Ambition versus Gradualism in Disinflation Horizons under Bounded Rationality: The Case of Chile [1]

Ali J. Al-Eyd [2] **Meral Karasulu [3]**

March 2008

NIESR DISUCSSION PAPER 308

This paper uses a stylized New Keynesian Model to examine alternative disinflation strategies under optimal monetary policy conditions with bounded rationality. The model is calibrated for Chile and presents some policy tradeoffs not necessarily captured under the full rational expectations solution. Under rational expectations and in the absence of nominal inertia, the optimal policy suggests an ambitious disinflation horizon as expectations of inflation are revised immediately. However, under full adaptive learning, a distinct policy trade-off emerges between ambition and gradualism. We also find that when expectations exhibit a degree of imperfect knowledge it is optimal to increase the strength of the policy response relative to that of the perfect knowledge solution.

JEL Classification Numbers: E52
Keywords: Monetary policy, inflation targeting, disinflation, optimal control, rational expectations, adaptive learning

Author's E-Mail Address: ali.aleyd@citi.com and mkarasulu@imf.org

---

[1] Address for correspondence: Ali Al-Eyd, National Institute of Economic and Social Research, 2 Dean Trench St., Smith Square, London SW1P 3HE, UK. The views presented here are those of the authors and do not necessarily reflect those of the National Institute of Economic and Social Research, Citigroup, or the International Monetary Fund. The authors would like to thank Stephen Hall for valuable comments and discussion. They also thank Ray Barrell, Patricia Brenner, Brian Henry, Kent Matthews and Martin Weale for useful comments. All remaining errors are those of the authors.

[2] Citigroup and National Institute of Economic and Social Research

[3] International Monetary Fund

I. INTRODUCTION

An important policy question for a central bank seeking to disinflate from high to moderate levels of inflation concerns the horizon over which this is feasible and least costly. The adjustment costs depend on several factors, including the credibility of the new policy regime and the degree of wage and price indexation; both of which depend on the adjustment in private sector expectations.

During a period of hyperinflation, Sargent (1986) favours a rapid and seemingly costless disinflation. This reasoning follows on the premise that prior periods of high and volatile inflation discourage long-term wage and price contracts, so that with a sufficiently credible policy reform in place it is possible to achieve rapid adjustments of prices. However, during periods of moderate inflation, wage and price agreements are more likely to be prolonged and overlapped with each other entrenching price expectations at current levels of inflation. In such circumstances, King (1996) advocates a more gradual timetable for disinflation.

The "appropriate" speed of disinflation depends on the credibility of policies and private sector expectations. From a modelling perspective, the assumption of full rational expectations implies that a new and untried policy regime is greeted with instant credibility since private expectations quickly converge on those of the authorities. Clearly, this may not be the most appropriate modelling framework for an economy with a history of high and variable inflation. Another and perhaps more reasonable approach would be to allow for agents who *learn* through economic outcomes about the newly adopted policy objectives and the implications that these have for pricing mechanisms.

This paper explores the merits of *ambition* versus *gradualism* in the setting of inflation targets in a small open economy using a stylized New Keynesian Model under alternative expectation sturctures. The analysis is conducted in an optimal setting where policy weights for a standard inflation forecast-based rule are derived with respect to alternative disinflation strategies, central bank preference settings, and expectational structures. This approach permits us to analyze standard variance frontiers and to comment on the desirability of each strategy. In doing so, our analysis highlights the implications for optimal policy weights and choice of disinflation strategy arising from the various expectational structures employed.

Our study is similar to Orphanides and Williams (2004) who develop a simple model of the US economy and examine, among other things, the role that imperfect knowledge – in the form of *adaptive learning* – about the structure of the economy plays in efficient formulation of monetary policy. However, here we use an open-economy framework to include the transmission mechanism of the real exchange rate and examine the efficient formulation of monetary policy over the disinflation horizon. We view the alternative model structures given by adaptive learning and rational expectations as complementary approaches to the investigation of the disinflation horizon.

This model is calibrated for Chile. We have chosen to use Chile as a case since it was the first emerging market economy to adopt an inflation targeting regime at moderately high levels of inflation. In 1990, the Chilean monetary authorities announced the first numerical target for inflation and successfully, but gradually, disinflated over an 11 year horizon to around 3 percent at the end of 2001. The case of Chile provides a relevant example to assess the implications of rational expectations and adaptive learning on the efficiency of policy rules and the speed of disinflation. Chile adopted inflation targeting after the failure of a short-lived exchange-rate based stabilization program in 1990, which created significant uncertainty regarding the resolve and ability of the newly independent central bank to reign in inflation under the untested and new IT regime. The Chilean monetary authorities were, however, successful in reigning in inflation over the course of the 1990s where the approach to disinflation was gradualist by design. Their success raises a few natural questions: Could Chile have pursued a more *ambitious* disinflation regime and what would be the economic costs of doing so? And, importantly, what role (if any) do expectational dynamics play in the choice of disinflation strategy?

The paper proceeds as follows. Section 2 provides a brief overview of inflation targeting in Chile. Section 3 develops a small New Keynesian Model and presents the central bank preference settings. Section 4 discusses the alternative expectations structures employed in this study. Section 5 explores the alternative disinflation strategies explored here and derives optimal monetary policy settings with respect to each. Section 6 investigates the policy tradeoffs associated with these exercises, while Section 7 examines the implications for the optimal policy weights resulting from differing degrees of nominal inertia in the structural model. Finally, Section 8 concludes.

II. Inflation Targeting in Chile

Chile's current inflation targeting framework grew out of a varied policy mix dating back to the mid 1970s when the authorities were struggling with hyperinflationary pressures. Throughout most of the latter half of the twentieth century, Chile, like most Latin American countries, experienced high and volatile levels of inflation. The inability to adhere to sound macroeconomic policies, coupled with external shocks, contributed to high and variable inflation rates ranging from as little as 10 percent in the 1960s to as much as 606 percent in 1973 (Figure 1). The threat of hyperinflation in the early 1970s helped refocus efforts on a stabilization program.[4] This concerted policy effort began with deep structural reforms as well as tight fiscal and monetary policies, but ended in an exchange rate based stabilization program beginning in 1979.

[4] The inconsistent macroeconomic policy framework prior to the 1970s resulted from several factors including political and social unrest, mismanaged policies, and balance of payments crises (Garcia (2000)).

Figure 1: Inflation in Chile 1974-2001

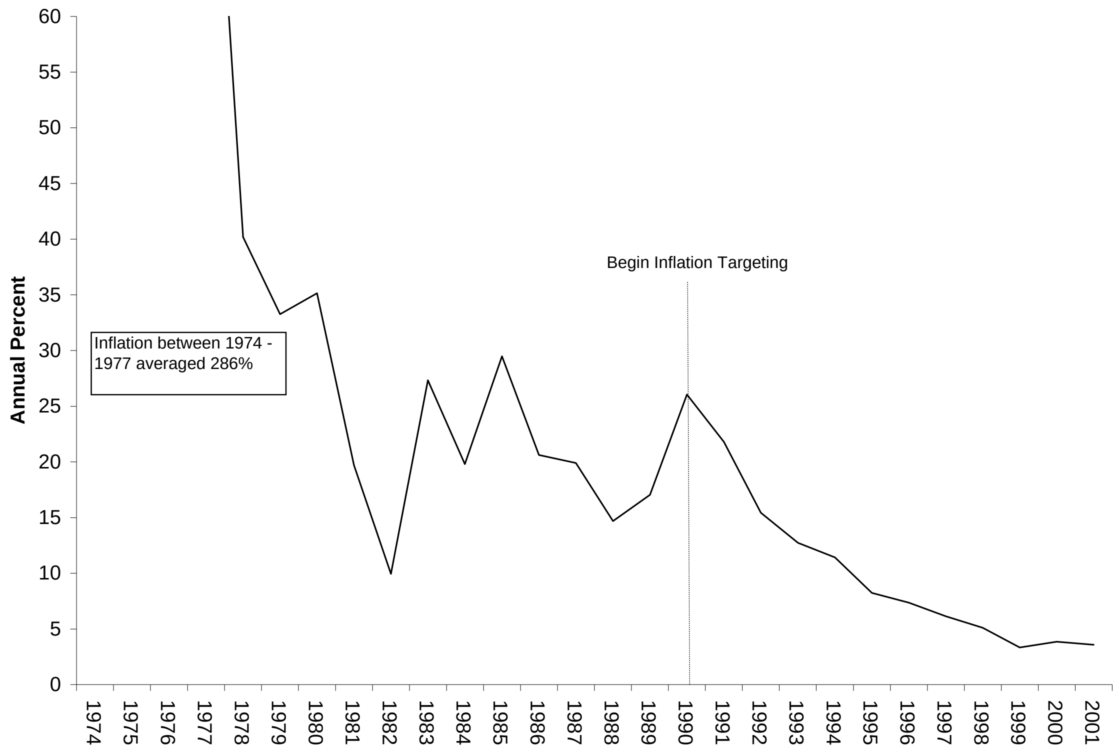

The failure of austere policy measures and structural reforms to rein in inflation can be attributed to several factors, including widespread wage and price indexation, the persistence of inflationary expectations, and adverse external shocks (Garcia (2000)). The introduction of a fixed exchange rate, however, provided only temporary relief as inflationary pressures began to build up following large capital inflows, which funded an unsustainable domestic expansion and put upward pressure on the real exchange rate, ultimately forcing the authorities to abandon the parity. The peso was floated in 1982 in the midst of a balance of payments and banking crisis.

