEC's profit as a share of output. The relationship however is highly non-linear and a further reduction of the capacity of oil producers outside OPEC results in a much more dramatic increase in the equilibrium price of oil and a larger output loss in the US.

Figure 2.4 shows the sensitivity of the results to the elasticity of oil in output. Keeping the capacity of non-OPEC producers constant, an increase of the oil elasticity to 0.1 raises the market share of OPEC. As a result, the oil price jumps to 57% over marginal cost and the US output gap widens to 5%.

2.4 Dynamic Properties of the Model

We solve the model numerically by first-order Taylor approximation of the decision rules around the deterministic steady-state with zero inflation (following Blanchard and

Kahn (1980)).[15] This section reports some of the more interesting dynamic features of the economy under our preferred calibration.

Figures 2.5, 2.7, 2.9 and 2.11 in Appendix 2.C show the impulse-response functions for several variables of interest. The signs of the shocks are chosen so that all impulses result in an increase in the oil price on impact. The figures plot the efficient allocation (denoted by the superscript "e"); the natural allocation (denoted by "n"; it coincides with the actual evolution under a policy of full price stability); and the actual evolution of the relevant variables with nominal rigidity and under the benchmark policy rule.

To help clarify the intuition, the bottom-right panel of the figures shows three GDP gap measures: the actual (or welfare-relevant GDP gap, denoted by Y_{gap}), the natural GDP gap (denoted by Y^n_{gap}), and the "sticky price GDP gap" (denoted by Y^s_{gap}), defined as the distance between the actual and the natural level of GDP.

2.4.1 US technology shock

We begin with a typical (one-standard-deviation) positive shock to US total factor productivity in figure 2.5. Consider first the efficient allocation. As is standard in RBC models, the efficient level of output rises (in our case by 0.74%). Since OPEC acts competitively and there is no change in the marginal cost of oil production, the oil price remains constant. Because there is no change in the price, the supply of the fringe stays fixed as well. With OPEC as the marginal oil producer, all of the additional oil demand is met by a rise in OPEC's supply, which raises OPEC's market share.

15 Solving the model by second-order approximation yields virtually identical impulse-response functions.

Now let's turn to the natural evolution and compare it to the efficient one. In response to the positive TFP shock, dominant OPEC raises its oil supply, while engineering a slight increase in the oil price markup.[16] This is a consequence of profit maximization subject to downward-sloping demand: since OPEC's profit is the product of the oil price markup and oil output, in the face of stronger US demand for oil due to oil's enhanced productivity, it is optimal to increase both profit factors. As figure 2.5 shows, this requires that OPEC increase its supply by a slightly smaller fraction of steady-state output than if it operated as a perfect competitor.[17]

Due to the oil price rise, the supply of non-OPEC increases as well, albeit by less than OPEC. OPEC's market share rises, consistent with the increase in the oil price markup as per equations (2.82) and (2.84). Natural output in the US increases by slightly less than the efficient amount because of the inefficient response of natural oil supply. Quantitatively, however, the natural GDP gap moves very little in response to a US technology shock. This suggests that, with respect to US TFP shocks, a policy aimed at full price stability would almost stabilize the GDP gap.

Finally, consider the actual allocation with nominal rigidity and given the benchmark policy rule (2.10). Inflation falls by around 30 basis points (annualized), while output increases by 0.61% — less than the efficient increase. As it turns out, most of the inefficiency in response to the US TFP shock stems from the suboptimality of the benchmark policy rule. This can be seen from the bottom-right panel, in which nearly all of the 13 basis points fall (that is, widening) of the GDP gap induced by the shock

[16] The latter can be seen as the difference between the natural and the efficient response of the oil price.

[17] Figure 2.6 illustrates this in the case of linear demand. If OPEC operated as a perfect competitor, an increase in demand would move it from point A to point A' where marginal cost crosses the new oil demand schedule. The oil price remains unchanged and all adjustment falls on oil supply. Since instead OPEC is a profit-maximizing monopolist, marginal revenue shifts out by less than the oil demand schedule. As a result, both oil output and the oil price rise as OPEC moves from point B to point B'.

is attributable to the fall of the "sticky price GDP gap" (Y^s_{gap}). Hence, there is almost no tradeoff between inflation and GDP gap stabilization. Compared to the benchmark Taylor-type rule, a positive technology shock calls for a somewhat more aggressive interest rate *reduction* even if the shock is associated with a slightly rising oil price.

2.4.2 Oil technology shock

We next discuss the responses to a one-standard-deviation negative shock to oil productivity shown in figure 2.7. Again, we focus in turn on the evolution of the efficient, the natural, and the actual allocations.

First, because a negative oil technology shock is a positive marginal cost shock for the oil industry, the efficient level of oil supply falls while the efficient oil price rises (by 12%). Since oil is an intermediate input, the efficiency of final goods production is also affected, so that the first-best level of output declines by 0.65%. The supply of the fringe remains constant because the oil price rise is completely offset by the increase in the marginal cost of oil production. As a result, OPEC's share declines in response to the shock.

In the natural equilibrium, since marginal revenue is steeper than the demand curve, OPEC's oil price markup decreases, meaning that the natural oil price rise (around 9%) is less than the efficient increase (of 12%).[18] Similarly, the fall in OPEC's output (as a fraction of steady-state) is less than the efficient decline. Because of the decrease of the oil price markup, non-OPEC supply falls by around 3%, while OPEC's market share declines by around 3 percentage points, shadowing the movement of the oil price markup.

[18] Figure 2.8 illustrates this in the case of linear demand.

Actual US GDP falls by around 0.4%, which is less than the efficient decline of 0.65%.[19] As a result, the rise in inflation by 20 basis points is accompanied by an increase (that is, narrowing) of the GDP gap by around 25 basis points. In contrast to the previous shock, however, this time much of the GDP gap movement is “natural” in the sense that it is attributed more to the temporary fall in the oil price markup than to sticky prices.

The part of the GDP gap due to sticky prices can be stabilized better by raising the nominal rate more aggressively than the benchmark rule (2.10) prescribes. In fact, a policy of full price stability would bring the response of the GDP gap down to that of the natural GDP gap (a 19 basis points rise, instead of 25), which is unambiguously welfare-improving compared to the benchmark rule. But, clearly, a policy of full price stability is not optimal either, as it is not able to fully stabilize the GDP gap, and in general results in excessive GDP gap variation. In order to stabilize the GDP gap more, the central bank would have to allow some amount of deflation. In other words, the optimal rule would seek to strike a balance between stabilizing prices and stabilizing the GDP gap. From the point of view of rule (2.10), in response to a negative oil technology shock which raises the oil price, the central bank should *raise* the nominal rate by more than what the benchmark rule prescribes (but not by so much as to cause excessive deflation).

---

[19] This output response is in the ballpark of empirical estimates of the response of US GDP to an "oil price shock"; admittedly, uncertainty about this empirical response is an order-of-magnitude large: according to Bernanke et al. (2004) and IMF(2005) a 10% increase in the oil price leads to a 0.10% to 0.20% drop in US GDP after 1 to 2 years. On the other extreme, Rotemberg and Woodford (1997) and Finn (2000) argue that the effect is as large as a 2.5% drop in GDP after 5 to 7 quarters.

2.4.3 Fringe capacity shock

In third place we analyze the effects of a one-standard-deviation negative shock to the total capacity of competitive fringe producers.[20]

First notice in figure 2.9 that this shock has no effect on the efficient oil price or on the first-best level of output (the latter can be seen also in expression (2.73) in which the fringe shock does not appear). The reason is that, unlike the oil technology disturbance, the fringe shock does not affect the efficiency of oil production. The latter in turn is related to the fact that in the efficient equilibrium, and for the allowed size of oil demand and fringe shocks, the aggregate oil supply curve is flat at the marginal cost of OPEC. Since OPEC can supply any amount of oil at that price, shocks to fringe capacity are of no relevance for the marginal cost of oil production and as a consequence do not affect the efficient level of output.

Turning to the natural allocation, a negative fringe shock decreases non-OPEC supply by 7.3% and raises OPEC's market share by around 2.6 percentage points. By (2.81) the effective demand for OPEC oil becomes less price-elastic, which implies that the profit-maximizing oil price increases by around 2.7%. OPEC's output increases by less than the decrease in non-OPEC supply, and as a consequence total oil production declines. The resulting drop in US output (by around 0.15%), coupled with the constancy of the efficient level of output, translates into a fall (that is, widening) of the natural GDP gap by 15 basis points.

The actual allocation for this shock almost coincides with the natural one. The GDP gap fall is by 14 basis points and it is accompanied by a rise in inflation by 3 basis

[20] Alternatively, one could think of the negative fringe shock as a positive demand shock from the rest of the world (e.g. China), where demand is postulated to decrease linearly in the price.

points. Importantly, virtually all of the GDP gap fall is due to imperfect competition in the oil sector and as such cannot be stabilized through a policy of price stability. In fact, any attempt to stabilize the GDP gap in this case would have to come at the cost of increasing inflation. Hence, with respect to this shock, optimal monetary policy would involve a welfare-relevant trade-off between inflation and GDP gap stabilization. With respect to the benchmark rule, the central bank should either *raise* or *lower* the nominal interest rate, depending on the relative benefit of inflation versus GDP gap stabilization.

Finally, notice in passing that this shock creates a negative conditional correlation between OPEC and non-OPEC oil supply. This negative correlation features importantly in the data throughout the 1980s when non-OPEC oil production (especially that of UK, Norway, Russia and Mexico) took off, while OPEC's output was essentially halved (see figure 2.10).

2.4.4 Monetary policy shock

Finally, we illustrate the monetary transmission mechanism by tracing out the effects of a monetary policy shock in figure 2.11. The efficient and the natural allocations are of course unaffected by this type of shock.

In terms of the actual allocation, in response to an unexpected 25 basis points interest rate cut, US output (+0.25%) and inflation (+45 bp) both rise as is standard in the New Keynesian model. OPEC responds to the rise in demand by raising its output (+1.3%) while engineering an increase in the oil price markup (+0.2%). The supply of the competitive fringe increases by the same proportion as the oil price markup. This is less than OPEC's supply rise and OPEC's share increases, in line with the oil price markup rise. Since the efficient and the natural levels of output remain constant, the shock results in an inefficient rise (narrowing) of the GDP gap by 25 bp, all of which is attributable to sticky prices. Monetary policy in this model has a strong influence on the actual evolution of output and prices and can be used as an effective tool to offset the real disturbances causing inefficient fluctuations in welfare-relevant variables.

2.4.5 Summary and policy implications

Table 2.2 summarizes the conditional correlation of the oil price with US output, the GDP gap, inflation and the oil price markup (or OPEC's share), induced by each of the four shocks under the benchmark monetary policy rule. In addition, the last column of the table sums up the policy implications of each type of shock, relative to the prescription of the benchmark policy rule.

The table shows that the oil price could be positively or negatively correlated with the GDP gap and inflation depending on the source of the shock. A somewhat surprising finding, perhaps, is that conditional on an oil technology shock, the oil price is positively correlated with the GDP gap. As mentioned earlier, the reason is that conditional on this shock, the oil price is negatively related to the oil price markup. In contrast, the oil price is negatively correlated with the GDP gap if the shock is due to an unexpected change in non-OPEC capacity.

Related to the above, the policy implications of an oil price change depend crucially on the underlying source of the shock. In particular, an oil price increase due to a negative oil technology shock calls for a somewhat *higher* interest rate vis-a-vis the benchmark, since this type of shock lowers the efficient level of output while imperfect competition in the oil market and price stickiness prevent actual output from falling sufficiently. As we saw in section 2.4.2, a typical negative oil technology shock which raises the oil price by 9% results in a 3% decrease in the oil price markup. Because of the relatively small share of oil in output, this translates into a 25 basis points increase in the GDP gap (and a 20 bp rise in inflation). If the central bank were to offset completely the effect of the shock on the GDP gap, it would have to raise the nominal rate by roughly 25 basis points above the benchmark policy rule.

In contrast, an oil price increase associated with a negative fringe shock may well require a *lower* interest rate with respect to the benchmark. This is because the efficient level of output remains unaffected, while actual output falls as OPEC uses the opportunity to raise the oil price markup. In particular, a typical fringe shock raises the oil price markup by 3%, which translates into a 20 basis points decrease in the GDP gap. Therefore, if the central bank wants to offset completely the effect of the shock on the GDP gap, it would have to lower the interest rate by around 20 basis points relative to the benchmark rule. Of course, in both scenarios, there is no reason why the central bank should want to completely insulate the GDP gap from the shock, since that would generate below target inflation (deflation) in the former case, and above target inflation in the latter.

Lastly, if the oil price rise is caused by a rise in technology (and oil productivity) in the US, the interest rate should be set *lower* than the benchmark rule for a reason independent of the oil price movement. Namely, the interest rate smoothing of rule (2.10) prevents it from offsetting the GDP gap and inflation fall due to the shock in the presence of nominal rigidities. Unlike the previous two disturbances, for this shock the tradeoff between inflation and GDP gap stabilization is quantitatively small.

Shock		Cond. correlation Y	$\tilde{Y}$	Π	υ	Desirable deviation from benchmark rule (2.10) in response to an oil price rise
Z	p_o^Z	$-$	$+$	$+$	$-$	$R\uparrow$ to stabilize $\tilde{Y}$, tradeoff for deflation
Ω	p_o^Ω	$-$	$-$	$+$	$+$	$R\downarrow\uparrow$, traditional $\Pi-\tilde{Y}$ tradeoff
A	p_o^A	$+$	$-$	$-$	$+$	$R\downarrow$ to stabilize Π, virtually no tradeoff
R	p_o^R	$+$	$+$	$+$	$+$	

Table 2.2. Oil price correlations and policy implications conditional on shock

2.4.6 A note on Taylor-type reaction to the oil price

Taylor (1993)-type rules are often advocated as useful guidelines for policy on the basis of their simplicity and good performance (in terms of implied welfare loss) in the standard sticky price model. In its simplest form, in the context of the New Keynesian model, a Taylor rule prescribes that the central bank should adjust sufficiently the interest rate in response to variations in inflation and the welfare-relevant output gap. In fact, as already discussed, in the standard New Keynesian model stabilizing inflation is equivalent to stabilizing the output gap and hence the latter term can be dropped from the rule. But in the absence of "divine coincidence" of monetary policy objectives, as in this model, the presence of the output gap in the rule is justified as it would result in superior performance in general compared to a rule which reacts to inflation only.

Unlike inflation, though, the output gap is an unobservable variable, making a rule which reacts to it less useful as a policy guide. In our context, it may be interesting to know whether there is an observable variable, perhaps the oil price or its change, which is a good substitute for the output gap. Indeed, to the extent that some inflation-targeting central banks target not "core" but "headline" inflation, which includes the price of energy, a Taylor type rule would implicitly react to energy price changes proportionately to the share of energy in CPI. What can we say about the advisability of a Taylor rule reacting to the oil price on the basis of our findings?

From our discussion in the previous section it is already clear that a mechanical Taylor-type reaction to the oil price regardless of the source of the shock is not likely to be very useful, and might even be harmful. The reason is that, as witnessed in table 2.2, the correlation of the oil price with the output gap can be either positive or negative conditional on the type of the shock. As a result, the unconditional correlation between the oil price and the output gap can be quite weak (–0.11 under our benchmark calibration).

As shown in section 2.2.6, it is instead the oil price *markup* which enters unambiguously in the expression for the output gap. And while the oil price markup may be difficult to come by in practice because of the lack of reliable estimates of OPEC's marginal costs, according to our model it should be highly positively correlated with OPEC's market share, a variable which is more directly observable. In this sense, rather than removing energy prices from the "headline" consumer price index to obtain an index of "core" inflation, our analysis suggests treating the oil price markup (or OPEC's market share) as an independent target variable.

2.4.7 Variance decomposition

To assess the relative importance of the four sources of fluctuations in our model, in table 2.3 we show the variance decomposition for several key variables, along with their unconditional standard deviations. Clearly, these statistics are sensitive to our calibration of the shock processes.

In particular, under our baseline calibration, US technology shocks account for around 40% of the volatility of inflation, 68% of the volatility of output, but only 3% of the volatility of the welfare-relevant output gap. Oil technology shocks are responsible for around 16% of the volatility of inflation, 26% of the volatility of output, and as much as 44% of the volatility of the output gap. Fringe shocks contribute only 1% of the volatility of inflation and 5% of the volatility of output, but as much as 44% of the volatility of the output gap. And monetary policy shocks are responsible for 44% of the volatility of inflation, 1% of the volatility of output, and 8% of the volatility of the output gap.

Not surprisingly, US output, inflation and the interest rate can be explained to a large extent by the US-originating technology and monetary policy shocks. Still, as much as 31% of US output variance and 17% of US inflation volatility can be accounted for by the combined contribution of oil technology and fringe shocks. More importantly, these two shocks together contribute close to 89% of the variance in the welfare-relevant output gap. Since these are precisely the shocks that make monetary policy interesting (in the sense of inducing a meaningful policy tradeoff), the fact that they account for much of the output gap and inflation variability confirms that the lack of a policy tradeoff in the standard New Keynesian model is just a coincidence.

