Stabilisation policy, rational expectations and price-level versus inflation targeting: a survey

Michael Hatcher (University of Glasgow)
Patrick Minford (Cardiff University and CEPR)

15 June 2014

Abstract

We survey literature comparing inflation targeting (IT) and price-level targeting (PT) as macroeconomic stabilisation policies. Our focus is on New Keynesian models and areas that have seen significant developments since Ambler's (2009) survey: optimal monetary policy; the zero lower bound; financial frictions; and transition costs of adopting a PT regime. Ambler's conclusion that PT improves social welfare in New Keynesian models is fairly robust, but we note an interesting split in the literature: PT consistently outperforms IT in models where policymakers commit to simple Taylor-type rules, but results in favour of PT when policymakers minimise loss functions are overturned with small deviations from the baseline model. Since the beneficial effects of PT appear to hang on the joint assumption that agents are rational and the economy New Keynesian, we discuss survey and experimental evidence on rational expectations and the applied macro literature on the empirical performance of New Keynesian models. Overall, the evidence is not clear-cut, but we note that New Keynesian models can pass formal statistical tests against macro data and that models with rational expectations outperform those with behavioural expectations (i.e. heuristics) in direct statistical tests. We therefore argue that policymakers should continue to pay attention to PT.

Keywords: inflation targeting, price-level targeting, rational expectations, stabilisation policy.

1   Introduction

The past two decades have seen a resurgence of interest in optimal monetary policy, including research on potential alternatives to inflation targeting (IT). In this survey, we focus on an alternative policy that is considered a serious contender to IT: price-level targeting (PT). The main difference between IT and PT is that, under a PT regime, the central bank aims to stabilise the aggregate price level around a predetermined target price path. Hence, for example, if there is an inflationary shock that takes the price level above the

target price path, below-average inflation will be required in the future in order to return the price level to target. This contrasts with an IT regime where inflation stabilisation is the goal of policy and above average inflation today would be followed by average (i.e. on-target) inflation tomorrow. In other words, 'bygones are bygones' under IT, whereas past deviations from target are corrected under a PT regime. We focus on PT because it has been the subject of considerable debate in recent years – both in academia and central banks. Most notably, the Bank of Canada recently published a detailed review of the costs and benefits of PT which argues that the potential benefits do not justify the costs and risks associated with a change in regime (see Bank of Canada, 2011). The Bank of Canada expressed concerns, in particular, about whether agents would understand the workings of a PT regime and take it into account in their future expectations in accordance with the standard modelling assumption of rational expectations.

In this survey, we shed more light on the likely costs and benefits of PT by reviewing recent theoretical and empirical evidence, and we also engage with the Bank of Canada's conclusions. Our survey begins by discussing recent developments in the monetary policy literature before turning to the macro stabilisation literature on IT versus PT. We focus on New Keynesian models and important developments since Ambler's (2009) survey: optimal monetary policy under commitment and discretion; the zero lower bound on nominal interest rates; financial frictions; and the costs of transition to a PT regime. Some of these recent developments are initial attempts to address issues raised by the financial crisis. The main theoretical finding of the survey is that Ambler's original conclusion that PT can improve social welfare in New Keynesian models remains intact, but the list of caveats is somewhat longer. In particular, we note an interesting split in the literature: PT consistently outperforms IT in New Keynesian models where policymakers commit to simple Taylor-type rules, but results in favour of PT which are derived from models where central banks minimise loss functions can be overturned with small deviations from the baseline model. In addition, the rational expectations assumption is crucial when considering the costs of transition to a PT regime. These findings encouraged us to focus on the assumptions on which the beneficial effects of PT hang: are economic agents rational and the economy New Keynesian?[1]

With this question in mind, the second half of the survey reviews recent experimental and survey literature on rational expectations, and the applied macro literature on New Keynesian models. Both of these areas have seen important developments in recent years and they offer different insights. Both experiments and surveys tend to reject the strict rational expectations hypothesis, but the experimental literature provides some support for rational expectations as an approximation, while survey evidence should not be interpreted as conclusive statistical evidence against rational expectations (for reasons that we discuss). On the other hand, the applied macro literature initiated by Christiano et al.

---

[1] Other useful surveys of price-level targeting include Cournède and Moccero (2009) and Kahn (2009). Neither Ambler (2009) or these papers survey empirical evidence on rational expectations and New Keynesian models.

(2005) is broadly supportive of New Keynesian models with rational expectations but should be treated with caution because it does not amount to a formal statistical test that accepts or rejects the model. There is, however, a more recent strand of literature that has formally tested New Keynesian models using indirect inference. We explain this approach before discussing some of the main results from this literature. Overall, the evidence on rational expectations and New Keynesian models is not clear-cut, but we note that New Keynesian models can pass formal statistical tests against macro data and that models with rational expectations outperform those with behavioural expectations (i.e. heuristics) in direct statistical tests.[2] These findings suggest that the benefits of PT might be higher than envisaged by many policymakers, including the Bank of Canada. We therefore argue that policymakers should continue to pay attention to PT in the future.

2 Background: New Keynesian models and the Bank of Canada Review

In this section we first set out the baseline New Keynesian model which provides the foundation for the macro stabilisation literature on IT versus PT. We then briefly discuss the literature on sticky prices in relation to the Calvo (1983) price-setting mechanism that is central in New Keynesian models. Finally, we provide an overview of the main conclusions of the Bank of Canada's 2011 Review of PT to provide context for our detailed survey of the literature that follows.

