Optimal Monetary Policy with Downward Nominal Wage Rigidity

A thesis submitted in fulfilment of the requirements for a

Bachelor of Economics (Honours)

By: Ajay Murugasu

520437814

Supervised by: Associate Professor Chris Gibbs, Professor Mariano Kulish

School of Economics

The University of Sydney, NSW

STATEMENT OF ORIGINALITY

I hereby declare that this submission is my own work and to the best of my knowledge it contains no material previously published or written by another person. Nor does it contain any material which has been accepted for the award of any other degree or diploma at the University of Sydney or at any other educational institution, except where due acknowledgment is made in this thesis.

Any contributions made to the research by others with whom I have had the benefit of working at the University of Sydney is explicitly acknowledged. I also declare that the intellectual content of this study is the product of my own work and research, except to the extent that assistance from others in the project's conception and design is acknowledged.

Acknowledgements

I would first like to thank my supervisors, Associate Professor Christopher Gibbs and Professor Mariano Kulish, for their support and guidance throughout this project. Their advice and feedback have been invaluable. On Chris's recommendation, and with the support and encouragement of both supervisors, I was fortunate to be awarded the Brian Gray APRA Scholarship. I am grateful to the Australian Prudential Regulation Authority (APRA) and the Reserve Bank of Australia (RBA) for their generosity in administering this scholarship.

Finally, to my family—thank you for your unwavering support, patience, and belief. Without your sacrifices and encouragement, none of this would have been possible.

Abstract

I investigate optimal monetary policy design in the presence of downward nominal wage rigidities (DNWR), a defining structural feature of the Australian economy. Embedding an occasionally binding wage inflation floor into a New Keynesian model, I compare the performance of optimal policy under commitment, discretion, and a Taylor rule under productivity and demand shocks. The analysis reveals that optimal credible policy takes the form of a target criterion adjusting wage inflation in response to changes in the output gap. Under positive productivity shocks, commitment achieves the lowest central bank losses through expectations management, maintaining parameter invariance across constraint states. The wage floor improves measured performance by providing automatic nominal anchoring. Discretion suffers severe deterioration as period-by-period optimisation creates destabilising expectational feedback loops. A simple Taylor Rule outperforms discretionary policy when the constraint binds. Under negative demand shocks, both Commitment and Discretion exploit the divine coincidence to achieve zero losses, while the Taylor Rule fails to do so regardless of calibration.

Keywords: optimal monetary policy, downward nominal wage rigidity, occasionally binding constraints, wage inflation targeting

Contents

1 Introduction

When wages are unable to adjust downward, the economy experiences what is known as downward nominal wage rigidity (DNWR). This restriction alters how the labour market absorbs shocks and, in turn, how the broader economy responds to them. This thesis examines the implications of downward nominal wage rigidity for monetary policy, focusing on how policy should be optimally designed when wages cannot fall.

Canonical optimal policy frameworks, from Clarida, Galí and Gertler (1999) to Woodford (2003), typically model nominal rigidities as symmetric adjustment frictions. However, these frictions do not generate the asymmetry inherent in DNWR. Faced with a shock that requires a lower real wage, a standard model with only symmetric rigidities permits gradual nominal wage declines. Under a binding DNWR constraint, this adjustment channel is blocked, forcing the central bank to tolerate a larger and more persistent output gap creating a policy problem absent from symmetric frameworks.

This thesis makes three primary contributions. First, it isolates DNWR as the sole nominal friction in an otherwise flexible-price New Keynesian model, thus exposing the pure mechanism through which the wage floor alters monetary trade-offs. Secondly, it implements an occasionally binding wage inflation constraint, reflecting the institutional structure of wage bargaining and generating nonlinear, state-dependent dynamics absent from the standard optimal monetary policy problem. Third, it quantifies the welfare improvement of optimal policy under commitment to that under discretion and a Taylor Rule across different macroeconomic shocks.

The results demonstrate that optimal commitment policy under this wage inflation constraint becomes inherently history- and state-dependent. When the constraint is slack, policy follows a target criterion that adjusts wage inflation in response to changes in the output gap. When the floor binds, restricted in its ability to target wage inflation, policy under commitment smooths adjustment inter-temporally by managing expectations. By committing to correct past deviations, a credible central bank anchors wage-setters' behaviour, mitigating the severity of the trade-off between wage inflation and the output gap.

Further analysis reveals a clear hierarchy of policy performance. Optimal policy under commitment consistently achieves superior outcomes, managing expectations effectively when the wage floor binds. Under positive productivity shocks, which create a trade-off between wage inflation and the output gap, the lack of commitment leads to significantly higher volatility under discretionary policy. When the constraint binds, a standard Taylor Rule is found to outperform discretion, as the rule's fixed coefficients prevent the expectational instability that discretionary re-optimisation triggers. Under negative demand shocks, however, both commitment and discretion can exploit the 'divine coincidence' to achieve near-perfect stabilisation whereas the Taylor Rule is unable to do so.

Taken together, these findings demonstrate that DNWR modelled as a wage inflation floor has large implications for optimal monetary policy design. When the constraint binds, the ability to make credible commitments about future policy is paramount for anchoring expectations and avoiding destabilising volatility. Conversely, when the constraint is slack, the capacity for flexible, state-contingent responses is sufficient. The presence of DNWR thus creates a regime-switching environment where no single policy framework dominates unconditionally. The relative merits of each regime are shock-dependent.

This thesis proceeds as follows. Section 2 documents the empirical and institutional foundations of DNWR in Australia. Section 3 situates this thesis within the literature on optimal policy. Section 4 presents the methodological framework. Section 5 reports analytical and numerical results for positive productivity and negative demand shocks. Section 6 discusses extensions and avenues for future research. Section 7 concludes.

2 Motivation

This section first establishes DNWR as a structural feature of the Australian labour market. It then argues that the architecture of Australia's wage-setting system, anchored in legal, contractual, and behavioural prohibitions against noinal cuts, renders the smooth, cost-based representations of wage stickiness found in canonical New Keynesian frameworks analytically misleading. Instead, a hard, occasionally binding constraint is required to capture the nonlinear and state-dependent dynamics that monetary policy must confront in practice.

2.1 Evidence of DNWR in Australia

Microeconomic evidence reveals pronounced asymmetry in the distribution of wage changes, with nominal reductions occurring infrequently while wage freezes and increases remain commonplace.

Figure 1 illustrates three distinctive features of this asymmetry. First, wage freezes cluster sharply at exactly zero—a spike whose frequency far exceeds predictions from any symmetric adjustment process. Second, the distribution's left tail is conspicuously thin, indicating that nominal wage cuts are exceptionally uncommon. Bishop and Cassidy (2017) found that such reductions affected fewer than 5 percent of job-stayers, with their incidence declining steadily from the 1990s through 2015. Third, the distribution exhibits pronounced positive skewness, with the vast majority of workers experiencing either frozen or rising wages.

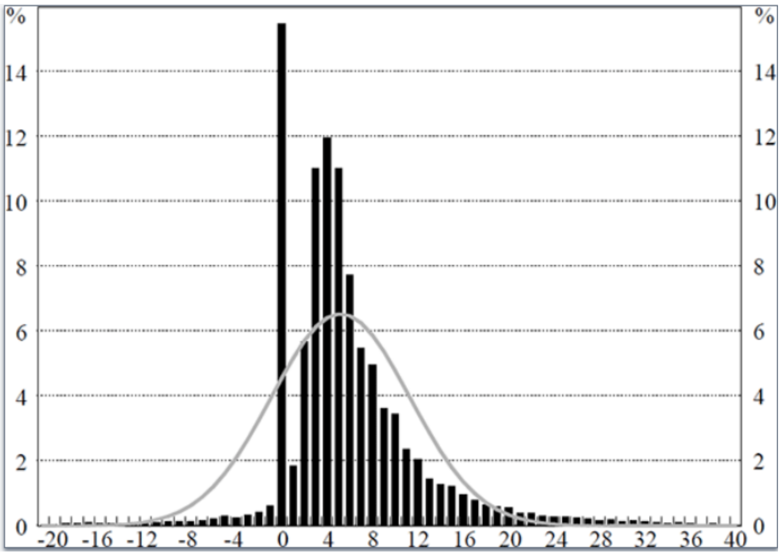

Figure 1: Distribution of Annual Wage Changes.

This rigidity has proven durable across diverse economic conditions, persisting through both the low-inflation, low-growth environment of the 2010s and the COVID-19 recession. Analysing wage dynamics during 2013-2015, when nominal wage growth averaged approximately 2 percent annually, Jacobs and Rush (2015) reveal that firms responded to weak demand primarily through wage freezes and reduced hiring rather than cutting incumbent workers' pay. Even amid subdued aggregate growth and labour market slack, nominal reductions remained scarce—behavior inconsistent with flexible adjustment and instead suggestive of an effective lower bound on individual wage changes.

The COVID-19 crisis provides particularly compelling evidence of the persistence of DNWR. Bishop and Lancaster (2023) document that over 40 percent of Australian households reported unchanged wages during 2020-2022, despite unprecedented economic uncertainty and sectoral disruptions. Among continuously employed individuals, Andrews et al. (2022) show that adjustment occurred predominantly through reduced hours, wage freezes, and foregone bonuses rather than nominal cuts. Further corroboration comes from Buchanan and Kells (2024), who find that New South Wales' temporary 2.5 percent public sector wage cap generated real wage declines through inflation erosion rather than nominal reductions, underscoring institutional and behavioural resistance to downward adjustment.

Collectively, this evidence establishes downward nominal wage rigidity as a structural feature of Australian wage determination rather than a transitory phenomenon. The asymmetry is marked, persistent across varying economic environments, and resilient to even the most severe macroeconomic shocks.

2.2 Justification of a Wage Inflation Floor

Extensive empirical evidence of DNWR justifies its inclusion in the optimal monetary policy problem. This subsection argues that DNWR should be modelled as a hard occasionally binding constraint, to reflect the institutional architecture of Australia's three-tiered wage-setting system.

The modern awards system, covering approximately 20 percent of Australian employees, establishes legally binding minimum terms and conditions determined by the Fair Work Commission through Annual Wage Reviews. Critically, the Commission has never implemented nominal reductions in these reviews. During periods of economic weakness, the FWC may freeze award wages or limit increases, but nominal cuts are institutionally precluded. Pol (2020) documents this consistent pattern, while Watts and Mitchell (1990) trace the norm against nominal reductions to earlier wage-fixing tribunals. This creates a legally binding floor at zero nominal growth for a substantial share of the workforce.

Enterprise bargaining agreements, covering approximately 40 percent of employees, reinforce this rigidity through contractual mechanisms. EBAs are multi-year agreements negotiated between employers and unions or employee representatives, typically specifying scheduled annual increases of 2-3 percent over three-to-four year terms. These agreements commonly include inflation indexation clauses and explicitly prohibit nominal reductions during their duration. Bishop and Chan (2019) emphasise that even during renegotiation following expiry, nominal cuts remain exceptional—union involvement and fairness norms create powerful resistance to reducing established wage levels. The contractual nature of EBAs thus generates a binding constraint rather than merely elevated adjustment costs.

Individual employment contracts, covering the remaining 40 percent of workers, are shaped by award and EBA norms despite lacking formal collective determination. Employers face behavioural and reputational constraints against nominal pay reductions. International evidence from Fougère et al. (2018) and Castellanos et al. (2004) demonstrates that fairness considerations, concerns about employee morale and retention, and social comparisons create strong reluctance to cut nominal wages even when legally permissible. In Australia, the prevalence of award-based minima and EBA-determined benchmarks establishes implicit norms extending to individually contracted workers.

Taken together, these institutional features render nominal wage reductions practically infeasible for the vast majority of the workforce. Empirical evidence for this can be seen in the behaviour of the national Wage Price Index (WPI), the ABS' preferred measure of pure nominal wage change. Since the series began, economy-wide nominal wage growth has never fallen below zero establishing a clear, empirically observable floor in aggregate wage data. The key question for monetary policy design is thus how a wage inflation floor, reflecting this institutional reality, shapes optimal policy.

3 Literature Review

3.1 Foundations of Optimal Policy

Modern theory of optimal monetary policy evolved from the recognition that policy credibility and expectations management are central to macroeconomic stability. Kydland and Prescott

(1977) demonstrated that commitment to policy rules dominates discretionary optimisation due to time inconsistency- optimal plans announced today become suboptimal to follow tomorrow, destroying their credibility. Barro and Gordon (1982) applied this insight to demonstrate how discretionary central banks generate inflation bias by attempting to exploit the short-run Phillips curve despite rational expectations.

Through the 1980s and 1990s, these dynamic consistency insights were gradually integrated with emerging New Keynesian microfoundations, particularly Calvo (1983)'s staggered price adjustment and Rotemberg (1982)'s monopolisitc competition framework. This culminated in the modern optimal policy framework summarized by Clarida, Galí and Gertler (1999), hereafter CGG, who unified rational expectations, nominal rigidities, and dynamic optimisation into a tractable model for analysing optimal monetary policy.

The CGG framework demonstrated that when prices adjust sluggishly through Calvo-style contracts, policymakers face a short-run trade-off between stabilising inflation and output gaps. Under committment, the optimal policy rule exhibits history dependence: the central bank promises to correct past inflation deviations, anchoring expectations and improving current outcomes. This inter-temporal dimension generates superior outcomes compared to discretionary period-by-period optimisation, where inability to credibly commit eliminates expectations management. The CGG approach's elegance and analytical tractability made it foundational for the subsequent literature on both theoretical and applied policy design.

Building on CGG, Woodford (1999), King and Wolman (1999) showed that with only price rigidity, stabilising inflation simultaneously stabilises the output gap in what was later termed the 'divine coincidence' by Blanchard and Gali (2007). However, Erceg, Henderson and Levin (2000), henceforth referred to as EHL, broke this symmetry by introducing nominal wage rigidity alongside price stickiness. They found that when both frictions coexist, the policymaker must balance three competing objectives- stabilising price inflation, wage inflation, and the output gap- meaning the divine coincidence no longer holds forcing the policymaker to tolerate volatility.

3.2 Monetary Policy Incorporating DNWR

Despite the empirical prominence of asymmetric wage adjustment, the EHL model and its descendants treat upward and downward wage movements symmetrically. This abstraction obscures the institutional and behavioural asymmetries documented in Section 2. Downward nominal wage rigidity has been recognised as important for monetary policy since Tobin (1972) which showed that positive inflation "greases the wheels" of labor markets by facilitating real wage adjustment when nominal cuts prove difficult. The literature on monetary policy, however, diverges methodologically in modelling DNWR.

One branch employs asymmetric adjustment costs while maintaining continuous optimisation. Kim and Ruge-Murcia (2011) introduce higher penalties for wage cuts than for increases, producing a kinked Philips curve that preserves tractability within standard DSGE methods. Benigno and Ricci (2011) similarly show that asymmetric rigidities yield nonlinear Philips curves where inflation reacts disproportionately to positive versus negative outpu gaps. These models, however, capture the difficulty of downward adjustment but not its impossibility. The friction is smooth rather than discrete. Consequently, they fail to capture the wage inflation floor empirically observed in Australia.

An alternative branch models DNWR as a hard, occasionally binding constraint. Schmitt-Grohé and Uribe (2016) impose a zero lower bound on nominal wages within a flexible-price RBC model, showing that adverse shocks requiring real wage cuts generate involuntary unemployment. They formalise Tobin's 'grease the wheels' mechanism, finding that small positive inflation targets mitigate welfare losses by relaxing the binding frequency of the constraint. Evans (2018) applies the same constraint to a medium-scale New Keynesian model with price and wage stickiness, finding that under Taylor rules, DNWR induces strong asymmetries. Wage floors bind under negative demand shocks, amplifying unemployment persistence, while upward flexibility yields smooth expansions. His results confirm that modest positive inflation (around 0.75-1 % annually) can improve welfare by reducing binding frequency and severity.

3.3 Contribution

Despite the advances in the literature, several gaps remain unaddressed. Most existing work analyses policy within models featuring both price and wage rigidities simultaneously, confounding the distinct effects of each friction. While papers like Evans (2018) examine simple rules, comprehensive characterisation of optimal commitment policy under wage floors using the loss function framework from CGG remains unexplored.