Another way of seeing this is by observing that under the benchmark policy rule the bulk of the volatility of the actual output gap (std 93 basis points) is due to fluctuations in the natural output gap (std 81 bp). Indeed, the correlation between these two output gap measures is around 0.95. In contrast, the correlation between the natural output gap and the sticky price output gap (std 29 bp) is +0.26. In other words, monetary regime (2.10) which targets only inflation misses on the opportunity to stabilize the welfare-relevant output gap by countering fluctuations in the natural output gap (caused by OPEC's time-varying market power), through *opposite* movements in the sticky price output gap (which would entail a negative correlation between the two).

		Std	Variance due to			
			A	Z	Ω	R
US output	Y	0.76	67.50	26.13	5.36	1.01
Output gap	$\tilde{Y}$	0.93	2.94	44.24	44.44	8.38
Natural output gap	$\tilde{Y}^n$	0.81	0.15	44.62	55.23	0.00
Sticky price output gap	$\tilde{Y}^s$	0.29	19.15	7.28	0.24	73.33
Inflation	Π	0.63	39.79	15.72	0.76	43.73
Interest rate	R	0.68	63.20	24.54	1.61	10.65

Table 2.3. Variance decomposition

Note: for inflation and the interest rate "std" is annualized; for US output "std" is the standard deviation (in percentage points) of the quarterly growth rate of output.

2.5 Sensitivity Analysis

In this section we report the sensitivity of our main findings to the elasticity of oil in production as well as to the monetary policy regime in place.

2.5.1 The elasticity of oil in production

Expression (2.92) for the output gap and the expressions for κ_ν and κ_{mc} suggest that the elasticity of oil in final goods production is likely to be an important parameter affecting the model's dynamics. At the same time there is evidence that, at least in the US, this parameter has declined, so that today the oil share in GDP is much smaller than what it used to be three decades ago. To test the extent to which the macroeconomic effects of oil sector shocks depend on this elasticity, we recompute our model with a twice larger oil share, by reducing the share of labor to 0.6 and the share of capital to 0.3.

We find that the impact of oil sector shocks on the US economy approximately doubles with respect to the baseline. In particular, the impact of a typical oil technology shock that raises the oil price by 9% is now a 0.75% drop in US output, a rise (narrowing) of the output gap by 55 basis points, and an increase in inflation by 40 basis points. The impact of a typical fringe shock which increases the oil price by 2.5% is a 0.25% drop in US output, a corresponding fall (widening) of the output gap by 25 basis points and a rise in inflation by 5 basis points.

A larger oil share amplifies the responses of US output and inflation also to monetary policy shocks. The overall effect is that doubling the oil share increases the unconditional volatility of US inflation by around 25%, of output by 41%, and of the output gap by 94% with respect to the baseline. These volatility effects are substantial and point to the possibility that reduced dependence of the US economy on oil may have played an important role in the pronounced decline in US inflation and output volatility since the mid 1980-s (a phenomenon dubbed by some economists as the "Great Moderation", e.g. McConnell and Perez-Quiros (2000)).

2.5.2 Monetary policy

Table 2.4 summarizes the stabilization properties of several monetary policy regimes in terms of the implied volatility of US welfare-relevant variables, as well as the impact responses to oil sector shocks (normalized to produce the same 10% increase in the oil price). The alternative monetary policies considered include the benchmark rule (2.10); full price stability, $\Pi_t = 1$; constant nominal interest rate, $R_t = 1/\beta$; rule (2.10) with $\phi_\pi = 2$ and without interest-rate smoothing, $\phi_R = 0$; rule (2.10) with $\phi_\pi = 2$ without smoothing and with (the optimal) oil price reaction $\phi_o = -0.02$; and rule (2.10) with smoothing and with (a sub-optimal) oil price reaction, $\phi_o = +0.04$. In what follows, we discuss briefly three of these monetary policy regimes.

Constant interest rate policy

How would the economy evolve in the wake of an "oil shock" if the interest rate did not react to any endogenous variable, but instead remained constant? To answer this question we simulate our model under the assumption that the central bank follows a constant nominal interest rate policy.[21]

We find that this rule amplifies dramatically the effects of oil sector shocks on the US economy. In particular, the impact of an oil technology shock which raises the oil price by 10% is an increase in inflation by 2 percentage points — a response which is ten times larger compared to the benchmark policy! US output *increases* by 0.45%, raising (narrowing) the output gap by more than a full percentage point — four times more than the benchmark policy! And in response to a negative fringe capacity shock which raises the oil price by the same 10%, US output (and the output gap) falls by 4% (4 percentage points for the gap), while inflation *falls* by more than 7 percentage points!

The reason for this very different impact of oil sector shocks is that a constant nominal interest rate policy implies that any movements in expected inflation (including those induced by oil sector developments) translate one-for-one to opposite movements in the ex-ante real interest rate, with the usual consequences for output demand and inflation. For instance, in response to a negative oil technology shock which lowers the efficient level of output, the dominant oil firm optimally commits to reducing future oil supply, inducing a rise in expected inflation. With a constant nominal rate, this lowers

[21] In our model, the endogeneity of the oil price implies that the Blanchard and Khan (1980) conditions for local determinacy of the solution are satisfied even under a constant interest rate policy. See chapter 3 and figure 3.4 for details.

the ex-ante real interest rate and stimulates US activity so that instead of falling, output actually increases. The latter boosts temporarily oil demand and mitigates the negative impact of the shock on the dominant oil supplier's profits. Thus, in the absence of active monetary policy, the pursuit of profit-smoothing on behalf of the dominant oil firm comes at the cost of higher volatility in the oil importer. As a result, the unconditional output volatility increases by 55%, output gap volatility doubles, and inflation volatility increases by a factor of 4.7 with respect to the benchmark policy rule!

Optimal uniform reaction to oil price changes

In section 4.6 we discussed the reasons why a uniform Taylor-type reaction to the oil price is not likely to improve significantly on the benchmark rule. To quantify the extent to which it might help, we compute the optimal uniform reaction to oil price changes, conditional on fixing the long-run reaction coefficient on inflation to its baseline value, and considering the cases with and without interest rate smoothing. To find the optimal coefficient, we approximate the solution of our model to second order and evaluate directly the expected welfare of the US consumer, conditional on the economy starting in the deterministic steady-state.

In the case with interest rate smoothing, the optimal uniform reaction to oil price changes is virtually zero and the welfare gain with respect to the benchmark rule is negligible. We then set the interest rate smoothing parameter to zero while maintaining the same long-run inflation response coefficient. This removes the dependence of the nominal interest rate on oil price and CPI inflation which occurred in the more distant past. We find that in this case, the expected welfare of the US consumer is maximized for a value of the reaction coefficient on the oil price $\phi_o \approx -0.02$.

The specific value of ϕ_o is not very interesting since it is clearly sensitive to the calibration (the relative size of the shocks) of our model. In particular, the optimal reaction should induce more efficient responses to the shocks which fall more strongly on welfare-relevant variables. However, the gain in expected welfare under this rule vis-a-vis the same rule with $\phi_o = 0$ is quite modest — equivalent to a permanent rise in consumption of only 0.02% (or around $1.8 billion per year based on US consumption expenditure in 2006).

Sub-optimal uniform reaction to oil price changes

If the optimal uniform reaction does not improve significantly on the performance of the benchmark policy, how harmful can a sub-optimal Taylor-type reaction to the oil price be (assuming a plausible response coefficient)? Let us suppose that the monetary authority chooses a contemporaneous reaction coefficient to oil price inflation $\phi_o = 0.04$ keeping all other parameters constant (that is, a long run inflation reaction of 2).

In response to a negative oil technology shock which raises the oil price by 10%, the nominal interest rate increases by around 125 basis points. As a result output falls by 1.3% and inflation falls by 90 basis points. Importantly, US output falls by *more* than the efficient decrease *widening* the output gap by around 50 basis points (contrary to the output gap *narrowing* by around 25 bp under the benchmark rule). And in response to a negative fringe shock which raises the oil price by 10%, output falls by 1.5% which widens the output gap by 150 basis points (compared to the 50 bp widening under the benchmark rule), at the same time as inflation decreases by around 140 basis points. Therefore, this policy is clearly destabilizing, throwing the economy into an unnecessary recession in response to oil sector shocks which raise significantly the oil price.

	Benchmark	$\Pi_t = 1$	$R_t = \beta^{-1}$	$\phi_R = 0$	$\phi_o = -.02$	$\phi_o = .04$
Unconditional standard deviation						
Output gap	0.93	0.81	1.85	0.83	0.85	1.11
Inflation	0.63	0	2.99	0.43	0.42	1.28
Interest rate	0.68	0.45	0	0.87	0.89	1.70
Impact responses to an oil tech. shock that raises the oil price by 10%						
Output	-0.43	-0.53	0.45	-0.52	-0.36	-1.29
Output gap	0.28	0.21	1.10	0.21	0.36	-0.51
Inflation	0.21	0	2.01	0.11	0.20	-0.90
Interest rate	0.08	0.11	0	0.21	-0.40	1.25
Impact responses to a fringe shock that raises the oil price by 10%						
Output (gap)	-0.49	-0.53	-4.10	-0.54	-0.37	-1.49
Inflation	0.11	0	-7.77	0.03	0.16	-1.39
Interest rate	0.04	0.05	0	0.06	-0.49	1.05

Table 2.4. Stabilization properties of alternative policy rules

Note: output (%); output gap (percentage points); inflation and interest rate (pp annualized)

To sum up, we find that the monetary policy *regime* in place in the US plays an important role for the behavior of the oil sector and the way in which oil sector shocks are transmitted to the US economy.

2.6 Conclusion

Killian (2006) argues that the economics profession should move beyond studying the effects of changes in the real price of oil and address the problem of identifying the structural shocks underlying such changes. Only then can economists make the next step of evaluating alternative policies in response to the fundamental shocks. Our model is an attempt in that direction, demonstrating how oil technology and fringe capacity shocks in the oil producing part of the world, combined with monetary policy and TFP shocks in the oil importing region, are transmitted to the price of oil in a world oil market dominated by OPEC. At the same time, and conditional on the monetary policy *regime* in place, each of these shocks affects through different channels the evolution of macroeconomic variables relevant for the oil importer, and as a consequence has distinct policy implications.

Unlike previous studies of the link between oil and the macroeconomy, we model explicitly OPEC as a dominant oil supplier with a fringe of competitive oil producers. This implies that, in equilibrium, the supply of oil fluctuates around an inefficiently low level, reflected in a positive oil price markup and a negative output gap in the US. Importantly, shocks in this environment induce inefficient variation of the oil price markup, and create a meaningful tradeoff between inflation and output gap stabilization — a feature which many central bankers perceive as realistic, but which is absent from the standard monetary policy model.

We are aware that by assuming a frictionless labor market we may be understating the efficiency costs of oil sector shocks. Moreover, our analysis ignores several potentially important aspects of the oil industry: the fact that oil is a storable commodity, which is actively traded on futures markets, and the long gestation lags in adding productive capacity, to name a few. By making oil supply less responsive in the short-run, the latter in particular may be relevant for explaining the puzzlingly high volatility of the oil price relative to oil output. At the same time, we may be omitting other important shocks, for example precautionary demand shifts due to fears about future oil availability (Killian, 2006). We must leave for future research the analysis of some of these issues in an appropriately modified framework.

2.A First Order Conditions of OPEC's Problem

Let $s_o \equiv 1 - \alpha_1 - \alpha_2$ denote the share of oil in GDP. The following are the first order conditions of Ramsey-type problem (2.54) solved by the dominant oil exporter:

$$
\begin{aligned}
0 &= 1/O_t - (\lambda_{1t} + \lambda_{7t})p_{ot} + \lambda_{9t}s_o\frac{Q_t\Delta_t}{O_t + X_t} && (O_t)\\
0 &= -\lambda_{1t} + \lambda_{2t}E_t\left[\frac{\beta R_t}{C_{t+1}\Pi_{t+1}}\right] - \lambda_{2t-1}\frac{R_{t-1}C_{t-1}}{C_t^2\Pi_t} \\
&\quad +\lambda_{3t}E_t\left[\beta\theta\Pi_{t+1}^{\epsilon-1}D_{t+1} - D_t\right] + \lambda_{4t}E_t\left[\beta\theta\Pi_{t+1}^{\epsilon}N_{t+1} - N_t\right] + \lambda_{10t}L_t^{\psi} && (C_t)\\
0 &= \lambda_{1t} + \lambda_{3t} + \lambda_{4t}\mu mc_t + \lambda_{7t}s_o mc_t\Delta_t - \lambda_{8t}\alpha mc_t\Delta_t - \lambda_{9t}\Delta_t && (Q_t)\\
0 &= \lambda_{3t-1}\theta C_{t-1}\Pi_t^{\epsilon-1} - \lambda_{3t}C_t + \lambda_{5t}\left(1-\theta\right)\left(\epsilon-1\right)N_t^{1-\epsilon}D_t^{\epsilon-2} \\
&\quad +\lambda_{6t}(1-\theta)\epsilon N_t^{-\epsilon}D_t^{\epsilon-1} && (D_t)\\
0 &= \lambda_{4t-1}\theta C_{t-1}\Pi_t^{\epsilon} - \lambda_{4t}C_t + \lambda_{5t}\left(1-\theta\right)\left(1-\epsilon\right)N_t^{-\epsilon}D_t^{\epsilon-1} \\
&\quad -\lambda_{6t}(1-\theta)\epsilon N_t^{-\epsilon-1}D_t^{\epsilon} && (N_t)\\
0 &= \lambda_{8t}w_t + \lambda_{9t}\alpha_1 Q_t\Delta_t/L_t + \lambda_{10t}C_t\psi L_t^{\psi-1} && (L_t)\\
0 &= \frac{1}{p_{ot} - Z_t^{-1}} - (O_t + 2X_t)(\lambda_{1t} + \lambda_{7t}) + \lambda_{9t}s_o\frac{Q_t\Delta_t}{O_t + X_t}\Omega_t Z_t \\
&\quad +\lambda_{11t}\phi_o\bar{R}^{1-\phi_R}R_{t-1}^{\phi_R}\left[\frac{\Pi_t}{\bar{\Pi}}\right]^{\phi_\pi}p_{ot}^{\phi_o-1}p_{ot-1}^{-\phi_o} \\
&\quad -\beta E_t\left[\lambda_{11t+1}\bar{R}^{1-\phi_R}R_{t-1}^{\phi_R}\left(\frac{\Pi_{t+1}}{\bar{\Pi}}\right)^{\phi_\pi}p_{ot+1}^{\phi_o}\phi_o p_{ot}^{-\phi_o-1}\right] && (p_{ot})\\
0 &= -\lambda_{2t-1}R_{t-1}\frac{C_{t-1}}{C_t}\Pi_t^{-2} + \lambda_{3t-1}\theta\left(\epsilon-1\right)C_{t-1}\Pi_t^{\epsilon-2}D_t \\
&\quad +\lambda_{4t-1}\theta\epsilon C_{t-1}\Pi_t^{\epsilon-1}N_t + \lambda_{5t}\left(\epsilon-1\right)\theta\Pi_t^{\epsilon-2} \\
&\quad +\lambda_{6t}\theta\epsilon\Pi_t^{\epsilon-1}\Delta_{t-1} + \lambda_{11t}\bar{R}^{1-\phi_R}R_{t-1}^{\phi_R}\phi_\pi\Pi_t^{\phi_\pi-1}\bar{\Pi}^{-\phi_\pi}\left[\frac{p_{ot}}{p_{ot-1}}\right]^{\phi_o} && (\Pi_t)\\
0 &= E_t\left[\lambda_{6t+1}\beta\theta\Pi_{t+1}^{\epsilon}\right] - \lambda_{6t} + \lambda_{7t}s_o mc_t Q_t - \lambda_{8t}\alpha_1 mc_t Q_t - \lambda_{9t}Q_t && (\Delta_t)\\
0 &= \lambda_{4t}\mu Q_t + \lambda_{7t}(1-\alpha_1-\alpha_2)Q_t\Delta_t - \lambda_{8t}\alpha_1 Q_t\Delta_t && (mc_t)
\end{aligned}
$$

$$\begin{aligned}
0 &= \lambda_{2t}\beta E_t\left[\frac{C_t}{C_{t+1}\Pi_{t+1}}\right]-\lambda_{11t} \\
&+E_t\left[\lambda_{11t+1}e^{r_t}\beta\phi_R\left(\frac{R_t}{\bar{R}}\right)^{\phi_R-1}\left(\frac{\Pi_{t+1}}{\bar{\Pi}}\right)^{\phi_\pi}\left(\frac{p_{ot+1}}{p_{ot}}\right)^{\phi_o}\right] && (R_t) \\
0 &= \lambda_{8t}L_t-\lambda_{10t} && (w_t)
\end{aligned}$$

2.B GDP Gap Derivation and Proof of Proposition 3

GDP Gap Derivation

We denote the natural (flexible-price) level of variables with the superscript n, the efficient level of variables with the superscript e, and steady-state variables with an upper bar. We define the production gap $\tilde{q}_t$ as the (log) difference between actual production Q_t and its efficient level Q_t^e.