2.1 The baseline New Keynesian model

The distinguishing feature of New Keynesian models is that monopolistically-competitive firms set prices optimally to maximise profits, subject to constraints on how frequently they can re-set prices or how costly it is to do so. In the baseline specification due to Calvo (1983), each firm can change its price with a constant probability, so the interval between price changes is a random variable. Profit maximisation by the firms that are able to re-optimise leads to a first-order condition that relates current price to a mark-up on marginal cost and expected future price. Log-linearising this first-order condition around a zero-inflation steady-state and aggregating across firms leads to the baseline New Keynesian Phillips curve in which economy-wide inflation π_t depends on the output gap x_t and expected future inflation:

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

where β is the discount factor of the representative household (the sole owner of firms), E_t is the conditional expectations operator, $\kappa > 0$ is the slope of the

---

[2] We provide a formal definition of behavioural expectations in Section 2.3. We discuss in Sections 2.3 and 5.1 how these expectations relate to alternative approaches in the literature.

Phillips curve, and u_t is a cost-push shock that follows an AR(1) process.[3]

Equation (1) states that inflation depends positively on the output gap and (rationally) expected future inflation. The main way in which the New Keynesian Phillips curve differs from a more traditional neoclassical Phillips curve is that inflation depends on expected future inflation ($E_t\pi_{t+1}$), and not expected current inflation ($E_{t-1}\pi_t$). This feature of the model is plausible because there is substantial empirical evidence that anticipated changes in monetary policy have real effects. It also implies that inflation depends on the expected discounted stream of future output gaps and cost-push shocks, so that managing expectations about the future is crucial for current inflation control.

The demand side of the model is standard: consumers maximise utility by choosing bond holdings optimally, giving an Euler equation for each household. If we aggregate across households we get a single consumption Euler equation in which aggregate consumption is a function of the real interest rate and expected future consumption. Log-linearising this equation and imposing the market-clearing condition that consumption equals output minus government spending, we get the IS curve in the baseline New Keynesian model:

$$x_t = E_t x_{t+1} - \sigma(R_t - E_t\pi_{t+1}) + g_t \tag{2}$$

where $\sigma > 0$ is the intertemporal elasticity of substitution, x_t is the output gap, R_t is the nominal interest rate, and g_t is a government spending shock that follows an AR(1) process.

Like the New Keynesian Phillips curve, the IS curve is forward-looking: it emphasises the importance of expectations about the future in the determination of current outcomes. In particular, the current output gap rises with the expected future output gap (due to consumers' desire to smooth consumption) and falls when the short-term real interest rate, $R_t - E_t\pi_{t+1}$, is increased.

In addition to the optimising microfoundations of the New Keynesian model, an important feature is that a social loss function in inflation and output gap variations can be derived as an approximation to the utility function of the representative household; see e.g. Walsh (2010, Ch. 8.6). As a result, reseachers can conduct social welfare analyses which are internally consistent with the model itself and broadly consistent with the stated objectives of real-world central banks. The approximate social loss function in the baseline New Keynesian model can be represented as follows:

$$L_t = E_t \sum_{j=0}^{\infty} \beta^j \left(\pi_{t+j}^2 + \alpha x_{t+j}^2 \right) \tag{3}$$

---

[3] The cost-push shock is appended to this equation because it introduces a meaningful trade-off between inflation and output gap volatility. Alternatively, a Phillips curve that includes a cost-push shock can be derived directly by allowing for exogenous fluctuations in firms' demand elasticities or exogenous fluctuations in taxes (see Steinsson, 2003).

where $\alpha > 0$ is the relative weight on output gap variations, which depends upon the parameters of the model.

This social loss function has two important roles. First, Equation (3) can be used as a basis for comparing social welfare across alternative monetary policy regimes like IT and PT; see Vestin (2006). Second, by minimising Equation (3) subject to (1) by choosing inflation and the output gap directly, we can derive first-order condition that implements the optimal monetary policy. This condition provides insight into the principles of effective policymaking and allows us to analyse the optimal response of interest rates to shocks that hit the economy; this can be done by substituting the first-order condition into Equation (2). The optimal responses of inflation and the output gap can also be 'backed out' in a similar manner (see e.g. Clarida et al., 1999).

The approach of deriving optimal policies has been popular in the theoretical literature. This requires the researcher to make an assumption about whether the central bank re-optimises on a period-by-period basis ('discretion') or can make binding promises about the entire future path of monetary policy ('commitment'). For instance, with the baseline New Keynesian model and the assumption of commitment, Evans and Honkapohja (2006) show that the optimal commitment policy and the *implied* interest rate rule are as follows:

$$\pi_t = -\frac{\alpha}{\kappa}(x_t - x_{t-1}) \tag{4}$$

$$R_t = c_1 x_{t-1} + c_2 g_t + c_3 u_t \tag{5}$$

where c_1, c_2 and c_3 are complicated functions of the model parameters.

An alternative to the optimal policy approach is to 'close' the model with an instrument rule for the nominal interest rate. The most commonly used interest rate rule is a log-linear Taylor rule (Taylor, 1993):

$$R_t = c_\pi \pi_t + c_x x_t \tag{6}$$

where c_π, $c_x > 0$ are the reaction coefficients on inflation and the output gap, respectively; we treat the inflation target and steady state interest rate as zero here and in what follows.

As we can see by comparing Equations (5) and (6), Taylor rules are generally suboptimal policies – i.e. they will not minimise Equation (3). Nevertheless, Taylor rules have become a popular way of modelling monetary policy because they are easy to communicate; more robust across alternative models of the economy than optimal policies (see Taylor, 1999); and a useful way of describing the behaviour of central banks in estimated general equilibrium models.

2.2 Calvo contracts and micro evidence on sticky prices

There is considerable empirical evidence suggesting that prices are 'sticky' in nominal terms as assumed by Calvo (1983) price-setting. In a seminal paper, Bils and Klenow (2004) conclude that average price spells in the US last between 3 and 4 months, though subsequent results in Nakamura and Steinsson (2008) that exclude sales point to longer price spells of 7 to 11 months – a result which is more consistent with early empirical work on price stickiness in the US. Dhyne et al. (2006) find similar results for the Euro Area as part of research conducted by the Inflation Persistence Network at the European Central Bank (ECB): the average duration of price spells is 10 to 13 months. Consistent with these findings, the Calvo reset probability is typically calibrated to imply average price spells of between 2 and 4 quarters in New Keynesian models.