The economic recovery from the crisis was accompanied by rising inflation that reached nearly 26 percent by 1985, more than double the figure than when the peso was floated. As a response, the government adopted a new framework supported by a commitment to fiscal restraint with an additional monetary component aimed at maintaining a competitive exchange rate. These policies worked as inflation began to fall over the course of the decade and the economy emerged from recession.

By the end of the 1980s, however, solid GDP growth brought with it a new round of inflationary pressures and it was clear that further policy measures would be needed. The move to an independent central bank in October 1989 was crucial in devising the policies that followed. Under Law 18,840 of the Constitutional Organic Act the Central Bank was granted unprecedented autonomy of both goal and instrument independence (see Central Bank of Chile (2002)).

The Central Bank's newly acquired autonomy yielded a degree of credibility and paved the way for a more aggressive policy framework in which explicit targets for inflation were announced, laying down the basic building blocks of the formal inflation targeting framework currently in place in Chile.[5] The first official numerical target for the consumer price index (CPI) was announced in 1990 when inflation was around 25 percent. The target horizon was set for one year and a future goal of 3.5 percent was to be achieved over a medium term horizon. These features provided a *transparent* structure for which the authorities could be held accountable. Table 1 below describes the evolution of inflation targets in Chile over the 1990s, which culminated in the current symmetric target range of 2 to 4 percent.

Table 1: Evolution of Inflation Targets in Chile 1991 – 2001*

Year	Inflation Target
1991	18
1992	15-20
1993	10-12
1994	9-11
1995	8
1996	6.5
1997	5.5
1998	4.5
1999	4.3
200	3.5
2001	2-4

* Source: Schaechter et al. (2000). Targets between 1996-2000 include an unspecified +/- range.

This new framework for monetary policy contained some unique features. First, the benefits of external competitiveness resulting from earlier policies signalled that the authorities would seek to maintain some control over the exchange rate. Second, due to many years of high and persistent inflation Chile had a high degree of wage and price indexation. These features in part dictated that the disinflation program would have to be *gradualist* in nature since a rapid disinflation could have adverse real effects. And, finally, the authorities had to establish a tough and credible stance in order to dislodge these entrenched expectations.

The new policy climate could therefore be described as one in which the objectives of the authorities encompassed a degree of control over the exchange rate, while maintaining monetary autonomy, and pursuing the benefits of capital inflows – a policy mix bordering on the "impossible trinity". Until 1999 Chile maintained this policy mix which was anchored on an inflation target, but also had another secondary objective which was pursued in an opportunistic fashion.

[5] An interesting fact noted in Morande (2001) is that, with the announcement of this new policy framework, Chile had actually adopted what is now referred to as "inflation targeting" yet this lingo had not been formalized at the time.

This second objective came in the form of a (soft) target for the current account deficit, which according to Morande (2001) ranged between 2 to 4 percent through the middle of the 1990s and increased to 5 percent by 1998. Three important mechanisms were in place to support this target. First, the authorities maintained a wide crawling exchange rate band in order to mitigate large swings in the exchange rate and also to provide an additional anchor for expectations. The band was used "opportunistically" to exploit the primary goal of price stability. This asymmetric approach enabled the authorities to adjust the band when, for example, exchange rate appreciation threatened the upper limit. In this way, the authorities prevented large swings in the exchange rate, while simultaneously guarding their inflation target.

Second, substantial capital inflows enabled the authorities to accumulate sizeable foreign exchange reserves which were used to support the exchange rate arrangement. Third, the imposition of capital controls allowed the authorities to maintain a degree of control over both the exchange rate and monetary policy. The interaction of these supporting mechanisms suggests that the use of the exchange rate was an essential tool in the fight against inflation. Indeed, Garcia (2000) states that up to one third of the disinflation over the 1990 to 1997 period can be attributed to innovations in the exchange rate band.

Speculative pressures built in 1998 as the fallout from the Asian crisis and a deteriorating fiscal position called into question the authorities' commitment to the peso peg. The sensibility of the varied policy mix, which had been working well over the course of the 1990s, also came into question. The authorities embarked on an asymmetric policy response to external pressures, which included the maintenance of the peg through the sale of foreign exchange reserves and periods of tight monetary policy. The defence of the peg lasted almost a year. In 1999 the authorities abandoned the asymmetric approach and its supporting mechanisms (exchange rate band and capital controls) for a fully-fledged symmetric IT framework with a constant target range of 2 to 4 percent.

III. A Small Open Economy Model

We adopt an open-economy version of the New Keynesian Model[6] similar to those presented in Svensson (1997b) and Batini and Nelson (2000) and augment the common underlying framework with alternative expectational structures. The model includes relations for aggregate demand, short run aggregate supply (Phillips curve), and a real uncovered interest parity condition:

$$y_t = \alpha_1 y_{t-1} + \alpha_2 \hat{E}_t y_{t+1} - \alpha_3 r_{t-1} + \alpha_4 q_{t-1} + \varepsilon_t^d \,, \tag{1}$$

$$\pi_t = \beta_1 \pi_{t-1} + (1-\beta_1)\hat{E}_t \pi_{t+1} + \beta_2 y_{t-1} + \beta_3 q_t + \varepsilon_t^s \,, \tag{2}$$

$$q_t = \hat{E}_t q_{t+1} - \gamma_1 (r_t - r_r^f) + \upsilon_t \,. \tag{3}$$

[6] See Clarida et al. (1999) and Henry and Pagan (2004) for respective theoretical and empirical assessments of the standard New Keynesian Model.

where, y_t is real output at time *t*, r_t is the domestic real interest rate, π_t is the inflation rate and q_t is the real exchange rate (the nominal exchange rate s_t in (4) below is defined as the domestic price of foreign currency so that a rise in s_t is a depreciation).[7] All parameters in (1) - (3) are positive and all variables, except the interest rate, are in logs and are taken as the difference from equilibrium levels (given here as the baseline data). $\varepsilon_t^d \sim iid(0, \sigma^2_{\varepsilon^d})$ denotes the demand shock, and the supply shock is given by $\varepsilon_t^s \sim iid(0, \sigma^2_{\varepsilon^s})$, with their covariance given by $c_{\varepsilon^d \varepsilon^s}$.

The expectations are denoted by $\hat{E}_t$ which need not be fully rational and the ' ^ ' symbol emphasizes this point. A full rational expectations solution is indicated by E_t without a ' ^ ' to alert the reader to the places where the basic structure may be altered to distinguish the expectations assumption. In order to avoid issues arising from simultaneity, we adopt the conventional assumption that expectation at time *t*, $\hat{E}_t$ are based on information held in period *t-1* (see Evans and Honkaphohja (2001)).

Equation (1) is the aggregate demand or IS relationship denoting output as a function of expected output, one period lagged real interest rate, the real exchange rate, and a demand shock. This is an expectations augmented IS function of the kind derived in McCallum and Nelson (1999) where $\hat{E}_t y_{t+1}$ captures expected future output. Values for α_1 and α_2 are set at 0.7 and 0.2, respectively, and reflect a small degree of forward-looking behaviour. We calibrate α_3 and α_4 to be 0.1 and 0.033, respectively, as in Batini and Nelson (2000). This calibration produces a ratio for the interest rate to exchange rate of approximately 3 in line with the conventional rule of thumb value reported in Ball (1998).

Equation (2) is an expectations augmented open-economy Phillips curve with both a forward and backward-looking term in inflation. This specification is set to capture the effects of inflation inertia which is governed by $\beta_1 = 0.7$ suggesting a reasonable degree of forward looking behaviour, while we set $\beta_2 = 0.2$. The term in the level of the real exchange rate provides a direct channel for pass-through of import prices on to domestic inflation. As discussed above, this is an important transmission channel for both policy and external events in a small open-economy and we have chosen $\beta_3 = 0.05$ in line with other studies.