From the final-goods production function (2.42) and the oil share condition (2.39), we obtain the following two equations,

$$\frac{\Delta_t Q_t}{Q_t^e} = \left(\frac{L_t}{\bar{L}}\right)^{\alpha_1}\left(\frac{O_{dt}}{O_{dt}^e}\right)^{1-\alpha_1-\alpha_2}, \tag{2.94}$$

$$\frac{p_{ot}O_{dt}}{p_{ot}^e O_{dt}^e} = \frac{mc_t\Delta_t}{\mu^{-1}}\frac{Q_t}{Q_t^e}. \tag{2.95}$$

Eliminating the O_{dt}/O_{dt}^e from the above equations and taking log-deviations (denoted by hats) yields,

$$\tilde{q}_t = \frac{\alpha_1}{\alpha_1+\alpha_2}\hat{l}_t + \frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}\left(\hat{mc}_t+\hat{\Delta}_t\right) - \frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}\left(\hat{p}_{ot}-\hat{p}_{ot}^e\right). \tag{2.96}$$

Taking a log-linear approximation of (2.67), we obtain,

$$(1+\psi)\hat{l}_t = \mu[\mu-1+\alpha_1+\alpha_2]^{-1}\hat{mc}_t. \tag{2.97}$$

Combining it with expression (2.96) above, dropping the constant and the higher order term related to Δ_t, we obtain the following first-order approximation for the production gap,

$$\tilde{q}_t \simeq \kappa^q_{mc}\hat{mc}_t - \kappa_\nu \hat{\nu}_t, \tag{2.98}$$

where $\hat{\nu}_t = \hat{p}_{ot} - \hat{p}^e_{ot}$, and

$$\kappa^q_{mc} \equiv \frac{\mu\alpha_1 + (1-\alpha_1-\alpha_2)(1+\psi)(\mu - 1 + \alpha_1 + \alpha_2)}{(1+\psi)(\mu - 1 + \alpha_1 + \alpha_2)(\alpha_1 + \alpha_2)} \tag{2.99}$$

$$\kappa_\nu \equiv \frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}. \tag{2.100}$$

To obtain the GDP gap $\tilde{y}_t$, defined as the log-difference between U.S. *value added* and its efficient level, we use the condition

$$Y_t = C_t = [1 - (1-\alpha_1-\alpha_2)mc_t\Delta_t]Q_t, \tag{2.101}$$

which implies

$$\tilde{y}_t = \tilde{q}_t - \frac{1-\alpha_1-\alpha_2}{\mu - 1 + \alpha_1 + \alpha_2}\hat{mc}_t. \tag{2.102}$$

Combining the above expression with the production gap (2.98), we obtain

$$\tilde{y}_t = \kappa_{mc}\hat{mc}_t - \kappa_\nu \hat{\nu}_t, \tag{2.103}$$

where

$$\kappa_{mc} \equiv \kappa^q_{mc} - \frac{1-\alpha_1-\alpha_2}{\mu - 1 + \alpha_1 + \alpha_2} = \frac{\mu\alpha_1 + (1-\alpha_1-\alpha_2)(1+\psi)(\mu-1)}{(1+\psi)(\mu - 1 + \alpha_1 + \alpha_2)(\alpha_1 + \alpha_2)}. \tag{2.104}$$

Proof of proposition 3

With exogenous or competitive oil price, $p^e_{ot} = Z_t^{-1}$, the model can be written as follows: the production function is given by,

$$Q_t = A_t L_t^{\alpha_1} \bar{K}^{\alpha_2} O_{dt}^{1-\alpha_1-\alpha_2}/\Delta_t. \tag{2.105}$$

Combining (2.40) with (2.22) we obtain,

$$C_t L_t^{1+\psi} = \alpha_1 mc_t Q_t \Delta_t. \tag{2.106}$$

And combining (2.39) and (2.106) yields,

$$p^e_{ot} O_{dt} = \frac{1 - \alpha_1 - \alpha_2}{\alpha_1} C_t L_t^{1+\psi}. \tag{2.107}$$

Resource constraint,

$$C_t = Q_t - p^e_{ot} O_{dt}. \tag{2.108}$$

The above four equations, together with (2.58), (2.59), (2.60), (2.61), describe fully the behavior of the private sector. A benevolent monetary policy maker who wants to maximize the welfare of the representative US household would solve the following Lagrangian:

$$\begin{aligned}
&\max \; E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \log C_t - \frac{L_t^{1+\psi}}{1+\psi} \right. \qquad (2.109)\\
&+\lambda_{1t} \left[Q_t - p^e_{ot} O_{dt} - C_t \right] \\
&+\lambda_{2t} \left[\frac{1 - \alpha_1 - \alpha_2}{\alpha_1} C_t L_t^{1+\psi} - p^e_{ot} O_{dt} \right] \\
&+\lambda_{3t} \left[A_t L_t^{\alpha_1} \bar{K}^{\alpha_2} O_{dt}^{1-\alpha_1-\alpha_2} / \Delta_t - Q_t \right] \\
&+\lambda_{4t} \left[\alpha_1 mc_t Q_t \Delta_t - C_t L_t^{1+\psi} \right] \\
&+\lambda_{5t} \left[Q_t + \beta\theta C_t E_t \left[\Pi_{t+1}^{\epsilon-1} D_{t+1} \right] - C_t D_t \right] \\
&+\lambda_{6t} \left[\mu mc_t Q_t + \beta\theta C_t E_t \left(\Pi_{t+1}^{\epsilon} N_{t+1} \right) - C_t N_t \right] \\
&+\lambda_{7t} \left[\theta \Pi_t^{\epsilon-1} + (1-\theta) \left(\frac{N_t}{D_t} \right)^{1-\epsilon} - 1 \right] \\
&\left. +\lambda_{8t} \left[\theta \Pi_t^{\epsilon} \Delta_{t-1} + (1-\theta) \left(\frac{1 - \theta \Pi_t^{\epsilon-1}}{1-\theta} \right)^{\frac{\epsilon}{\epsilon-1}} - \Delta_t \right] \right\}
\end{aligned}$$

System of first-order conditions:

$$1/C_t - \lambda_{1t} + \frac{p^e_{ot}O_{dt}}{C_t}\lambda_{2t} - L_t^{1+\psi}\lambda_{4t} - \lambda_{5t}Q_t/C_t - \lambda_{6t}\mu mc_t Q_t/C_t = 0 \qquad (C_t)$$

$$-L_t^{\psi} + (1+\psi)\lambda_{2t}\frac{p^e_{ot}O_{dt}}{L_t} + \alpha_1\lambda_{3t}Q_t/L_t - (1+\psi)C_t L_t^{\psi}\lambda_{4t} = 0 \qquad (L_t)$$

$$p^e_{ot}O_{dt}(\lambda_{1t} + \lambda_{2t}) - (1-\alpha_1-\alpha_2)Q_t\lambda_{3t} = 0 \qquad (O_{dt})$$

$$\lambda_{1t} - \lambda_{3t} + \alpha_1 mc_t\Delta_t\lambda_{4t} + \lambda_{5t} + \mu mc_t\lambda_{6t} = 0 \qquad (Q_t)$$

$$\alpha_1\Delta_t\lambda_{4t} + \mu\lambda_{6t} = 0 \qquad (mc_t)$$

$$\alpha_1 mc_t Q_t\lambda_{4t} + \lambda_{8t+1}\beta\theta\Pi_{t+1}^{\epsilon} - \lambda_{8t} - \frac{Q_t}{\Delta_t}\lambda_{3t} = 0 \qquad (\Delta_t)$$

$$\lambda_{5t-1}\theta C_{t-1}\Pi_t^{\epsilon-1} - C_t\lambda_{5t} - \lambda_{7t}(1-\theta)(1-\epsilon)N_t^{1-\epsilon}D_t^{\epsilon-2} = 0 \qquad (D_t)$$

$$\lambda_{6t-1}\theta C_{t-1}\Pi_t^{\epsilon} - C_t\lambda_{6t} + \lambda_{7t}(1-\theta)(1-\epsilon)N_t^{-\epsilon}D_t^{\epsilon-1} = 0 \qquad (N_t)$$

$$\lambda_{5t-1}(\epsilon-1)\Pi_t^{\epsilon-2}D_tC_{t-1} + \lambda_{6t-1}\epsilon\Pi_t^{\epsilon-1}N_tC_{t-1} + \lambda_{7t}(\epsilon-1)\Pi_t^{\epsilon-2} + \qquad (\Pi_t)$$
$$+\lambda_{8t}\epsilon\left[\Pi_t^{\epsilon-1}\Delta_{t-1} - \left(\frac{1-\theta\Pi_t^{\epsilon-1}}{1-\theta}\right)^{\frac{1}{\epsilon-1}}\Pi_t^{\epsilon-2}\right] = 0$$

In what follows we guess and verify that zero inflation in each period is a solution.

From (2.60), our guess $\Pi_t = 1$ implies that $N_t = D_t$. This, from (2.58) and (2.59) yields $mc_t = \mu^{-1}$ (the price markup is constant). In addition, from (2.61) and starting with $\Delta_{-1} = 0$, we have $\Delta_t = 1$ (there is no price dispersion). Substituting $\Delta_t = 1$ and $mc_t = \mu^{-1}$ in (2.67) we obtain

$$L_t = \left[\frac{\alpha_1}{\mu - (1-\alpha_1-\alpha_2)}\right]^{\frac{1}{1+\psi}} = \bar{L}, \qquad (2.110)$$

which is equal to the efficient level of labor effort established in section 2.2.5. Rewriting (2.39) using the above results,

$$p^e_{ot}O_{dt} = (1-\alpha_1-\alpha_2)Q_t/\mu, \qquad (2.111)$$

and substituting O_{dt} from above equation into (2.105), we obtain

$$Q_t = \left[A_t Z_t^{1-\alpha_1-\alpha_2} \bar{L}^{\alpha_1} \bar{K}^{\alpha_2}\right]^{\frac{1}{\alpha_1+\alpha_2}} \left[(1-\alpha_1-\alpha_2)\mu^{-1}\right]^{\frac{1-\alpha_1-\alpha_2}{\alpha_1+\alpha_2}} = Q_t^e \tag{2.112}$$

which is equal to the efficient level of output derived in section 2.2.5. Moreover,

$$C_t = \left[1-(1-\alpha_1-\alpha_2)\mu^{-1}\right] Q_t = C_t^e \tag{2.113}$$

corresponds to the efficient level of consumption. All other real endogenous variables can be expressed similarly in terms of $Q_t = Q_t^e$. Thus, a policy of full price stability replicates the real allocation attained in the efficient (first-best) equilibrium.

Using the above, we can rewrite the conditions that have to be satisfied by the Lagrange multipliers as follows:

$$\begin{aligned}
\lambda_{3t} &= \frac{1}{(\alpha_1+\alpha_2)Q_t} && (2.114)\\
\lambda_{5t} &= \lambda_{3t} - \lambda_{1t} && (2.115)\\
\lambda_{6t} &= -\lambda_{4t}\alpha_1/\mu && (2.116)\\
p_{ot}^e O_{dt}\left(\lambda_{2t} - \frac{\alpha_1}{1-\alpha_1-\alpha_2}\lambda_{4t}\right) &= \frac{\bar{L}^{1+\psi}}{(1+\psi)} - \frac{\alpha_1}{(1+\psi)(\alpha_1+\alpha_2)} && (2.117)\\
\alpha_1 Q_t \lambda_{4t}/\mu + \lambda_{8t+1}\beta\theta - \lambda_{8t} &= \frac{1}{\alpha_1+\alpha_2} && (2.118)\\
\lambda_{5t-1}\theta C_{t-1} - C_t\lambda_{5t} &= \lambda_{7t}(1-\theta)(1-\epsilon)/D && (2.119)\\
\lambda_{6t-1}\theta C_{t-1} - C_t\lambda_{6t} &= -\lambda_{7t}(1-\theta)(1-\epsilon)/D && (2.120)\\
\lambda_{5t-1}(\epsilon-1)C_{t-1} + \lambda_{6t-1}\epsilon C_{t-1} &= \lambda_{7t}(1-\epsilon)/D && (2.121)
\end{aligned}$$

A recursive solution consistent with the "timeless perspective" requires to set $\lambda_{5,-1} + \lambda_{6,-1} = 0$:

$$\begin{aligned}
\lambda_{5t} &= -\lambda_{6t} && (2.122)\\
C_t\lambda_{6t} &= C_{t-1}\lambda_{6t-1} && (2.123)\\
\lambda_{7t} &= \frac{D}{1-\epsilon}C_t\lambda_{6t} && (2.124)\\
\lambda_{4t} &= -\frac{\mu}{\alpha_1}\lambda_{6t} && (2.125)\\
(1-\beta\theta)^{-1}\lambda_{8t} &= -1/(\alpha_1+\alpha_2) - \frac{Q}{C}C_t\lambda_{6t} && (2.126)\\
\lambda_{1t} &= \lambda_{3t} - \lambda_{5t}. && (2.127)
\end{aligned}$$

2.C Figures

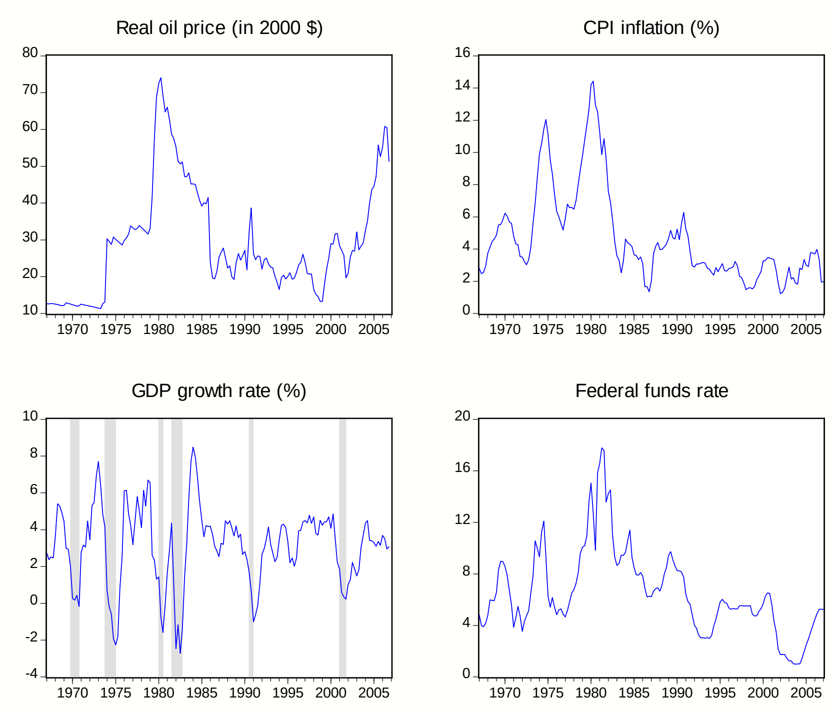

Fig. 2.1. Price of oil, inflation, growth and Fed rate

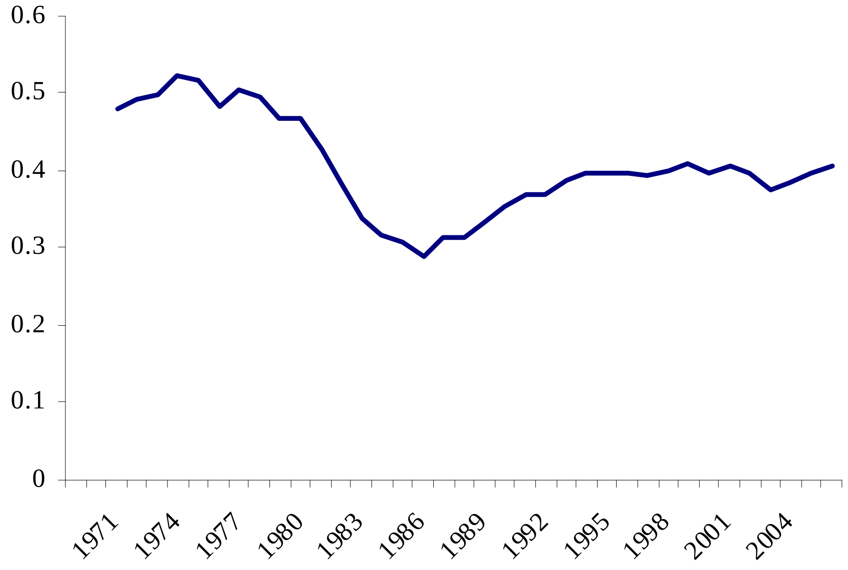

Fig. 2.2. OPEC's world market share

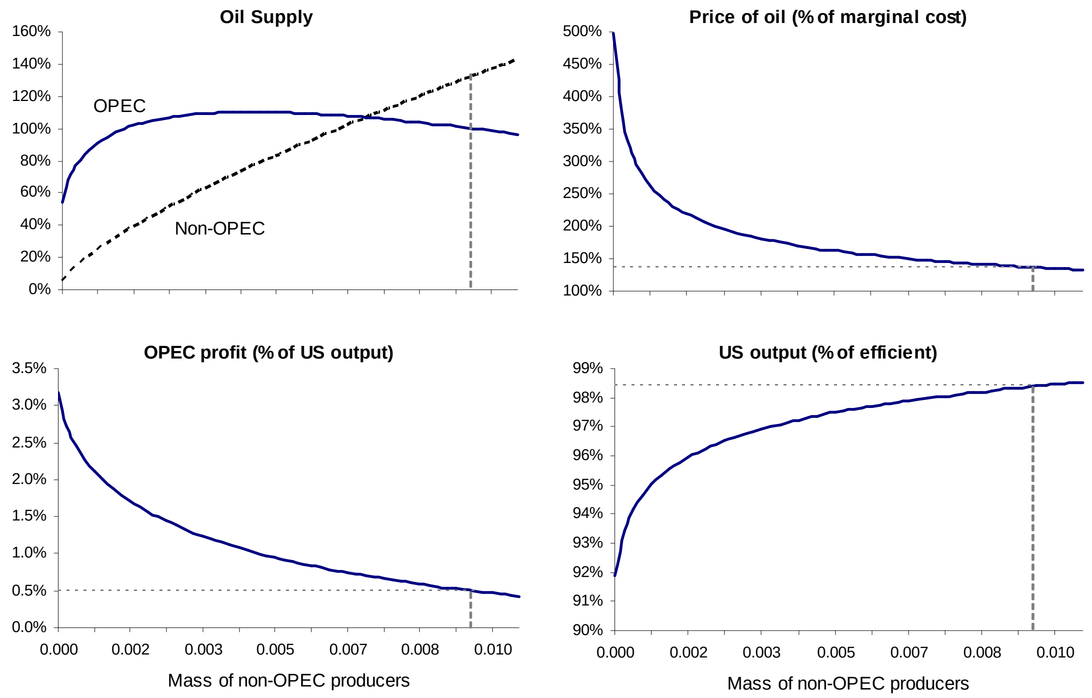

Fig. 2.3. Comparative statics: mass of non-OPEC producers ($\bar{\Omega}$)