This thesis makes three distinct contributions.

First, it isolates the DNWR friction by assuming flexible prices, providing a clean benchmark for wage rigidity's policy implications without the confounding effect of price rigidity. Abstracting from price stickiness enables clear identification of policy trade-offs attributable solely to the wage floor. In EHL, the central bank balances both rigidities simultaneously. Focusing exclusively on wage dynamics in the presence of a constraint clarifies what optimal policy strives to achieve when confronting DNWR alone, providing a benchmark for more complex models.

Second, this thesis extends New Keynesian optimal policy analysis to accommodate occasionally binding constraints, rather than asymmetric adjustment costs, capturing institutional realities —a class of problems increasingly recognized as important (Brassil, Gibbs and Ryan 2025) but methodologically challenging. The constraint creates state-dependence: the economy operates under different regimes depending on if the constraint binds, enabling analysis of how optimal policy adjusts with and without the constraint binding- a mechanism absent from smooth cost specifications like Erceg et al. (2000).

Third, this thesis quantifies the welfare gains from commitment and compares it to discretionary policy and wage inflation Taylor rules under both productivity and demand shocks. The welfare comparison reveals how commitment's ability to exploit forward-looking behaviour, already established in CGG, varies across different shocks and if the constraint affects commitment's ability to do so which remains under-explored in the literature.

4 Methodology

This section develops the framework for analysing optimal monetary policy under downward nominal wage rigidity. I first lay out the model environment, then formalise the wage floor, define the policy regimes under consideration, and finally detail the solution method and evaluation methodology.

4.1 Model Environment

I construct a small-scale New Keynesian DSGE model with flexible prices and a hard floor on nominal wage inflation as an occasionally binding constraint.

4.1.1 Firms

The production side features a continuum of differentiated intermediate goods indexed by $i \in [0, 1]$, aggregated by perfectly competitive final goods producers into a single final consumption good Y_t. This monopolistic competition structure, standard in New Keynesian models following Dixit and Stiglitz (1977), generates endogenous price markups.

Final goods producers aggregate intermediate varieties using constant elasticity of substitution (CES) technology:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} \, di \right)^{\frac{\epsilon}{\epsilon-1}} \tag{1}$$

where $\epsilon > 1$ measures the elasticity of substitution between varieties. Cost minimization yields demand for each variety:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t \tag{2}$$

where $P_t = \left(\int_0^1 P_t(i)^{1-\epsilon} \, di \right)^{\frac{1}{1-\epsilon}}$ is the aggregate price index.

Each intermediate firm produces output using linear technology in labor:

$$Y_t(i) = A_t N_t(i) \tag{3}$$

where A_t is aggregate TFP and $N_t(i)$ is labor employed. Firms hire labor competitively at nominal wage W_t.

Under flexible prices, profit maximization yields:

$$P_t(i) = \mu \frac{W_t}{A_t} \tag{4}$$

where $\mu = \frac{\epsilon}{\epsilon - 1} > 1$ is the constant markup over nominal marginal cost. Symmetric equilibrium implies all firms choose identical prices $P_t(i) = P_t$.

This flexible-price assumption is critical for isolating wage rigidity effects. As Erceg, Henderson and Levin (2000) demonstrate, combining sticky prices and wages confounds their distinct contributions to policy trade-offs. By abstracting from price stickiness, we identify cleanly how the wage floor alone alters optimal policy.

4.1.2 Households

The representative household maximizes expected lifetime utility by choosing consumption (C_t), labor supply (N_t), nominal wages (W_t), and one-period nominal bonds (B_t), subject to downward nominal wage rigidity as an occasionally binding constraint:

$$\max_{\{C_t, N_t, W_t, B_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, N_t)$$

$$\text{subject to} \quad P_t C_t + Q_t B_t \leq B_{t-1} + W_t N_t + D_t - T_t^{\text{adj}}, \quad \forall t$$

$$\pi_t^w \geq \bar{\pi}^w, \quad \forall t$$

where Q_t is the bond price, D_t represents dividends, and T_t^{adj} denotes Rotemberg (1982) wage adjustment costs:

$$T_t^{\text{adj}} = \frac{\phi_w}{2} \left(\frac{W_t}{W_{t-1}} - 1 \right)^2 P_t Y_t \tag{5}$$

where $\phi_w > 0$ governs adjustment cost magnitude and $\pi_t^w \equiv \frac{W_t}{W_{t-1}} - 1$.

I adopt Rotemberg costs over Calvo wage-setting for three reasons: analytical tractability (closed-form first-order conditions without additional state variables), symmetric adjustment costs providing a neutral baseline, and natural connection to the quadratic loss function framework.

The instantaneous utility function is:

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

where $\sigma > 0$ measures inverse intertemporal elasticity of substitution and $\psi > 0$ is inverse Frisch elasticity of labor supply.

Deriving the Dynamic IS Curve The household's Euler equation yields:

$$Q_t = \beta E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{P_t}{P_{t+1}} \right] \tag{6}$$

Log-linearizing around zero-inflation steady state with $i_t \equiv -\ln Q_t$:

$$c_t = E_t c_{t+1} - \frac{1}{\sigma}(i_t - E_t \pi_{t+1} - \rho) \tag{7}$$

where lowercase variables denote log-deviations and $\rho \equiv -\ln \beta$.

With flexible prices, optimal price-setting from equation (4) implies:

$$\pi_t = \pi_t^w - \Delta a_t \tag{8}$$

Substituting and imposing market clearing $y_t = c_t$, then defining the output gap $x_t \equiv y_t - y_t^n$ and subtracting the natural-output Euler equation yields:

$$x_t = E_t x_{t+1} - \frac{1}{\sigma}(i_t - E_t \pi_{t+1}^w - r_t^n) \tag{9}$$

where r_t^n is the natural rate of interest.

Deriving the Wage Phillips Curve with DNWR The household's optimization problem, incorporating Rotemberg adjustment costs and the occasionally binding constraint $\pi_t^w \geq \bar{\pi}^w$, leads to a wage-setting equation that can be expressed in log-linearized form as a piecewise

Wage Phillips Curve:

$$\pi_t^w = \begin{cases} \beta E_t[\hat{\pi}_{t+1}^w] + \kappa_w \left(\widehat{MRS}_t - \hat{w}_t \right) & \text{if } \pi_t^w > \bar{\pi}^w \quad \text{(Slack)} \\ \bar{\pi}^w & \text{if } \pi_t^w = \bar{\pi}^w \quad \text{(Binding)} \end{cases} \tag{10}$$

where $\kappa_w \equiv \frac{(\sigma+\psi)(1-\beta)}{\phi_w}$ measures the sensitivity of wage inflation to labor market conditions, $\widehat{MRS}_t = \sigma c_t + \psi n_t$ is the marginal rate of substitution, and $\hat{w}_t = w_t - p_t$ represents real wages.

The equilibrium can ultimately be characterised as $\pi_t^w = \beta E_t[\pi_{t+1}^w] + \kappa_w \left(\widehat{MRS}_t - \hat{w}_t \right) + \lambda_t$. The multiplier λ_t represents the shadow value of the constraint- when the constraint is slack, $\lambda_t = 0$ and when it binds, $\lambda_t = \bar{\pi}^w - \beta E_t[\pi_{t+1}^w] - \kappa_w \left(\widehat{MRS}_t - \hat{w}_t \right)$.

The two regimes operate as follows:

- When the constraint is slack ($\pi_t^w > \bar{\pi}^w$), standard forward-looking wage Phillips curve dynamics prevail. Wage inflation responds to expected future wage inflation and the gap between workers' desired wage (based on their marginal rate of substitution) and the current real wage.

- When the constraint is binding ($\pi_t^w = \bar{\pi}^w$), wage inflation is exogenously pinned at the floor. This creates the state-dependent nonlinearity central to DNWR's effects on monetary transmission.

4.2 Wage Floor Constraint

The wage inflation floor $\pi_t^w \geq \bar{\pi}^w$ represents institutional barriers to nominal wage cuts documented in Section 2, creating a hard, occasionally binding constraint.

Computationally, we implement the floor at $\bar{\pi}^w = -0.02$ (negative two percent annually). Since the model is linearised around a steady state with positive wage inflation, this value represents a deviation from steady-state wage inflation rather than an absolute level. With a steady-state annual wage inflation of approximately two percent, this floor corresponds to a binding constraint at zero nominal wage growth, consistent with the empirical evidence detailed in Section 2.

4.3 Policy Regimes

4.3.1 Optimal Policy Framework

Two methodological approaches dominate optimal policy analysis. The loss function approach, most notably used in CGG, posits a central bank objective quadratic in inflation and output gap deviations from steady state and derives optimal policy by minimising expected losses subject to equilibrium constraints. This framework offers substantial tractability. Linear quadratic methods yield closed-form solutions and interpretable rules connecting naturally to Taylor rule literature. Woodford (2003) demonstrates that appropriate quadratic specifications approximate household welfare through second-order expansions around steady state. The primary limitation is potential ad-hoc specification if not microfounded. Quadratic forms penalize deviations symmetrically, potentially missing asymmetric welfare costs- though this is less concerning when the rigidity itself creates asymmetry.

The Ramsey approach, employed by Khan, King and Wolman (2003) and Erceg et al. (2000), directly maximizes representative household utility subject to equilibrium constraints. Fully micro-founded, it captures all welfare-relevant trade-offs, representing the constrained social optimium. However, computational demands are substantial, particularly with non-linear constraints, and solutions are often state-contingent and implicit rather than simple targeting rules.

This thesis adopts the loss function approach for three reasons. First, the occasionally binding wage floor requires Occbin methodology, adopted from Guerrieri and Iacoviello (2015), integrating naturally with linear-quadratic frameworks but becoming more complex under fully non-linear Ramsey optimization. Second, direct comparison with wage inflation Taylor rules requires a common loss metric. Third, quadratic losses in wage inflation and the output gap provide transparent policy objectives whilst still approximating welfare. While full Ramsey optimization would provide complete welfare accounting, the loss function method delivers analytical clarity and computational feasability essential for analyzing occasionally binding constraints.

4.3.2 Optimal Policy Under Commitment

Optimal policy minimises the central bank's quadratic loss function subject to equilibrium constraints including the wage floor. Following Clarida, Galí and Gertler (1999):

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left[(\pi_t^w)^2 + \alpha x_t^2 \right] \tag{11}$$

where $\alpha > 0$ weights output gap stabilization relative to wage inflation stabilization. This quadratic specification derives from second-order Taylor approximation to household utility around zero-inflation steady state (Woodford 2003).

The central bank commits at $t = 0$ to a complete state-contingent policy plan minimizing expected losses subject to the wage Phillips curve and wage floor constraint. Commitment means honoring this plan even when ex-post incentives favor deviation—the ability to bind future policy creates credibility influencing forward-looking expectations.

Taking first-order conditions yields the optimal policy target criterion:

$$\pi_t^w = -\frac{\alpha}{\kappa_w}(x_t - x_{t-1}) \quad \text{when } \lambda_t = 0 \tag{12}$$

This target criterion exhibits history dependence that is fundamental to optimal policy under commitment. The dependence on output gap changes rather than levels is a product of the inter-temporal structure of the welfare loss function. Since wage inflation costs appear in each period while the Phillips curve links current inflation to expected future output gaps, the planner optimally 'leans against the wind' by offsetting the accumulated deviations, creating endogenous policy inertia wherein past shocks continue to influence current stabilization efforts even after the shocks themselves have dissipated.

The optimality of this history-dependent rule operates through an expectations management channel. By committing to undo past output gap expansions through subsequent contractions (and vice-versa), the central bank makes agents internalise future policy responses. This commitment anchors wage inflation expectations today because forward-looking agents recognise that any current deviation will trigger offsetting adjustments. The result is that the credible promise of future correction reduces the need for large contemporaneous responses, thus reducing the policy trade-off between output and wage inflation stabilisation.

4.3.3 Optimal Policy Under Discretion

Under discretion, the central bank lacks the ability to commit to future actions and instead reoptimizes each period taking private sector expectations as given. Whereas commitment allows the policymaker to shape expectations by credible future decisions, discretionary policy treates expectations as predetermined state variables (Kydland and Prescott 1977).

The discretionary policymaker treats $E_t\pi^w_{t+1}$ as exogenous, yielding the target criterion:

$$\pi^w_t = -\frac{\alpha}{\kappa_w}x_t \quad \text{when } \lambda_t = 0 \tag{13}$$

The absence of x_{t-1} compared to under commitment reveals the time-inconsistency problem inherent in discretionary policy. Because the central bank cannot commit to future actions, it has no mechanism to influence how today's decisions affect tomorrow's expectations.

4.3.4 Wage Inflation Taylor Rule

For comparison, I analyze a conventional Taylor rule where the nominal interest rate responds mechanically to observed deviations:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)(\phi_\pi \pi^w_t + \phi_y x_t) + \varepsilon^i_t \tag{14}$$

where $\rho_i \in [0, 1)$ captures interest rate smoothing, $\phi_\pi > 1$ and $\phi_y \geq 0$ are response coefficients, and ε^i_t is a monetary policy shock.

This rule is backward-looking, time-invariant, and has fixed coefficients. Unlike optimal policy, the rule cannot promise future responses that anchor expectations beyond what is implied by its fixed coefficients.

When the wage floor binds, the rule continues prescribing policy based on constrained observed wage inflation, providing misleading guidance since $\pi^w_t = \bar{\pi}^w$ conveys no information about underlying fundamentals when the constraint is active.

4.4 Rational Expectations Solution with Occassionally Binding Constraints

4.4.1 Occasionally Binding Constraints Methodology

I solve the model using the OccBin toolkit (Guerrieri and Iacoviello 2015), designed for occasionally binding constraints in linear rational expectations models. The method exploits piecewise-linearity: when the constraint is slack, the model is linear; when binding, wage inflation is fixed at the floor. OccBin determines endogenously when transitions occur, accounting for agents' rational expectations about future regime switches.

The algorithm proceeds iteratively. Given initial conditions and shocks, OccBin conjectures regime sequences (binding or slack each period), solves the piecewise-linear model forward, checks consistency between conjectured regimes and model-implied wage inflation, and iterates until convergence where regime assignments match dynamics.

This handles forward-looking expectations under regime uncertainty. Agents form expectations knowing the constraint may bind or remain slack depending on future states. OccBin assumes perfect foresight about regime transitions within each iteration, performing well for moderate shocks and finite binding spells.

4.4.2 Alternative Solution Methods and Justification

While OccBin is our core method, alternative approaches exist with distinct trade-offs:

Global non-linear methods (Smolyak collocation, projection methods) approximate policy functions globally, handling large shocks and capturing non-linearities far from steady state. However, computational costs grow exponentially with state variables, implementation is complex for forward-looking models, and the wage floor creates localized non-linearity that OccBin's piecewise approach handles efficiently. For our small-scale model with moderate shocks, first-order accuracy is reasonable and global methods would provide limited gains at substantial computational cost.

Higher-order perturbation (second/third order) captures precautionary motives and risk premia while maintaining low computational costs. However, standard perturbation cannot handle inequality constraints—the wage floor creates non-differentiability requiring discrete regime-switching. The symmetric nature of perturbation misses the fundamental asymmetry: wage increases occur freely while decreases face a hard floor.

Other piecewise-linear methods exist (e.g., Holden 2016), but OccBin is chosen for extensive validation in monetary policy applications with occasionally binding constraints, efficient handling of forward-looking expectations with regime transitions, well-documented MATLAB/Dynare implementation, and widespread adoption facilitating comparability with existing studies.

Given the model's tractable scale, moderate shock calibrations, discrete regime-switching nature of the constraint, and computational feasibility requirements, OccBin provides optimal accuracy-feasibility balance. Global methods would offer marginal accuracy improvements at substantial cost, while higher-order perturbation cannot handle the inequality constraint without regime modeling.

4.4.3 Calibration Strategy

Model parameters are calibrated to Australian wage-setting institutions, observed rigidities, and standard quarterly relationships.