The factors responsible for nominal price rigidities remain somewhat of a mystery, however. Under Calvo price-setting these factors are not modelled: firms are literally unable to change their prices in some periods, and the probability of being able to change price in any given period is independent of the date at which prices were last re-set. These are clearly not realistic assumptions, but they make the model highly tractable. And because the group of firms that is able to change price is selected randomly in each period, the Calvo specification implies that, over a given period of time, some firms will be able to change their prices more frequently than others, as we see in the data (e.g. Bils and Klenow, 2004). Therefore, the Calvo model captures some important features of real-world price-setting within a highly tractable framework.[4]

Calvo price-setting has, however, been criticised along several lines. It is well known, for example, that the assumption of a constant reset probability means that the Calvo model cannot explain why firms reset their prices more frequently when inflation is high. In addition, some authors have noted that the standard New Keynesian model with Calvo price-setting implies that moderate to high levels of trend inflation have implausibly large effects on the deterministic steady-state and stochastic mean levels of output (Ascari, 2004; Amano et al., 2007). Several alternative models of nominal rigidity have been proposed in the literature, including state-dependent pricing models where firms change prices more frequently when it is more profitable to do so (e.g. Golosov and Lucas, 2007) and generalised Taylor models where there are several sectors with Taylor (1980) contracts of different durations (e.g. Dixon and Le Bihan, 2012).[5]

Most notably, the recent papers by Midrigan (2011) and Karadi and Reiff (2012) consider state-dependent pricing models in which heterogenous firms face 'menu costs' of price adjustment. Midrigan (2011) shows that a menu-cost model that can match a wide array of micro price-setting facts implies aggregate real effects of monetary policy shocks which are quantitatively similar to those in

---

[4] An alternative to Calvo price-setting is the Taylor model (Taylor, 1980) of overlapping fixed-duration price contracts, but this specification is less convenient because it requires as many groups of price-setters as price-spell durations.

[5] For a useful discussion of state-dependent pricing models, see Walsh (2010, Ch. 6.2). Both Taylor models and Calvo models are examples of time-dependent pricing models.

the Calvo model. However, Karadi and Reiff (2012) point out that this result is overturned for relatively large business cycle shocks because, in state-dependent models, the fraction of firms changing price rises sharply as the magnitude of shocks is increased, making the aggregate price level somewhat more flexible. By contrast, the fraction of firms changing price remains constant in the Calvo model. Using data on value-added tax changes in Hungary, Karadi and Reiff show that the price level reacted flexibly and asymmetrically to large positive and negative shocks. They note that a menu-cost model similar to Midrigan (2011) can replicate both of these facts, whereas the Calvo model cannot, though it does remain a reasonable approximation of the menu-cost model for standard business cycle shocks in an economy with low inflation.

In summary, state-dependent models of price-setting can match the micro data somewhat better than the standard Calvo model, but the aggregate implications of the two models appear to be similar in normal economic circumstances. Here we limit our attention to the New Keynesian model with Calvo price-setting because this enables us to concentrate (i) on what the literature has to say about whether central banks should drop an inflation target where 'bygones are bygones' in favour of a price-level target,[6] and (ii) on models that have been subjected to formal statistical tests and not rejected by them; see Sections 5.2 to 5.4. However, since we cannot be sure that our conclusions about PT would carry over to alternative models of price-setting, the reader should bear in mind the caveats of the Calvo model mentioned in this section. We save a fuller discussion of the aggregate implications of New Keynesian models until Section 5, where we survey several papers from the applied macro literature.[7]

2.3   Bank of Canada Review of Price-Level Targeting

In 2006 the Bank of Canada announced its intention to study the costs and benefits of PT, with a view to potentially changing its monetary policy mandate at the next renewal date in 2011 (see Bank of Canada, 2006). This made it the first central bank in recent history to seriously entertain the possibility of PT regime. The Bank of Canada subsequently published a Review of the costs and benefits of PT that included several policy recommendations, including the reasons behind its decision to stick with an IT mandate until at least 2016 (see Bank of Canada, 2011). In this section, we discuss these conclusions in order to provide context for our survey of the recent literature that follows.

Since the only episode of PT in history was in Sweden during the Great Depression (see Berg and Jonung, 1999), the Bank of Canada Review focused on

---

[6]It should be noted that results regarding the relative merits of PT have not yet been established in state-dependent pricing models due to the additional complexities involved in working with these models. We are aware of only a small number of studies that assess the merits of PT in alternative time-dependent pricing models (e.g. Wolman, 2005).

[7]We note in Section 5.1 that price indexation – which helps New Keynesian models to match macro data – is at odds with the micro evidence on price-setting cited above. Importantly, however, the test of rational versus behavioural expectations that we consider in Section 5.4 takes place in a New Keynesian model where price indexation is absent. The rational expectations version of the model is not rejected by this formal statistical test.

model simulations and laboratory experiments, rather than empirical evidence. It argues that although there are "modest, but economically significant, potential gains from PT," realising these gains would be difficult in practice because the "models assume that agents are forward-looking, fully conversant with the implications of PT and trust policy-makers to live up to their commitments" (Bank of Canada 2011, p. 14).[8] In other words, the Bank expressed doubts about the ability of households and firms to understand the workings of PT and argues that the regime might lack credibility in the short-term. The Bank of Canada summarised its overall conclusion on PT as follows:

> Given the current state of knowledge, the potential benefits of PT in.increasing long-term certainty about the price level and.providing greater short-term macroeconomic stability, relative to the current IT framework, do not clearly outweigh the costs and risks associated with real-world expectations and credibility falling short of the model ideal. (Bank of Canada 2011, p. 20)