The real uncovered interest rate parity (RUIP) (3) develops the link between the real exchange rate and the real interest rate. In line with Svensson (1997b), we have chosen to represent the interest parity condition in real terms to provide a more stable long run solution. The real interest rate differential governs the movements in the real exchange

[7] Foreign variables are denoted by superscript 'f' where data for the United States are used to proxy the foreign country.

rate where we have calibrated a correlation of one, $\gamma_r = 1$, in line with long run theoretical considerations. The shock υ_t is meant to act as a risk premium capturing the effects of unobservables – e.g. market sentiment on the exchange rate – and is assumed $\upsilon_t \sim iid(0, \sigma_v^2)$.

The annual data for Chile are from the IFS, and hence one period equals one year. The lag structure is meant to reflect realistic channels for the transmission of both monetary policy and shocks and is similar to that found in the small open economy model of Ball (1998). There are two channels for the transmission of policy to affect inflation; the interest rate and the real exchange rate channel. A monetary contraction through the nominal interest rate takes two years to affect inflation via the traditional aggregate demand channel. The adjustment in the real interest rate is felt with one year lag in the aggregate demand equation and takes a further year to feed through to inflation in the Phillips curve. The real exchange rate channel is meant to capture the features of a small open economy where the transmission channel to inflation is more direct for both policy and external shocks. Schaechter et al. (2000) report that for Chile the monetary policy transmission lag from the exchange rate and wage shocks to inflation has slowed from a rapid pace to approximately 1-2 years in recent times. Additionally, Caputo (2004) reports that real exchange rate shocks have an impact on inflation between 11-16 months in Chile subsequent to the shock. Our dynamic specification is congruent with the transmission lags reported in these studies.

The real exchange rate, q_t, is defined as:

$$q_t = s_t + p_t^f - p_t \tag{4}$$

where, p_t is the price level of domestically produced goods, p_t^f is the foreign price level, and s_t is the nominal exchange rate. The short term real interest rate, r_t, follows from the Fisher identity:

$$r_t = i_t - \hat{E}_t \pi_{t+1} \tag{5}$$

where, i_t is the domestic nominal interest rate and, $\hat{E}_t \pi_{t+1}$ is the private sector's expectation of inflation. The foreign short term real interest rate, r_t^f is defined as

$$r_t^f = i_t^f - E_t \pi_{t+1}^f, \tag{6}$$

where, i_t^f is the foreign domestic nominal interest rate and $E_t \pi_{t+1}$ is the rational expectation of inflation.[8]

The existence of transmission lags in monetary policy implies that, in practice, IT is invariably forward-looking and requires that the monetary authorities incorporate forecasts of future inflation into today's policy decisions. This forward-looking approach to policy has given rise to the notion of "inflation-forecast based" (IFB) targeting in both optimal and simple monetary policy rules (see Hall and Nixon (1997), Batini and

[8] The foreign real interest rate is denoted with rational expectations as a simplification and does not alter the results obtained under the adaptive learning or hybrid models.

Haldane (1998), Svensson and Woodford (1999), Batini and Nelson (2000), among others, and Svensson (1997a)). IFB rules rely on the official inflation forecast of the monetary authorities and, therefore, are seen to encompass all of the required information for setting policy. Specifically, Batini and Haldane (1998) find that IFB rules in a rational expectations setting are "output, lag, and information encompassing."

The model is closed with an IFB rule similar to that found in Batini and Haldane (1998).[9] Here we specify the policy rule in terms of the nominal interest rate and modify it to respond to the deviation of a one-period-ahead inflation forecast from the inflation target:

$$i_t = \delta_\pi \left(\hat{E}_t \pi_{t+1} - \pi^*_{t+1} \right), \tag{7}$$

where, $\delta_\pi > 0$ is the policy weight, and π^*_{t+1} is the inflation target at time *t+1* (to correspond with inflation expectations at that time). The time subscript on π^*_{t+1} is intentional and facilitates a time varying target for inflation. However, this need not be the case since one can set $\pi^*_{t+1} = 0$ (or to some other constant target) without any loss in generality.[10]

We employ the standard assumption that the instrument of the central bank is the nominal interest rate, i_t, and that the authorities have preferences over a Barro and Gordon (1983) loss function

$$L_t = \lambda_\pi \left(\pi_t - \pi^*_t \right)^2 + \lambda_y y_t^2, \tag{8}$$

where the fixed preference weights (λ_π and λ_y) are associated with inflation and output, and π^*_t corresponds to the disinflation strategy, or target in (7). Optimal policy is determined by deriving the interest rate which minimizes the intertemporal function[11]

$$E_t \sum_{\tau=0}^{\infty} \theta^\tau L_{t+\tau} \tag{9}$$

subject to the dynamics of the economy described above and where $0 < \theta \leq 1$ is the discount factor.[12] By alternating the preference weights in (8) we examine three central bank types; (i) a "Hawk" central bank focused mostly on minimizing deviations of inflation from its target; (ii) a "Balanced" central bank that values inflation and output variations equally, and, finally (iii) a "Dove" central bank that is more concerned with output variations than inflation.

---

[9] In addition, Batini and Haldane (1998) allow for some interest smoothing and forward-looking behavior with the inclusion of the short term *ex ante* lagged real interest rate and the equilibrium value of the real interest rate.

[10] The policy rule, (7), is simulated in first differences and, therefore, leaves the rule determinate in the absence of an intercept term.

[11] Model solution and all optimization/simulations are conducted using CEF Modelling Suite (see CEF (2000)).

[12] We assume that the authorities do not discount the future, so we set $\theta = 1$. However, it is noted that Henry et al. (2006) show that policy outcomes (specifically, inflation and output variances), and hence social welfare, can depend significantly on the degree to which policymakers discount the future.

Table 2: Central Bank's Objective Function

Objective	*Description*	*Weights*
L1	Hawk	$\lambda_\pi = 0.9, \quad \lambda_y = 0.1$
L2	Balanced	$\lambda_\pi = 0.5, \quad \lambda_y = 0.5$
L3	Dove	$\lambda_\pi = 0.1, \quad \lambda_y = 0.9$

IV. LEARNING AND EXPECTATIONS STRUCTURES

The assumption of adaptive learning stands in contrast to that of rational expectations (or model consistent expectations) where agents are assumed to know the full structure of the economy and its parameters at all times. Under rational expectations, a switch in regime to, say, IT is fully credible and, with no nominal inertia, will lead to expectations of inflation to be revised immediately and for output to remain at its natural rate (Hall et al. (2000)). Therefore, the policy prescription under rational expectations with no nominal inertia would be for the central bank to move immediately to the long-run inflation target.

Here we propose an alternative specification for expectations which represents a complete dynamic system enabling agents to learn *adaptively* using a least squares algorithm based on the Kalman filter (see Hall et al. (2000) and Garratt and Hall (1995), for a similar set-up). This specification takes into account the uncertainties inherent in a new policy regime and the central bank's willingness to fully support the new framework. Under adaptive learning, agents are unaware of the true structure of the economy as it would exist under rational expectations and must *learn* about it over time using a 'reasonable rule of thumb' or 'forecast rule'. It is in this spirit that agents are *boundedly rational*. Therefore, agents employ their own 'forecast rule' (or *perceived law of motion* (PLM)) to obtain estimates of the unobserved state variables, or the *actual law of motion* (ALM) of the economy. Using their PLM, agents optimally update their expectations of future endogenous variables each period subject to observed data.[13] The system, therefore, is self-referential and we draw on the properties inherent in the E-stability principle whereby small forecasting errors made relative to the rational solution (ALM) are corrected over time.[14]

[13] As is common in the learning literature, studies, such as Orphanides and Williams (2004), employ an algorithm based on a weighted form of *recursive least squares* to update the parameters of the PLM (e.g. Marcet and Sargent (1989a,b)). The Appendix shows that the Kalman filter can be identified as a more general form of this procedure allowing for the application of the convergence criteria as set out in Marcet and Sargent.

[14] More formally, a fixed point of the mapping (or learning process) between the PLM and the ALM represents a rational expectations (or E-stable) solution, and one which is not dependent upon an arbitrary terminal condition (see Evans and Honkapohja (2001) for details). More loosely, if the parameters of the PLM cease to change then this can be taken as evidence of E-stability since, in this case, by virtue of the learning algorithm the expectations error term is zero; the outcome is equal to the expectation and so the system has converged.