Table 1: Model Parameter Calibration

Symbol	Parameter	Value	Reasoning
α	Weight on output gap in loss function	0.25	Standard
β	Discount Factor	0.99	Standard
σ	Inverse Intertemporal Elasticity	1.0	Standard
φ	Inverse Frisch Elasticity	5	Gali (2015)
ϕ_π	Taylor Rule Weight on Inflation	1.5	Gali (2015)
ϕ_y	Taylor Rule Weight on Output Gap	0.5	Gali (2015)
ρ	Time Preference Rate	0.01	$-\log(\beta)$
u^n	Natural Rate of Unemployment	5%	Gali (2015)
γ	DNWR Parameter	0.99	Schmitt-Grohé and Uribe (2016)
κ_w	Slope of Wage Phillips Curve	0.1	Model-specific
ρ_u	Persistence of Cost-push Shock	0.8	Standard
ρ_{rn}	Persistence of Natural Rate Shock	0.8	Standard
$\bar{\pi}^w$	Wage Inflation Floor	-0.02	-2% deviation from steady state

4.5 Welfare Evaluation

Welfare is evaluated by computing losses from deterministic impulse responses to unit shocks, comparing outcomes when the wage floor can bind versus when effectively removed. Throughout this analysis, welfare is measured by the central bank's quadratic loss function rather than derived household utility, a distinction that becomes particularly salient under binding constraints where the loss function abstracts from the welfare costs of the nominal rigidity itself.

Loss Function Specification Welfare is measured by the central bank's quadratic loss function:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left[(\pi_t^w)^2 + \alpha x_t^2 \right] \tag{15}$$

This derives from second-order Taylor approximation to household utility (Woodford 2003), capturing primary policy trade-offs while providing transparent welfare comparisons.

Computational Implementation Welfare evaluation proceeds through deterministic impulse response analysis:

1. **Constrained case**: Set wage floor at $\bar{\pi}^w = -0.02$ and compute impulse responses to unit shock ($\varepsilon_t^r = 0.3$ or $\varepsilon_t^u = 0.3$) over 50 periods. OccBin determines endogenously when the floor binds, generating regime-switching paths $\{\pi_t^w, x_t, i_t, \lambda_t\}_{t=1}^{50}$.

2. **Unconstrained case**: Set floor to $\bar{\pi}^w = -999$ ensuring the constraint never binds. Compute responses to identical shock, yielding unconstrained dynamics where wage inflation adjusts freely.

3. **Period loss**: Compute period-by-period losses for both cases:

$$L_t = \frac{1}{2} \left[(\pi_t^w)^2 + \alpha x_t^2 \right], \quad t = 1, 2, ..., 50$$

4. **Welfare via backward recursion**: Starting from $W_{50} = L_{50}$, recursively compute:

$$W_t = L_t + \beta W_{t+1}, \quad t = 49, 48, ..., 1$$

The initial period welfare W_1 represents total expected discounted loss.

5. **Central bank loss comparison**: Compute constraint cost as:

$$\text{Welfare Cost} = 100 \times \frac{W_1^{\text{constrained}} - W_1^{\text{unconstrained}}}{W_1^{\text{unconstrained}}}$$

6. **Cross-regime comparison**: Repeat for discretion and Taylor rule, yielding four welfare measures enabling quantification of: (a) welfare gain from optimal commitment versus discretion versus Taylor rule, and (b) welfare cost of the constraint within each regime.

This deterministic approach provides transparent insight into how optimal policy manages constraint binding for specific shock types, with OccBin's regime-switching solution most reliable for deterministic paths where transitions can be precisely identified.

4.6 Equilibrium Conditions

$$\text{Resource Constraint:} \quad x_t = c_t \tag{16}$$

$$\text{Production Function:} \quad x_t = n_t \tag{17}$$

$$\text{Euler Equation:} \quad c_t = E_t c_{t+1} - \frac{1}{\sigma}(i_t - E_t \pi_{t+1} - r_t^n) \tag{18}$$

$$\text{Wage New Keynesian Phillips Curve:} \quad \pi_t^w = \beta E_t[\hat{\pi}_{t+1}^w] + \kappa_w(\widehat{MRS}_t - \hat{w}_t) + \lambda_t \tag{19}$$

$$\text{Wage Inflation Floor:} \quad \pi_t^w \geq \bar{\pi}^w \tag{20}$$

$$\text{Complementary Slackness:} \quad \lambda_t \geq 0, \quad \lambda_t(\pi_t^w - \bar{\pi}^w) = 0 \tag{21}$$

$$\text{Price-Wage Relationship:} \quad \pi_t = \pi_t^w - (u_t - u_{t-1}) \tag{22}$$

$$\text{Real Wage Identity:} \quad \omega_t = w_t - p_t \tag{23}$$

$$\text{Labor Force:} \quad w_t - p_t = \sigma c_t + \varphi l_t \tag{24}$$

$$\text{Wage Inflation Definition:} \quad \pi_t^w = w_t - w_{t-1} \tag{25}$$

$$\text{Price Inflation Definition:} \quad \pi_t = p_t - p_{t-1} \tag{26}$$

$$\text{TFP Shock Process:} \quad u_t = \rho_u u_{t-1} + \varepsilon_t^u \tag{27}$$

$$\text{Natural Rate Process:} \quad r_t^n = \rho_r r_{t-1}^n + \varepsilon_t^r \tag{28}$$

To close the system, monetary policy is defined by one of three regimes:

Commitment: $\pi_t^w = -\frac{\alpha}{\kappa_w}(x_t - x_{t-1})$ when $\lambda_t = 0$, otherwise $\pi_t^w = \bar{\pi}^w$.

Discretion: $\pi_t^w = -\frac{\alpha}{\kappa_w}x_t$ when $\lambda_t = 0$, otherwise $\pi_t^w = \bar{\pi}^w$.

Taylor Rule: $i_t = \rho_i i_{t-1} + (1 - \rho_i)(\phi_\pi \pi_t^w + \phi_y x_t)$.

OccBin solves this system, determining endogenously when the constraint binds given the shocks that hit the economy.

5 Results

This section presents the core analytical results. Through a series of impulse response and welfare analyses, it will demonstrate how the presence of a hard wage floor alters the conduct and effectiveness of optimal monetary policy with a Taylor Rule as comparison.

5.1 Positive TFP Shock

I begin by analysing the economy's response under three different regimes to a positive productivity shock u of magnitude 0.3 and standard deviation 0.01.

5.1.1 Optimal Commitment Policy

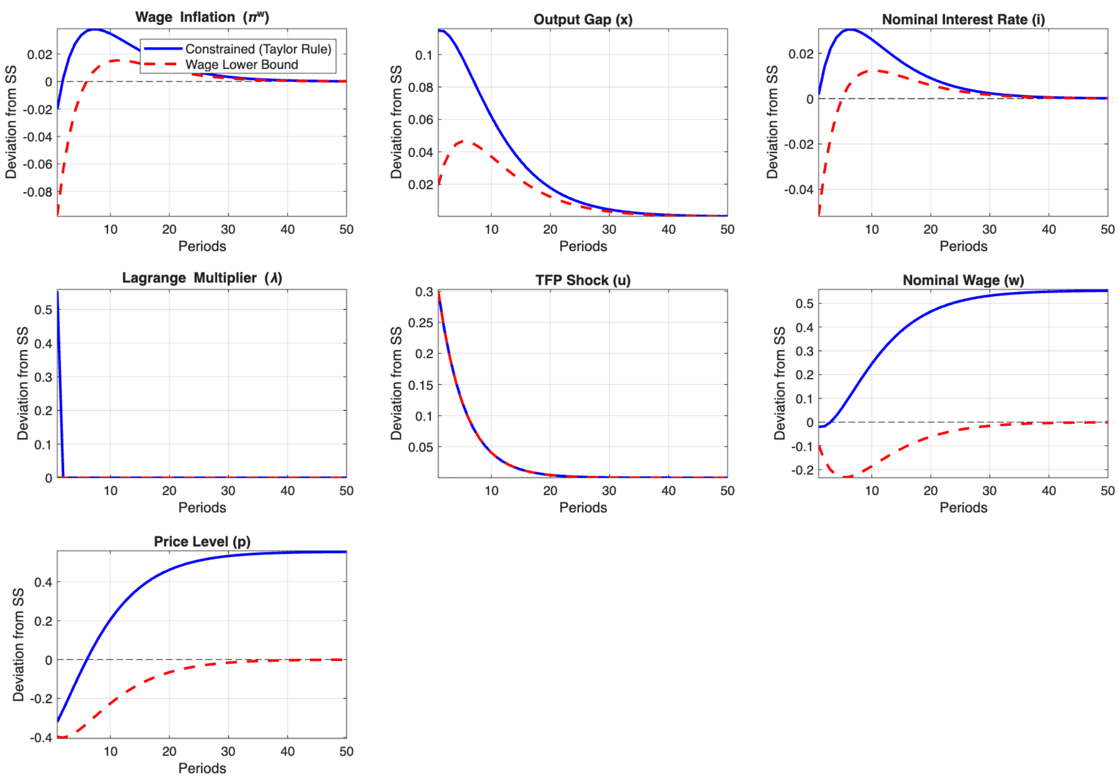

Figure 2: Impulse responses of variables with optimal commitment policy under a positive productivity shock. The shock follows an AR(1) with $\rho_u = 0.5$ and $\sigma = 0.1$

Unconstrained Case A positive productivity shock raises the marginal product of labor, requiring real wages to rise. With sticky nominal wages, this adjustment can occur through two channels: nominal wage increases or price level decreases. Sticky wages prevent immediate nominal adjustment, so flexible prices fall to facilitate the real wage rise through deflation. This creates the central policy trade-off: stabilizing wage inflation, which falls as productivity gains reduce marginal costs, against managing the output gap, which rises as the economy operates above potential during adjustment.

The hump-shaped output gap response reveals commitment's optimal strategy. Rather than immediately closing the gap, the central bank engineers a temporary boom to compress wage markups and facilitate alignment between labor supply and productivity-enhanced demand without excessive wage deflation. This spreads adjustment burden across both margins over time, minimizing present discounted distortions. The optimal policy takes the form of a target criterion $\pi_t^w + \frac{\alpha}{\kappa_w}(x_t - x_{t-1}) = 0$, where wage inflation responds to changes in the output gap rather than its level. This structure creates history-dependence: the output gap's past values influence current wage inflation targets, generating the inertial hump-shaped path that spreads adjustment costs inter-temporally.

This operates through expectations management. By credibly promising future accommodation, the central bank influences current wage setting through the forward-looking Phillips curve. Wage setters moderate current adjustment in anticipation of future policy support, allowing gradual adjustment along both nominal and real margins simultaneously. The target criterion's dependence on output gap changes means past promises constrain current policy, which shapes future expectations, creating smoothing unavailable to policies that re-optimize each period. Nominal interest rates rise initially to prevent deflationary shocks from reducing real rates excessively, then gradually return to steady state.

Constrained Case The constraint eliminates the nominal wage adjustment channel, forcing all equilibration through the output gap. While productivity gains still require real wages to rise, this can no longer occur through efficient combination of nominal wage adjustment and price changes. Instead, the output gap must expand sufficiently to raise the marginal product of labour enough to justify the required real wage at the constrained nominal wage.

The output gap nearly triples and exhibits monotonic decline rather than a hump. This reflects the necessity of using quantities as the sole equilibrating mechanism. Without wage deflation, the central bank cannot spread adjustment across margins. The economy must move further along the labor demand curve through output expansion to reach equilibrium where higher marginal product justifies the real wage the constraint prevents from adjusting nominally. The larger output gap represents resource misallocation rather than optimal smoothing.

Commitment's strategic power erodes substantially. In the unconstrained case, promises of future accommodation moderate current wage setting through the forward-looking Phillips curve. When the constraint binds, this channel breaks down. Wage setters cannot respond by moderating demands because the constraint prevents downward adjustment regardless of expectations. Expectations management becomes irrelevant when the constraint makes it economically infeasible.

The target criterion makes this transparent. When unconstrained, $\pi_t^w + \frac{\alpha}{\kappa_w}(x_t - x_{t-1}) = 0$ balances wage inflation against output gap changes. When constrained, $\frac{\lambda_t}{\kappa_w} = \pi_t^w + \frac{\alpha}{\kappa_w}(x_t - x_{t-1})$ where $\lambda_t > 0$ represents the shadow value of relaxing the constraint. With π_t^w at its lower bound, the positive shadow value must be offset entirely by larger $\frac{\alpha}{\kappa_w}(x_t - x_{t-1})$, requiring substantially larger output movements. This shadow value quantifies the economic cost of losing one adjustment margin.

The permanent nominal wage and price increase illustrates different adjustment mechanics. Unconstrained deflation facilitates temporary real wage adjustment with nominal wages returning to steady state. When the constraint prevents this deflation, nominal wages must permanently increase to achieve the same real wage path. Larger wage (and price) inflation in later periods compensates for constrained earlier deflation, generating larger cumulative nominal distortions and reflecting the inefficiency of forced quantity-only adjustment.

5.1.2 Optimal Discretionary Policy

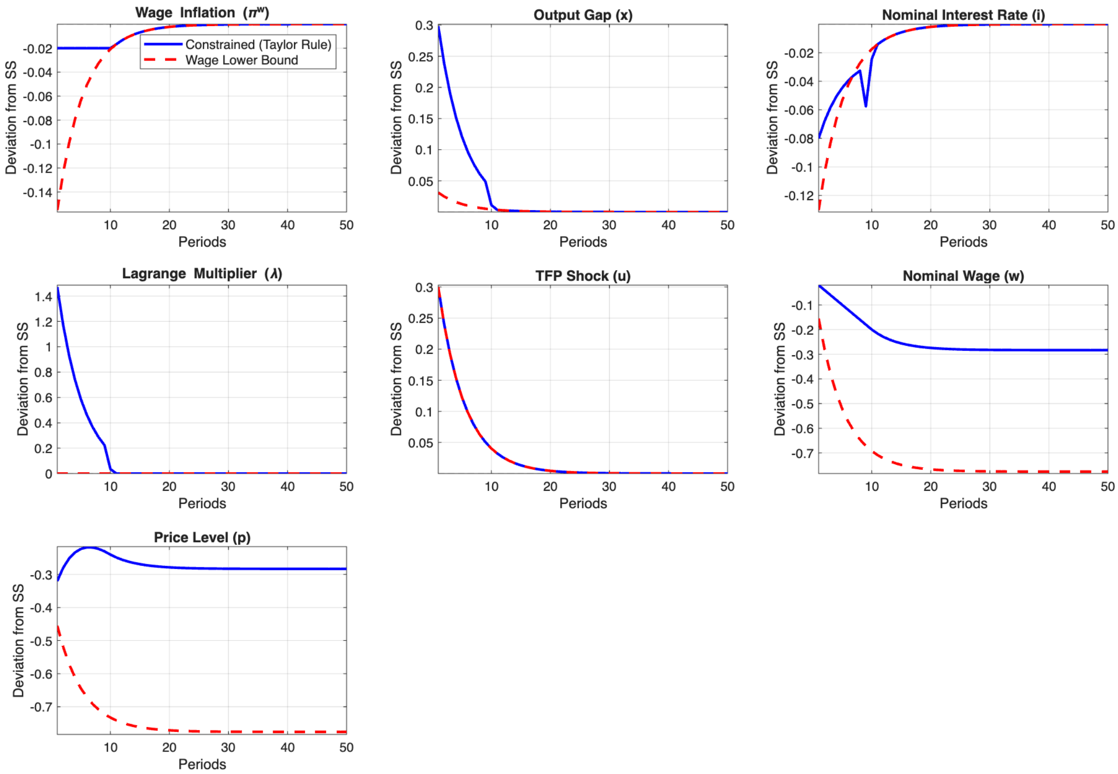

Figure 3: Impulse responses of variables with optimal discretionary policy under a positive productivity shock. The shock follows an AR(1) with $\rho_u = 0.5$ and $\sigma = 0.1$

Unconstrained Case The output gap surges immediately to its peak, far more aggressively than under commitment, then monotonically decreases. Nominal interest rates drop sharply below zero initially before returning to steady state. Prices fall due to reduced marginal costs. Nominal wages decline permanently below steady state. Wage inflation drops deeply negative initially as the output gap surges, then gradually rises but never overshoots above zero.