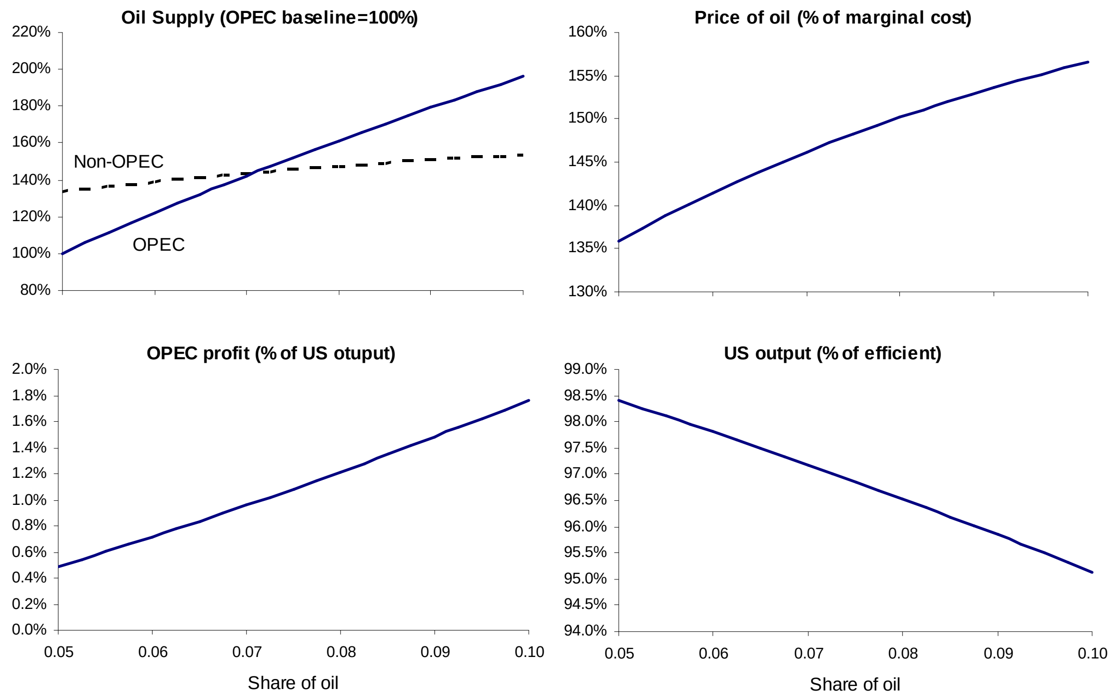

Fig. 2.4. Comparative statics: elasticity of oil in production (keeping the mass of non-OPEC constant)

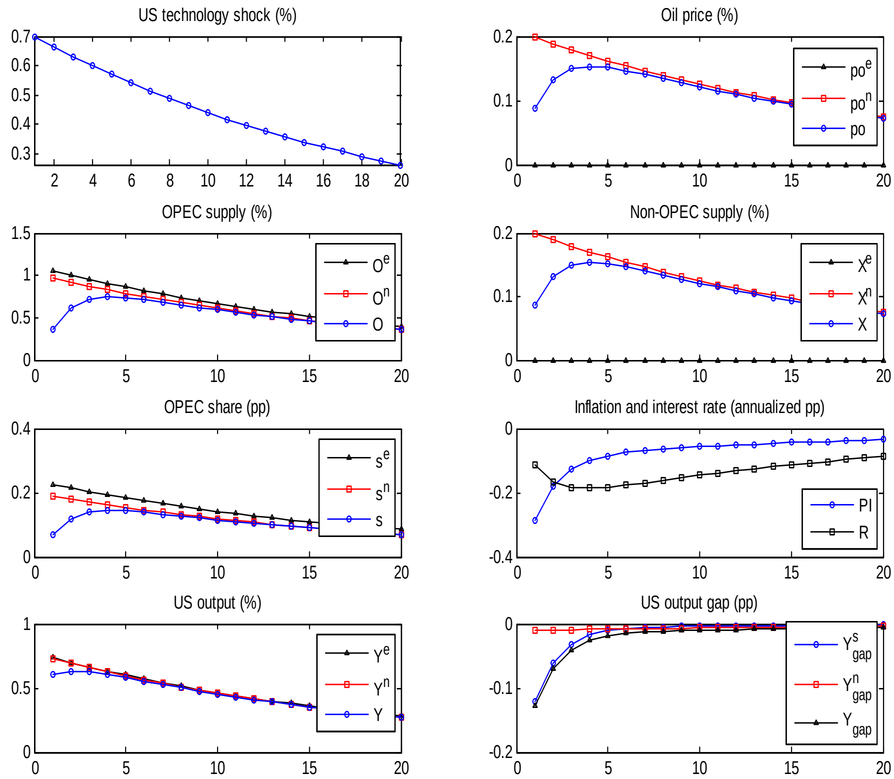

Fig. 2.5. Responses to a positive US TFP shock

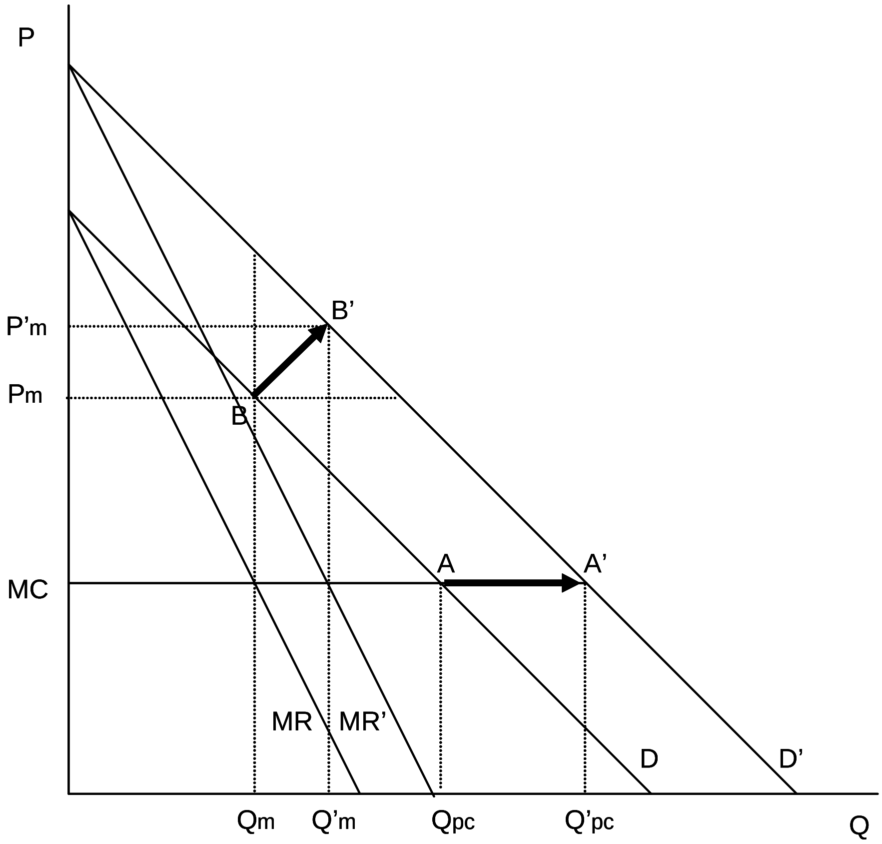

Fig. 2.6. Competitive and monopoly response to a positive demand shock.

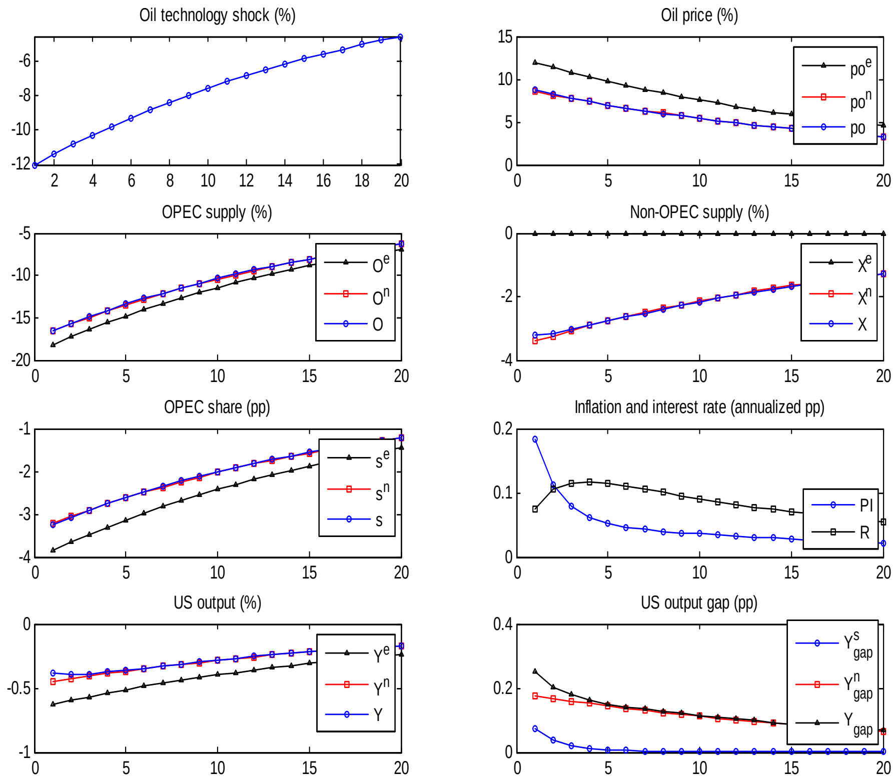

Fig. 2.7. Responses to a negative oil technology shock

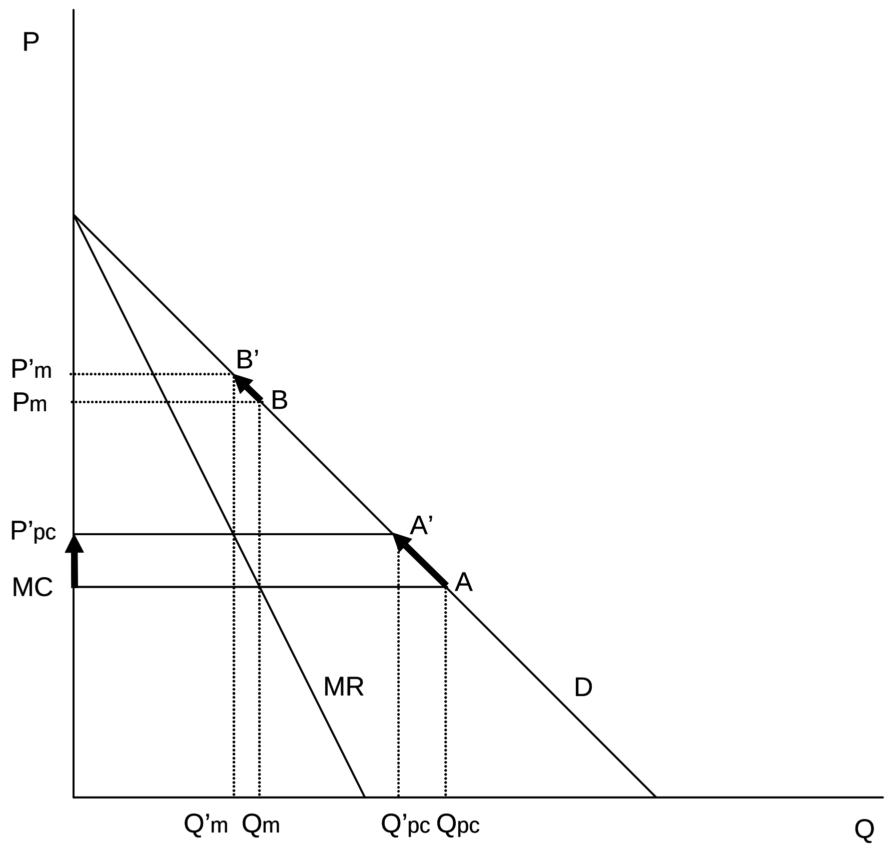

Fig. 2.8. Competitive and monopoly response to an oil technology shock.

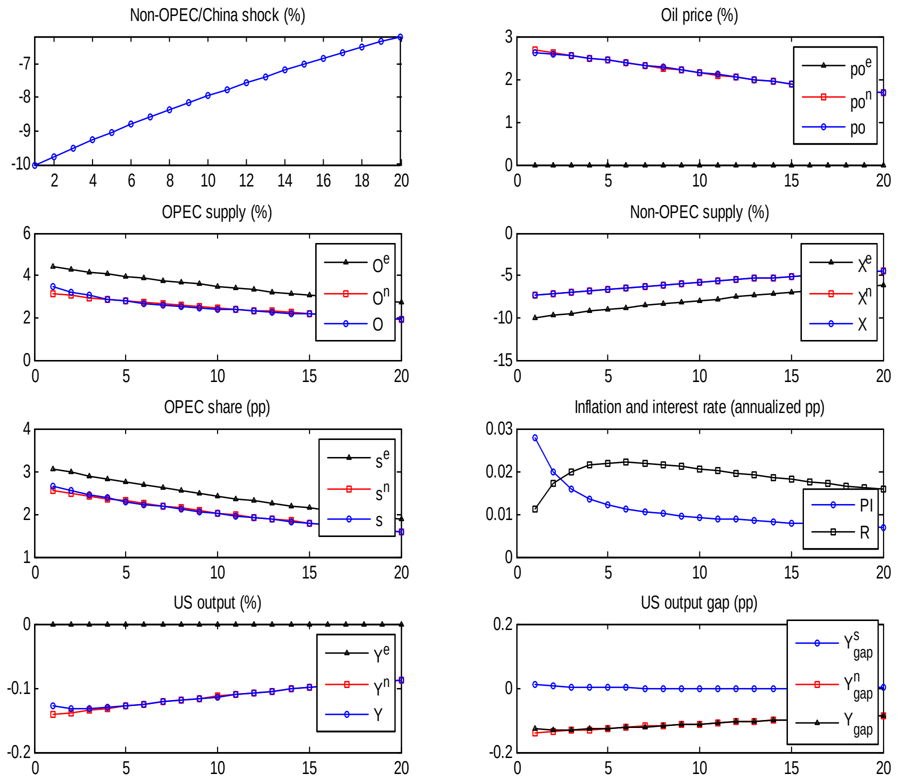

Fig. 2.9. Responses to a negative fringe capacity shock

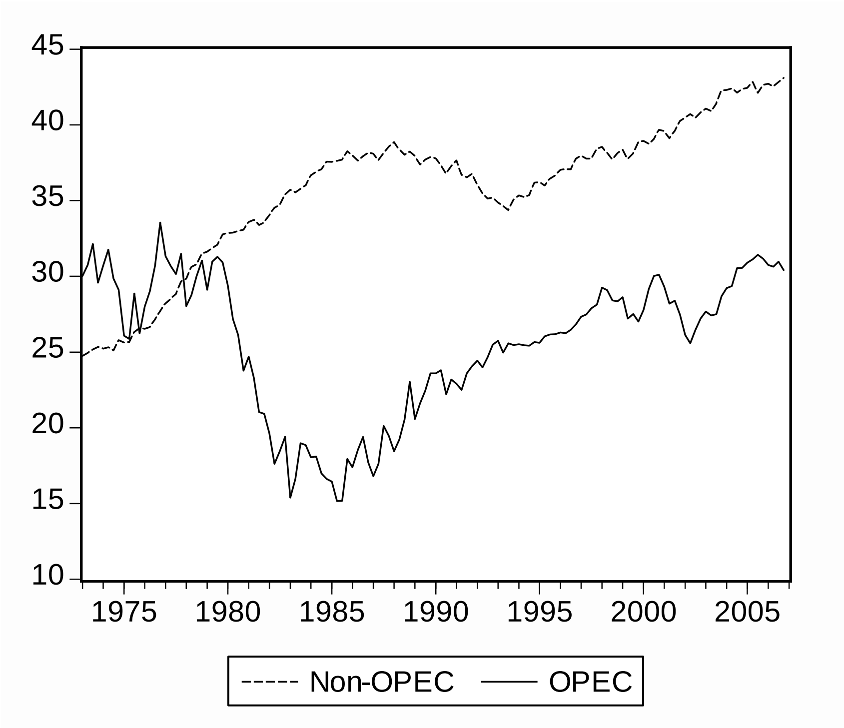

Fig. 2.10. OPEC and non-OPEC oil supply (million barrels per day)