The 'costs and risks' of real-world expectations falling short of the model ideal have been highlighted in the early literature on PT. For example, Lebow et al. (1992) and Haldane and Salmon (1995) find that PT raises both output and inflation volatility under adaptive expectations, which are purely backward-looking.[9] Thus, the rational expectations assumption is central to the costs and benefits of PT relative to IT. For this reason, we review empirical evidence on rational expectations and New Keynesian models in the second half of this survey. We also consider in some detail a macro test of rational expectations against behavioural expectations, an alternative type of expectation formation that differs non-trivially from rational expectations. In particular, agents with behavioural expectations choose relative weights on extrapolative (i.e. backward-looking) forecasts and steady-state forecasts based on their past success.[10] The test of expectations we consider is relevant to the IT versus PT debate because, as emphasised by the Bank of Canada quote above, there is a widespread view among policymakers that private agents have expectations which are (at least partly) backward-looking, which would imply poor performance of PT.

On the theoretical side, we review recent research that sheds more light on the transition costs associated with PT and the extent to which important results in the PT literature are robust to deviations from the baseline New Keynesian model. We turn first to the theoretical literature and then to the empirical evidence on rational expectations and New Keynesian models.

---

[8] To ensure consistency with the abbreviations used in this survey, price-level targeting was abbreviated PT and not as PLT as in Bank of Canada (2011).

[9] For a more detailed discussion of this early literature, see Section 2 of Ambler (2009).

[10] The seminal paper in this literature is Brock and Hommes (1997). In Section 5.1 we discuss the relationship between these expectations and alternatives in the macro literature.

3 Macro stabilisation literature

As Ambler (2009) makes clear, PT was initially motivated as a way of providing the economy with long-term price stability. Nevertheless, most of the recent interest in PT comes from its implications for short run macro stabilisation, that is, economic stability at business cycle frequencies. In fact, the recent resurgence of interest in PT has gone hand-in-hand with the rise in popularity of New Keynesian models, since there are welfare gains from 'history dependence' in a forward-looking, rational expectations environment (Woodford, 2003). History dependence means that monetary policy should respond systematically to past economic conditions, as well current and expected future economic outcomes. PT is history dependent because the central bank promises to offset deviations from the target price path in the future. As a result, PT will tend to produce superior stabilisation outcomes to policies that lack history dependence.

In this section we first discuss the main findings from the macro stabilisation literature on IT versus PT, as surveyed by Ambler (2009). We then turn to more recent contributions that shed light on robustness to extensions of the baseline New Keynesian model. For ease of reference, we split the discussion that follows into four subsections: discretion versus commitment; the zero lower bound; financial frictions; and the costs of transition from IT to PT.

3.1 Discretion versus commitment

In a seminal paper, Kydland and Prescott (1977) showed that the distinction between discretion and commitment is crucial in rational expectations models. Under discretion, the policymaker optimises on a period-by-period basis. As a result, it cannot make binding promises about the future and must take future expectations as given. This inability to commit implies suboptimal outcomes because policy cannot influence future expectations and so effectively has fewer instruments available to achieve its objectives. A classic example of a suboptimal outcome under discretion is the 'inflation bias' problem highlighted by Barro and Gordon (1983), whereby an equilibrium with steady-state inflation above the socially-optimal level results from a desire to push output above the natural rate. An efficient equilibrium outcome can be restored in various ways which involve delegating the central bank a loss function that differs from the true social loss function (Rogoff, 1985; Walsh, 1995; Svensson, 1999).

It is important to note that discretion leads to suboptimal outcomes in dynamic models with rational expectations, even if the central bank's output target is *not* over-ambitious (the assumption made in almost all of the recent literature, and in all the papers we review here). Svensson (1997) and Clarida et al. (1999) dub this phenomenon 'stabilisation bias' because it implies that the central bank's response to shocks that hit the economy is suboptimal, so that inflation and output gaps deviate from their first-best outcomes. As with the inflation bias, stabilisation bias implies that delegating monetary policy to a central bank without IT preferences could improve social outcomes.

3.1.1 Discretion

Svensson (1999) was the first to formally investigate whether delegating monetary policy to a central bank with a PT loss function could improve stabilisation outcomes under rational expectations. Using a model with a Lucas-type (or Neoclassical) Phillips curve, he showed that delegating PT preferences to the central bank delivers a 'free lunch' result: inflation variability is lower for any given level of output gap variability if output is sufficiently persistent. In effect, PT eliminates some of the useless discretion present under IT because the central bank is required to undo current deviations from its target price path in the future, so that discretionary behaviour becomes more costly. The crucial distinction is that 'bygones are bygones' under IT, whereas PT makes the objective function of a discretionary policymaker depend in part upon past outcomes, so that discretion carries a penalty which is absent under IT.

Although Svensson's model features rational expectations, it is subject to two important limitations. First, the economy is described by a Phillips curve which performs poorly empirically because it implies that only unanticipated changes in policy have real effects. Second, the social loss function Svensson uses cannot be derived from the utility function of a representative agent. As noted in Section 2.1, the New Keynesian model can address both these criticisms.

Vestin (2006) assesses the performance of PT in the baseline New Keynesian model. He shows that the free lunch result remains intact. Intuitively, since firms in New Keynesian models set current prices as a function of expected future prices, the extent to which the aggregate price level rises in response to an inflationary shock depends on the impact of policy upon price-setters expectations about future inflation. Under IT, inflation expectations are effectively fixed on the inflation target because 'bygones are bygones'. Under PT, by contrast, firms expect a rise in current prices to be followed by a contraction in demand and lower future prices, in order to return the price level to its target path. As a result, firms raise their prices less in response to inflationary (i.e. cost-push) shocks under a PT regime, implying lower inflation variability for any given level of output gap variability.