This greater volatility reveals discretion's fundamental weakness: inability to make credible promises about future policy. Without credibility, the central bank cannot spread adjustment costs intertemporally through expectations management. Instead, it front-loads the entire response, engineering an immediate aggressive output expansion to compress wage markups as much as possible in the current period. When private agents do not believe promises of future accommodation, the only way to influence their behavior is through current economic conditions.

The central bank must create actual current demand stimulus rather than managing expected future stimulus.

The target criterion makes this explicit. Under discretion, wage inflation responds to the output gap level, not its change as under commitment. This reflects absence of history-dependence: each period, the policymaker re-optimizes given current conditions alone, ignoring how today's decisions affect tomorrow's expectations. This creates amplified short-run volatility and prevents the gradual adjustment path commitment achieves. The permanently depressed nominal wage level also shows that without ability to commit to future paths, the policy fails to orchestrate the nominal wage trajectory that returns to steady state under commitment.

The interest rate response contrasts sharply with commitment. Discretion cuts rates aggressively because it cannot influence expectations through credible future promises and must instead stimulate actual current demand to moderate wage deflation. Under commitment, the central bank raises rates while promising future accommodation, exploiting the forward-looking Phillips curve to stabilise current inflation through expectations rather than current output manipulation. This difference quantifies the value of credibility: commitment achieves similar wage stabilisation with less output volatility by substituting expectations management for demand stimulus.

Constrained Case When the constraint binds, the output gap increases substantially more than under commitment and returns to steady state only once the constraint stops binding. Nominal interest rates decrease by less initially than unconstrained because a similar cut would overstimulate demand without aiding necessary real wage adjustment when nominal wages cannot fall. A 'hitch' occurs in the last few periods before the constraint stops binding. Expected wage inflation jumps discontinuously from the floor to a value determined by the Phillips cure. This discontinuous change in $E_t\pi^w_{t+1}$ requires a sharp compensating adjustment in the nominal rate through the IS curve to maintain equilibrium as discretion lacks future credibility. Nominal wages decrease less than unconstrained, then stabilize permanently below steady state. Wage inflation hits the lower bound immediately and remains there for multiple periods, then returns gradually to steady state.

The constraint exposes discretion's critical fragility through expectational feedback. Under discretion, the target criterion becomes $\frac{\lambda_t}{\kappa_w} = \pi_t^w + \frac{\alpha}{\kappa_w} x_t$. The output gap enters in levels, not changes, reflecting absence of history-dependence. With $\lambda_t > 0$ and π_t^w fixed at the floor, the bank must tolerate an elevated output gap level each period to satisfy this condition. This creates a vicious cycle: the large output gap represents aggressive easing to stimulate demand. Private agents observe this aggressive easing each period and rationally increase inflation expectations. Through the forward-looking Phillips curve, higher expected inflation increases current wage pressure, making the constraint bind more severely in subsequent periods. This requires even larger output gaps, propagating the cycle.

Under commitment, the central bank promises future accommodation that moderates current wage setting, allowing the constraint to act as a beneficial nominal anchor. Under discretion, aggressive current easing without credible future promises destabilises expectations, turning the constraint into a source of instability. Each period the bank re-optimises given the current binding constraint, attempting to stimulate demand to moderate wage pressure. But without ability to commit to future policy paths, these attempts only raise expected inflation, worsening the constraint bind. The prolonged periods at the wage floor quantify this destabilization: discretion forces the constraint to bind repeatedly because its lack of credibility prevents the expectations management that would allow smoother adjustment.

5.1.3 Wage Inflation Taylor Rule

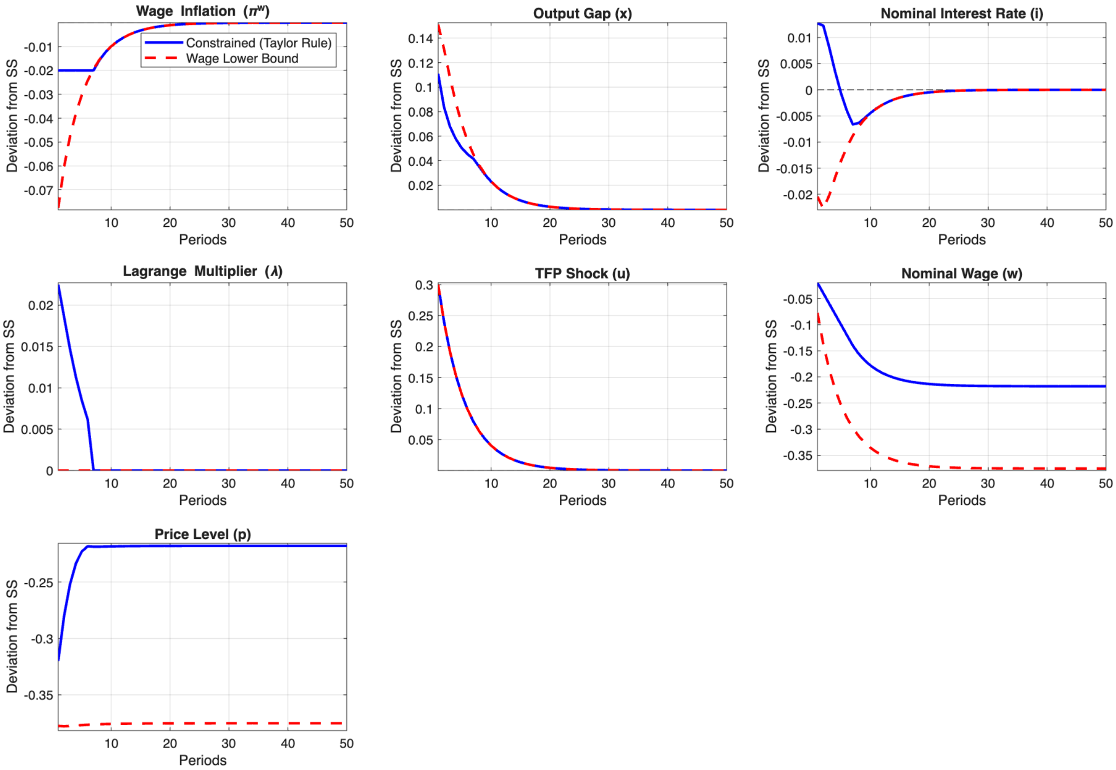

Figure 4: Impulse responses of variables with a wage inflation Taylor Rule given by $i_t = \rho_i i_{t-1} + (1 - \rho_i)(\phi_\pi \pi_t^w + \phi_y x_t)$ where $\rho_i = 0.8$, $\phi_\pi = 1.5$, $\phi_y = 0.5$ under a positive productivity shock. The shock follows an AR(1) with $\rho_u = 0.5$ and $\sigma = 0.1$.

Unconstrained Case The output gap rises substantially, exceeding even discretionary policy, without commitment's hump-shaped pattern. The nominal interest rate falls below zero initially and declines further in a negative hump before returning to steady state. Nominal wages decrease permanently. Prices drop initially then partially reflate, stabilizing below steady state. Wage inflation decreases below zero and returns gradually to steady state.

The Taylor Rule's inferior performance reflects its mechanical nature. The rule observes falling wage inflation and rising output gap, then applies fixed coefficients $\phi_\pi = 1.5$ and $\phi_y = 0.5$ to determine the interest rate response. This mechanical prescription cuts rates aggressively, stimulating demand without regard for whether this efficiently balances nominal and real stabilization for productivity shocks. The rule cannot access either commitment's expectations management or discretion's state-contingent re-optimization.

Despite lacking commitment's forward guidance, discretion performs substantially better because it re-optimizes each period, choosing precisely how much output expansion efficiently compresses wage markups given current conditions. The discretionary policymaker solves for the optimal output-inflation trade-off in real time. The Taylor Rule mechanically applies predetermined weights, producing larger output expansion than optimal. The initial price drop reflects reduced marginal costs, facilitating real wage adjustment. The rate cut stimulates demand, pushing prices partially back up, though they stabilize below steady state as the shock dissipates. This reveals the rule's core limitation: fixed coefficients calibrated for general conditions cannot adapt to the specific trade-offs productivity shocks create between wage deflation and real adjustment.

Constrained Case The output gap increases but substantially less than under discretion when constrained. The nominal interest rate initially rises above zero before declining below zero and eventually returning to steady state. Nominal wages decrease permanently but less than unconstrained. Prices decrease initially then partially reflate. Wage inflation hits the lower bound for similar duration to discretion before returning gradually to zero.

The initial interest rate rise reveals how the constraint alters mechanical rule behaviour. In $i_t = \rho_i i_{t-1} + (1 - \rho_i)(\phi_\pi \pi_t^w + \phi_y x_t)$, when wage inflation hits the floor at -0.02, the $\phi_\pi \pi_t^w$ term provides less downward pull on rates than when wage inflation falls freely. Meanwhile, the rising output gap pushes rates upward through $\phi_y x_t$. The mechanical rule prescribes rate increases initially, not because of strategic calculation but because the fixed formula produces this result when one argument is constrained. As the output gap begins declining in subsequent periods, this feeds into the output term, eventually pulling rates below zero.

The constraint binds for similar duration to discretion, but the critical difference lies in the output gap magnitude. Under discretion, the constrained output gap actually exceeds the Taylor Rule's substantially. This reversal occurs through expectational feedback. Period-by-period re-optimisation means the discretionary policymaker repeatedly observes the binding constraint and targets aggressive negative wage inflation through large output gaps. This creates larger output gaps each period, raising inflation expectations and making the constraint bind more severely in a self-reinforcing cycle.

The Taylor Rule avoids this destabilisation precisely through its inflexibility. The rule cannot re-optimise each period in response to the binding constraint. It mechanically applies the same fixed coefficients regardless of whether the constraint binds. This prevents the optimisation trap: the rule cannot repeatedly attempt to induce wage deflation through output expansion because it is not solving an optimisation problem. When the constraint prevents wages from falling, the Taylor Rule simply continues applying its predetermined formula. The wage floor acts as an external brake on the volatility the rule would otherwise generate, limiting output gap fluctuations that dominate the rule's loss function. The rule's inability to distinguish between constrained and unconstrained states becomes an advantage: it cannot fight against the rigidity because it does not recognise the rigidity exists. This mechanical simplicity prevents the sophisticated but destabilising re-optimisation that undermines discretion under constraints.

5.1.4 Comparison

Optimal Commitment achieves superior performance through expectations management. By credibly promising future accommodation, it moderates current wage setting through the forward-looking Phillips curve, allowing gradual adjustment across both nominal and real margins. This intertemporal smoothing minimises both the magnitude and duration of constraint binding. The hump-shaped output gap reflects efficient spreading of adjustment costs over time rather than front-loading distortions. Discretion when unconstrained outperforms the Taylor Rule despite lacking commitment's forward guidance because period-by-period re-optimisation allows state-contingent balancing of the wage inflation-output gap trade-off. The discretionary policymaker solves for the optimal response given current conditions, whereas the Taylor Rule mechanically applies fixed coefficients calibrated for general conditions. However, when the constraint binds, discretion's re-optimisation prescribes aggressive easing to stimulate demand raises inflation expectations, which through the Phillips curve increases wage pressure, making the constraint bind more severely and requiring even larger output gaps. This creates a destabilising feedback loop absent under commitment.

The Taylor Rule's inflexibility prevents this optimisation trap. Unable to re-optimise in response to the binding constraint, the rule simply continues applying predetermined coefficients.

The wage floor acts as an external brake on the volatility the rule would otherwise generate, limiting output gap fluctuations. The rule cannot fight against the rigidity because it does not recognise the rigidity exists, allowing the constraint to perform its stabilising function unopposed. This demonstrates that under productivity shocks, sophisticated but non-credible optimisation can be worse than mechanical simplicity when constraints bind.

These findings extend Schmitt-Grohé and Uribe (2016), confirming commitment's superiority in New Keynesian models while revealing why alternatives fail differently. Discretion lacks the credibility essential for navigating constraints through expectations management. The Taylor Rule's effectiveness relative to discretion when constrained validates that policy design quality matters more than optimization sophistication. The modest positive wage inflation during commitment's adjustment phase supports Tobin's (1972) "grease the wheels" hypothesis: moderate inflation facilitates relative wage adjustments under downward rigidity. However, substantial productivity shocks reveal even optimal inflation management cannot fully eliminate asymmetric wage rigidity's real costs, making credible commitment frameworks essential for economies with strong institutional wage rigidities.

5.2 Negative Demand Shock

Next, I analyse the economy's response under the three different regimes to a negative demand shock d of magnitude -0.3 and standard deviation 0.01.

5.2.1 Optimal Policy

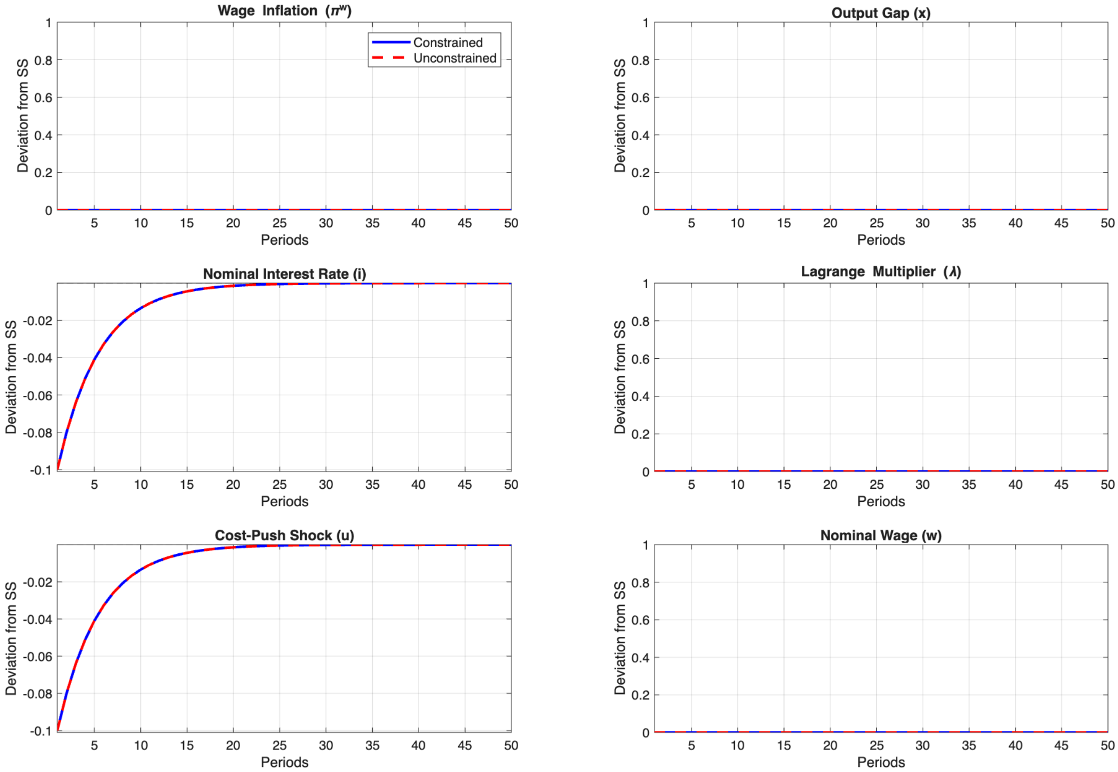

Figure 5: Impulse responses of variables with both optimal commitment and discretionary policy under a positive productivity shock. The shock follows an AR(1) with $\rho_d = 0.5$ and $\sigma = 0.1$

Optimal commitment and discretionary policy generate identical impulse responses in both constrained and unconstrained cases. These results demonstrate the 'divine coincidence' result detailed in Section 3. A negative demand shock moves wage inflation and the output gap in the same direction. The shock opens a negative output gap, which feeds through the wage NKPC, generating dis-inflationary pressure. The central bank thus faces no trade-off and simply adjusts the nominal interest rate to perfectly offset the shock's impact on aggregate demand through the Euler equation. Positive supply shocks, by contrast, create deflationary pressure through lower marginal costs yet raises output above potential, forcing the central bank to choose between stabilising wage inflation and closing the output gap, explaining the importance of credibility as seen in Section 5.1 above. The constraint never binds because policy ensures there is no deviation of wage inflation from steady state.