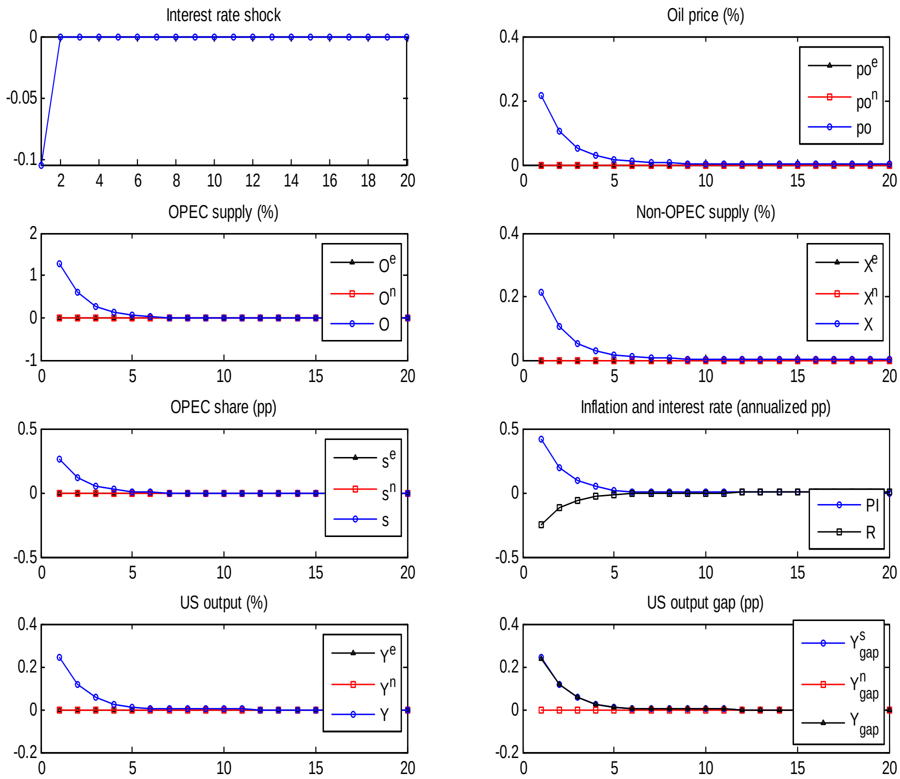

Fig. 2.11. Responses to a negative interest rate shock

Chapter 3
Oil and the Great Moderation

3.1 Introduction

Hamilton (1983) noticed that all but one US recessions since World War II were preceded by increases in the price of crude oil, suggesting that exogenous oil shocks were responsible for much of the post-war volatility of US GDP growth (see Figure 3.1). Other authors found similar evidence of a link between oil price rises and US inflation, as well as a link between oil price fluctuations and both output growth and inflation in other industrialized countries (e.g. Darby (1982), Burbidge and Harrison (1984)).

The relevance of oil as a source of macroeconomic fluctuations was established as conventional wisdom at least until Hooker (1999) pointed to a break in the oil price–GDP relationship and Hooker (2002) found a parallel break in the oil price–inflation relationship around 1981.[1] This break date roughly coincides with (but precedes) the beginning of a period of remarkable macroeconomic stability, dubbed by some economists as the “Great Moderation”, and reflected in a dramatic decline in the volatility (and persistence) of key macro variables in a number of industrialized economies, including the US (see Table 3.1 and figures 3.2 and 3.3).[2]

[1] Specifically, Hooker (1999) found that two widely used transformations of the oil price do not Granger cause output in the post-1980 period, while Hooker (2002) identified a structural break in core US inflation Phillips curves such that oil prices contributed substantially to core inflation before 1981 but since that time the pass-through has been negligible.

[2] The "Great Moderation" was noticed by Kim and Nelson (1999) and McConnel and Perez-Quiros (2000) and its beginning is usually dated around 1984.

Since evidence suggests that the moderation is spread across a number of countries[3], and oil supply shocks are likely to affect many oil-importing countries in a similar way, a reduction in oil sector volatility or a dampening of the transmission (or "pass-through") of that volatility to the rest of the world economy is a natural candidate (perhaps working alongside other factors) for explaining the rise of macroeconomic stability in the advanced world. One possibility is that oil supply shocks have simply become smaller in the period after 1984; at the same time, diversification towards less oil-intensive sectors and increased energy efficiency (reducing the share of oil in GDP) may have diminished the importance of oil supply shocks.

We asses the extent to which the macroeconomic moderation in the US can be accounted for by changes in oil shocks and the oil share, by performing counterfactual simulations based on Bayesian estimation of the model developed in Chapter 2 for the periods before and after 1984. In doing so, we nest two popular explanations for the Great Moderation: (1) "good luck" in the form of a favorable change in the distribution of TFP and other (non-oil) real shocks, as claimed for example by Ahmed, Levin and Wilson (2002) and Stock and Watson (2002); and (2) an improvement in the conduct of monetary policy, as argued by Clarida, Gali and Gertler (2000) and Boivin and Giannoni (2003).[4]

We find that oil played a non-trivial role in the moderation. In particular, the reduction of the oil share alone can explain around one third of the inflation moderation, and 13% of the GDP growth moderation. In turn, oil sector shocks alone can account

---

3 Cecchetti et. al. (2006) find evidence of moderation in 16 out of 25 industrialized countries and Stock and Watson (2003) report similar evidence for 6 of the G-7 countries; on the other hand, see Canova et. al. (2007) for evidence that the moderation has been more of an Anglo-Saxon phenomenon.

4 We do not control for other possible explanations involving structural changes in private sector behavior, such as better inventory management (McConnell and Perez-Quiros, 2000), or financial innovation (Dynan, Elmendorf and Sichel, 2005).

for 7% of the growth moderation and 11% of the inflation moderation. Still, consistent with other studies, we find that the dominant role in the reduction of macroeconomic volatility was played by non-oil real shocks and by monetary policy. In particular, smaller TFP shocks account for around two-thirds of GDP growth moderation, while better monetary policy alone can explain about two-thirds of the inflation moderation.

Related to the above, we find evidence that the inflation-output gap tradeoff has become more benign after 1984 due to the smaller share of oil in GDP. As a result, oil sector shocks have become less important for US macroeconomic fluctuations relative to US-originating shocks to TFP, preferences and policy.

The rest of the paper is organized as follows. The next section puts our work in the context of the related literature; section 3.3 presents some stylized volatility facts; section 3.4 sketches a log-linearized version of the oil pricing model developed in Chapter 2 and illustrates how different factors could cause moderation; section 3.5 covers the data and estimation methodology; section 3.6 describes our priors and the estimation results; section 3.7 contains counterfactual analysis decomposing the volatility moderation into contributions by each factor, and discusses the implied changes in the Phillips curve; section 3.8 relates our results to those of the literature and the last section concludes.

3.2 Related Literature

Our paper is related to several distinct lines of research. One is the empirical literature on the link between oil and the macroeconomy starting with Darby (1982) and Hamilton (1983). Bernanke, Gertler and Watson (1997) challenged the finding of Hamilton

(1983), documenting that essentially all U.S. recessions in the postwar period were preceded by both oil price increases as well as a tightening of monetary policy. Using a modified VAR methodology they found that the systematic monetary policy response to inflation (presumably caused by oil price increases) accounted for the bulk of the depressing effects of oil price shocks on the real economy. What is more, Barsky and Killian (2001) and Killian (2005) argued that even the major oil price increases in the 1970s were not an essential part of the mechanism that generated stagflation, and that the latter is attributable instead to monetary factors. Unlike these studies, our analysis is based on a structural model featuring optimal oil price setting, estimated with Bayesian methods. This allows us to disentangle the contribution of policy from the effects of oil shocks and the oil share without running into the Lucas critique.

Another strand of research deals with theoretical models of the link between oil and the macroeconomy. Some of the more recent contributions include Kim and Loungani (1992), Rotemberg and Woodford (1996), Finn (1995, 2000), Leduc and Sill (2004), and Carlstrom and Fuerst (2005). While these studies differ in the way oil is employed in the economy (as a consumption good, as a standard productive input, or as a factor linked to capital utilization), and in the implications of oil shocks, they all share the assumption that the oil price (or oil supply) is exogenous, and hence unrelated to any economic fundamentals. This is not only unappealing from a theoretical point of view as pointed out by Killian (2007), and inconsistent with the evidence presented in Killian (2007), Mabro (1998), and Hamilton (1983).[5] The issue is that with an exogenous (or a perfectly competitive) oil sector, and absent any real rigidities (e.g.

[5] When testing the null hypothesis that the oil price is not Granger-caused collectively by US output, unemployment, inflation, wages, money and import prices, Hamilton (1983) obtained a rejection at the 6% significance level.

real wage rigidities as in Blanchard and Gali, 2007), there is no meaningful trade-off between inflation and output gap stabilization, implying that full price stability is optimal even in the face of oil sector shocks. The fact that inflation in the 1970s was highly volatile suggests that either policy was very far from optimal, or that indeed there was an important policy trade-off. Different from the existing contributions, our model features a dominant oil exporter that charges an endogenously varying (optimal) oil price markup, which enters the Phillips curve as a "cost-push" term and induces a trade-off between the output gap and inflation (see Chapter 2).

Finally, our paper is related to the literature on the Great Moderation, starting with Nelson (1999) and McConnell and Perez-Quiros (2000). With some simplification, most of the explanations for the rise in macroeconomic stability can be classified into three broad categories: (a) "good practices", that is, changes in private sector behavior unrelated to stabilization policy, for instance improved inventory management (McConnell and Perez-Quiros, 2000) or financial innovation (Dynan, Elmendorf and Sichel, 2005); (b) "good policy", notably better monetary policy as argued by Clarida, Gali and Gertler (2000), Boivin and Giannoni (2003), and Gali and Gambetti (2007); and (c) "good luck", meaning a favorable shift in the distribution of real shocks, as in Ahmed, Levin and Wilson (2002), Stock and Watson (2002), and Justiniano and Primiceri (2006). Explanations of "good luck" in particular often give smaller oil shocks as an example (e.g. Summers, 2005).[6]

---

6 Not all studies fit the above classification. For example, Canova et. al. (2007) claim that it is impossible to account for both the Great Inflation of the 1970s and the strong output growth of the 1990s with a single explanation. Using a different approach, Canova (2007) similarly finds that the fall in variances of output and inflation had different causes, and suggests that the quest for a single explanation is likely to be misplaced. See section 8 for more on this.

Our framework allows us to separate oil from non-oil factors, while nesting the "better policy" and "smaller non-oil shocks" explanations. In this respect, our work is most closely related to Leduc and Sill (2007) who assess the role played by monetary policy relative to TFP and oil shocks in the Great Moderation. The main advantage of our approach lies in modelling the oil sector from optimizing first principles rather than assuming an exogenous process for oil supply. Another difference is that we estimate most of the model's parameters separately for each sample with Bayesian techniques which allows us to fit better the volatility reduction facts compared to Leduc and Sill who calibrate their model. In addition, compared to their paper, we put special focus on the role played by the oil share and not only on oil shocks.

3.3 Volatility Reduction Facts

Table 3.1 shows the standard deviations of three quarterly US macro series: GDP growth, deflator inflation, and the federal funds rate, for two subsamples, pre- and post-1984. "The Great Moderation" refers to the pronounced decline in the volatility of these (and other) macro variables in the post-1984 sample. In particular, the volatility of GDP growth declined by about 55%, of inflation by 60%, and of the nominal interest rate by 30%. For comparison, the last row of the table shows the standard deviation of the quarterly percentage change in the real price of oil. While there is a reduction in its volatility by 20%, this is somewhat less pronounced than for the other three variables.

Clearly, the volatility reduction facts reported in Table 3.1 are not insensitive to the choice of break year. Different studies have estimated different break dates for the different variables, but usually they lie in the range around 1982 to 1986. Redoing the calculations with 1982:I as the break date, we obtain volatility reductions of 45%, 57%, 20%, and 25%, respectively. And doing the same with 1986:I, we obtained 54%, 62%, 36%, and 13%. While the differences are non-trivial, by and large all three calculations show a similar picture.

The aim of this paper is to evaluate empirically the contribution of oil sector volatility and transmission, and compare it with alternative explanations for the volatility reduction (better monetary policy and non-oil related "good luck"). While the Great Moderation is sometimes associated also with a reduction in the persistence of macro variables (e.g. Canova et. al. 2007), we do not attempt to explain this phenomenon or attribute it to the various factors.

	Standard deviation (x 100)		Volatility
	1965:I – 1983:IV	1984:I – 2006:IV	reduction
GDP growth	1.126	0.508	55%
Inflation	0.609	0.244	60%
Interest rate	0.847	0.583	31%
Real oil price	16.33	12.99	20%

Table 3.1. US volatility reduction since 1984

3.4 The Log-Linearized Model

We base our empirical analysis on the full version of the model developed in Chapter 2. Here we sketch a compact representation, expressed in terms of log-deviations from the efficient equilibrium. For equal treatment of the household sector with the other four types of agents (firms, monetary authority, OPEC and non-OPEC), we include a shock to the time discount factor as an additional source of aggregate demand disturbance. This improves the fit of our model in terms of matching the volatility and moderation facts reported in Table 3.1. In addition, we allow for monetary policy to react to the output gap besides inflation, which in our model is an appropriate objective for a central bank concerned with the welfare of the representative household.

3.4.1 Dynamic IS curve

Log-linearizing the consumer's Euler equation, replacing consumption with final goods value added (that is, GDP), and casting the resulting expression in deviation from the efficient allocation, we obtain,

$$\hat{y}_t = E_t \hat{y}_{t+1} - \left(\hat{\imath}_t - E_t \pi_{t+1} - \hat{rr}_t^e\right), \tag{3.1}$$

where $\hat{y}_t = y_t - y_t^e$ is the (log) distance between actual value added and its efficient level (we refer to it as the "output gap" for simplicity)

The IS curve thus relates the current output gap positively to its expected future level, and negatively to the distance between the ex-ante real interest rate $\hat{\imath}_t - E_t \pi_{t+1}$ and the efficient real interest rate $\hat{rr}_t^e$. The latter is defined as the expected growth rate of efficient GDP, and in equilibrium is given by the expression,

$$\hat{rr}_t^e = (1 - \rho_b)\, \hat{b}_t - \left(\frac{1 - \rho_a}{1 - s_o}\right) \hat{a}_t - \left(\frac{s_o\,(1 - \rho_z)}{1 - s_o}\right) \hat{z}_t, \tag{3.2}$$

which depends negatively on shocks to technology $\hat{a}_t$ (in final goods) and $\hat{z}_t$ (in the oil sector), and positively on the shock to the discount factor $\hat{b}_t$, where s_o[7] is the share of oil in GDP. The driving variables $\hat{a}_t$, $\hat{z}_t$, and $\hat{b}_t$ are assumed to follow independent stationary AR(1) processes,

$$\hat{a}_t = \rho_a \hat{a}_{t-1} + \epsilon_t^a, \tag{3.3}$$

$$\hat{z}_t = \rho_z \hat{z}_{t-1} + \epsilon_t^z, \tag{3.4}$$

$$\hat{b}_t = \rho_b \hat{b}_{t-1} + \epsilon_t^b, \tag{3.5}$$

[7] In the notation of Chapter 2, $s_o \equiv 1 - \alpha_1 - \alpha_2$.

where

$$\hat{a}_t \equiv \log(A_t), \tag{3.6}$$

$$\hat{z}_t \equiv \log(Z_t), \tag{3.7}$$

$$\hat{b}_t \equiv \log(\beta_t) - \log(\beta), \tag{3.8}$$

ρ_a, $\rho_{z,}$ and ρ_b are shock persistence parameters, and ϵ_t^a, ϵ_t^z, and ϵ_t^b are $i.i.d.$ innovations to US total factor productivity, oil technology, and the time discount factor, all mean zero and with standard deviations σ_a, σ_z, and σ_b respectively.