Garratt and Hall (1995) and Evans and Honkapohja (2001) among others note that there is no theoretical basis for modelling the learning behaviour of agents. Therefore, we are left to form a guess as to how agents learn and in order to remain consistent with the endogenous learning literature we adopt simple forecast rules based on past values of relevant variables[15]

$$\hat{E}_t \pi_{t+1} = \rho_{0,t} + \rho_{1,t}\pi_{t-1} + \rho_{2,t}\pi^*_{t-1}, \tag{10}$$

$$\hat{E}_t q_{t+1} = \rho_{3,t} + \rho_{4,t} q_{t-1}, \tag{11}$$

$$\hat{E}_t y_{t+1} = \rho_{5,t} + \rho_{6,t} y_{t-1}. \tag{12}$$

These PLMs are used to form one period ahead expectations of each respective variable based on information held at time *t* where $\rho_{0t} \ldots \rho_{6t}$ are a set of time-varying parameters each evolving according to the following process

$$\rho_{i,t} = \alpha_i \rho_{i,t-1} + a_{i,t}, \tag{13}$$

where $i = 0,...,6$ and $a_{i,t}$ is an *iid* error term. The PLM given by (10) incorporates the inflation target so that when $\rho_{1,t}$ is equal to zero we have the rational expectations solution where the inflation expectation equals the policy target.

By varying the expectational assumptions in both the structural equations and the policy rule we are able to examine the policy implications arising from three different scenarios. The first two models represent fully rational (RE) and fully adaptive learning (AL) agents, respectively, where expectational assumptions are consistent throughout both in the policy rule and the structural model. The third case we consider combines a fully rational policy rule and adaptive learning in the structural model to yield a hybrid (HYB) model. The idea behind the HYB model is to capture the notion that policymakers may (think they) have perfect knowledge, but agents in the economy are slow to adjust as in the King (1996) sense. Therefore, a perfectly rational instrument must feed back from a learning structure. Table 3 summarizes these model structures.

Table 3: Expectations Structures

Model	***Structural Equations***	***Policy Rule***
Rational Expectations (RE)	RE	RE
Adaptive Learning (AL)	AL	AL
Hybrid (HYB)	AL	RE

[15] More complicated rules, possibly containing the reduced form of the whole system can be employed. However, Beeby et al. (2002) conduct a study across a set of learning rules and conclude that making a rule more complicated does not necessarily improve its learning performance.

In adopting the adaptive learning structure, we seek to operationalize the sentiments of King (1996) who provides a conceptual basis and framework for its application in exploring simple monetary policy rules and the ideas surrounding the speed of disinflation. The author states that the optimal speed of disinflation in countries attempting to change from a regime of moderate or high inflation to that of price stability is subject to how rapidly expectations of inflation adapt to the change in regime. King attributes this link to the fact that, following a period of prolonged inflation, agents' expectations do not immediately adjust since the new regime is untried and, hence, cannot be considered credible. In this respect, the author concludes that the costs to disinflation, in terms of lost output from wage and price indexation, increase more than proportionally with the speed of disinflation because "[agents'] learning takes time"(see p.58).[16] King suggests that treating expectations under endogenous learning is a more natural assumption than the extreme view of rapid adjustment in expectations (e.g. rational expectations) or that of slow adjustment (e.g. adaptive expectations).

Taken in the context of IT frameworks in emerging economies, the combination of rapidly adjusting expectations (RE) and optimal policy rules derived from the full state vector on a disinflation path do not take into account that the new policy framework will not have full policy credibility. By its nature, full rational expectations requires that expectations adjust immediately (to within a white noise error) to the announced inflation target and therefore provides no scope for the (gradual) establishment of credibility and adjustment to the desired inflation target that we witness in practice. As such, a more plausible combination would consist of forward-looking policy rules, as in Batini and Haldane (1998) and Batini and Nelson (2000), and endogenous learning, as described in King (1996) and applied in Hall et al. (2000) and Orphanides and Williams (2004).

V. DISINFLATION AND OPTIMAL POLICY

In examining alternative speeds of disinflation, we take the actual *gradualist* approach (GRAD) pursued by the Central Bank of Chile over the period 1991 – 2001 as our benchmark and specify two alternative disinflation strategies relative to this actual inflation target path: (i) *ambitious* target path (AMB2) and (ii) *gradualist* target path (GRAD2). The former is defined as minus 2 percent relative to GRAD, while the latter is defined as plus 2 percent relative to GRAD, with all disinflation strategies converging on the authorities' long run target of 3 percent.

Optimal parameter values for ζ_{π} are derived for the IFB rule (7) using each version of the model (RE, AL, and HYB) and the three specifications (L1, L2, L3) for the central bank's objective function (9). The analysis proceeds by embedding (7) into each of the model structures and minimizing the respective loss function subject to the structural model, the goal weights, and the respective disinflation path to derive an optimal policy weight(s).

---

[16] Insert is our addition.

Table 4 lists the optimal policy weights from the simulations.[17] First, by construction, the inflation Hawk displays the strongest policy response followed by the Balanced and Dove policymakers, respectively, regardless of expectational assumptions and the speed of disinflation. Second, the most ambitious disinflation strategy (AMB2) elicits the strongest policy response followed by the gradualist (GRAD) and most gradualist (GRAD2) disinflation strategies, irrespective of expectational assumptions and loss preferences of the central bank. These results are quite intuitive given the nature of the exercise where *a priori* a positive correlation is expected between (i) the policymakers' aversion to inflation and the optimal policy weights, and between (ii) the ambitiousness of the disinflation strategy and the optimal policy weights.

Table 4: Optimal Policy Weights for ζ_{π}

	AMB2	**GRAD**	**GRAD2**
Rational expectations			
Hawk	1.546	1.534	1.512
Balanced	1.496	1.485	1.466
Dove	1.341	1.333	1.318
Adaptive learning			
Hawk	1.651	1.623	1.579
Balanced	1.624	1.598	1.556
Dove	1.566	1.543	1.505
Hybrid expectations			
Hawk	1.465	1.444	1.411
Balanced	1.447	1.427	1.397
Dove	1.405	1.388	1.361

A third observation is that the adaptive learning model yields optimal policy weights which are strictly greater than the corresponding values obtained with the rational expectations and hybrid expectations models. This result is quite strong and is similar to that reported by Orphanides and Williams (2004). When expectations incorporate a degree of imperfect knowledge it is optimal to increase the strength of the policy response relative to that of the perfect information solution. Since the lag structure inherent in the learning mechanism dictates that policy responds with a delay it is plausible that a stronger policy response is required to control inflation. Under rational expectations there is no lag in the policy response and so the instrument adjusts appropriately in the current period.

Fourth, the optimal response is always stronger under rational expectations than under hybrid expectations when there is a non-negligible weight on inflation in the central

[17] That all optimal weights are greater than one implies a desirable property analogous to the Taylor principle (highlighted by Bullard and Mitra (2002) in a learning framework) whereby the policy weight attached to deviations of inflation from its target is greater than one to ensure that the nominal interest rate responds more than one-for-one to changes in inflation. This ensures that the real interest rate does not fall when inflation rises.

bank's loss function (i.e. Hawk or Balanced). However, when there is little preference weight associated with inflation (i.e. Dove) then the optimal hybrid expectations policy response is stronger than that of the rational expectations model.

The former result suggests that when the authorities are perfectly rational and care about inflation volatility but face a learning structure then it is optimal to moderate the policy response over the disinflation horizon. This finding stands in contrast to that under the full adaptive learning model above where the optimal learning response is always stronger than the fully rational response. The reason for this departure in policy under the hybrid expectations structure is that the authorities are now aware that the combination of expectational rigidities and a forceful disinflation may lead to an overshooting of the target and hence destabilizing cycles in economic activity; whereas when both the authorities and the private sector are learning it is optimal to have a stronger policy response, perhaps helping the authorities to establish a degree of policy credibility in an uncertain environment.

The latter result suggests that when the central bank is a dove then the optimal policy response under the hybrid model there is a need for a stronger policy response to gain control of inflationary expectations over the disinflation horizon. Without sufficient policy control, expectations of inflation will become decoupled from those of the authorities resulting in unstable dynamics. Under full rational expectations, this additional control is not required since expectations automatically align themselves with those of the authorities; therefore, a slightly weaker policy response is optimal against the preference settings of the authorities.

It is useful to recall that the open economy framework employed here includes an additional transmission channel for monetary policy through the real exchange rate to affect inflation and output, and will therefore also affect the optimal policy response. In the model as the target for inflation is lowered the real exchange rate begins to appreciate reinforcing the downward dynamics of inflation. Reflecting a relatively rapid pass-through from exchange rate shocks to inflation, the real exchange rate affects inflation in the Phillips curve contemporaneously. In the full rational expectations model, the real exchange rate will "jump" downwards (appreciate) immediately in response to the inflation target shock and perhaps moderates the size of the optimal policy response. In the full adaptive learning model, the adjustment in the real exchange rate is felt with a lagged effect and perhaps amplifies the optimal policy response.