5.2.2 Wage Inflation Taylor Rule

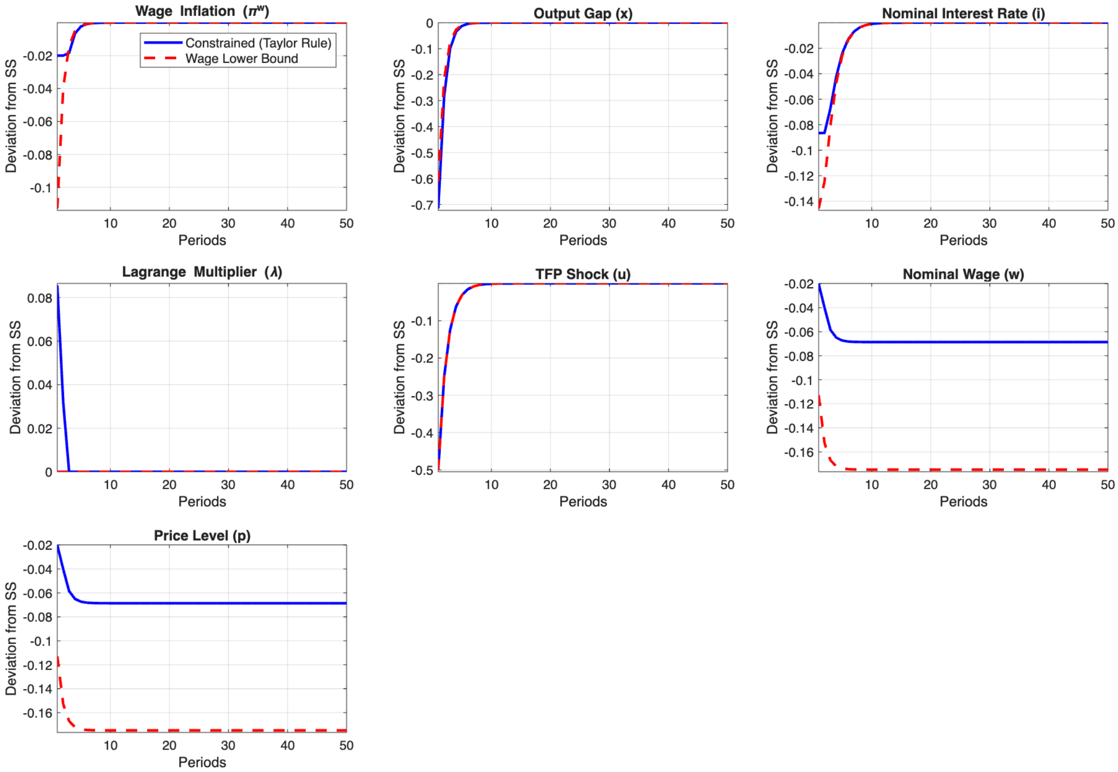

Figure 6: Impulse responses of variables with a wage inflation Taylor Rule given by $i_t = \rho_i i_{t-1} + (1-\rho_i)(\phi_\pi \pi_t^w + \phi_y x_t)$ where $\rho_i = 0.8$, $\phi_\pi = 1.5$, $\phi_y = 0.5$ under a negative demand shock. The shock follows an AR(1) with $\rho_d = 0.5$ and $\sigma = 0.1$.

Unconstrained Case The negative demand shock reduces aggregate consumption and output below potential, creating disinflationary pressure. Unlike optimal policies that exploit the divine coincidence to achieve zero losses by perfectly offsetting the natural rate movement, the Taylor Rule's fixed coefficients prescribe insufficient easing. The output gap falls on impact and returns monotonically to steady state. The nominal interest rate follows a similar path, cutting rates but not by enough to restore output to potential.

The Taylor Rule's mechanical formula $i_t = \rho_i i_{t-1} + (1-\rho_i)(\phi_\pi \pi_t^w + \phi_y x_t)$ cannot achieve the 'divine coincidence'. The rule responds to observed inflation and output deviations with predetermined weights, but these weights do not correspond to the precise adjustment needed to eliminate both gaps completely. The negative output gap drives wage inflation downward through the Phillips curve. Reduced marginal costs generate persistent deflationary dynamics that feed back into expectations. Cumulative periods of negative wage inflation cause nominal wages to decline permanently, as do prices. The permanently lower nominal wage and price levels quantify the Taylor Rule's failure to provide adequate stimulus, leaving real variables persistently below potential during the adjustment process.

Constrained Case When the constraint binds, the Taylor Rule's already inadequate response worsens significantly. The output gap decreases by more than in the unconstrained case and takes longer to return to steady state. The nominal interest rate decreases by less than unconstrained, creating an even larger gap between actual and required policy accommodation. Wage inflation hits the lower bound for two periods before gradually returning to steady state. Nominal wages and prices decrease permanently but by different amounts than without the constraint.

The constraint amplifies the Taylor Rule's failure through a distinct channel than under productivity shocks. With demand shocks, no trade-off exists between stabilization objectives—the constraint provides no beneficial nominal anchoring as it does under supply shocks. When wage inflation hits the floor, the rule's inflation term $\phi_\pi \pi_t^w$ generates less downward pressure on the policy rate since measured inflation cannot fall below the constraint. The rule therefore eases less aggressively than its formula would otherwise prescribe, providing even less stimulus when more is needed.

However, what matters for aggregate demand is the real interest rate, not the nominal rate. The slightly less negative nominal rate in the constrained case is partially offset by less negative expected inflation (since the constraint prevents wage inflation from falling as much), keeping real rates somewhat similar. This explains why output gaps do not differ as dramatically as might be expected. Nevertheless, the constraint still worsens outcomes because it forces more adjustment through quantities when the economy already lacks adequate demand stimulus. The permanently lower nominal wage and price levels are larger in magnitude than optimal, reflecting cumulative deflationary pressure the inadequate policy response failed to prevent. Importantly, the constraint binds for fewer periods than under equivalent-magnitude productivity shocks. This asymmetry reveals the fundamental difference between shock types: productivity shocks create genuine trade-offs where the constraint can be beneficial, while demand shocks do not and the constraint impedes adjustment.

5.2.3 Comparison

Optimal policy- both under commitment and discretion- achieves the ’divine coincidence’ detailed in Section 3, perfectly offsetting the shock by adjusting the nominal rate as the Euler equation demands. The Taylor rule lacks optimal policy’s flexibility and its formula instead prescribes a drop in interest rates insufficient to offset the demand shock, thus causing a negative output gap and consequently wage deflation through the Wage NKPC, propagating through expectations.

5.3 Central Bank Loss Comparison

This section evaluates the loss implications of alternative regimes in response to productivity and demand shocks of different sizes ranging from 10 to 50 %, both in unconstrained and constrained environments.

5.3.1 Positive TFP Shock

Table 2: Unconstrained Central Bank Losses by Shock Size (No Wage Floor)

Shock	Commitment	Discretion	Taylor Rule	Disc-Comm	Taylor-Comm
0.10	0.000536	0.001364	0.001410	+0.000828	+0.000874
0.15	0.001207	0.003071	0.003172	+0.001863	+0.001966
0.20	0.002214	0.005458	0.005640	+0.003213	+0.003426
0.25	0.003351	0.008528	0.008814	+0.005176	+0.005463
0.30	0.004826	0.012281	0.012700	+0.007454	+0.007875
0.35	0.006587	0.016714	0.017151	+0.010146	+0.010563
0.40	0.008579	0.021831	0.022566	+0.013252	+0.013987
0.45	0.010858	0.027630	0.028555	+0.016772	+0.017696
0.50	0.013406	0.034111	0.035252	+0.020706	+0.021847

Table 3: Loss Decomposition by Component (Unconstrained Case)

	Commitment				**Optimal Discretion**				**Taylor Rule**			
Shock	Inflation	Output	(%) π	(%) y	Inflation	Output	(%) π	(%) y	Inflation	Output	(%) π	(%) y
0.10	0.000526	0.000010	98.1	1.9	0.001138	0.000027	97.7	2.3	0.000010	0.001310	0.8	99.2
0.15	0.001119	0.000236	82.6	17.4	0.003098	0.000062	98.0	2.0	0.000224	0.002948	7.1	92.9
0.20	0.002129	0.000421	83.5	16.5	0.005351	0.000071	98.7	1.3	0.000394	0.005214	7.0	93.0
0.25	0.003285	0.000657	83.3	16.7	0.008361	0.000162	98.1	1.9	0.000653	0.008189	7.4	92.6
0.30	0.004731	0.000963	83.1	16.9	0.012063	0.000478	96.2	3.8	0.000937	0.011753	7.4	92.6
0.35	0.006440	0.001288	83.3	16.7	0.016367	0.000410	97.6	2.4	0.001224	0.016053	7.1	92.9
0.40	0.008413	0.001682	83.4	16.6	0.021406	0.000425	98.1	1.9	0.001569	0.021009	6.9	93.1
0.45	0.010645	0.002192	82.9	17.1	0.027088	0.000545	98.0	2.0	0.002020	0.026535	7.1	92.9
0.50	0.013142	0.002626	83.4	16.6	0.033442	0.000669	98.0	2.0	0.002494	0.032758	7.1	92.9

Unconstrained Central Bank Loss Table 2 establishes a clear hierarchy of central bank losses. Commitment dominates with losses roughly 2.5 times lower than either Discretion or the Taylor Rule. At 10 percent shocks, Commitment generates loss of 0.000536 versus 0.001364 for Discretion and 0.001410 for the Taylor Rule. At 50 percent shocks, these become 0.013406, 0.034111, and 0.035252. The Commitment-Discretion gap grows disproportionately from 0.000828 at 10 percent to 0.020706 at 50 percent, a 25-fold increase for a 5-fold shock increase, demonstrating that inability to manage expectations becomes increasingly costly as required adjustment intensifies. Discretion and the Taylor Rule generate nearly identical losses.

Table 3 reveals the source through loss composition. Commitment concentrates losses on inflation variance, accounting for 98 percent at 10 percent and 83 percent thereafter. Output variance remains minimal throughout. By credibly promising future accommodation, Commitment stabilizes current wage inflation through the forward-looking Phillips curve without large output movements. Wage setters moderate demands today because they internalize expected policy support, allowing adjustment with minimal distortions.

Discretion achieves similar composition with inflation losses at 96 to 98 percent, but absolute magnitudes are consistently 2 to 3 times larger despite comparable output volatility. At 10 percent, Discretion's inflation loss of 0.001138 exceeds Commitment's 0.000526 by 116 percent. Without credible commitment to future accommodation, private agents do not moderate current demands. Discretion operates entirely through contemporaneous trade-offs, lacking the temporal dimension that makes Commitment efficient.

The Taylor Rule exhibits inverted composition. Output variance dominates at 93 to 99 percent; inflation variance never exceeds 7 percent. As a fixed coefficient rule responding only to current observables, it cannot access expectations management or state-contingent re-optimization. The predetermined coefficients do not optimally balance nominal and real stabilization, generating substantial real volatility while tolerating large inflation fluctuations and forcing adjustment burden onto the real economy.

Table 4: Constrained Central Bank Losses by Shock Size (With Wage Floor)

Shock	Commitment	Discretion	Taylor Rule	Disc-Comm	Taylor-Comm
0.10	0.000319	0.002618	0.001410	+0.002298	+0.001091
0.15	0.000708	0.005771	0.003079	+0.004993	+0.002370
0.20	0.001367	0.010304	0.004789	+0.008667	+0.003423
0.25	0.002293	0.015173	0.006858	+0.012880	+0.004562
0.30	0.003496	0.022456	0.009283	+0.018959	+0.005788
0.35	0.004965	0.030534	0.012687	+0.025569	+0.007104
0.40	0.006706	0.039674	0.017077	+0.032968	+0.008499
0.45	0.008716	0.054051	0.018703	+0.041735	+0.009988
0.50	0.010996	0.062284	0.022556	+0.051288	+0.011559

Table 5: Cost of Constraint on Central Bank Loss

Shock	Commitment	Discretion	Taylor	Least Costly	Disc-Comm	Taylor-Comm
0.10	−0.000217	+0.001253	+0.000000	Commitment	+0.001470	+0.000217
0.15	−0.000498	+0.002631	−0.000095	Commitment	+0.003129	+0.000404
0.20	−0.000778	+0.004576	−0.000851	Taylor	+0.005354	−0.000073
0.25	−0.001015	+0.007090	−0.001549	Taylor	+0.008145	−0.000090
0.30	−0.001335	+0.010171	−0.003047	Taylor	+0.011501	−0.002076
0.35	−0.001603	+0.013820	−0.005205	Taylor	+0.015423	−0.003601
0.40	−0.001874	+0.018036	−0.007357	Taylor	+0.019910	−0.005483
0.45	−0.002142	+0.022826	−0.009851	Taylor	+0.024946	−0.007708
0.50	−0.002409	+0.028173	−0.012968	Taylor	+0.030582	−0.010828

Constrained Central Bank Loss The constraint reshapes the loss hierarchy. Table 4 shows Commitment maintains dominance while the suboptimal ranking reverses. At 10 percent, Discretion's constrained loss of 0.002618 exceeds the Taylor Rule's unchanged 0.001410 by 86 percent. At 50 percent, Discretion reaches 0.062284 while the Taylor Rule achieves 0.022556, making Discretion 176 percent worse. The constrained Commitment-Discretion gap expands from 0.002298 to 0.051288, amplifying the unconstrained gaps by 177 percent and 148 percent respectively. The constraint discriminates between policies that adapt strategically and those that cannot.

Table 5 reveals asymmetric constraint effects. Commitment experiences negative costs across all shocks, reaching 0.002409 at 50 percent. The Taylor Rule exhibits negative costs from 20 percent onward, reaching 0.012968 at 50 percent. Discretion suffers positive costs growing nonlinearly from 0.001253 to 0.028173, a 22-fold increase for a 5-fold shock increase. This accelerating pattern distinguishes Discretion from other regimes.

Commitment's lower loss under the constraint occurs because the wage floor prevents deflationary pressure, reducing required stabilization. The constraint performs part of the nominal anchoring function automatically. When wages cannot fall, measured wage inflation volatility declines. Commitment accepts greater real volatility to respect the rigidity, concentrating losses on the output gap. The ability to redirect stabilization toward available margins allows navigation of the constraint with limited cost increases. The negative constraint cost indicates the automatic nominal anchoring complements Commitment's strategic framework, improving measured performance relative to the unconstrained benchmark where more active intervention would be required.

The Taylor Rule's lower loss stems from its inability to distinguish between states. The rule applies fixed coefficients regardless of whether the floor binds. When the constraint prevents wage declines, it acts as a beneficial brake on excessive volatility the rule would otherwise generate through rigid interest rate responses. The wage floor automatically limits wage inflation fluctuations that dominate the loss function, reducing measured loss without strategic adjustment. The rule cannot exploit the constraint strategically but equally cannot fall into the trap of discretionary optimisation. Its simplicity becomes advantageous when the constraint disciplines policy automatically, preventing the rule from fighting against the rigidity.

Discretion's higher loss reflects destabilising expectational feedback. Each period, the discretionary policymaker observes the binding constraint and eases aggressively to reduce the output gap. However, period-by-period optimisation without commitment ability causes private agents to rationally increase inflation expectations. Through the Phillips curve, higher expected inflation increases current wage pressure, making the constraint bind more severely and requiring larger output movements. Further easing propagates a self-reinforcing cycle. Each attempt to improve the current situation through re-optimisation undermines future stabilisation by destabilising expectations, causing both loss components to deteriorate. The gap between Discretion and the Taylor Rule grows from 0.001208 to 0.039728, a 33-fold expansion, demonstrating that re-optimisation instability intensifies disproportionately as required adjustment increases.

5.3.2 Negative Demand Shock

Optimal Policy Both Optimal Commitment and Optimal Discretion achieve zero losses under negative demand shocks across all shock magnitudes. This perfect stabilisation reflects the divine coincidence in New Keynesian models without cost-push disturbances. Demand shocks create no inherent trade-off between inflation and output gap stabilization. The central bank can adjust the nominal interest rate to exactly offset the demand disturbance, maintaining both wage inflation and the output gap at their steady state values throughout the transition. With no deviations in either loss function argument, central bank losses remain precisely zero regardless of shock size or constraint presence. The wage floor never binds under optimal policy responding to demand shocks because the policy prevents any deflationary pressure from emerging. This stands in sharp contrast to productivity shocks, where the tension between stabilising wages pushed downward by increased efficiency and maintaining output near potential creates unavoidable welfare losses that differ substantially across policy regimes.