Notice that the observable GDP growth rate is given by

$$\Delta y_t = \Delta \hat{y}_t + \Delta y_t^e. \tag{3.9}$$

3.4.2 Phillips curve

Aggregating the optimal staggered price-setting decision of final goods firms, we obtain the following first-order approximation to the dynamics of inflation around the deterministic steady-state with zero inflation,

$$\pi_t = \beta E_t \pi_{t+1} + (1 - s_o)\lambda \hat{y}_t + s_o \lambda \hat{\nu}_t, \tag{3.10}$$

where π_t denotes inflation, $\hat{y}_t$ the output gap, $\hat{\nu}_t \equiv \hat{p}_{ot}^* - \hat{p}_{ot}^e$ is the optimal oil price markup (determined below), β is the (mean) time discount factor; and parameter λ is related to the structural parameters of the underlying model as follows,

$$\lambda = \frac{(1+\psi)(\mu - s_o)(1-\theta)(1-\beta\theta)}{[\mu\alpha_1 + (\mu - 1)(1+\psi)s_o]\theta}, \tag{3.11}$$

where ψ is the inverse of the Frisch labor supply elasticity, μ is the average markup in the final goods sector, $1 - \theta$ is the frequency of price adjustment, and α_1 is the labor share in final goods production.

Notice that the oil price markup enters the Phillips curve like a "cost-push" term. Namely, a rise in the oil price markup leads to a rise in inflation and/or a negative output gap, implying a trade-off between the two policy objectives. This is in contrast with the case of perfect competition in the oil sector (or exogenous oil price), in which oil price shifts are necessarily associated with an opposite movement in the efficient level of output and imply no tension between inflation and output gap stabilization (for more details see Chapter 2).

Iterating the Phillips curve forward, we obtain the expression,

$$\pi_t = \lambda \sum_{k=0}^{\infty} \beta^k E_t \left[(1 - s_o)\hat{y}_{t+k} + s_o \hat{\nu}_{t+k} \right] \tag{3.12}$$

which shows that inflation is a weighted average of current and expected future output gaps and oil price markups.

3.4.3 Monetary policy

The central bank follows a Taylor-type rule of the form,

$$\hat{\imath}_t = \phi_i \hat{\imath}_{t-1} + (1 - \phi_i) \left(\phi_\pi \pi_t + \phi_y \hat{y}_t \right) + \hat{r}_t, \tag{3.13}$$

where π_t is inflation, $\hat{y}_t$ is the output gap, $\hat{r}_t$ is a zero mean $i.i.d.$ monetary policy shock, and ϕ_i, ϕ_π and ϕ_y are policy reaction coefficients.

3.4.4 Oil sector

In Chapter 2 we model OPEC as a dominant firm which seeks to maximize the welfare of its owner, internalizing the effect of its pricing decision on global output and oil demand. Operating alongside a competitive fringe of price-taking oil suppliers, the dominant oil exporter sells its output to an oil importing country (the US), which uses it to produce final goods.

A first-order approximation of the optimal oil price setting rule of the dominant oil supplier takes the form,

$$\hat{\nu}_t = \boldsymbol{\gamma}' \begin{bmatrix} \hat{y}_{t-1} \\ \hat{\imath}_{t-1} \\ \hat{\boldsymbol{\lambda}}_{t-1} \\ \hat{\boldsymbol{\xi}}_t \end{bmatrix}, \tag{3.14}$$

in which past period's value added, $\hat{y}_{t-1}$, and nominal interest rate, $\hat{\imath}_{t-1}$, are state variables; $\hat{\boldsymbol{\lambda}}_{t-1}$ is a vector of co-state variables (Lagrange multipliers) capturing the effect of commitment; $\hat{\boldsymbol{\xi}}_t = \left[\hat{a}_t, \hat{b}_t, \hat{r}_t, \hat{z}_t, \hat{\omega}_t\right]'$ is a vector of exogenous states; and $\boldsymbol{\gamma}$ is a vector of non-linear functions of the structural parameters of the model.

Competitive fringe producers seek to maximize profits while taking the oil price as given. In equilibrium, competitive fringe output $\hat{x}_t$ is an increasing function of the oil price $\hat{p}^*_{ot}$, oil technology $\hat{z}_t$, and the shock to fringe capacity $\hat{\omega}_t$,

$$\hat{x}_t = \hat{p}^*_{ot} + \hat{z}_t + \hat{\omega}_t. \tag{3.15}$$

The total capacity of the competitive fringe is assumed to follow a stationary AR(1) process with persistence ρ_ω,

$$\hat{\omega}_t = \rho_\omega \hat{\omega}_t + \epsilon^\omega_t, \tag{3.16}$$

where $\hat{\omega}_t \equiv \log\left(\Omega_t/\bar{\Omega}\right)$ and ϵ^ω_t is $i.i.d.$ with mean zero and standard deviation σ_ω.

3.4.5 What factors could cause moderation?

We illustrate how different factors may contribute to the moderation based on the above model.

The first explanation is that the distribution of real disturbances hitting the economy has changed so that real shocks have become "smaller" on average. Notice that smaller real shocks would reduce the volatility of $\hat{rr}^e_t$, while smaller oil sector shocks in particular are likely to diminish the variance of the oil price markup, $\hat{\nu}_t$. Since these are the two driving variables in our model, for any given interest rate rule and oil share, the volatility of output, inflation and the interest rate would be reduced.

An alternative (or complementary) explanation has to do with better monetary policy. This includes smaller monetary surprises ($\hat{r}_t$ shocks), as well as a more stabilizing policy rule. Smaller monetary shocks reduce the volatility of the interest rate, which is transmitted through the IS and Phillips curves to actual output and inflation. At the same time stronger systematic reaction of the policy instrument to inflation and output deviations from target results in better stabilization of these variables over the cycle.[8]

Finally, part of the moderation may be due to the fact that oil – perhaps once an important source of volatility – now accounts for a smaller fraction of GDP compared to the past. The latter can be due to increased energy efficiency and diversification away from oil-intensive sectors.

The oil share affects the volatility of $\hat{rr}^e_t$ and hence of the output gap, and at the same time it affects the pass-through coefficient on the oil price markup in the Phillips curve. Other things equal, a smaller oil share is likely to reduce the volatility of output and the pass-through from the oil price to inflation.

To see how the oil share affects the inflation–output gap tradeoff, notice that a policy of strict price stability ($\pi_t = 0$) implies

$$\hat{y}_t = -\frac{s_o}{1 - s_o}\hat{\nu}_t, \tag{3.17}$$

[8] Strictly speaking, stronger reaction to the output gap would result in better alignment of output with its efficient level, which need not imply a smaller volatility of the *growth rate* of output, especially if real shocks are large.

while a policy aimed at strict output gap stability ($\hat{y}_t = 0$) implies,

$$\pi_t = s_o \lambda \sum_{k=0}^{\infty} \beta^k E_t \hat{\nu}_{t+k}. \tag{3.18}$$

In both cases the extent to which stabilizing one variable induces inefficient fluctuations in the other is a function of the oil share. Finally, the oil share affects the elasticity of demand for OPEC's oil and thus the volatility of the oil price markup.

3.5 Data and Methodology

We asses the extent to which the macroeconomic moderation in the US can be accounted for by changes in oil shocks and the oil share by performing counterfactual simulations based on Bayesian estimation of the model of Chapter 2 for the periods pre- and post-1984. Our estimation methodology is in the spirit of Smets and Wouters (2003), Gali and Rabanal (2005), and An and Shorfheide (2005). The observable variables (the moderation of which we want to explain) are US GDP growth, inflation, the nominal interest rate, and the percentage change of the real price of oil. Quarterly data on real GDP, the GDP deflator, the Federal Funds rate and the West Texas Intermediate oil price from 1965:I to 2006:IV are taken from FRED II.[9] GDP growth and inflation are computed as quarterly percentage changes of real GDP and the GDP deflator[10]; the nominal interest rate is converted to quarterly frequency to render it consistent with the model; and the oil price is detrended by the GDP deflator and cast in quarterly percentage changes. The resulting series are demeaned prior to estimation.

Since our model is meant to describe the behavior of OPEC, we start the sample in 1965 which marks the year in which OPEC based their Secretariat in Vienna. Be-

[9] The original series names are GDPC96, GDPDEF, FEDFUNDS and OILPRICE.

[10] Our model makes no difference between GDP deflator and CPI inflation.

fore that the international oil industry was dominated by seven major oil companies of Anglo-Saxon origin, known as the "Seven Sisters". Of these five belonged to the US (Esso, Mobil, Chevron, Texaco and Gulf), one to the UK (BP), and one was Anglo-Dutch (Shell). Even though OPEC was created in 1960, in the first few years of its existence its activities were of a low-profile nature, as it set out its objectives, established a secretariat, and engaged in negotiations with the oil companies.[11] Thus, throughout the period 1959-1964 the nominal oil price remained unchanged at just below 3$ a barrel.

The sample is split in 1984:I. This corresponds to the estimated break in US output volatility by McConnell and Perez-Quiros (2000), Cecchetti et al (2006) and others. A break in inflation volatility was found around that date as well (Kahn, McConnell, and Perez-Quiros, 2002); a break in the oil – GDP link (Hooker, 1999) and the oil – inflation relationship (Hooker, 2002) was identified around 1981; and a break in the conduct of monetary policy around 1979–1982 (Gali and Gertler, 2000).

We fix several parameters of the model based on historical averages over the full sample (as in the case of the time discount factor), or on values which are standard in the literature (as with the elasticity of substitution among final goods). These calibrated parameter values are given in Appendix 3.A.

The elasticity of oil in production is calibrated separately for each sub-sample based on the average nominal expenditure on oil as a share of nominal GDP, that is,

$$s_o = \sum_t \frac{(\text{barrels of oil consumed in the US})_t \text{ x } (\text{\$ per barrel})_t}{(\text{nominal GDP})_t}, \tag{3.19}$$

where t runs from 1965 to 1983 in the first sample and from 1984 to 2006 in the second. This yields a value of 0.036 for the first period and 0.022 for the second, which we fix

11 Source: OPEC

prior to estimation.[12] The reason we choose to calibrate the oil share in this way rather than letting the estimation procedure tell us about its distribution is that we do not expect the variables we use in the estimation to be informative about this parameter. Instead, we use a formula for the oil share which is consistent with our model, and for which we have accurate data.

The above procedure leaves us with fourteen parameters to estimate: the frequency of price adjustment (θ), the Frisch labor supply elasticity (ψ), the parameters of the monetary policy rule (ϕ_i, ϕ_π, ϕ_y), the shocks' autoregressive parameters (ρ_a, ρ_b, ρ_z, ρ_ω) and standard deviations of the innovations (σ_a, σ_b, σ_z, σ_ω, σ_r).

We approximate our model to first-order and solve it with a standard method for linear rational expectations models (e.g. Sims 2002, and Klein, 2000). Given the state-space representation, we use the Kalman filter to evaluate the likelihood of the four observable variables. From Bayes' rule the posterior density function is proportional to the product of the likelihood and the prior density of the parameters. We use a random walk Metropolis-Hastings algorithm to obtain 5 chains of 50000 draws from the posterior distribution. We choose a scale for the jumping distribution in the MH algorithm which yields an acceptance rate of around 30%. The posterior distributions are obtained by discarding the first half of the draws from each chain.

Once we obtain the estimates for each sample period, we perform counterfactual simulations isolating the effect of a change in a single factor (e.g. the oil share) on the volatility moderation.[13]

---

[12] We do this by setting the share of labor to 0.634 in the first sample, and to 0.648 in the second, while keeping the share of capital fixed at 0.33.

[13] We do not model any transition dynamics; Canova and Gambetti (2007) propose an alternative method of performing counterfactual simulations based on re-estimating all the model's parameters conditional on the chosen counterfactual value for any given parameter. We stick to the more standard approach, followed e.g. by Stock and Watson (2003), treating our parameters as behavioral and thus independent

3.6 Priors and Estimation Results

3.6.1 Choice of priors

The first four columns of tables 3.2a and 3.2b show the assumed prior densities for the parameters whose posterior distributions we want to characterize. For each parameter we use the same prior density in both samples, except for the parameter on inflation in the monetary policy rule. For this parameter we assume a normal (1.5, 0.5) distribution in the second sample, but a gamma prior with mean 1.1 and a standard deviation of 0.5 in the first sample. Following Lubik and Schorfheide (2004) and Justiniano and Primiceri (2007), this assigns roughly equal probability on the inflation coefficient being either less or greater than one, while restricting it to be positive.[14]

We should stress that the conditions for local determinacy of equilibria in our model are not the standard ones. In particular, $\phi_\pi > 1$ is not a necessary condition for local uniqueness, and indeed there is a large region of determinacy for values of ϕ_π sufficiently below 1 (see Figure 3.4). The reason is that, different from the standard three equation New Keyenesian framework, in our model the Phillips curve includes an additional term – the oil price markup – which responds (optimally) to other endogenous variables, and in particular to the past output gap. This explains why we can solve and estimate our model for values of ϕ_π below 1.

For the other parameters of the monetary policy rule we use normal prior densities in both samples. For the price adjustment probability we assume a beta prior with mean 0.75 and standard deviation of 0.1. For the inverse Frisch labor supply elasticity we as-

---

of the experiment.

[14] The estimation results turn out to be very similar if instead we assume the same normal prior density for the coefficient on inflation in both samples.

sume a gamma prior with mean 1 and standard deviation of 0.5. The autocorrelation coefficients of the shocks are assumed to be beta with mean 0.9 and standard deviation of 0.05. And for the standard deviation of the innovations we assume an inverted gamma distribution (which ensures non-negativity) and use prior information from the calibrated model in Chapter 2 to specify the mean.

3.6.2 Estimation results

Comparing the two sets of estimated posterior modes in tables 3.2a and 3.2b we notice several important parameter shifts. First, the mode of the inflation coefficient of the monetary policy rule is larger in the second sample, implying that monetary policy was reacting more strongly to inflation compared to the first period. At the same time, the estimated standard deviation of the interest rate innovation in the pre-1984 sample is more than double that in the post-1984 sample, suggesting that policy was more erratic in the first period.

Secondly, the mode of the Calvo parameter governing the frequency of price adjustment is smaller in the post-1984 period suggesting that prices have become more flexible.

Third, there is evidence of changes in the volatility (and persistence) of real shocks. In particular, the volatility of the US technology innovation is cut by half in the post-1984 period, while preference shocks have become more persistent. Finally, oil sector shocks (especially oil technology shocks) have become smaller in the latter period.

Para-	Prior distribution				Posterior distribution		
meter	Density and domain		Mean	Std	Mean	Std	Mode
θ	Beta	$[0,1)$	0.75	0.100	0.649	0.076	0.627
ψ	Gamma	$\mathbb{R}^+$	1.00	0.500	1.097	0.397	0.901
ϕ_i	Normal	$\mathbb{R}$	0.60	0.100	0.557	0.075	0.543
ϕ_π	Gamma	$\mathbb{R}^+$	1.10	0.500	1.887	0.292	2.096
ϕ_y	Normal	$\mathbb{R}$	0.50	0.125	0.596	0.105	0.586
ρ_a	Beta	$[0,1)$	0.90	0.050	0.957	0.015	0.974
ρ_b	Beta	$[0,1)$	0.90	0.050	0.883	0.035	0.894
ρ_z	Beta	$[0,1)$	0.90	0.050	0.933	0.026	0.940
ρ_ω	Beta	$[0,1)$	0.90	0.050	0.931	0.024	0.947
$100\sigma_a$	Inv. Gamma	$\mathbb{R}^+$	0.70	∞	1.220	0.098	1.180
$100\sigma_b$	Inv. Gamma	$\mathbb{R}^+$	0.70	∞	2.170	0.480	1.900
$100\sigma_z$	Inv. Gamma	$\mathbb{R}^+$	10.0	∞	18.27	1.870	18.59
$100\sigma_\omega$	Inv. Gamma	$\mathbb{R}^+$	10.0	∞	31.74	5.300	28.64
$100\sigma_r$	Inv. Gamma	$\mathbb{R}^+$	0.10	∞	0.430	0.053	0.430

Table 3.2a. Prior and posterior distributions, 1965 –1983

Table 3.3 shows that the estimated model does quite a good job at matching the second moments and the post-1984 volatility reduction of the variables of interest. To be precise, the model slightly overestimates the volatility of GDP growth and inflation in both periods but matches quite well the post-1984 reduction in volatility of these variables. The moderation of the nominal interest rate is somewhat overestimated but the volatility and moderation of the oil price is matched pretty well.