VI. TRADE-OFF FRONTIERS

The variation in the optimal policy weights described above suggests that the costs to disinflation will depend on the expectational assumptions employed. In line with the vast policy rule literature stemming from the seminal work of Taylor (1993), we derive variance trade-off frontiers for output and inflation to illustrate these costs. In particular, sample variances associated with each optimal policy simulation (Table 4) are calculated for output and inflation permitting us to examine the costs to disinflation associated with alternative policy preferences as well as with alternative disinflation strategies (Table A1 in the Appendix details the variances).

Figures 2-4 show output and inflation frontiers grouped by disinflation strategy, illustrating the trade-offs associated with alternative policy preferences as measured against each disinflation strategy. Figures 5-7 show the same frontiers grouped by preference function and illustrate the trade-offs associated with alternative disinflation strategies as measured against each policy preference.

An initial observation from the first group of frontiers (Figures 2-4) is that each expectations structure displays a "clustering" in their respective outcomes. For example, all of the AL trade-offs frontiers imply higher inflation volatility than output volatility. The most favourable cluster comes from the HYB model which appears to balance the weak inflation control of the AL model and the weak output control of the RE model. Within each respective cluster the inflation volatility is lower the more hawkish the central banker. These findings reflect the ordering of the optimal policy weights in Table 4.

Figure 2: AMB2 disinflation strategy under different policy preferences and expectations

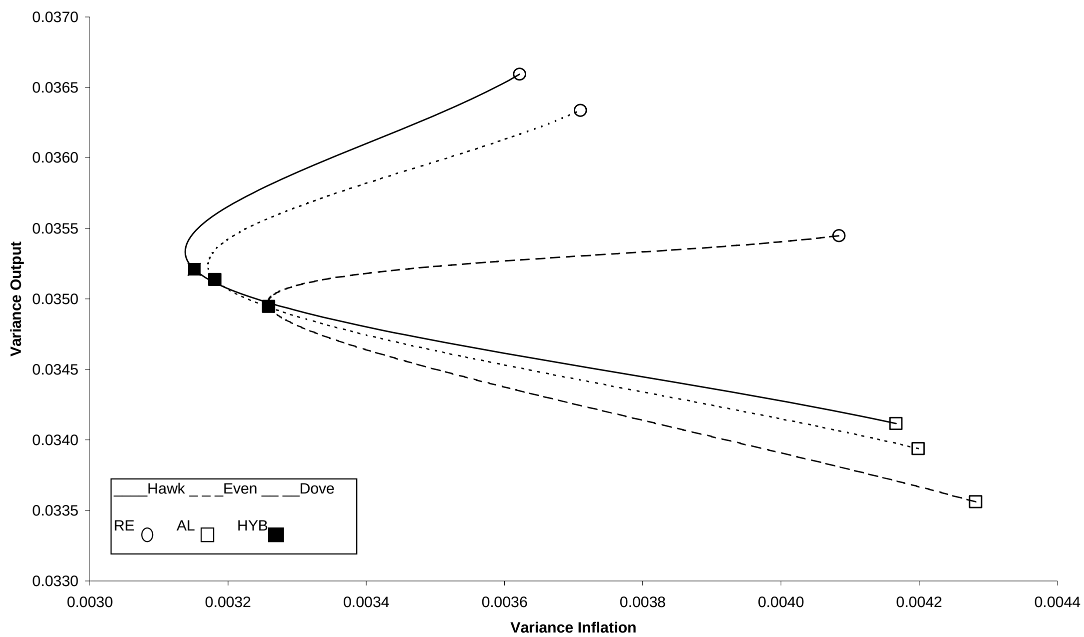

Figure 3: GRAD disinflation strategy under different policy preferences and expectations

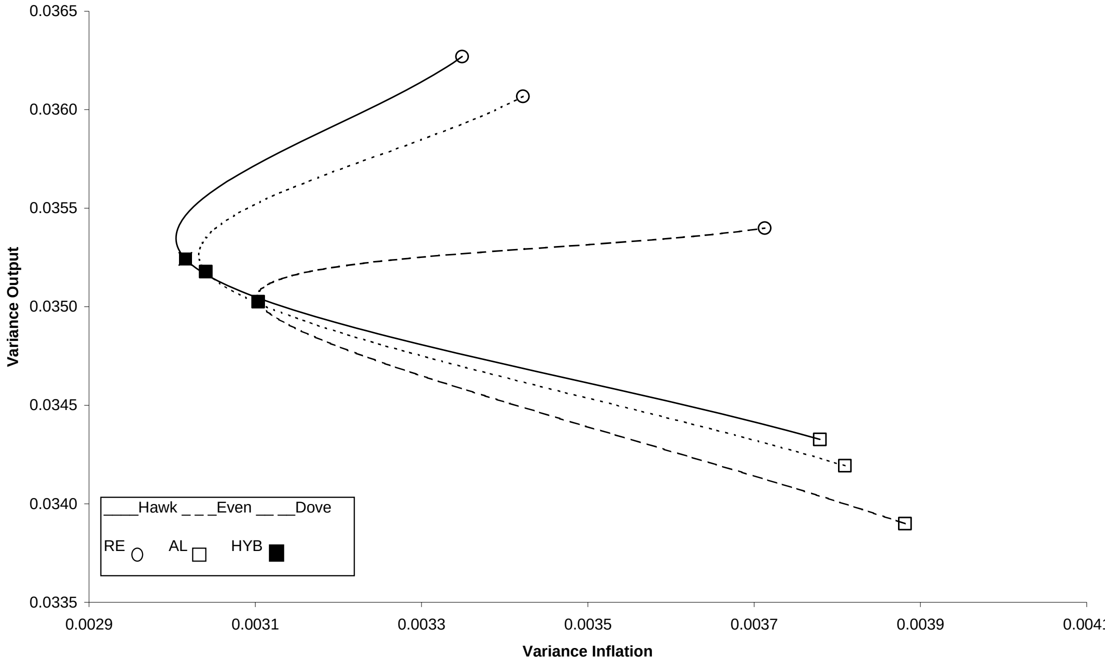

Figure 4: GRAD2 disinflation strategy under different policy preferences and expectations

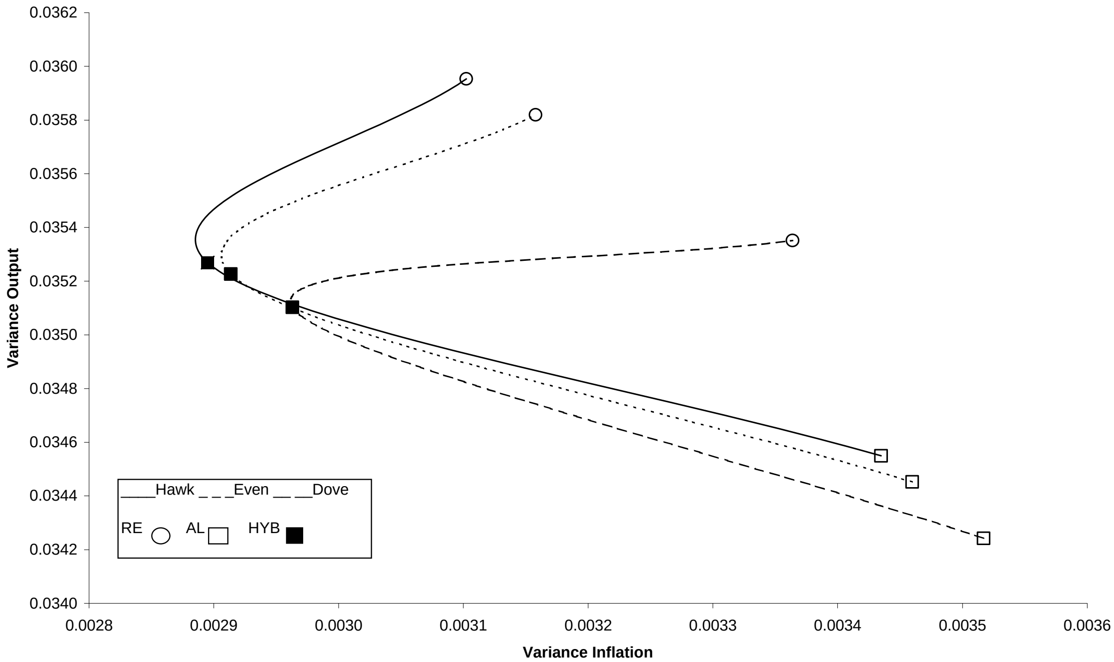

Turning now to the tradeoffs associated with alternative speeds of disinflation (Figures 5-7) some key findings emerge. First, there is a positive correlation between the ambitiousness of the strategy and the volatility of inflation. The more ambitious the disinflation strategy the further to the east its frontier appears.[18] This result is congruent across expectation structures and central bank preferences and suggests that, in terms of inflation volatility, it pays to be more gradual. Although, this finding may not be too surprising since more ambitious – and hence stronger optimal – policy may cause inflation to fall quickly and possibly overshoot its long run target leading to greater volatility. Therefore, the interesting question becomes: is there any gain in terms of output stability from a more ambitious disinflation?