Taylor Rule The Taylor Rule's performance under negative demand shocks reveals different constraint interactions than under productivity shocks. Tables 6 and 7 present central bank losses and decomposition for demand shocks ranging from negative 10 to negative 50 percent. While optimal policies achieve zero losses through the divine coincidence, the Taylor Rule generates substantial welfare costs systematically amplified by the constraint despite minimal binding frequency.

Table 6: Total Central Bank Losses by Demand Shock Size (Taylor Rule)

Shock	Unconstrained	Constrained	Central Bank Loss	Cost (%)	W_Con/W_Unc
−0.10	0.00398	0.00410	0.00012	3.02	1.030
−0.15	0.00896	0.01004	0.00108	12.10	1.121
−0.20	0.01593	0.01869	0.00276	17.31	1.173
−0.25	0.02490	0.03008	0.00518	20.82	1.208
−0.30	0.03585	0.04422	0.00837	23.36	1.234
−0.35	0.04880	0.06110	0.01230	25.22	1.252
−0.40	0.06373	0.08071	0.01698	26.64	1.266
−0.45	0.08066	0.10305	0.02239	27.76	1.278
−0.50	0.09958	0.12813	0.02854	28.66	1.287

Table 6 shows the Taylor Rule generates losses scaling quadratically with shock magnitude, rising from 0.00398 at negative 10 percent to 0.09958 at negative 50 percent. While optimal policy achieves the divine coincidence by adjusting interest rates to exactly offset natural rate movements, the Taylor Rule's fixed coefficients prescribe insufficient monetary easing. Table 7 reveals output gap variance dominates throughout, contributing 94 percent of total losses. At negative 10 percent, output losses of 0.00373 dwarf inflation losses of 0.00025. At negative 50 percent, output losses reach 0.09321 while inflation losses remain modest at 0.00637. The rule's inadequate easing fails to stimulate aggregate demand sufficiently, leaving output persistently below potential while deflationary dynamics build.

Under the wage floor constraint, Table 6 shows losses systematically exceed unconstrained levels. At negative 10 percent, the constrained loss of 0.00410 represents a 3.02 percent deterioration. This gap accelerates to 28.66 percent at negative 50 percent. The ratio of constrained to unconstrained loss rises monotonically from 1.030 to 1.287.

Table 7: Loss Decomposition by Component (Demand Shocks, Taylor Rule)

	Unconstrained			Constrained			Additional Loss	
Shock	Inflation	Output	Total	Inflation	Output	Total	$\Delta(\pi^w)$	$\Delta(x)$
−0.10	0.00025	0.00373	0.0040	0.00020	0.00390	0.0041	−0.00005	+0.00017
−0.15	0.00057	0.00838	0.0090	0.00020	0.00984	0.0100	−0.00037	+0.00146
−0.20	0.00102	0.01491	0.0159	0.00020	0.01849	0.0187	−0.00082	+0.00358
−0.25	0.00159	0.02330	0.0249	0.00020	0.02988	0.0301	−0.00139	+0.00657
−0.30	0.00229	0.03356	0.0358	0.00020	0.04402	0.0442	−0.00209	+0.01046
−0.35	0.00312	0.04567	0.0488	0.00020	0.06090	0.0611	−0.00292	+0.01523
−0.40	0.00408	0.05966	0.0637	0.00020	0.08051	0.0807	−0.00388	+0.02085
−0.45	0.00516	0.07550	0.0807	0.00020	0.10285	0.1031	−0.00496	+0.02735
−0.50	0.00637	0.09321	0.0996	0.00020	0.12793	0.1281	−0.00617	+0.03471

Table 7 exposes the mechanism. In the constrained case, inflation losses collapse to 0.00020 regardless of shock size as the wage floor prevents measured inflation from falling below zero, while output losses explode from 0.00390 to 0.12793. Additional losses show inflation declining by 0.00005 to 0.00617, but output losses increasing by 0.00017 to 0.03471. This asymmetric pattern reveals the constraint forces more adjustment through the output gap even when not binding. Private agents know wages cannot fall easily if severe deflation emerges, affecting wage setting and employment throughout the transition. This anticipation shifts the Phillips curve trade-off, requiring larger output movements to generate disinflation. The Taylor Rule, responding mechanically to observed variables, cannot recognize how the constraint's shadow has altered transmission.

This deterioration contrasts starkly with productivity shocks. Supply shocks push wages downward directly through reduced marginal costs, and the constraint acts as a beneficial brake on deflationary pressure, reducing losses. Demand shocks create no inherent stabilization trade-off but generate asymmetries where downward wage flexibility is impaired while upward flexibility remains. The constraint provides no stabilizing benefit and instead worsens losses.

The loss pattern in Table 6 quantifies this. At negative 10 percent, the modest shadow effect generates only 3.02 percent loss. At negative 50 percent, despite identical binding frequency, the loss reaches 28.66 percent. This superlinear relationship demonstrates constraint effects operate primarily through expectational channels. Each marginal shock increase amplifies the shadow effect disproportionately, as larger potential deflations make the constraint's presence more salient even when rarely binding. The Taylor Rule's simplicity prevents discretionary optimisation traps under supply shocks but becomes a liability under demand shocks as it is unable to produce the 'divine coincidence'.

5.4 Optimal Parameters

The preceding loss analysis established a clear performance hierarchy across policy regimes and demonstrated how binding constraints reshape this ranking. However, understanding which policy performs best under given conditions is distinct from understanding how to design that policy optimally. This section examines the optimal parameter configurations for each regime across both constrained and unconstrained environments.

Tables 8 through 13 present the results of comprehensive grid searches over the parameter spaces for Optimal Commitment, Optimal Discretion, and the Taylor Rule. For Commitment and Discretion, the search spans wage flexibility parameters and output gap preferences. For the Taylor Rule, the search covers interest rate smoothing, inflation responsiveness, and output gap responsiveness. Each regime's optimal configuration is identified for both the presence and absence of the wage inflation constraint.

5.4.1 Positive TFP Shock

Table 8: Optimal Commitment Parameters and Central Bank Loss (TFP Shocks)

	Constrained	Unconstrained
Optimal κ_w	0.1	0.1
Optimal α	0.1	0.1
Minimum Central Bank Loss	0.001385	0.001892

Table 9: Optimal Commitment Central Bank Loss under $u = 0.3$

κ_w	α	γ	$W_{\text{Constrained}}$	$W_{\text{Unconstrained}}$	$W_{\text{Difference}}$	W_{Ratio}
0.1	0.1	1	0.0013847	0.0018923	-0.00050762	0.73174
0.1	0.3	3	0.0026897	0.0037323	-0.0010426	0.72065
0.1	0.9	9	0.0045523	0.0061974	-0.0016451	0.73455
0.3	0.1	0.33333	0.0029312	0.0033315	-0.00010035	0.96669
0.3	0.3	1	0.007203	0.0074617	-0.00025871	0.96533
0.3	0.9	3	0.0159	0.017031	-0.0011303	0.93363
0.5	0.1	0.2	0.0035458	0.0035597	-0.00001388	0.9961
0.5	0.3	0.6	0.0093253	0.0089781	0.00034718	1.0387
0.5	0.9	1.8	0.022803	0.022335	0.00046746	1.0209
0.7	0.1	0.14286	0.0038636	0.0038783	-0.00001465	0.99622
0.7	0.3	0.42857	0.010499	0.010018	0.0004816	1.0481
0.7	0.9	1.2857	0.027102	0.025034	0.0020684	1.0826
0.9	0.1	0.11111	0.0040524	0.0040702	-0.00001781	0.99562
0.9	0.3	0.33333	0.011226	0.010796	0.00042984	1.0398
0.9	0.9	1	0.029967	0.027284	0.0026829	1.0983

Optimal Commitment Optimal Commitment exhibits optimal parameter invariance across constraint regimes. Table 8 shows the loss-minimising configuration remains $\kappa_w^* = 0.1$, $\alpha^* = 0.1$ whether the wage floor is present or absent. The constrained loss of 0.001385 is 26.8 percent lower than the unconstrained loss of 0.001892, generating a loss of 0.000508. This improvement under the constraint confirms that credible forward-looking policy can exploit nominal rigidities to enhance stabilization outcomes.

The optimal calibration favoring low wage flexibility and minimal output weight reflects Commitment's strategic advantage. With $\kappa_w^* = 0.1$, wages adjust gradually, creating a stable nominal anchor that amplifies the power of expectations management. With $\alpha^* = 0.1$, the policy signals overwhelming priority for nominal wage inflation stability, which itself becomes a stabilisation tool through its effect on wage-setter expectations. By credibly promising future accommodation when productivity shocks push wages downward, the central bank prevents deflationary spirals before they emerge, allowing the constraint to act as a beneficial brake on wage inflation.

Table 9 reveals the cost of the constraint varies systematically with parameter choices. At the optimum, the constraint reduces losses by 0.000508. This improvement intensifies as output weight increases. At $\alpha = 0.3$, the constraint benefit reaches 0.001043, and at $\alpha = 0.9$, it reaches 0.001645. Higher output weights force more aggressive real stabilization in the unconstrained case, creating volatility that the wage floor helps discipline in the constrained case. The constraint prevents the policy from tolerating excessive wage deflation in pursuit of output stabilization, automatically enforcing nominal anchoring that complements Commitment's expectations management.

Increasing wage flexibility while maintaining low output weight diminishes but preserves the constraint's loss benefit. At $\kappa_w = 0.3$, the benefit falls to 0.000100, and at $\kappa_w = 0.9$, it reaches 0.000178. Greater wage flexibility reduces the frequency and severity of constraint binding, limiting its stabilizing effect. Only the extreme combination of high wage flexibility and high output weight generates positive constraint costs. At $\kappa_w = 0.9$ and $\alpha = 0.9$, the constraint imposes a loss of 0.002683, as the policy's aggressive real stabilization repeatedly collides with the wage floor, creating destabilizing feedback despite Commitment's expectations management capacity.

The parameter invariance and loss reduction under constraints demonstrate Commitment's robustness. The same calibration that minimizes losses when constraints bind also minimizes losses when they do not, because the policy's framework already incorporates optimal responses to all contingencies through expectations management. The constraint's benefit at the optimum occurs because the wage floor prevents deflationary pressures that would otherwise require active policy offset, effectively performing part of the nominal anchoring function automatically. Commitment's credible promises about future accommodation interact constructively with this automatic anchoring, using the constraint to enhance rather than impede stabilization.

Table 10: Optimal Discretion Parameters and Central Bank Loss (TFP Shocks)

	Constrained	Unconstrained
Optimal κ_w	0.1	0.9
Optimal α	0.1	0.1
Minimum Central Bank Loss	0.004650	0.004846

Table 11: Optimal Discretion Central Bank Loss under $u = 0.3$

κ_w	α	γ	$W_{\text{Constrained}}$	$W_{\text{Unconstrained}}$	$W_{\text{Difference}}$	W_{Ratio}
0.1	0.1	1	0.0046501	0.0066255	-0.0019754	0.70185
0.1	0.3	3	0.01355	0.010887	0.0026633	1.2446
0.1	0.9	9	0.040251	0.013375	0.026877	3.0095
0.3	0.1	0.33333	0.0046833	0.0066794	-0.0019961	0.70106
0.3	0.3	1	0.01365	0.0235	-0.0098498	0.58806
0.3	0.9	3	0.04055	0.05963	-0.01908	0.68083
0.5	0.1	0.2	0.00469	0.0055095	-0.00081952	0.85125
0.5	0.3	0.6	0.01367	0.020353	-0.0066827	0.67166
0.5	0.9	1.8	0.04061	0.070638	-0.030028	0.5749
0.7	0.1	0.14286	0.0046929	0.0050544	-0.00036159	0.92846
0.7	0.3	0.42857	0.013679	0.017789	-0.0041107	0.76892
0.7	0.9	1.2857	0.040636	0.0661	-0.025464	0.61476
0.9	0.1	0.11111	0.0046944	0.0048459	-0.00015149	0.96874
0.9	0.3	0.33333	0.013683	0.016335	-0.002652	0.83765
0.9	0.9	1	0.04065	0.060115	-0.019465	0.67621

Optimal Discretion Optimal Discretion exhibits parameter sensitivity that contrasts sharply with Commitment's invariance. Table 10 shows the loss-minimizing configuration shifts between constraint regimes. Under the constraint, the optimum occurs at $\kappa_w^* = 0.1$, $\alpha^* = 0.1$ with loss of 0.004650. Without the constraint, the optimum shifts to $\kappa_w^* = 0.9$, $\alpha^* = 0.1$ with loss of 0.004846. This parameter dependence reveals that discretionary policy cannot maintain optimal calibration across different economic states, requiring adjustment based on whether constraints are expected to bind.

The constrained optimum favoring minimal wage flexibility differs from the unconstrained optimum favoring maximal flexibility. Low κ_w in the constrained case limits the frequency of wage adjustment, reducing opportunities for the constraint to bind and creating a disciplining mechanism that partially compensates for Discretion's inability to commit. High κ_w in the unconstrained case allows rapid wage adjustment without risk of binding, permitting the discretionary policymaker to respond more flexibly to shocks. The optimal output weight remains at 0.1 across both regimes, indicating that strong nominal anchoring through the loss function provides some stability regardless of constraint presence.

Table 11 reveals the constraint generates improvements across most parameter combinations, with losses ranging from 0.000152 at high wage flexibility to 0.025449 at the extreme configuration of maximum wage flexibility and output weight. At the constrained optimum, the constraint improves outcomes by 0.001975 relative to the same parameters in the unconstrained environment. This pattern indicates the wage floor acts as a beneficial commitment device for discretionary policy, preventing the excessive wage deflation that myopic period-by-period optimization would otherwise tolerate.

The reduction in losses intensifies as wage flexibility increases while maintaining low output weight. At $\kappa_w = 0.3$, the constraint benefit reaches 0.002596, and at $\kappa_w = 0.5$, it reaches 0.001441. Greater wage flexibility without the constraint allows discretionary policy to generate more frequent adjustments that accumulate into larger nominal instability. The wage floor interrupts this dynamic, forcing stabilization through alternative channels. The relationship becomes non-monotonic at very high flexibility, with the benefit declining to 0.000152 at $\kappa_w = 0.9$, as the constraint binds less frequently and provides diminishing disciplining effect.

Increasing output weight amplifies the constraint's benefit. At $\alpha = 0.3$ with moderate wage flexibility, the constraint benefit reaches 0.004111, and at the extreme configuration of $\kappa_w = 0.9$ and $\alpha = 0.9$, it reaches 0.025449. Higher output weights induce more aggressive real stabilization under discretion, creating volatility that the wage floor helps contain. The constraint prevents the policy from sacrificing nominal stability excessively in pursuit of output objectives, automatically imposing discipline that discretionary optimization cannot self-impose.

The divergence between constrained and unconstrained optimal parameters quantifies Discretion's design fragility. The unconstrained optimum at high wage flexibility performs poorly when the constraint is present, generating losses of 0.004694 compared to the constrained optimum of 0.004650. Conversely, the constrained optimum at low wage flexibility performs worse without the constraint, generating losses of 0.006626 compared to the unconstrained optimum of 0.004846. A discretionary policymaker must forecast which state will prevail and calibrate accordingly, or accept suboptimal performance when the forecast proves incorrect.

Comparing Tables 8 and 10 demonstrates the gap in performance between the regimes. At their respective optima, Commitment achieves constrained loss of 0.001385 while Discretion reaches 0.004650, making Discretion 236 percent worse despite both selecting identical parameters in the constrained case. In the unconstrained case, Commitment's loss of 0.001892 compares to Discretion's 0.004846, a 156 percent performance gap. This persistent disadvantage quantifies the cost of being unable to manage expectations. Even when the constraint acts as a beneficial commitment device for Discretion, reducing its costs relative to the unconstrained case, the regime cannot approach Commitment's strategic superiority. The constraint helps discipline discretionary policy but cannot substitute for genuine forward-looking credibility.