The free lunch result also remains intact when the firms who are unable to re-optimise choose to index their prices to past inflation, except in the special case of full indexation (Røisland, 2006; Gaspar et al. 2007).[11] Indexing prices to past inflation leads to the 'hybrid New Keynesian Phillips curve' where inflation is also related to lagged inflation:

$$\pi_t - \gamma\pi_{t-1} = \beta E_t(\pi_{t+1} - \gamma\pi_t) + \kappa x_t + u_t \tag{7}$$

---

[11] Røisland (2006) shows that discretion will replicate the optimal commitment policy when the central bank is delegated a period loss function $(p_t - \gamma p_{t-1})^2 + \tilde{\alpha}x_t^2$, where p_t is the price level (in logs) and $\tilde{\alpha}$ is the modified weight on output gap variations. Vestin's (2006) result that PT is optimal arises as a special case when there is zero indexation (i.e. $\gamma = 0$), while IT is optimal under full indexation (i.e. $\gamma = 1$). When $0 < \gamma < 1$, the optimal loss function under discretion is usually interpreted as a 'hybrid inflation/price-level target', but Gaspar et al. (2007, pp. 19-20) argue that this is effectively PT with a longer target horizon.

where $0 \leq \gamma \leq 1$ represents the degree of indexation to past inflation, and $\pi_t - \gamma\pi_{t-1}$ is the 'quasi-difference' in inflation.

Although this hybrid Phillips curve performs better against macro data than the standard New Keynesian Phillips curve, it is difficult to justify because price-setters are assumed to index their prices in a purely backward-looking (and hence non-rational) manner. The assumption of price indexation is also strongly at odds with the empirical evidence on price-setting cited in Section 2.2. Galí and Gertler (1999) have shown that one can also justify a partially backward-looking New Keynesian Phillips curve by having some fraction of price-setters follow a simple rule-of-thumb when setting prices, but this assumption can be criticised on similar grounds. In this case, the results in Nessén and Vestin (2005) show that PT will dominate IT only if the fraction of rule-of-thumb price-setters remains sufficiently small.[12]

Blake et al. (2011) tackle a somewhat different issue. In particular, since there are multiple rational expectations equilibria under discretion when the baseline New Keynesian model is augmented with capital accumulation (see Blake and Kirsanova, 2012), the lowest welfare equilibrium under delegation to a PT central bank may be inferior to the highest welfare equilibrium attainable under IT. Consequently, in the absence of a mechanism for selecting one equilibrium over another, it is generally not possible to predict whether PT will dominate IT. Finally, Masson and Shukayev (2012) show that if the central bank operates under discretion and the public believes that there is the possibility it will rebase the price-level target in response to large shocks, there can be multiple equilibria even in the absence of capital accumulation. As a result, the economy could end up in a low credibility equilibrium where output volatility is increased and the beneficial effects of PT are reduced or even reversed.

3.1.2 Commitment

The standard argument made in favour of discretion is that, in practice, "no major central bank makes any kind of binding commitment over the course of its future monetary policy" (Clarida et al. 1999, p. 1671). It is not obvious, however, that discretion provides a better description of the real-world behavior of central banks than commitment, because it implies that avoidable policy mistakes are repeated ad infinitum. Moreover, as is now familiar in the context of formal 'forward guidance', central banks do adopt contingent rules in which there are binding forward targets for inflation and output that commit the central bank except in respect of its contemporary response to shocks.[13]

---

[12] Under the baseline calibration of Nessén and Vestin (2005), PT outperforms IT under discretion if the fraction of rule-of-thumb price-setters is less than 0.68 (see Figure 6 of their paper and the surrouding discussion, pp. 853–55).

[13] Since 2008, explicit forward guidance has also been used by several central banks, including the Bank of Canada, the ECB, the Bank of England, the Federal Reserve, and the Bank of Japan. For example, the Federal Reserve announced in August 2011 that economic conditions "are likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013"; see the Press Releases section of the Federal Reserve's website.

Commitment solutions are also of interest because they provide guidance on optimal monetary policy. The commitment-based literature on PT has two different strands: the first focuses on optimal monetary policy under commitment, and the second on commitments to simple Taylor-type interest rate rules.

Optimal commitment policies. If commitment is possible, then the social loss function, Equation (3), will be minimised by giving the central bank an IT loss function. Few economists believe that central banks can implement fully optimal commitment policies in practice, but it is nevertheless useful to study these policies to identify circumstances where delegating an alternative loss function such as PT may, or may not, improve social outcomes. This section dicusses recent literature in this regard.

One well-known result in relation to the baseline New Keynesian model is that optimal commitment implies 'price stationarity' and 'history dependence' (Clarida et al., 1999; Woodford, 2003). As noted by Vestin (2006), PT outperforms IT under discretion precisely because delegating a PT loss function to the central bank makes the price level stationary and policy history dependent, hence moving the discretionary solution closer to the optimal commitment policy.[14] By contrast, a discretionary IT regime implies base-level drift in the price level and lacks history dependence because 'bygones are bygones'. History dependence is a general feature of optimal policy in forward-looking models; it is therefore robust across alternative New Keynesian models.

The price stationarity result is not similarly robust, however: minor modifications of the baseline New Keynesian model can overturn the result that a stationary price level is optimal. For instance, in the case of the hybrid New Keynesian Phillips curve, Røisland (2006) and Gaspar et al. (2007) have shown that if all firms index their prices to past inflation (i.e. $\gamma = 1$), then the optimal commitment policy implies base-level drift in the price level. Steinsson (2003) showed that the same conclusion holds when some fraction of price-setters are backward-looking and follow a simple rule-of-thumb when setting prices, as in Galí and Gertler (1999). In addition, Levin et al. (2010) show that optimal policy involves considerable base-level drift in response to contractionary demand shocks when the zero lower bound on interest rates is a binding constraint.