This highlights the second point, namely that there is a clear positive trade-off for output stability under the AL model, but a clear negative trade-off under the RE model; the trade-off under the HYB model is similar to that under the AL model though it is less pronounced. These findings imply that there is a strict deterrent to ambitious disinflation under the full rational expectations model; the more ambitious the disinflation the worse off the economy is in terms of both inflation and output variation.[19] Under full adaptive learning, however, the policymakers face a real choice: an ambitious disinflation strategy is less costly in terms of output, but it must be measured against increasing inflation volatility.

As discussed above, expectations adjust with a lagged effect in the adaptive learning model. And given that rigidities in the adjustment of expectations can give rise to large output costs during the disinflation process it makes sense that a more ambitious and stronger policy response will be associated with less output volatility. The reason is that a stronger policy response under learning helps to align expectations with those of the authorities more quickly reducing the costs of adjustment.[20] However, under full rational expectations, there are no inherent rigidities to adjustment in expectations and so inflation adjusts (perhaps) too quickly under strong policy. Therefore, given the transmission mechanisms of monetary policy, one may find that inflation is more likely to overshoot under more ambitious policies leading to increased output volatility in the RE model.

These findings may help shed some empirical light on the gradualist approach to disinflation which was actually pursued by the Chilean authorities. Under the assumption of fully rational agents and fully credible policy a gradualist policy dominates with lower

---

[18] For example, this is seen in Figure 5 by the solid line "AMB2 frontier" and the corresponding trade-off points for each expectations structure.

[19] It can be seen by the flattening of the tops of the curves in Figures 5-7 that the negative implications for output stability in the RE model are greater under Hawk preferences and less under Dove preferences. However, this is clearly the result of the optimization routine whereby volatility in output is punished in the latter but not in the former.

[20] Or as Orphanides and Williams (2004) state, policies that fail to maintain a tight control over inflation are prone to episodes in which the public's expectations of inflation become uncoupled from the policy objective.

output and inflation variability. However, both the full learning and hybrid expectations models suggest that there is a real policy trade-off for the authorities to explore where a more ambitious disinflation strategy would elicit greater inflation volatility, but would be rewarded with less output volatility.