Table 12: Optimal Taylor Rule Parameters and Central Bank Loss (TFP Shocks)

	Constrained	Unconstrained
Optimal ρ_i	0.1	0.5
Optimal ϕ_π	1.0	1.0
Optimal ϕ_y	1.0	0.5
Minimum Central Bank Loss	0.000854	0.007876

Table 13: Taylor Rule Central Bank Loss under $u = 0.3$

Parameters			**Constrained**	**Unconstrained**
ρ_i	ϕ_π	ϕ_y	$W_{\text{Constrained}}$	$W_{\text{Unconstrained}}$
0.1	1.0	1.0	0.000854	0.009828
0.5	1.0	1.0	0.001392	0.009190
1.0	1.0	1.0	0.004115	0.011016
0.1	1.5	1.0	0.001135	0.008879
0.5	1.5	1.0	0.001354	0.008156
1.0	1.5	1.0	0.011051	0.013146
0.1	2.0	1.0	0.001466	0.009050
0.5	2.0	1.0	0.001718	0.008294
1.0	2.0	1.0	0.018268	0.018268

Taylor Rule The optimal Taylor Rule design exhibits environment-specific requirements distinct from both Commitment's invariance and Discretion's parameter sensitivity. Table 12 shows optimal parameters differ substantially between constraint regimes. When the wage floor is present, the optimal configuration is $\rho_i^* = 0.1$, $\phi_\pi^* = 1.0$, and $\phi_y^* = 1.0$, yielding loss of 0.000854. In the unconstrained case, the optimum shifts to $\rho_i^* = 0.5$, $\phi_\pi^* = 1.0$, and $\phi_y^* = 0.5$, with loss of 0.007876.

When constraints bind, the optimal rule favours minimal smoothing and aggressive output response. Low smoothing allows immediate policy adjustment, preventing delayed responses that amplify real volatility when wages cannot adjust downward. High output response of 1.0 reflects that the constraint forces all equilibration through quantities, requiring proportional policy response. In the unconstrained environment, optimal smoothing rises to 0.5 and output response falls to 0.5, as the economy's flexibility across adjustment margins permits gradual policy changes and reduced real stabilisation. The inflation response remains at 1.0 across both environments, satisfying the Taylor principle requirement for determinacy.

Table 13 exposes miscalibration risks. In the constrained environment, increasing smoothing from 0.1 to 0.5 raises loss by 63 percent, and to 1.0 raises loss nearly fivefold. High smoothing prevents rapid adjustments necessary when constraints eliminate nominal flexibility. The combination of high smoothing and aggressive inflation response proves particularly damaging, with losses reaching 0.018268 at $\rho_i = 1.0$ and $\phi_\pi = 2.0$. In the unconstrained case, reducing smoothing from the optimal 0.5 to 0.1 increases loss by 25 percent, indicating some smoothing becomes valuable when multiple adjustment channels exist.

This state dependence reveals the limitation of smple rules. Unlike Commitment, which maintains optimal design across environments through expectations management, the Taylor Rule requires explicit recalibration based on expected constraint status. The constrained-optimal rule achieves remarkably low loss of 0.000854, outperforming both Commitment's 0.001385 and Discretion's 0.004650. However, this configuration generates losses of 0.009828 unconstrained, substantially worse than the unconstrained optimum of 0.007876 and inferior to both Commitment's 0.001892 and Discretion's 0.004846. In the constrained case, the optimally-calibrated Taylor Rule performs 38 percent better than Commitment and 445 percent better than Discretion. In the unconstrained case, this ranking inverts completely, with the Taylor Rule's loss exceeding Commitment's by 316 percent and Discretion's by 63 percent. The policymaker must forecast which state will prevail and calibrate accordingly, or accept suboptimal performance across states.

5.4.2 Negative Demand Shock

Optimal Policy Both Optimal Commitment and Optimal Discretion achieve zero losses across all shock magnitudes and parameter configurations, reflecting the divine coincidence in New Keynesian models without cost-push disturbances. When policy faces no stabilization trade-offs, calibration of preference weights and structural parameters becomes immaterial. This contrasts with productivity shocks, where parameter choice critically determines welfare through its effect on navigating the trade-off between stabilizing wage deflation and maintaining output near potential. Because optimal policy prevents any deflationary pressure from emerging, the constraint never binds. The distinction between commitment and discretion vanishes entirely, as the absence of expectational trade-offs eliminates any advantage to forward-looking policy. Both regimes achieve identical zero losses, confirming that the commitment premium observed under productivity shocks stems specifically from managing expectations when stabilization trade-offs exist, not from general superiority across all shock types.

Table 14: Optimal Taylor Rule Parameters and Central Bank Loss (Demand Shocks)

	Constrained	Unconstrained
Optimal ρ_i	0.1	0.1
Optimal ϕ_π	2.0	2.0
Optimal ϕ_y	1.0	1.0
Minimum Central Bank Loss	0.026781	0.022185

Table 15: Taylor Rule Central Bank Loss under $d = -0.3$

Parameters			**Constrained**	**Unconstrained**
ρ_i	ϕ_π	ϕ_y	$W_{\text{Constrained}}$	$W_{\text{Unconstrained}}$
0.1	1.0	1.0	0.029737	0.028516
0.1	1.5	1.0	0.028222	0.025040
0.1	2.0	1.0	0.026781	0.022185
0.5	1.0	1.0	0.041480	0.038030
0.5	1.5	1.0	0.040121	0.035428
0.5	2.0	1.0	0.049294	0.041010
0.1	1.0	0.5	0.064141	0.058010
0.1	1.5	0.5	0.060226	0.048106
0.1	2.0	0.5	0.056515	0.040620

Taylor Rule The optimal Taylor Rule design for demand shocks exhibits parameter invariance different from productivity shocks. Table 14 shows the loss-minimizing configuration is identical across constrained and unconstrained environments: $\rho_i^* = 0.1$, $\phi_\pi^* = 2.0$, and $\phi_y^* = 1.0$. The constrained loss of 0.026781 exceeds the unconstrained loss of 0.022185, generating a cost of 0.004596. This deterioration under constraints opposes the improvement the Taylor Rule experiences under productivity shocks.

The invariant optimal configuration reflects the symmetric nature of demand shocks. Unlike productivity shocks that push wages downward directly through reduced marginal costs, demand shocks operate through aggregate demand channels that create no inherent directional bias. The constraint's presence or absence does not fundamentally alter which policy aggressiveness is required. Minimal smoothing of 0.1 allows rapid policy adjustment across both environments. Aggressive inflation response of 2.0, double the productivity shock optimum of 1.0, reflects the need to combat deflationary pressure more forcefully when the Taylor Rule cannot achieve the divine coincidence that optimal policies deliver. High output response of 1.0 maintains real stabilization across both constraint states.

Table 15 reveals that losses under demand shocks substantially exceed those under productivity shocks across all parameter combinations. The optimal constrained loss of 0.026781 is 31 times larger than the productivity shock constrained optimum of 0.000854. This difference quantifies the cost of the Taylor Rule's inability to achieve perfect stabilization when no inherent policy trade-offs exist. While optimal policies exploit the divine coincidence to deliver zero losses, the simple rule's fixed coefficients prescribe inadequate easing, generating persistent output gaps and deflationary dynamics.

The cost of the constraint remains positive across all examined parameter combinations for demand shocks, unlike productivity shocks where negative costs emerged for most calibrations. At the optimum, the constraint increases losses by 20.7 percent. This systematic deterioration stems from the constraint's shadow effect on transmission. Even rare binding alters private sector wage-setting and employment decisions throughout the transition, shifting the Phillips curve and requiring larger output movements to generate disinflation. The Taylor Rule cannot recognize or adjust for how the constraint has altered the instrument-target relationship, delivering identical responses regardless of transmission changes.

The parameter invariance under demand shocks simplifies Taylor Rule design relative to productivity shocks but cannot overcome the fundamental limitation. A single calibration performs optimally across both constraint states, eliminating the forecast dependence required under supply disturbances. However, this unified optimum delivers losses orders of magnitude worse than optimal policies achieve. The simple rule's inability to preemptively offset natural rate movements dooms it to generating substantial costs under demand shocks regardless of calibration quality. The contrast with productivity shocks, where careful state-contingent design allowed the Taylor Rule to nearly match or exceed Optimal Discretion's performance, demonstrates that the value of simplicity depends on the shock type.

5.4.3 Design Principles for Constrained Environments

The optimal parameter analysis reveals fundamental distinctions in policy design quality that depend on shock type. Under positive TFP shocks, Optimal Commitment demonstrates robustness through parameter invariance. The same configuration minimises losses whether constraints bind or not, as credible forward-looking policy manages expectations optimally across all contingencies. The Taylor Rule exhibits state-dependent design requirements, with optimal calibration varying systematically between constrained and unconstrained environments. A rule optimised for constrained conditions achieves strong performance, surpassing even Optimal Discretion, but performs poorly in unconstrained states. Optimal Discretion prove fragile, with complex, non-monotonic loss landscapes highly sensitive to parameters and structural conditions. The interaction between myopic optimisation and binding constraints produces explosive instability, with constrained losses exceeding the optimally-calibrated Taylor Rule by more than fivefold despite Discretion's theoretical advantage of continuous re-optimisation.

Under negative demand shocks, the results change. Both Optimal Commitment and Optimal Discretion achieve zero losses across all parameter configurations through the divine coincidence, rendering calibration choices immaterial and eliminating any distinction between forward-looking and period by period optimisation. The Taylor Rule exhibits parameter invariance across constraint states, maintaining identical optimal calibration whether the wage floor is present or absent. However, this unified optimum delivers losses orders of magnitude worse than optimal policies, as the simple rule cannot achieve the perfect natural rate offset that eliminates all stabilisation costs. The constraint systematically deteriorates Taylor Rule performance through shadow effects on transmission, contrasting sharply with the improvements the rule experiences under productivity shocks.

These findings extend established results in the optimal monetary policy literature while revealing new shock-type dependencies. The dominance of commitment over discretion under productivity shock constraints confirms EHL, who demonstrate that credible commitment becomes increasingly valuable when nominal rigidities bind. The ranking reversal between discretion and simple rules parallels Schmitt-Grohe and Uribe (2016), who show that well-designed instrument rules can outperform discretionary optimisation under occasionally binding constraints.

Most distinctively, this analysis demonstrates that optimal design principles depend fundamentally on shock type. Under productivity shocks creating inherent stabilisation trade-offs, commitment's expectations management delivers decisive advantages, discretion's myopic re-optimisation generates instability through expectational feedback, and simple rules require sophisticated state-contingent calibration but can achieve strong performance. Under demand shocks eliminating stabilisation trade-offs, optimal policies exploit the divine coincidence for perfect stabilisation regardless of commitment capacity, while simple rules suffer fundamental limitations from their inability to preemptively offset natural rate movements. This shock-type dependence of design principles provides new insight beyond the literature's traditional focus on time-consistency problems, revealing that the value of commitment, the danger of discretion, and the requirements for Taylor rule design all vary systematically with the nature of disturbances facing the economy.

5.5 Limitations

An important limitation must be acknowledged. The central bank loss function provides a tractable metric for comparing policy performance but does not directly measure household welfare, particularly when the wage floor binds. The quadratic loss function penalises squared deviations of wage inflation and the output gap, assuming these deviations matter symmetrically and that welfare costs can be captured through second-order approximation around steady state.

The negative constraint costs for Commitment and the Taylor Rule illustrate this divergence. These regimes achieve lower measured central bank losses when the wage floor is present because the constraint automatically limits measured wage inflation volatility. However, this does not necessarily imply households are better off, as the constraint forces adjustment through employment rather than nominal wages, potentially creating larger real disruptions.

Despite this limitation, the central bank loss function remains valuable for comparing regime performance. It provides a consistent metric across all regimes, allowing systematic evaluation of how different policy frameworks navigate the constraint. The ranking reveals genuine insights about expectations management, commitment value, and the interaction between policy sophistication and binding rigidities. However, the quantitative magnitudes should be interpreted as measures of performance relative to the central bank's stabilisation objectives rather than precise measures of household welfare. A complete welfare analysis would require explicit micro-foundation of the constraint, careful accounting of distributional effects, and validation that the quadratic approximation remains adequate when the constraint binds.

6 Extensions and Future Research

The most natural extension incorporates price stickiness alongside wage rigidity, creating a dual-friction model. While flexible prices enable clean identification of DNWR's policy implications, real-world central banks face both constraints simultaneously. Dual rigidity creates additional trade-offs: stabilizing price inflation may conflict with stabilizing wage inflation when the wage floor binds. This would reveal whether commitment's welfare advantage persists when policy must balance competing objectives, or whether managing price dynamics alongside constrained wage dynamics fundamentally alters optimal design. The flexible-price baseline may overstate DNWR costs by permitting instantaneous real adjustment, or understate them if sticky prices propagate wage rigidity distortions more persistently.

Open-economy dynamics provide a second valuable extension. For small open economies like Australia, exchange rate flexibility offers an additional adjustment channel: real depreciation can substitute for real wage cuts that nominal floors prevent. However, exchange rate pass-through complicates transmission, potentially amplifying inflation persistence under wage rigidity. An open-economy model would clarify whether optimal policy should exploit the exchange rate channel more aggressively under DNWR and how terms-of-trade shocks interact with domestic wage constraints.

Incorporating heterogeneous agents would capture distributional implications absent from the representative-agent framework. DNWR likely binds heterogeneously across worker types: low-wage workers, young entrants, and non-unionized sectors face more frequent constraint binding. When wages cannot adjust downward, employment bears the burden unequally. Aggregate welfare calculations conceal these distributional costs, potentially misrepresenting optimal policy if central bank objectives include employment inequality concerns. This proves particularly relevant for Australia where Modern Awards create heterogeneous wage floors across industries, generating differential constraint binding patterns.

Formal Bayesian estimation would strengthen quantitative conclusions by disciplining key parameters with Australian data. The calibration strategy relies on literature consensus, but critical parameters—particularly the wage Phillips curve slope and Rotemberg adjustment cost—remain empirically uncertain. Estimation would provide posterior distributions quantifying parameter uncertainty and enable formal model comparison: does incorporating the DNWR constraint improve empirical fit relative to standard New Keynesian models?

Finally, alternative constraint specifications merit exploration. Reality may feature time-varying floors adjusting with inflation expectations or institutional changes, state-dependent floors that tighten during expansions, or probabilistic constraints where wage cuts become increasingly costly but not literally impossible. Comparing optimal policy across these specifications would establish whether core findings are robust to constraint modeling details.

7 Conclusion

This thesis set out to determine how monetary policy should be optimally designed when DNWR acts as a binding constraint on wage inflation in the economy.

The results demonstrate that optimal commitment policy in this environment targets wage inflation based on changes in the output gap. The value of this commitment framework, relative to discretionary optimization or simple rules, depends on the nature of macroeconomic shocks facing the economy.

Under positive productivity shocks, commitment achieves the lowest losses for the policymaker through expectations management. The constraint improves measured outcomes by providing wage inflation anchoring that benefits credible forward-looking policy. Discretion suffers severe performance deterioration under the constraint, with losses growing disproportionately large as period-by-period optimisation destabilises expectations. A simple wage inflation Taylor Rule outperforms discretionary policy when the constraint binds as the constraint limits excessive volatility in wage inflation. Under negative demand shocks, both commitment and discretion exploit the 'divine coincidence' to achieve negligible losses regardless of constraint presence, as demand shocks create no inherent trade-off between inflation and output gap stabilisation. The Taylor Rule is unable to perfectly offset demand shocks through adjusting the nominal rate of interest by virtue of its fixed coefficients. Furthermore, the constraint worsens its performance as it limits the ability of the Taylor Rule to ease monetary policy.

Further analysis reveals that under positive productivity shocks, commitment exhibits complete parameter invariance across constraint states, while discretion requires different optimal calibrations depending on whether the constraint binds. The Taylor Rule also demands state-contingent design, with a rule calibrated for constrained conditions achieving remarkably low losses yet performing poorly in absence of the constraint. Under demand shocks, optimal policy parameters become irrelevant due to the divine coincidence whilst the Taylor Rule delivers non-negligible losses.