Para-meter	Prior distribution				Posterior distribution		
	Density and domain		Mean	Std	Mean	Std	Mode
θ	Beta	$[0,1)$	0.75	0.100	0.529	0.067	0.543
ψ	Gamma	$\mathbb{R}^+$	1.00	0.500	1.328	0.405	1.248
ϕ_i	Normal	$\mathbb{R}$	0.60	0.100	0.684	0.057	0.704
ϕ_π	Normal	$\mathbb{R}$	1.50	0.500	3.115	0.290	3.012
ϕ_y	Normal	$\mathbb{R}$	0.50	0.125	0.549	0.101	0.572
ρ_a	Beta	$[0,1)$	0.90	0.050	0.978	0.009	0.984
ρ_b	Beta	$[0,1)$	0.90	0.050	0.950	0.015	0.950
ρ_z	Beta	$[0,1)$	0.90	0.050	0.870	0.040	0.867
ρ_ω	Beta	$[0,1)$	0.90	0.050	0.948	0.021	0.955
$100\sigma_a$	Inv. Gamma	$\mathbb{R}^+$	0.70	∞	0.595	0.044	0.590
$100\sigma_b$	Inv. Gamma	$\mathbb{R}^+$	0.70	∞	2.108	0.512	1.880
$100\sigma_z$	Inv. Gamma	$\mathbb{R}^+$	10.0	∞	13.70	1.708	14.45
$100\sigma_\omega$	Inv. Gamma	$\mathbb{R}^+$	10.0	∞	28.70	5.134	25.63
$100\sigma_r$	Inv. Gamma	$\mathbb{R}^+$	0.10	∞	0.226	0.033	0.200

Table 3.2b. Prior and posterior distributions, 1984 – 2006

	1965-1983		1984-2006		Volat. reduction	
	Data	Model	Data	Model	Data	Model
GDP growth	1.126	1.381	0.508	0.669	55%	52%
Inflation	0.609	0.658	0.244	0.279	60%	58%
Interest rate	0.847	0.860	0.583	0.455	31%	47%
Real oil price	16.32	15.71	12.99	12.92	20%	18%

Table 3.3. Second moments of model and data

3.7 Implications

3.7.1 What accounts for the Moderation?

In this section we attribute the volatility reduction implied by the model (the last column of Table 3.3) to counterfactual changes in each factor in isolation, including: (1) the oil share; (2) monetary policy; (3) real shocks, including oil sector shocks and US shocks; (4) a shift in the frequency of price adjustment or (5) in the Frisch labor supply elasticity; and (6) the residual due to the interaction of all factors (this is just to say that the contributions of factors 1 to 5 are not linearly additive and do not sum up to explaining 100% of the moderation).

Table 3.4a presents the percentage reduction in volatility which would be achieved by a change in any single factor keeping the rest of the parameters at their pre-1984 values. Thus, had the oil share in the period 1965-1983 been at its post-1984 value (that is, 0.022 instead of 0.036), GDP growth would have been 6.5% less volatile, the nominal interest rate 10.7% less volatile, and inflation 23% less volatile, while the oil price would have been 0.4% more volatile. Expressed in percent of the actual reduction in the volatility of these variables (reproduced in the last column), the change in the oil share alone could explain around one tenth of the GDP growth moderation, a quarter of the interest rate moderation, and one third of the inflation moderation. This points to the oil share decline having played a considerable role in the moderation, especially of inflation.

By the same token, better monetary policy alone could explain around two thirds of the inflation moderation, a half of the interest rate moderation, but only 5% of the GDP growth moderation. And smaller real shocks explain around three quarters of the GDP growth moderation, one quarter of inflation moderation, and one third of the interest rate moderation. Smaller oil shocks in particular account for 7% of GDP growth moderation, 11% of inflation moderation, and all of the oil price moderation.

Table 3.4b shows that around two thirds of the GDP growth moderation can be accounted for by smaller TFP shocks alone, while smaller time preference shocks account for around one tenth of the inflation moderation and a quarter of the interest rate moderation. Interestingly, smaller oil technology shocks were responsible for the bulk of the oil price moderation, with fringe (or cartelization) shocks playing a smaller role.

Our findings ascribe to monetary policy quite a modest role in GDP *growth* moderation. This could be for several reasons. One is the proximity of our simple model to the RBC paradigm: apart from nominal price rigidities (with a Calvo parameter estimated at 0.63 in the first period) and imperfect competition in oil, our model features no other imperfections or real rigidities (e.g. as in Blanchard and Gali, 2007) that would raise the importance of the interest rate channel. Second, we assume that the central bank reacts to the output *gap* (and not to output *growth*), which in our model is a relevant target variable for a central bank concerned with the welfare of the representative household. Given this rule, better monetary policy does not necessarily imply smaller output fluctuations, especially if real disturbances are large. Third, the estimated reaction to the output gap is not much different across the two samples (it is the reaction to inflation which increases substantially in the second period), so even if the fluctuations of efficient output were not large, the post-1984 rule would not have stabilized output much better than the pre-1984 one.

The bottom line of this analysis is that the reduced oil share and smaller oil shocks have played a non-trivial role in the volatility reduction even if the other two factors – smaller TFP shocks and better monetary policy – have played the dominant role in the moderation of GDP and inflation respectively.

	Oil share	Monet. policy	Real shocks Oil	Real shocks US	Real shocks All	Calvo param	Frisch elast.	All factors
GDP growth	6.49	2.29	3.40	34.9	40.3	1.07	0.37	52
Inflation	23.0	40.2	6.06	6.43	13.0	-6.89	-2.28	58
Interest rate	10.7	26.5	1.22	15.8	17.3	-1.98	-0.71	47
Real oil price	-0.42	0.01	17.9	0.02	18.0	0.05	0.01	18

Table 3.4a. Percent moderation by factor[15]

	Real shock $\hat{a}$	Real shock $\hat{b}$	Real shock $\hat{z}$	Real shock $\hat{\omega}$	All factors
GDP growth	34.1	0.54	3.13	0.27	52
Inflation	0.88	5.50	4.64	1.37	58
Interest rate	2.98	12.4	0.52	0.70	47
Real oil price	0.02	0.00	14.8	2.65	18

Table 3.4b. Percent moderation by shock

3.7.2 Changes in the Phillips curve

Hooker (2002) finds evidence of a break in standard (backward-looking) core US inflation Phillips curves regressions, with oil price changes making a substantial contribution to core inflation before 1981 but little or no pass-through since that time. Similarly, estimating the standard New Keynesian model via maximum likelihood, Ireland (2004) finds that "cost push" shocks have become smaller since the 1980s.

Our model estimated with Bayesian techniques is in broad agreement with these claims. Indeed, it points to the decrease in the oil share as a likely cause for the improvement in the Phillips curve tradeoff as inflation and the output gap have become more aligned with each other and less sensitive to oil price markup fluctuations. In particular,

15 The numbers indicate by how much the volatility of a (row) variable would have been reduced by a change in a single (column) factor. A negative sign means that the factor alone would have raised volatility post-1984

the last column of Table 3.5 shows that conditional on a reduction of the oil share from 3.6% to 2.2% (and keeping all other factors unchanged), the coefficient on the oil price markup (the "cost push" term) in the Phillips curve is reduced by about 40%. In addition, thanks mostly to smaller oil sector shocks, the volatility of the oil price markup itself has decreased by around 15% in the period after 1984. This, together with increased price flexibility (in the form of higher λ) since the mid-1980s and a stronger transmission channel, has made it possible for monetary policy to stabilize better both the output gap and inflation. As a result, the volatility of inflation has declined by about 23%, while the volatility of the output gap has been cut by around 40%.

		1965-1983	1984-2006	Counterf s_o
Oil share	s_o	0.036	0.022	0.022
Common slope coeff.	λ	0.643	1.337	0.641
Oil price markup coeff.	$s_o\lambda$	0.023	0.029	0.014
Output gap coefficient	$(1-s_o)\lambda$	0.620	1.312	0.627
Oil price markup volatility	std$(\hat{\nu})$	29.35	25.08	28.77
Oil price markup persist.	$\rho_{\hat{\nu}}$	0.946	0.947	0.946
Output gap volatility	std$(\hat{y})$	1.072	0.562	0.678
Output gap persistence	$\rho_{\hat{y}}$	0.902	0.925	0.817

Table 3.5. Changes in the Phillips Curve

3.7.3 Changes in the relative importance of shocks

Tables 3.6a and 3.6b show the asymptotic variance decomposition[16] of the four variables of interest in the first and the second sample.

Notably, the last two columns of each table reveal that the contribution of oil sector shock to US GDP growth and inflation variability has been considerable, both

[16] This is obtained by solving the Lyapunov equation $\Sigma_y = A\Sigma_y A' + B\Sigma_u B'$ in Σ_y, the unconditional variance of y, where y_t is the solution to the linear rational expectations model of the form $y_t = Ay_{t-1} + Bu_t$. It is thus simply the decomposition of the unconditional variance of endogenous variables, given that shocks occur in every period from now to infinity.

before and after 1984. In particular, oil shocks ($\hat{z}$ and $\hat{\omega}$) contributed to around 17% of GDP growth volatility, and as much as 60% of inflation volatility and 32% of interest rate volatility in the period 1965–1983. Thereafter, oil shocks continued to account for around 17% of growth volatility, but were responsible for "only" 33% percent of inflation volatility and 14% of interest rate volatility. Interestingly, the shock to oil productivity turns out to be more important for GDP growth and oil price volatility, while the fringe shock is more relevant for the volatility of inflation and the interest rate.

	US shocks			Oil shocks	
	Real		Nom.		
	$\hat{a}$	$\hat{b}$	$\hat{r}$	$\hat{z}$	$\hat{\omega}$
GDP growth	75.6	1.38	5.76	14.4	2.84
Inflation	1.95	24.4	10.4	9.86	53.4
Interest rate	6.77	57.5	3.77	1.34	30.6
Real oil price	0.04	0.01	0.04	74.0	25.9

Table 3.6a. Variance decomposition, 1965 – 1983

	US shocks			Oil shocks	
	Real		Nom.		
	$\hat{a}$	$\hat{b}$	$\hat{r}$	$\hat{z}$	$\hat{\omega}$
GDP growth	79.0	0.76	3.53	12.3	4.39
Inflation	0.60	33.1	32.8	0.55	32.9
Interest rate	2.59	82.1	1.41	0.76	13.2
Real oil price	0.02	0.00	0.01	70.0	30.0

Table 3.6b. Variance decomposition, 1984 – 2006

Turning to US-originating disturbances, the shock to TFP ($\hat{a}$) which accounted for the bulk of GDP growth volatility pre-1984 has become even more important for GDP growth, and has decreased its impact on inflation, the interest rate and the oil price. The preference shock ($\hat{b}$) has become even more important for inflation and the interest rate, and less important for GDP growth; and the interest rate shock ($\hat{r}$) has increased its relative importance for inflation, but has become less relevant for GDP, the interest rate and the oil price.

3.8 Comparison of the Results with the Literature

Based on a calibrated model with exogenous oil supply, Leduc and Sill (2007) conclude that improved monetary policy can account for 45% of the decline in inflation volatility but only 5% to 10% of the reduction in output volatility, the bulk of which can be explained by smaller TFP shocks. In this regard our findings are similar to theirs: we find that better policy can explain around two thirds of the inflation moderation but only around 5% of the GDP growth moderation. However, our results are distinct when it comes to attribution of the Great moderation to oil shocks. While we find that smaller oil sector shocks would have contributed to 7% of GDP growth moderation and 11% of the inflation moderation, Leduc and Sill claim that oil quantity shocks would have *raised* volatility in the first period. This discrepancy is likely due to the fact that Leduc and Sill treat oil supply as exogenous[17], in contrast to our explicit modelling of the oil sector with a dominant player. In addition, we find that the reduced oil share can explain 13% of the GDP growth moderation, and one third of the inflation moderation, a question which is not addressed by Leduc and Sill.

Gali and Blanchard (2007) introduce real wage rigidities to generate an inflation-output gap trade-off. They demonstrate how a reduction in the oil share in consumption and production shifts inward the policy frontier and goes some way towards explaining the observed reduction in inflation and output volatility. Our model in comparison

[17] In fact they treat it as constant, apart from a few discrete jumps lasting for one quarter, as identified by Hamilton (2000).

generates a policy tradeoff by assuming imperfect competition in the oil market while ignoring real wage rigidities. We also attempt to quantify more precisely the contribution of each factor by estimating the model with Bayesian techniques and performing counterfactual simulations.

Canova (2007) investigates the causes of output and inflation moderation in the US by estimating the benchmark small scale New Keynesian model with Bayesian techniques over rolling samples. He finds instability in the posterior of the parameters describing private sector behavior, the coefficients of the policy rule, and the covariance structure of shocks. Canova further shows that even though changes in the parameters of the private sector are largest, they cannot account by themselves for the full decline in volatility of output and inflation, while changes in the parameters of the policy rule and the covariance of the shocks can. Our findings are similar to Canova in that the bulk of GDP growth moderation is attributed to changes in real shocks, while most of the inflation moderation is due to monetary policy. Yet we find that as much as a third of the inflation moderation and 10%-13% of the moderation of output is attributable to the smaller share of oil in GDP, which is not directly measurable in the benchmark New Keynesian model estimated by Canova.

Gali and Gambetti (2007) look for the sources of the Great Moderation using a VAR with time-varying coefficients and stochastic volatility. Their findings point to structural change, as opposed to just good luck, as an explanation. In particular, they show that a significant fraction of the observed changes in comovements and impulse-responses can be accounted for by stronger reaction of monetary policy to inflation, and an apparent end of short-run increasing returns to labor. On the other hand Canova and Gambetti (2007) using a VAR with time-varying coefficients identified through sign restrictions, find no evidence that there was an increase in the response of the interest rate to inflation, and overall conclude that monetary policy was only marginally responsible for the Great Moderation. Compared with these studies, our counterfactual analysis based on a structural DSGE model estimated with Bayesian techniques assigns an important role to monetary policy, especially in the moderation of inflation (and the nominal interest rate). At the same time we point to the non-trivial role played by oil, especially in the moderation of inflation and in the improvement of the inflation-output gap tradeoff.

3.9 Conclusions

We asses the extent to which the increased macroeconomic stability after 1984 can be accounted for by changes in oil shocks and the oil share by taking the model developed in Chapter 2 to the data with Bayesian techniques and performing counterfactual simulations. In doing so we nest two popular explanations for the Great Moderation, namely smaller non-oil shocks, and better monetary policy.

Our estimates indicate that oil played a non-trivial role in the volatility reduction. In particular, the reduced oil share alone can explain around one third of the inflation moderation, and 13% of the GDP growth moderation. Smaller oil sector shocks account for 7% of the growth moderation and 11% of the inflation moderation. Nevertheless, we find that smaller TFP shocks can explain two-thirds of the growth moderation, while better monetary policy alone can explain two-thirds of the inflation moderation.