Figure 5: Hawk policy preferences and all disinflation strategies

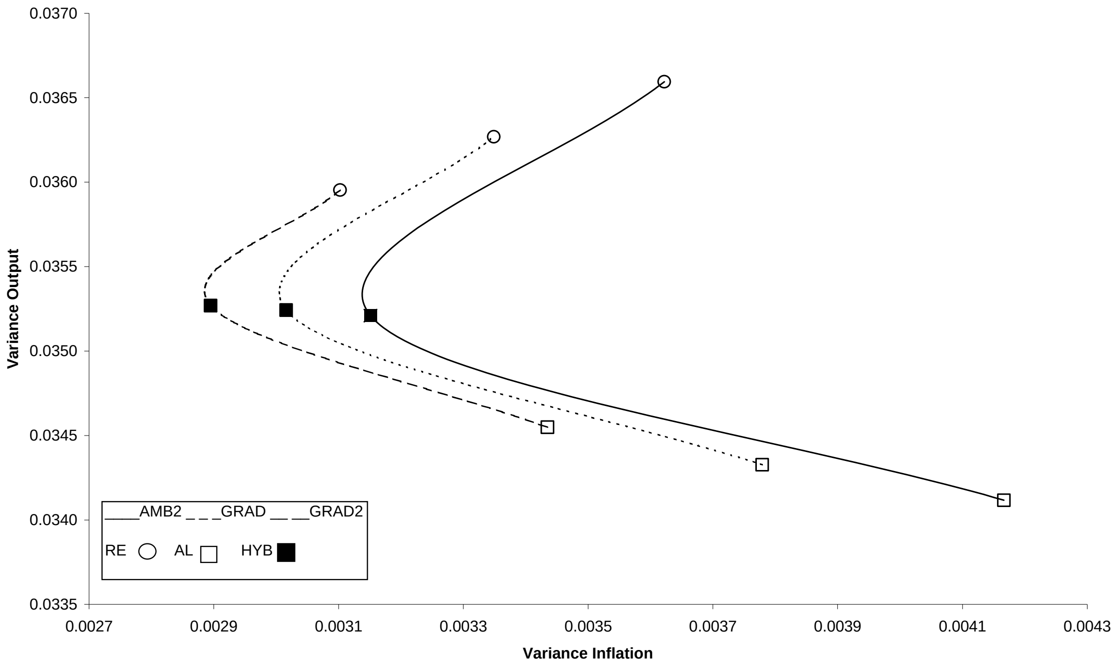

Figure 6: Balanced policy preferences and all disinflation strategies

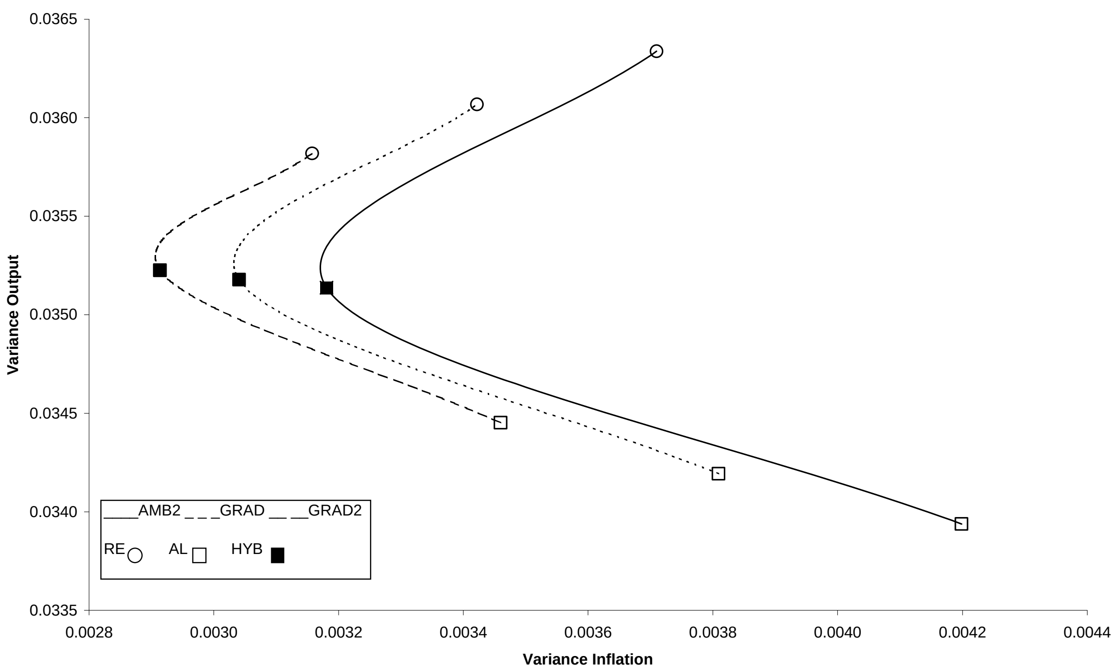

Figure 7: Dove policy preferences and all disinflation strategies

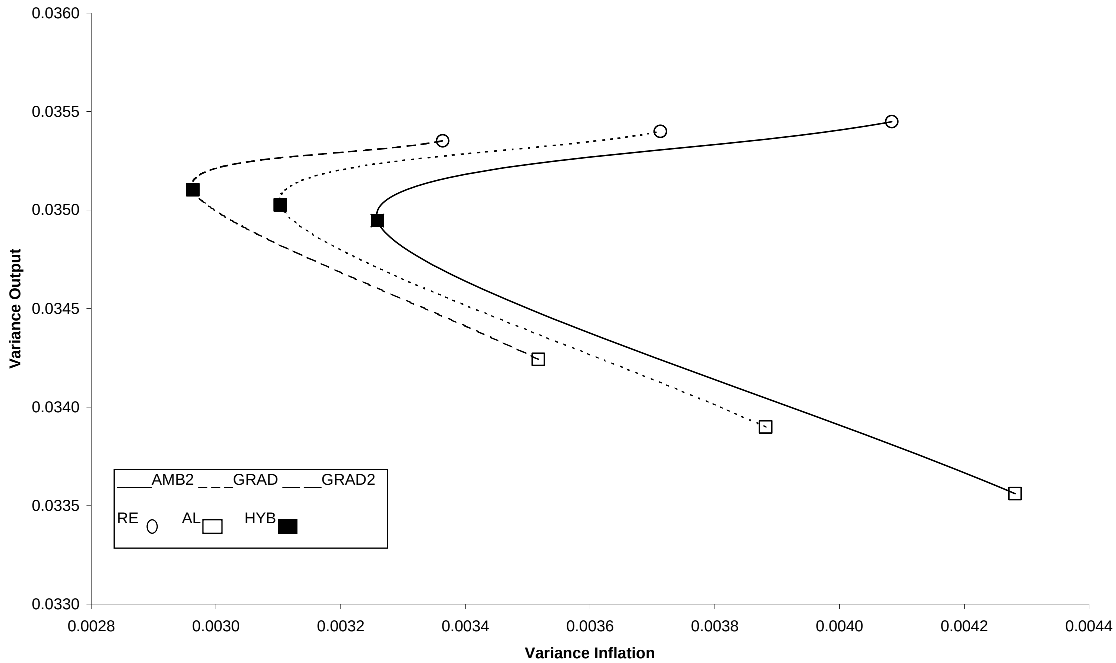

VII. SENSITIVITY ANALYSIS: NOMINAL INERTIA

The above results are in line with the literature which suggests that the costs to disinflation are to some extent dependent upon the adjustment in expectations (e.g. Sargent (1986) and King (1996)). Moreover, the use of adaptive learning adds an additional layer of dynamic adjustment in expected variables which, as we have found, gives rise to a stronger optimal policy response (e.g. in line with Orphanides and Williams (2004)). Therefore, it is instructive to assess the impact on these optimal policy weights from a change in the degree of nominal inertia as set out in the structural model.

Here we re-examine the optimal disinflation response by calculating new policy weights while changing the degree of nominal inertia. This is done by varying the coefficient β_1 on lagged inflation in the Phillips curve (2) where a smaller (larger) weight is associated with less (more) nominal inertia (.i.e. more (less) forward-looking).[21] The results are detailed in Table 5.

In both models, it is clear that, regardless of the central bank's preferences, a reduction in the degree of inflation persistence in the structural model elicits a stronger optimal policy response; the less is the nominal inertia, the smaller are the real costs associated with a faster disinflation, so the greater is the incentive to disinflate more quickly. This highlights the thrust behind the Sargent-type disinflation strategy where, for example,

[21] The focus here has been restricted to the extreme cases: (i) full rational expectations and full adaptive learning; (ii) Hawk and Dove preferences and (iii) AMB2 disinflation strategy.

under rational expectations the perfectly credible policy would emerge with no barrier to the instantaneous adjustment of prices (i.e. $\beta_1 = 1.0$). In the same vein, an increase in nominal inertia elicits a smaller optimal policy response to minimize the costs associated with disinflation. This is representative of the thrust behind the King-type disinflation strategy where the greater the rigidity the more costly it is to disinflate rapidly.

Table 5: Optimal policy weights for δ_π with changing nominal inertia under AMB2 disinflation strategy

	Rational Expectations		**Adaptive Learning**	
	Hawk	*Dove*	*Hawk*	*Dove*
β_1 =0.1	1.632	1.446	1.865	1.870
β_1 =0.5	1.578	1.377	1.726	1.710
β_1 =0.9	1.539	1.334	1.609	1.453

Note: Nominal inertia is governed by weight β_1 on lagged inflation in Phillips curve. The smaller the weight the less inertia.

The results derived here are in line with those obtained by Orphanides and Williams (2004) who find that the more forward-looking the inflation process the stronger the efficient policy response. In this regard, it is also interesting to note that in the adaptive learning model the more forward-looking the inflation process the weaker the distinction between a Hawk and a Dove becomes – little or no drag from inflation dynamics implies similar costs in terms of output.

VIII. CONCLUSION

This paper employs a stylized New Keynesian Model to examine alternative disinflation strategies in Chile under optimal monetary policy setting. The analysis is enriched through the adoption of alternative expectations structures. The assumption of adaptive learning provides an additional layer of dynamic interactions that help illuminate key policy trade-offs not necessarily captured under the full rational expectations solution. Given the uncertainty associated with the adoption of a new and untried policy regime, this brief departure from perfect knowledge is seen as a reasonable and complementary framework for the disinflation analysis conducted above.

Our analysis suggests that there is a strict deterrent to ambitious disinflation under full rational expectations since the volatility of inflation and output increase. However, under full adaptive learning, a policy trade-off emerges where the authorities must weigh the merits of a faster disinflation that is associated with less output volatility but increasing inflation volatility against a more gradual disinflation with the opposite effects.

Appendix

Table A1: Output and Inflation Variances

	Hawk		**Balanced**		**Dove**	
	π	y	π	y	π	y
AMB2						
RE	0.0042	0.0341	0.0042	0.0339	0.0043	0.0336
AL	0.0032	0.0352	0.0032	0.0351	0.0033	0.0349
HYB	0.0036	0.0366	0.0037	0.0363	0.0041	0.0354
GRAD						
RE	0.0038	0.0343	0.0038	0.0342	0.0039	0.0339
AL	0.0030	0.0352	0.0030	0.0352	0.0031	0.0350
HYB	0.0033	0.0363	0.0034	0.0361	0.0037	0.0354
GRAD2						
RE	0.0034	0.0345	0.0035	0.0345	0.0035	0.0342
AL	0.0029	0.0353	0.0029	0.0352	0.0030	0.0351
HYB	0.0031	0.0360	0.0032	0.0358	0.0034	0.0354

Uncovering the Kalman filter learning algorithm

Using the work of Garratt and Hall (1997) and Bullard (1992) this appendix describes how the general approach of Marcet and Sargent (1989b) can be rewritten to take the form of the Kalman filter learning algorithm. As such, we are able to draw upon the convergence criteria set out in Marcet and Sargent (1989a, b).

We begin by defining the state of the economy at time *t* as being characterized by the values of the elements of a $k \times 1$ vector z_t. Agents observe only a portion of the elements of z_t which are important for future decisions, hence future values of the unobserved variables in z_t should be forecast. Let $z_{1t} = e_1 z_t$ be a $k_1 \times 1$ subvector of z_t, where the subscript 1 indicates that future values of these variables must be predicted and e_1 is a matrix that selects from z_t. Let $z_{2t} = e_2 z_t$ be a $k_2 \times 1$ subvector of z_t, where the subscript 2 indicates that these are the variables agents use to predict z_{1t}. Therefore, agents are concerned only with

$$z_{1t} = f\left(z_{2t-1}\right), \tag{A1}$$

where $f(\cdot)$ is the forecast function. This specification assumes a lagged information structure, but contemporaneous information can be incorporated. The form of the law of motion (or the PLM) is dependant upon the forecast function $f(\cdot)$ used by agents and Marcet and Sargent (1989a, b) analyze

$$z_{1t} = \beta_t z_{2t-1} + \eta_t , \tag{A2}$$

where β_t is a $k_1 \times k_2$ matrix of parameters, and η_t is a $k_1 \times 1$ vector of noise such that $E\eta_t = 0$ and $E\eta_t \eta_t' = H$. Therefore, the actual law of motion is given by

$$\begin{bmatrix} z_{1t} \\ z_{1t}^c \end{bmatrix} = \begin{bmatrix} T_{11}(\beta_t) & T_{12}(\beta_t) \\ T_{21}(\beta_t) & T_{22}(\beta_t) \end{bmatrix} \begin{bmatrix} z_{2t-1}^c \\ z_{2t-1} \end{bmatrix} + \begin{bmatrix} V_1(\beta_t) \\ V_2(\beta_t) \end{bmatrix} u_t , \tag{A3}$$

where $T_{11}(\cdot)$ is $k \times (k_1 - k_2)$, $T_{12}(\cdot)$ is $k_1 \times k_2$, $[T_{21}(\beta_t) T_{22}(\beta_t)]$ is $(k - k_1) \times k$, $V_1(\cdot)$ is $k_1 \times m$, $V_2(\cdot)$ is $(k - k_1) \times m$ and u_t is an $m \times 1$ vector of white noise such that $Eu_t u_t' = U \forall t$. The superscript *c* indicates the compliment and for simplification we set $T_{11}(\beta_t) = 0$. Equation (A3) can be written as

$$z_t = T(\beta_t) z_{t-1} + V(\beta_t) u_t . \tag{A4}$$

The partitioned operators *T* and *V* are of conformable dimensions and are determined by the economic model employed. The PLMs given by (10)-(12) above are nested within the framework given by (A4) and the learning algorithm is any method employed to update the parameters of the time-varying matrix β_t (or the ρ_{it} given by our equation(s) (13)). In this study, we employ the Kalman filter updating algorithm. However, Marcet and Sargent (1989a, b) assume that agents use ordinary least squares to update β_t and, using Ljung and Söderström (1983), this can be written as

$$\overline{\beta}_t' = \beta_{t-1}' + (\alpha_t / t) R_t^{-1} \{ z_{2t-2} z_{2t-2}' [T_{12}(\beta_{t-1})' - \beta_{t-1}'] + z_{2t-2} u_{t-1}' V_1(\beta_{t-1})' \}, \tag{A5a}$$
$$\overline{R}_t' = R_{t-1} + (\alpha_t / t)[z_{2t-2} z_{2t-2}' - R_{t-1}/\alpha_1], \tag{A5b}$$

where $R_t = [1/(t-1)] \sum_{j=2}^{t} \alpha_j z_{2j-2} z_{2j-2}'$ is a $k \times k$ matrix and $\{\alpha_t\}_{t=0}^{\infty}$ is a positive sequence. The α_t sequence can be set to weight observations appropriately. The algorithm given by (A5a) and (A5b) is a version of weighted least squares and can be combined with a projection facility to create a learning mechanism (our equations (10-12) and (13)). Using Bullard (1992), (A5a) and (A5b) can be written as a Kalman filter.