The results demonstrate that when productivity shocks dominate, institutional frameworks supporting credible commitment deliver substantial value whilst discretionary frameworks generate severe instability which simple rules calibrated appropriately can outperform. When demand shocks dominate, however, priority shifts towards achieving the flexible policy responses required to offset the demand shock through changes in the nominal interest rate.

Further research should also incorporate price rigidities and heterogeneous agents. Deriving the household welfare function would validate whether the central bank loss patterns documented here correspond to genuine improvements in economic welfare or merely reflect better performance against the central bank's stabilisation objectives. Furthermore, Occbin assumes perfect foresight regarding the sequence of constraint regimes, abstracting from expectational uncertainty or learning dynamics that could emerge if agents were uncertain about the duration or recurrence of binding episodes. Future research could investigate the extent to which perfect foresight is relevant to my findings.

Ultimately, this thesis demonstrates that the presence of DNWR reshapes the design of optimal monetary policy in ways that depend critically on the economic environment. The value of credibility compared to simple rules and discretion varies systematically with shock type and constraint status. Recognising this state-dependence is essential for policymakers seeking to design robust frameworks that navigate institutional wage-setting realities effectively.

References

[1] Akerlof, G. A., Dickens, W. T., & Perry, G. L. (1996). The macroeconomics of low inflation. *Brookings Papers on Economic Activity*, *1996*(1), 1–76. https://doi.org/10.2307/2534646

[2] Barattieri, A., Basu, S., & Gottschalk, P. (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics*, *6*(1), 70–101. https://doi.org/10.1257/mac.6.1.70

[3] Benigno, P., & Ricci, L. A. (2011). The inflation-output trade-off with downward wage rigidities. *American Economic Review*, *101*(4), 1436–1466. https://doi.org/10.1257/aer.101.4.1436

[4] Bishop, J., & Cassidy, N. (2017). Insights into low wage growth in Australia. *RBA Bulletin*, March, 13–20. https://www.rba.gov.au/publications/bulletin/2017/mar/

[5] Bishop, J., & Lancaster, D. (2023). New timely indicators of wages growth. *RBA Bulletin*, September, 1–12. https://www.rba.gov.au/publications/bulletin/2023/sep/

[6] Brassil, A., Gibbs, C., & Ryan, N. (2025). Flexible average inflation targeting and the role of expectations. *RBA Research Discussion Paper*. Reserve Bank of Australia.

[7] Buchanan, J., & Kells, S. (2024). Why real wages in Australia have fallen. University of Sydney Business School.

[8] Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, *12*(3), 383–398. https://doi.org/10.1016/0304-3932(83)90060-0

[9] Clarida, R., Galí, J., & Gertler, M. (1999). The science of monetary policy: A new Keynesian perspective. *Journal of Economic Literature*, *37*(4), 1661–1707. https://doi.org/10.1257/jel.37.4.1661

[10] Coibion, O., Gorodnichenko, Y., & Weber, M. (2021). Labor markets during the COVID-19 crisis: A preliminary view. *Brookings Papers on Economic Activity*, *Spring*, 1–63. https://doi.org/10.1353/eca.2021.0008

[11] Daly, M. C., & Hobijn, B. (2014). Downward nominal wage rigidities bend the Phillips curve. *Journal of Money, Credit and Banking*, *46*(S2), 51–93. https://doi.org/10.1111/jmcb.12152

[12] Daly, M. C., Hobijn, B., Lucking, B., & Şahin, A. (2012). A search and matching approach to labor markets: Did the natural rate of unemployment rise? *Journal of Economic Perspectives*, *26*(3), 3–26. https://doi.org/10.1257/jep.26.3.3

[13] Elsby, M. W. L., Shin, D., & Solon, G. (2016). Wage adjustment in the Great Recession and other downturns: Evidence from the United States and Great Britain. *Journal of Labor Economics*, *34*(S1), S249–S291. https://doi.org/10.1086/682407

[14] Erceg, C. J., Henderson, D. W., & Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, *46*(2), 281–313. https://doi.org/10.1016/S0304-3932(00)00028-3

[15] Evans, C. (2018). *Optimal monetary policy with downward nominal wage rigidity* [Unpublished manuscript].

[16] Fahr, S., & Smets, F. (2010). Downward wage rigidities and optimal monetary policy in a monetary union. *Scandinavian Journal of Economics*, *112*(4), 812–840. https://doi.org/10.1111/j.1467-9442.2010.01630.x

[17] Friedman, M. (1968). The role of monetary policy. *American Economic Review*, *58*(1), 1–17. https://www.jstor.org/stable/1831652

[18] Galí, J. (2011). The return of the wage Phillips curve. *Journal of the European Economic Association*, *9*(3), 436–461. https://doi.org/10.1111/j.1542-4774.2011.01023.x

[19] Galí, J. (2015). *Monetary policy, inflation, and the business cycle: An introduction to the new Keynesian framework and its applications* (2nd ed.). Princeton University Press.

[20] Guerrieri, L., & Iacoviello, M. (2015). OccBin: A toolkit for solving models with occasionally binding constraints easily. *Journal of Monetary Economics*, *70*, 22–38. https://doi.org/10.1016/j.jmoneco.2014.08.005

[21] Jacobs, D., & Rush, A. (2015). Why is wage growth so low? *RBA Bulletin*, June, 9–18. https://www.rba.gov.au/publications/bulletin/2015/jun/

[22] Khan, A., King, R. G., & Wolman, A. L. (2003). Optimal monetary policy. *Review of Economic Studies*, *70*(4), 825–860. https://doi.org/10.1111/1467-937X.00269

[23] Kim, Y. B., & Ruge-Murcia, F. J. (2011). Monetary policy when wages are downward sticky: Friedman meets Tobin. *Journal of Economic Dynamics and Control*, *35*(12), 2064–2077. https://doi.org/10.1016/j.jedc.2011.08.005

[24] King, R. G., & Wolman, A. L. (1999). What should the monetary authority do when prices are sticky? In J. B. Taylor (Ed.), *Monetary policy rules* (pp. 349–398). University of Chicago Press.

[25] Lowe, P. (2019, June 18). *The labour market and monetary policy.* Speech delivered at the Anika Foundation Fundraising Lunch, Sydney. Reserve Bank of Australia. https://www.rba.gov.au/speeches/2019/sp-gov-2019-06-18.html

[26] Rotemberg, J. J. (1982). Sticky prices in the United States. *Journal of Political Economy*, *90*(6), 1187–1211. https://doi.org/10.1086/261117

[27] Schmitt-Grohé, S., & Uribe, M. (2016). Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy*, *124*(5), 1466–1514. https://doi.org/10.1086/688175

[28] The Australia Institute. (2021). *The wages crisis: Revisited.* Centre for Future Work.

[29] Tobin, J. (1972). Inflation and unemployment. *American Economic Review*, *62*(1), 1–18. https://www.jstor.org/stable/1821468

[30] Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy.* Princeton University Press.

A Appendix: Microfoundations

A.1 Final Goods Production

A perfectly competitive final goods firm combines a continuum of intermediate goods $Y_t(i)$ for $i \in [0,1]$ to produce aggregate output Y_t using CES technology:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon_p - 1}{\epsilon_p}} di \right)^{\frac{\epsilon_p}{\epsilon_p - 1}}, \quad \epsilon_p > 1 \tag{29}$$

Profit maximization yields the demand curve for intermediate good i:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon_p} Y_t \tag{30}$$

and the aggregate price index:

$$P_t = \left(\int_0^1 P_t(i)^{1-\epsilon_p} di \right)^{\frac{1}{1-\epsilon_p}} \tag{31}$$

A.2 Intermediate Goods Production

Intermediate firm i produces output using labor according to:

$$Y_t(i) = A_t N_t(i) \tag{32}$$

where A_t is aggregate productivity and $N_t(i)$ is a composite labor input assembled from differentiated labor types $j \in [0,1]$ via CES aggregation:

$$N_t(i) = \left(\int_0^1 N_t(i,j)^{\frac{\epsilon_w - 1}{\epsilon_w}} dj \right)^{\frac{\epsilon_w}{\epsilon_w - 1}}, \quad \epsilon_w > 1 \tag{33}$$

Cost minimization yields demand for labor type j:

$$N_t(i,j) = \left(\frac{W_t(j)}{W_t} \right)^{-\epsilon_w} N_t(i) \tag{34}$$

and the aggregate wage index:

$$W_t = \left(\int_0^1 W_t(j)^{1-\epsilon_w} dj \right)^{\frac{1}{1-\epsilon_w}} \tag{35}$$

Intermediate firm i maximizes profits $\Pi_t(i) = P_t(i)Y_t(i) - W_t N_t(i)$ subject to demand curve (30) and production function (32). With flexible prices, this yields the optimal pricing rule:

$$P_t(i) = \mu_p \frac{W_t}{A_t}, \quad \mu_p \equiv \frac{\epsilon_p}{\epsilon_p - 1} > 1 \tag{36}$$

In symmetric equilibrium where $P_t(i) = P_t$ for all i, the real wage is:

$$\frac{W_t}{P_t} = \frac{A_t}{\mu_p} \tag{37}$$

A.3 Households

A continuum of households indexed by $j \in [0, 1]$ each supply differentiated labor type j. Household j sets its nominal wage $W_t(j)$ and then supplies hours $N_t(j)$ according to firm demand. Aggregating demand across all intermediate firms gives economy-wide demand for labor type j:

$$N_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\epsilon_w} N_t \tag{38}$$

where $N_t \equiv \int_0^1 N_t(i)di$ is aggregate labor.

A.3.1 Household Optimization

Household j maximizes expected lifetime utility:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t(j)^{1-\sigma} - 1}{1-\sigma} - \frac{N_t(j)^{1+\psi}}{1+\psi} \right] \tag{39}$$

subject to the budget constraint:

$$P_t C_t(j) + Q_t B_t(j) \leq B_{t-1}(j) + W_t(j) N_t(j) + D_t(j) - T_t^{adj}(j) \tag{40}$$

and the Rotemberg wage adjustment cost:

$$T_t^{adj}(j) = \frac{\phi_w}{2}\left(\frac{W_t(j)}{W_{t-1}(j)} - 1\right)^2 P_t Y_t \tag{41}$$

where $C_t(j)$ is consumption, $B_t(j)$ are nominal bond holdings, Q_t is the bond price, $D_t(j)$ are dividends, $\beta \in (0,1)$ is the discount factor, $\sigma > 0$ is the inverse intertemporal elasticity of substitution, $\psi > 0$ is the inverse Frisch elasticity, and $\phi_w > 0$ governs wage adjustment costs.

Complete asset markets ensure perfect risk sharing, so $C_t(j) = C_t$ for all j. The first-order condition for consumption yields:

$$\lambda_t = \frac{C_t^{-\sigma}}{P_t} \tag{42}$$

The first-order condition for bonds yields the Euler equation:

$$Q_t = \beta E_t\left[\frac{\lambda_{t+1}}{\lambda_t}\frac{P_t}{P_{t+1}}\right] = \beta E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\sigma}\frac{1}{\Pi_{t+1}}\right] \tag{43}$$

where $\Pi_{t+1} \equiv P_{t+1}/P_t$ is gross inflation.

A.3.2 Wage Setting

The first-order condition for $W_t(j)$ accounts for how wage changes affect labor demand via equation (38) and create adjustment costs via equation (41). Defining wage inflation as $\pi_t^w(j) \equiv W_t(j)/W_{t-1}(j)$, the optimality condition is:

$$\begin{aligned}\epsilon_w N_t^{1+\psi}\left(\frac{W_t(j)}{W_t}\right)^{-\epsilon_w(1+\psi)-1}\frac{1}{W_t} + \lambda_t(1-\epsilon_w)W_t^{\epsilon_w}N_t W_t(j)^{-\epsilon_w} \\ - \lambda_t\phi_w(\pi_t^w(j)-1)\pi_t^w(j)P_tY_t\frac{1}{W_{t-1}(j)} - \beta E_t[\lambda_{t+1}\phi_w(\pi_{t+1}^w(j)-1)\pi_{t+1}^w(j)P_{t+1}Y_{t+1}]\frac{1}{W_t(j)} = 0\end{aligned} \tag{44}$$

In symmetric equilibrium where $W_t(j) = W_t$ and $N_t(j) = N_t$ for all j, this simplifies to:

$$\epsilon_w N_t^{1+\psi} + (1-\epsilon_w)\lambda_t W_t N_t - \lambda_t\phi_w(\pi_t^w - 1)\pi_t^w P_t Y_t - \beta E_t[\lambda_{t+1}\phi_w(\pi_{t+1}^w - 1)\pi_{t+1}^w P_{t+1}Y_{t+1}] = 0 \tag{45}$$

A.3.3 Steady State

In a zero-inflation steady state where $\pi^w = 1$ and all variables are constant, equation (45) reduces to:

$$\epsilon_w N^{1+\psi} + (1-\epsilon_w)\lambda W N = 0 \tag{46}$$

Using $\lambda = C^{-\sigma}/P = Y^{-\sigma}/P$ and the real wage $w \equiv W/P$, this yields:

$$w = \frac{\epsilon_w}{\epsilon_w - 1} N^{\psi} Y^{\sigma} = \mu_w \cdot MRS \tag{47}$$

where $\mu_w \equiv \epsilon_w/(\epsilon_w - 1) > 1$ is the wage markup and $MRS \equiv N^{\psi}Y^{\sigma}$ is the marginal rate of substitution between consumption and labor.

A.4 IS Curve

Log-linearizing the Euler equation (43) around steady state and using market clearing $C_t = Y_t$ yields:

$$y_t = E_t[y_{t+1}] - \frac{1}{\sigma}(i_t - E_t[\pi_{t+1}] - r_t^n) \tag{48}$$

where lowercase variables denote log deviations from steady state, i_t is the log deviation of the nominal interest rate, and r_t^n is the natural rate of interest.

Under flexible prices, the relationship between price and wage inflation is:

$$\pi_t = \pi_t^w - \Delta a_t \tag{49}$$

where $a_t \equiv \log A_t$ is log productivity. Substituting equation (49) into equation (48) and defining the output gap as $x_t \equiv y_t - y_t^n$ where $y_t^n = a_t$ is the natural level of output, we obtain:

$$x_t = E_t[x_{t+1}] - \frac{1}{\sigma}(i_t - E_t[\pi_{t+1}^w] - \tilde{r}_t^n) \tag{50}$$

where $\tilde{r}_t^n \equiv \rho + (\sigma - 1)E_t[\Delta a_{t+1}]$ is the modified natural rate and $\rho \equiv -\log \beta$ is the time preference rate.

A.5 Wage Phillips Curve

Log-linearizing equation (45) around the zero-inflation steady state yields the wage New Keynesian Phillips Curve. Define the marginal rate of substitution in log deviations as $\widehat{MRS}_t \equiv \psi \hat{N}_t + \sigma \hat{y}_t$ and the real wage in log deviations as $\hat{w}_t$.

The linearized wage-setting condition becomes:

$$\hat{\pi}_t^w = \beta E_t[\hat{\pi}_{t+1}^w] + \kappa_w(\widehat{MRS}_t - \hat{w}_t) \tag{51}$$

where $\kappa_w \equiv \epsilon_w N^{1+\psi}/(\phi_w Y^{1-\sigma}) > 0$ is the slope of the wage Phillips curve.

A.6 Downward Nominal Wage Rigidity

The downward nominal wage rigidity constraint imposes a floor on wage inflation:

$$\pi_t^w \geq \bar{\pi}^w \tag{52}$$

The complete wage inflation dynamics are characterized by the complementarity condition:

$$\hat{\pi}_t^w = \begin{cases} \beta E_t[\hat{\pi}_{t+1}^w] + \kappa_w(\widehat{MRS}_t - \hat{w}_t) & \text{if } \pi_t^w > \bar{\pi}^w \text{ (constraint does not bind)} \\ \hat{\bar{\pi}}^w & \text{if } \pi_t^w = \bar{\pi}^w \text{ (constraint binds)} \end{cases} \tag{53}$$

where $\hat{\bar{\pi}}^w$ is the log deviation of the wage floor from steady state.

This occasionally binding constraint creates state-dependent nonlinear dynamics. When the constraint binds, wage inflation is fixed at the floor regardless of economic conditions, severing the relationship between wage adjustment and the output gap established by the Phillips curve.