3.A Calibrated parameters

Quarterly discount factor	β	0.9926	Aver. annual real rate 3%
Steady-state markup	μ	1.15	Aver. markup 15%
Mean of non-OPEC capacity	$\bar{\Omega}$	0.004925	OPEC market share 40%
Inflation target	$\bar{\Pi}$	1	Optimal long-run inflation
Capital share	α_2	0.33	Aver. capital income share
Oil share, 1965-1983	s_o	0.036	Aver. oil income share
Oil share, 1984-2006	s_o	0.022	

Table 3.1. Calibrated parameters

3.B Figures

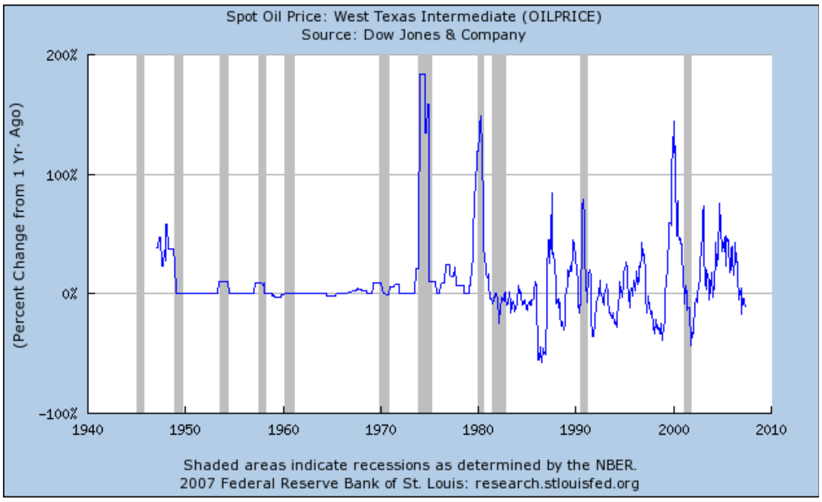

Fig. 3.1. Oil price and US recessions

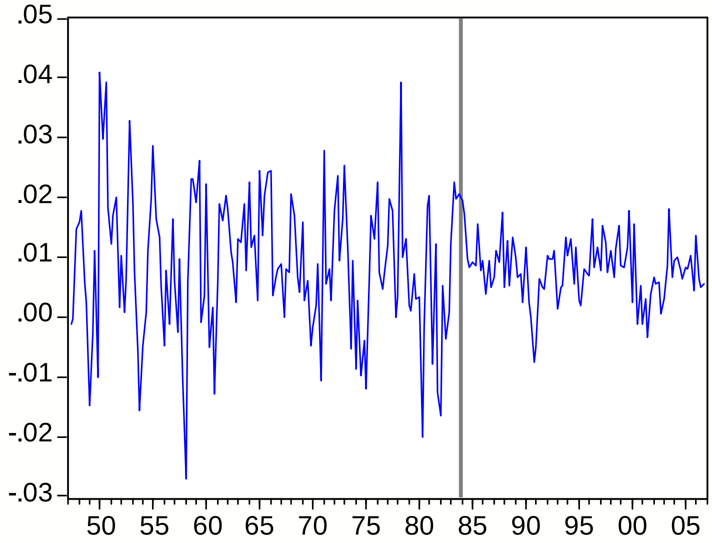

Fig. 3.2. US GDP growth moderation

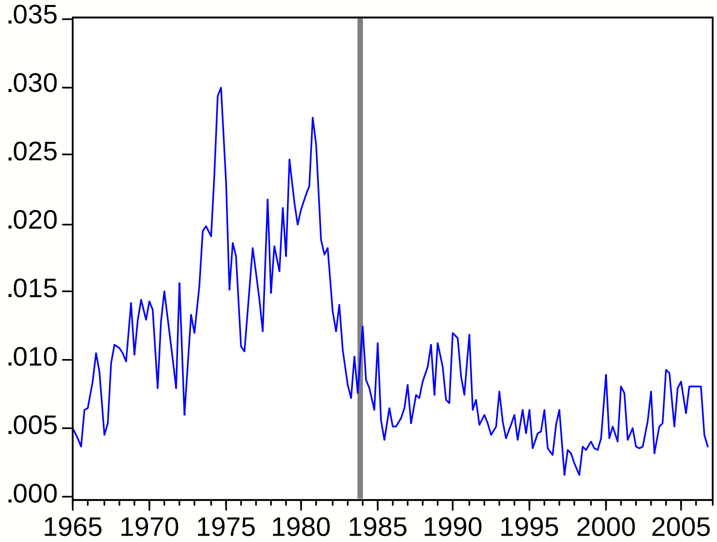

Fig. 3.3. US inflation moderation

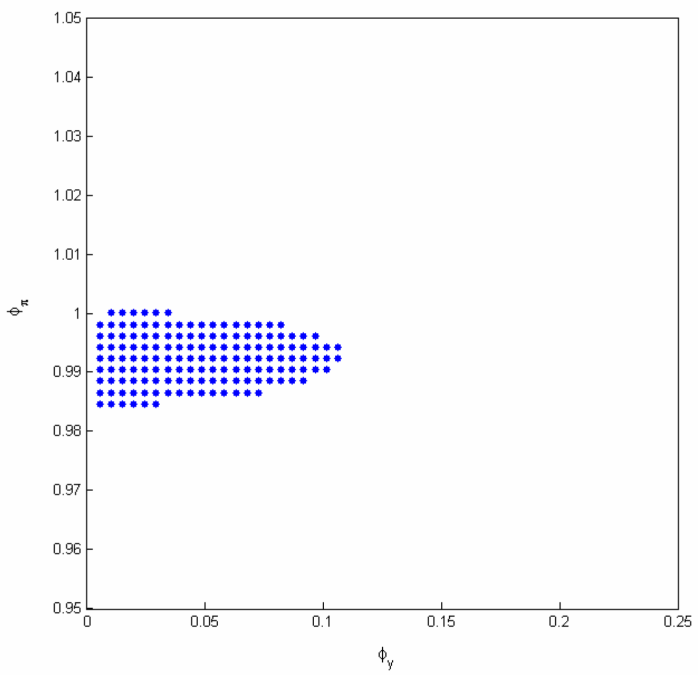

Fig. 3.4. Instability (dark) and determinacy (white) regions of the model

References

Adam, K. and Billi, R.: 2004a, Optimal monetary policy under commitment with a zero bound on nominal interest rates, *ECB Working Paper Series 377*, European Central Bank.

Adam, K. and Billi, R.: 2004b, Optimal monetary policy under discretion with a zero bound on nominal interest rates, *ECB Working Paper Series 380*, European Central Bank.

Adam, K. and Billi, R.: 2006, Optimal monetary policy under commitment with a zero bound on nominal interest rates, *Journal of Money, Credit and Banking* **38**, 1877–1905.

Adam, K. and Billi, R.: 2007, Optimal monetary policy under discretion with a zero bound on nominal interest rates, *Journal of Monetary Economics, forthcoming* .

Adelman, M. A. and Shahi, M.: 1989, Oil development-operating cost estimates, 1955-1985, *Energy Economics* **11**(1), 2–10.

Ahmed, S., Levin, A. and Wilson, B.: 2004, Recent US Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?, *Review of Economics and Statistics* **86**(3), 824–832.

An, S. and Schorfheide, F.: 2005, *Bayesian Analysis of DSGE Models*, Centre for Economic Policy Research.

Barsky, R. B. and Kilian, L.: 2001, Do we really know that oil caused the great stagflation? A monetary alternative, *NBER Working Papers 8389*, National Bureau of Economic Research, Inc.

Barsky, R. B. and Kilian, L.: 2004, Oil and the macroeconomy since the 1970s, *Journal of Economic Perspectives* **18**(4), 115–134.

Bernanke, Ben S., Gertler, Mark and Watson, Mark: 1997, Systematic monetary policy and the effects of oil price shocks, *Brookings Papers on Economic Activity* **1997**(1), 91–157.

Blanchard, O. and Galí, J.: 2007, The macroeconomic effects of oil price shocks: Why are the 2000s so different from the 1970s?, *Working Paper No.13368*, National Bureau of Economic Research, Inc.

Blanchard, O. and Khan, C.: 1980, The Solution of Linear Difference Models under Rational Expectations, *Econometrica* **48**, 1305–1311.

Blanchard, Olivier and Galí, Jordi: 2006, Real wage rigidities and the New Keynesian model, *Journal of Money, Credit, and Banking (forthcoming)* .

Boivin, J. and Giannoni, M.: 2006, Has Monetary Policy Become More Effective?, *The Review of Economics and Statistics* **88**(3), 445–462.

Burbidge, John and Harrison, Alan: 1984, Testing for the effects of oil-price rises using vector autoregressions, *International Economic Review* **25**(2), 459–484.

Calvo, G.: 1983, Staggered prices in a utility maximizing framework, *Journal of Monetary Economics* **12**(3), 383–398.

Canova, F.: 2007, Sources of structural changes in the US economy, *Mimeo*, Universitat Pompeu Fabra.

Canova, F. and Gambetti, L.: 2007, Structural changes in the US economy: is there a role for monetary policy, *Mimeo*, Universitat Pompeu Fabra and Universitat Autonoma de Barcelona.

Canova, F., Gambetti, L. and Pappa, E.: 2007, The structural dynamics of output growth and inflation: Some international evidence, *The Economic Journal* **117**, 167–191.

Carlstrom, C. T. and Fuerst, T. S.: 2005, Oil prices, monetary policy, and counterfactual experiments, *Working Paper 0510*, Federal Reserve Bank of Cleveland.

Cecchetti, S., Flores-Lagunes, A. and Krause, S.: 2006, Assessing the Sources of Changes in the Volatility of Real Growth, *NBER Working Papers 11946*, National Bureau of Economic Research, Inc.

Christiano, L., Eichenbaum, M. and Evans, C.: 2001, Nominal rigidities and the dynamic effects of a shock to monetary policy, *NBER Working Papers 8403*, National Bureau of Economic Research, Inc.

Christiano, L. and Fischer, J.: 2000, Algorithms for solving dynamic models with occasionally binding constraints, *Journal of Economic Dynamics and Control* **24**, 1179–1232.

Clarida, R., Gali, J. and Gertler, M.: 1999, The science of monetary policy: A New Keynesian perspective, *Journal of Economic Literature* **37**, 1661–1707.

Clarida, Richard, Gali, Jordi and Gertler, Mark: 2000, Monetary policy rules and macroeconomic stability: Evidence and some theory, *The Quarterly Journal of Economics* **115**(1), 147–180.

Coenen, G., Orphanides, A. and Wieland, V.: 2004, Price stability and monetary policy effectiveness when nominal rates are bounded at zero, *CEPR Discussion Paper Series 3892*, Centre for Economic Policy Research.

Cooley, T.: 1997, Calibrated models, *Oxford Review of Economic Policy* **13**(3), 55.

Dahl, C. and Yücel, M.: 1991, Testing Alternative Hypotheses of Oil Producer Behavior, *The Energy Journal* **12**(4), 117–138.

Darby, M. R.: 1982, The price of oil and world inflation and recession, *American Economic Review* **72**(4), 738–51.

Eggertsson, G. and Woodford, M.: 2003, The zero bound on interest rates and optimal monetary policy. Brookings Papers on Economic Activity.

EIA: 2007, Monthly energy review, *Technical report*, Energy Information Administration, US Department of Energy.

Finn, M.: 1995, Variance properties of Solow's productivity residual and their cyclical implications, *Journal of Economic Dynamics and Control* **19**(5), 1249–1281.

Finn, M.: 2000, Perfect Competition and the Effects of Energy Price Increases on Economic Activity, *Journal of Money, Credit and Banking* **32**(3), 400–416.

FRED II: 2007, Federal reserve economic data, *Database*, Federal Reserve Bank of St. Luis.

Galí, J.: 2003, *New Perspectives on Monetary Policy, Inflation, and the Business Cycle*, Vol. 3, Cambridge University Press, pp. 151–197.

Galí, J. and Gambetti, L.: 2007, On the sources of the great moderation, *Mimeo*, CREI and Universitat Pompeu Fabra.

Galí, J. and Rabanal, P.: 2005, Technology shocks and aggregate fluctuations: how well does the RBC model fit postwar US data?, *NBER Macroeconomics Annual* .

Gisser, Micha and Goodwin, Thomas H.: 1986, Crude oil and the macroeconomy: Tests of some popular notions: Note, *Journal of Money, Credit and Banking* **18**(1), 95–103.

Griffin, J. M.: 1985, OPEC behavior: A test of alternative hypotheses, *The American Economic Review* **75**(5), 954–963.

Hamilton, J.: 2000, What is an Oil Shock?, *Working Paper No.7755*, National Bureau of Economic Research, Inc.

Hamilton, J. D.: 1983, Oil and the macroeconomy since World War II, *The Journal of Political Economy* **91**(2), 228–248.

Hamilton, J. D.: 1996, This is what happened to the oil price-macroeconomy relationship, *The Journal of Monetary Economics* **38**(2), 215–220.

Hodrick, R. and Prescott, E.: 1981, Post-War US Cycles: An Empirical Investigation, *Journal of Money, Credit, and Banking* **19**, 1–16.

Hooker, M.: 1999, Oil and the Macroeconomy Revisited, *FEDS Working Paper No. 99-43*, The Federal Reserve Board.

Hooker, M. A.: 2002, Are oil shocks inflationary? Asymmetric and nonlinear specifications versus changes in regime, *Journal of Money, Credit and Banking* **34**(2), 540–561.

IFS: 2007, International financial statistics, *Database*, International Monetary Fund.

Ireland, P.: 2004, Technology Shocks in the New Keynesian Model, *Review of Economics and Statistics* **86**(4), 923–936.

Jones, C.: 1990, OPEC Behavior Under Falling Prices: Implications for Cartel Stability, *The Energy Journal* **11**(3), 117–129.

Jung, T., Teranishi, Y. and Watanabe, T.: 2005, Optimal monetary policy at the zero-interest-rate bound, *Journal of Money, Credit and Banking* **37**.

Justiniano, A. and Primiceri, G.: 2006, The Time Varying Volatility of Macroeconomic Fluctuations, *NBER Working Papers 9127*, National Bureau of Economic Research, Inc.

Kahn, J., McConnell, M. and Perez-Quiros, G.: 2002, On the Causes of the Increased Stability of the US Economy, *Federal Reserve Bank of New York Economic Policy Review* **8**(1), 183–202.

Kato, R. and Nishiyama, S.: 2005, Optimal monetary policy when interest rates are bounded at zero, *Journal of Economic Dynamics and Control* **29**.

Keynes, J.: 1936, *The General Theory of Employment, Interest and Money*, MacMillan, London.

Kilian, L.: 2005, The effects of exogenous oil supply shocks on output and inflation: Evidence from the G7 countries, *CEPR Discussion Papers 5404*, Centre for Economic Policy Research.

Kim, C. and Nelson, C.: 1999, Has the US Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle, *The Review of Economics and Statistics* **81**(4), 608–616.

Kim, In-Moo and Loungani, Prakash: 1992, The role of energy in real business cycle models, *The Journal of Monetary Economics* **29**(2), 173–189.

Klaeffling, M. and Lopez-Perez, V.: 2003, Inflation targets and the liquidity trap, *ECB Working Paper Series 272*, European Central Bank.

Klein, P.: 2000, Using the generalized schur form to solve a multivariate linear rational expectations model, *Journal of Economic Dynamics and Control* **24**(10), 1405–1423.

Krugman, P.: 1998, It's baaaack: Japan's slump and the return of the liquidity trap, *Brooking Papers on Economic Activity* **2**, 137–187.

Leduc, S. and Sill, K.: 2007, Monetary Policy, Oil Shocks, and TFP: Accounting for the Decline in US Volatility, *Working Paper*, Board of Governors of the Federal Reserve System.

Leduc, Sylvain and Sill, Keith: 2004, A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns, *The Journal of Monetary Economics* **51**(4), 781–808.

Lubik, T. and Schorfheide, F.: 2004, Testing for Indeterminacy: An Application to US Monetary Policy, *American Economic Review* **94**(1), 190–217.

Mabro, R.: 1998, OPEC behaviour 1960-1998: A review of the literature, *Journal of Energy Literature* **4**(1), 3–27.

Mankiw, N. G.: 2006, *Macroeconomics*, Worth Publishers New York.

McCallum, B.: 2000, Theoretical analysis regarding a zero lower bound on nominal interest rates, *Journal of Money, Credit and Banking* **32**(4), 870–904.

McConnell, M. and Perez-Quiros, G.: 2000, Output Fluctuations in the United States: What Has Changed Since the Early 1980's?, *The American Economic Review* **90**(5), 1464–1476.

Miranda, M. and Fackler, P.: 2001, *Applied Computational Economics*, MIT Press, Cambridge, MA.

Mork, K. A.: 1989, Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results, *The Journal of Political Economy* **97**(3), 740–744.

Nakov, A. and Pescatori, A.: 2007, Inflation-Output Gap Tradeoff with a Dominant Oil Supplier, *working Paper 0723*, Banco de España.

Prescott, E.: 1986, *Theory Ahead of Business Cycle Measurement*, Federal Reserve Bank of Minneapolis.

Rasche, R. and Tatom, J.: 1977, The Effects of the New Energy Regime on Economic Capacity, Production, and Prices, *Federal Reserve Bank of St. Louis Review* **59**(4), 2–12.

Reifschneider, D. and Williams, J.: 2000, Three lessons for monetary policy in a low inflation era, *Federal Reserve Bank of Boston Conference Proceedings* pp. 936–978.

Rotemberg, J. J. and Woodford, M.: 1996, Imperfect competition and the effects of energy price increases on economic activity, *Journal of Money, Credit and Banking* **28**(4), 550–77.

Rotemberg, J. and Woodford, M.: 1997, An optimization-based econometric framework for the evaluation of monetary policy, *NBER Macroeconomics Annual* **12**, 297–346.

Salant, S. W.: 1976, Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market, *The Journal of Political Economy* **84**(5), 1079–1094.

Schmitt-Grohe, S. and Uribe, M.: 2004, Optimal operational monetary policy in the Christiano-Eichenbaum-Evans model of the U.S. business cycle, *Working Paper No.10724*, National Bureau of Economic Research, Inc.

Sims, C.: 2002, Solving Linear Rational Expectations Models, *Computational Economics* **20**(1), 1–20.

Smets, F. and Wouters, R.: 2003, An estimated stochastic dynamic general equilibrium model of the euro area, *Working Paper*, European Central Bank.

Stock, J. and Watson, M.: 2002, Has the Business Cycle Changed and Why?, *NBER Working Papers 9127*, National Bureau of Economic Research, Inc.

Taylor, J.: 1993, Discretion versus policy rules in practice, *Carnegie-Rochester Conference Series on Public Policy* **39**, 195–214.

Wolman, A. L.: 1998, Staggered price setting and the zero bound on nominal interest rates, *Federal Reserve Bank of Richmond Economic Quaerterly* pp. 1–24.

Woodford, M.: 2003, *Interest and Prices*, Princeton University Press, Princeton, NJ.