If we ignore the projection facility and set the left-hand sides of (A5a) and (A5b) equal to β_t and R_t, respectively, then we can rewrite these equations in terms of a new matrix P_t. Since $z_{1t-1}' = z_{2t-2}' T_{12}(\beta_{t-1})' + u_{t-1}' V(\beta_{t-1})'$ then

$$\beta_t' = \beta_{t-1}' + (\alpha_t / t) R_{t-1}^{-1} z_{2t-2} [z_{1t-1}' z_{2t-2}' \beta_{t-1}'], \tag{A6a}$$
$$R_t = R_{t-1} + (\alpha_t / t)[z_{2t-2} z_{2t-2}' - R_{t-1}/\alpha_1]. \tag{A6b}$$

Now define $P_t = (1/t)R_{t-1}^{-1}$, so that we can write the algorithm as

$$\beta_t' = \beta_{t-1}' + \alpha_t P_t z_{2t-2}\left[z_{1t-1}' - z_{2t-2}'\beta_{t-1}'\right], \tag{A7a}$$
$$P_t^{-1} = P_{t-1}^{-1} + \alpha_t z_{2t-2} z_{2t-2}' . \tag{A7b}$$

Using the lemma $[W + XYZ]^{-1} = W^{-1} - W^{-1}X\left[ZW^{-1}X + Y^{-1}\right]^{-1} ZW^{-1}$ and by setting $W = P_{t-1}^{-1}$, $X = z_{2t-2}$, and $Z = z_{2t-2}'$ the following algorithm can be developed

$$\overline{\beta}_t' = \beta_{t-1}' + P_{t-1} z_{2t-2} f_{t-1}^{-1}\left[z_{1t-1}' - z_{2t-2}'\beta_{t-1}'\right], \tag{A8a}$$
$$\overline{P}_t = P_{t-1} - P_{t-1} z_{2t-2} z_{2t-2}' P_{t-1} f_{t-1}^{-1}, \tag{A8b}$$

where the scalar $f_{t-1} = (1/\alpha_t) + z_{2t-2}' P_{t-1} z_{2t-2}$. Equations (A8a) and (A8b) are now directly interpretable as standard Kalman filter updating equations where the term in brackets in the former equation is the expectations error and the first part of the expression for f_t is set equal to *W*. The single parameter α governs the rate of change for all parameters on the Kalman updating equations (our equation(s) (13)). In this study a general matrix is applied so that parameters may change at varying rates and it is in this sense that the Kalman filter applied here is more general than that used by Marcet and Sargent (1989a, b).

References

Ball, L. (1998), “Efficient Rules for Open Economies,” NBER Working Paper No. 6760, National Bureau of Economic and Research.

Barro, R. J., and D. B. Gordon (1983), “Rules, Discretion and Reputation in a Model of Monetary Policy,” *Journal of Monetary Economics*, 12, 101 – 121.

Batini, N., and A. G. Haldane (1998), “Forward-Looking Rules for Monetary Policy,” Working Paper No. 91, Bank of England.

Batini, N., and E. Nelson (2000), “Optimal Horizons for Inflation Targeting,” Working Paper No. 119, Bank of England.

Beeby, M., S. G. Hall, and S. G. B. Henry (2002), “Rational Expectations and Near Rational Alternatives: How Best to form Expectations,: CIM Discussion Paper 2002.01, Centre for International Macroeconomics, University of Oxford.

Bernanke, B., and S. Gilchrist (1998), “The Financial Accelerator in a Quantitative Business Cycle Framework,” NBER Working Paper No. 6455, National Bureau of Economic Research.

Bullard, J. (1992), “Time-varying parameters and nonconvergence to rational expectations under least squares learning,” *Economics Letters*, 40, 159-166.

Bullard, J., and K. Mitra (2002), “Learning about Monetary Policy Rules,” *Journal of Monetary Economics*, 49, 1105-1129.

Caputo, R. (2004), “Persistence and the role of exchange rate and interest rate inertia in monetary policy,” Central Bank of Chile, Working Paper No. 300.

CEF (2000), “CEF Modelling Suite,” www.cef-software.co.uk, Directors: Prof. Stephen G. Hall, Prof. Brian Henry, and Anne Shepard.

Central Bank of Chile (2002), “Constitutional Organic Act of the Central Bank of Chile: Law 18,840,” Publication, Central Bank of Chile.

Clarida, R., J. Gali, and M. Gertler (1999), “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37(4), 1661 – 1707.

Evans, G., and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*. Princeton University Press.

Evans, G., and S. Honkapohja (2003), “Adaptive Learning and Monetary Policy Design,” CEPR Discussion Paper No. 3962.

Garcia, P. S. (2000), “The Chilean Experience with Inflation Targeting,” in *Inflation Targeting Experiences*, Central Bank of the Republic of Turkey.

Garratt, A., and S. G. Hall (1995), “Model Consistent Learning and Regime Switching in the London Business School Model,” *Economic Modelling*, 12, 87 – 95.

_____ (1997), “E-equilibria and adaptive expectations: Output and inflation in the LBS model, *Journal of Economic Dynamics and Control*, 21, 1149-71.

Hall, S. G., S. G. B. Henry, and J. Nixon (2000), “Inflation Targeting: Delegation and Coordination of Monetary Policy,” CIM Discussion Paper No. 2000.08, Centre for International Macroeconomics, University of Oxford.

Hall, S. G., and J. Nixon (1997), “Controlling Inflation: Modelling Monetary Policy in the 1990s,” in *Macroeconomic Modelling in a Changing World*, eds. C. Allen, and S. G. Hall. John Wiley and Sons.

Henry, S. G. B., Satchi, M., and Vines, D. (2006), “The Effect of Discounting on Policy Choices in Inflation Targeting Regimes,” The *Economic Journal*, 116 (508), 266-282.

Henry, S. G. B., and Pagan, A. R. (2004), “The econometrics of the New Keynesian Policy Model: Introduction,” *Oxford Bulletin of Economics and Statistics*, 66, 581-608.

King, M. (1996), “How Should Central Banks Reduce Inflation: Conceptual Issues,” *Economic Review*, Fourth Quarter, pp. 25 – 52, Federal Reserve Bank of Kansas.

Marcet, A., and T. J. Sargent (1989a), “Convergence of Least Squares Learning in and Environment with Hidden State Variables and Private Information,” *Journal of Political Economy*, 97, 1306 – 22.

Marcet, A., and T. J. Sargent (1989b), “Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337 – 68.

Morande, F. G. (2001), “A Decand of Inflation Targeting in Chile: Developments, Lessons, and Challenges,” Working Paper No. 115, Central Bank of Chile.

Orphanides, A. and J. C. Williams (2004), “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation Targeting Debate*, Ben Bernanke and Michael Woodford (eds.), Chicago: University of Chicago Press.

Sargent, T. J. (1986), *Rational Expectations and Inflation*. New York, Harper and Row.

Schaechter, A. M., Stone, M. R., and Zelmer, (2000), “Adopting Inflation Targeting: Practical Issues for Emerging Market Economies,” IMF Occasional Paper 202, International Monetary Fund, Washington DC.

Svensson, L. E. O. (1997a), “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review*, 41, 1111 – 1146.

Svensson, L. E. O. (1997b), “Open-Economy Inflation Targeting,” *Journal of International Economics*,” 50, 155 – 83.

Svensson, L. E. O., and M. Woodford (1999), “Implementing Optimal Policy Through Inflation-Forecast Targeting,” Princeton University Press.

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