More recently, Amano et al. (2012) showed that base-level drift is optimal in the baseline New Keynesian model if the central bank and agents must make decisions before current shocks to the economy are observed. Intuitively, the benefits of price stationarity in the baseline case come from the expectation that policy will offset inflationary shocks, which in turn dampens the impact of shocks on current inflation. But if the central bank cannot observe current shocks to inflation, it cannot react directly to these shocks and so cannot influence price-setters' expectations in a favourable way. The benefits of price stationarity are therefore lost and it becomes optimal to treat inflationary shocks as bygones.

Lastly, Gerberding et al. (2012) consider a New Keynesian model with

---

[14] Vestin (2006) shows that PT can exactly replicate the optimal commitment policy in the baseline New Keynesian model if there is no persistence in the cost-push shock.

nominal rigidity in both the intermediate goods sector and the final goods sector and show that PT performs poorly. The reasoning is that PT makes the price level in the final goods sector stationary, whereas under the optimal policy there is considerable base-level drift in the price level after a sector-specific negative productivity shock. Moreover, PT permits somewhat more drift in the intermediate goods price level than is optimal, so it has additional welfare costs which would not arise if one of the sectors had flexible prices that could adjust freely with relative productivity without creating welfare losses (see Aoki, 2001). This trade-off between targeting one nominal rigidity at the expense of exacerbating the welfare costs of another was first highlighted by Erceg et al. (2000) in a model with sticky wages and prices.[15] While increasing the target horizon of policy lowers the welfare costs of PT by making this trade-off less severe, PT is clearly outperformed by IT for policy-relevant target horizons.

In short, history dependence is a robust feature of the optimal commitment policy in New Keynesian models, but price stationarity is not. This finding has implications for the design of the optimal central bank loss function under discretion since it narrows down the set of models in which PT is likely to outperform IT.

Commitment to Taylor rules. The second strand of commitment literature has assessed the performance of IT and PT Taylor rules in New Keynesian models. Simple interest rate rules are motivated by the argument that they may provide a better representation of the behaviour of central banks than fully optimal policies since they are easy to communicate, more robust across alternative models of the economy (see Taylor, 1999), and easier to implement as they do not imply a direct policy response to economic shocks which are difficult to observe in practice.[16] Whereas IT Taylor rules relate the nominal interest rate to a measure of the output gap and an inflation gap (like the original Taylor rule), PT Taylor rules instead respond to a price-level gap:[17]

$$R_t = c_p(p_t - p_t^*) + c_x x_t \tag{8}$$

where $c_p, c_x > 0$ are reaction coefficients, p_t is the log of the aggregate price level and p_t^* is the target price level (expressed in logs).

---

[15] Erceg et al. (2000) note that if there is more than one source of nominal rigidity in the economy, then it is optimal to target most forcefully the wage or price that is most rigid. In the baseline calibration of the Gerberding et al. (2012) model, consumer prices are only slightly more rigid than intermediate goods prices, so it is highly suboptimal to focus solely on stabilisation of consumer prices as under PT.

[16] For example, Equation (6) does not imply a direct response to economic shocks but Equation (5) does, and so does the implied optimal interest rate rule under discretion (see Clarida et al. 1999, p. 1672).

[17] When comparing regimes using Taylor rules, the reaction coefficients are typically chosen to minimise the social loss function obtainable under that rule. In addition, some papers in the literature allow for 'interest rate smoothing' through dependence of the nominal interest rate on its lagged value.

The Taylor rule literature includes several analyses of PT in open economies. An early contribution to this literature was Batini and Yates (2003). They show that the degree of openness of an economy is important for comparisons of IT and PT. In their model the real exchange rate enters the Phillips curve, so that the variability of inflation depends on fluctuations in the real exchange rate, with this channel becoming more important as the degree of openness of the economy is increased. On the one hand, PT can have a positive impact on inflation stabilisation because the uncovered interest parity (UIP) condition introduces an additional channel through which managing expectations matters for economic outcomes. On the other hand, PT could lead to a deterioration in economic stability relative to IT, because it makes interest rates somewhat more volatile, and this feeds back to greater real exchange rate volatility by the UIP condition. In ranking IT and PT regimes in welfare terms, it is the relative size of these two effects that matters.

Much of the subsequent research in this literature has been carried out by or in conjunction with the Bank of Canada. For instance, Coletti et al. (2008) compared IT and PT in the two-country IMF Global Economy Model (GEM), a medium-scale model designed to enable open-economy issues to be investigated within a microfounded framework that is suitable for policy analysis (see Laxton, 2008). This model contains several sources of nominal rigidity and is forward-looking; it can therefore be viewed as an extension of the New Keynesian model. Coletti et al. (2008) calibrated the model for Canada (with the US as the second country) and included tradables and non-tradables sectors which are subject to sector-specific shocks, giving the BoC-GEM model. They found that a PT Taylor rule outperforms an IT one in terms of inflation and output gap stability, primarily because shocks to the terms of trade strengthen the case for PT due to its role as a nominal anchor that stabilises the domestic price level.

PT also tends to outperform IT in other similar open economy models of Canada because the gains from PT due to extra stabilisation through the expectations channel outweigh the losses involved in responding to additional shocks in order to stabilise the aggregate price level; see Murchison (2010). It is worth noting, however, that dominance of PT over IT is not guaranteed when policymakers commit to Taylor rules: Coletti et al. (2012) consider an extended version of the BoC-GEM model with separate energy and non-energy sectors and find that PT improves stabilisation relative to IT in the case of non-energy commodity supply shocks, but not in response to energy commodity supply shocks and commodity demand shocks.

3.1.3 Summary

To summarise, PT tends to dominate IT in several different variants of open economy New Keynesian models when policymakers commit to simple Taylor-type rules. Hence, the main results highlighted by Giannoni (2010) in the context of the baseline New Keynesian model appear to carry over to open economy models. By contrast, the result that the optimal policy implies a stationary price level is sensitive to minor modifications of the New Keynesian model, as

is the performance of PT under discretion. These findings have relevance for the choice between IT and PT regimes because the real-world policies of central banks are probably best characterised either by discretion or commitment to simple instrument rules. We return to the argument that committing to PT Taylor rules is beneficial after discussing whether this result also holds in the presence of the zero lower bound and financial frictions.

3.2   The zero lower bound on nominal interest rates

In simple terms, the zero lower bound (ZLB) states that nominal interest rates cannot fall below zero in an economy where money is untaxed and can be stored without cost. Aided by advances in computing power, the macro stabilisation literature has recently incorporated the ZLB into New Keynesian models as an occasionally-binding constraint.[18]

Eggertsson and Woodford (2003) was the seminal contribution in this literature. They introduce the ZLB into the baseline New Keynesian model and show that the optimal commitment policy implies a state-contingent target price level at times when the ZLB is binding. Intuitively, by promising *additional* future inflation in response to deflationary pressure at the ZLB, policymakers can drive down (ex ante) real interest rates and provide extra stimulus to the output gap. Since this policy implies base-level drift in response to shocks that drive the economy to the ZLB, price level stationarity is no longer optimal, unlike in the baseline New Keynesian model. However, Eggertsson and Woodford show that the optimal commitment policy – which would be difficult to implement in practice – can be approximated by a simple rule that aims at a fixed target price level, whereas a standard IT rule performs poorly by comparison. The key to this result is that the real interest rate does not respond to the severity of a ZLB episode under the IT rule, because inflation expectations are fixed on the inflation target. By contrast, under the PT rule, the real interest rate falls as ZLB episodes worsen because agents rationally expect enough future inflation to return the price level to its fixed target. The stimulus to the output gap from lower real interest rates aids the economy's recovery from the ZLB.

Formal welfare analyses of optimal policy with an ocassionally-binding ZLB can be found in Adam and Billi (2006, 2007). The model in Adam and Billi (2006) follows Eggertsson and Woodford (2003), but they calibrate the model to US data using shocks from the Great Moderation period. In order to evaluate social welfare, these shocks are used to simulate the model. They find that zero nominal interest rates occur rather infrequently under the optimal commitment policy – only about one quarter in every 17 years, or an unconditional probability of 1.5 per cent. As a result, the additional welfare loss due to the ZLB is quite small at approximately 1 per cent of the welfare loss generated by sticky prices. Adam and Billi (2007) extend the analysis of optimal policy to the case of discretion. They find that the welfare losses imposed by discretion relative to optimal commitment increase by around two-thirds when the ZLB is

---

[18]For a detailed review of the early literature on the zero lower bound, see Yates (2004).

an occasionally-binding constraint, suggesting that the potential welfare gains from well-designed alternatives to IT are much larger when one accounts for the lower bound. Nakov (2008) takes up this theme. His contribution is to assess social welfare for a variety of *zero-truncated* Taylor rules, which he argues provide a more plausible representation of real-world policymaking than optimal policies. All the rules perform poorly compared to optimal commitment, with relative losses for PT and IT Taylor rules of 800 per cent and 1400 per cent respectively.[19] Crucially, however, constraining the IT-PT comparison to simple Taylor rules leaves intact the conclusion that PT is a significant improvement on IT in the presence of an ocassionally-binding ZLB.

Finally, an important recent contribution to the ZLB literature is Coibion et al. (2012). They also focus on a New Keynesian model with an occasionally-binding ZLB and Taylor rules, but their model is log-linearised around a non-zero trend inflation rate, so that the optimal rate of inflation can be studied. Under IT, moderate trend inflation is optimal in order to raise average nominal interest rates and hence reduce the probability of hitting the ZLB (and its associated welfare losses) Under PT, the likelihood of the ZLB being reached for any given trend inflation rate is much lower than under IT because inflation and output volatility are reduced, while the promise of above-average future inflation when at the lower bound means that the economy exits rapidly from ZLB episodes . As a result, it is optimal to set trend inflation close to zero under PT, so that the welfare gains of lower trend inflation can be realised while still keeping the overall welfare costs of the ZLB low.

The potential welfare gains from PT through this channel are quantitatively quite significant: raising the response to the price level in the Taylor rule from 0 to 0.25 implies a welfare gain equivalent to a permanent increase in aggregate consumption of 0.5 per cent.[20] This finding and the results of Nakov (2008) are further examples of PT dominating IT when policymakers commit to simple Taylor rules.

3.3   Financial frictions

The recent financial crisis has underlined the potential importance of financial frictions for the real economy. It is therefore important to establish whether PT dominates IT in models with these frictions which are subject to substantial financial disturbances. We discuss here some recent contributions in this area and their implications for the IT versus PT debate. We caution the reader, however, that this literature is at an early stage and presently relies on quite simple models of the financial sector.

---

[19] Nakov (2008) notes that the loss under IT is much larger if interest rate smoothing is not permitted. The reason is that interest rate smoothing implies history dependence, which is otherwise absent under IT.

[20] Positive trend inflation has three distinct costs in New Keynesian models; see Ambler (2007) for a useful discussion. The traditional welfare cost of inflation due to inflation acting as a tax on money holdings (see Bailey, 1956) is not one of them. It is therefore conceivable that the welfare gains attainable from lowering trend inflation under PT could be larger than estimated by Coibion et al. (2012).

Dib et al. (2008) investigated the impact of PT within a medium-scale open economy New Keynesian model whose parameters were estimated on Canadian data. The model is augmented with credit market imperfections as in Bernanke et al. (1999) and entrepreneurs who enter into one-period nominal debt contracts in order to finance investment. In this model, PT significantly lowers the distortion in the economy that arises due to nominal debt contracts, because inflation expectations are better stabilised than under an IT regime, which in turn lowers real interest rate variability. As a result, the welfare gains from PT are strengthened. More recently, Dib et al. (2013) use the same model to show the additional result that a PT Taylor rule outperforms an IT Taylor rule when the rules are extended to include a response to the external finance premium. Consequently, the case for PT remains intact even if the central bank can respond directly to financial variables in its policy rules.

Covas and Zhang (2010) also focus on the implications of financial market imperfections for comparisons of IT and PT, but they compare the two regimes in a New Keynesian model with imperfections in both debt and equity markets. The parameters of the model were again estimated for the Canadian economy. They find that PT outperforms IT in terms of stabilisation because the expectations channel means that inflation is better anchored under PT, so that it is less costly for the central bank to address financial market distortions through monetary policy. It should be noted, however, that the benefits of PT become smaller in this model as financial market imperfections are strengthened.

Lastly, Bailliu et al. (2012) study the interaction between macroprudential rules and monetary policy in a model with financial market imperfections which are again introduced through debt contracts as in Bernanke et al. (1999). The model is estimated on Canadian data and calibrated accordingly. Under optimised IT and PT Taylor rules, it is beneficial to respond to financial imbalances in both a Taylor rule and a macroprudential rule. Nevertheless, PT rules deliver substantial welfare gains over IT rules due to policy being history dependent.

These results provide further confirmation that augmenting New Keynesian models with simple financial frictions does not substantially alter the relative performance of IT and PT regimes, as well as additional evidence that PT tends to dominate IT when policymakers can commit to Taylor rules.

3.4 The transition from inflation to price-level targeting

The Bank of Canada's 2011 Review argues that the transition to PT could be costly because households and firms may not understand the workings of the new regime and ignore it when forming expectations. This could lead to a situation where PT initially has imperfect credibility and agents actively learn about the new regime.[21] Several studies in the literature have assessed the performance of PT in these circumstances. For instance, Gaspar et al. (2007) set up a New Keynesian model where PT initially has imperfect credibility and agents'

---

[21]Imperfect credibility is a situation where agents do not believe fully that PT will be implemented by policymakers.

expectations are determined by adaptive learning as in Evans and Honkapohja (2001). They find that an initial period of imperfect credibility and learning is sufficient to turn the net welfare gains from PT negative if agents are slow to learn.

Kryvtsov et al. (2008) model imperfect credibility as an exogenous process that converges towards perfect credibility over time. In the baseline New Keynesian model, it takes two-and-a-half years for the PT central bank to earn enough credibility to outperform IT, but this relatively short period is enough to ensure that IT dominates PT in net welfare terms. Cateau et al. (2009) extend the analysis of the transition from IT to PT using the Bank of Canada's main policy analysis model, ToTEM. They find results more favourable to PT: the long run welfare gains dominate the short run transition costs as long as the initial spell of imperfect credibility lasts less than 13 years. Finally, in relation to the lower bound on nominal interest rates, preliminary analysis by Cateau and Dorich (2012) suggests that the welfare gains of PT hinge crucially on the assumption that the new regime has a high level of crediblity.

Intuitively, PT performs poorly relative to IT in these studies because the expectational benefits of a price-level target are lost if agents are backward-looking or do not take the new regime fully into account in their expectations. We thus see again that the assumption of rational expectations is central to the costs and benefits of PT relative to IT. Moreover, the finding that initial periods of imperfect credibility and learning can overturn the welfare gains from PT arises because in many models the gains are relatively modest, albeit economically non-trivial. For instance, many of the studies in this survey that find in favour of PT conclude that the potential welfare gains would be equivalent to a permanent increase in aggregate consumption of less than 0.1 per cent.[22]

Hence, a crucial issue when evaluating the case for PT is the extent to which agents' expectations are forward-looking and take account of the new regime in the short-term before it has gained full credibility, since it is this which will determine whether or not welfare gains are realised in practice.

3.5   Summary

A key finding from the macro stabilisation literature is that PT can potentially improve upon IT in New Keynesian models. In these models, optimal policy implies price stationarity and history dependence – key features of a PT regime. Although the price stationarity result and the performance of PT under discretion are *not* robust, PT consistently outperforms IT when policymakers commit to simple Taylor rules. This result is robust in the sense that it holds for extensions of the baseline New Keynesian model, including open economies, the zero lower bound, and financial frictions, but the assumption that agents understand PT and incorporate it into their expectations is crucial. The beneficial effects

---

[22]See, for example, Dib et al. (2008), Covas and Zhang (2010), Bailliu et al. (2012), and Dib et al. (2013). An alternative measure of the welfare gain from PT is reported by Kryvtsov et al. (2008), but they note that "even under immediate perfect credibility, PT gives only a small welfare gain over IT" (p. 20).

of PT therefore seem to hang on the assumption that agents are rational and the economy New Keynesian. We therefore focus in the remainder of the survey on empirical evidence on rational expectations and New Keynesian models.

4   Survey and experimental evidence on rational expectations

As our discussion above makes clear, rational expectations are crucial in order for PT to have beneficial effects. It is therefore important to test this assumption scientifically. Broadly speaking, empirical evidence on rational expectations has come from two sources: studies investigating the performance of macro models with rational expectations, and surveys and experiments. This section focuses on the survey and experimental evidence..

4.1   Survey evidence on rational expectations

Surveys have been a popular way of testing the rational expectations assumption. On an intuitive level this popularity is hardly surprising: surveys give researchers the freedom to 'ask the right question' while also enabling them to collect expectations data at the disaggregated level at which household decisions are taken. Tests of rational expectations using individual-level survey data can also avoid biases that arise from aggregation or pooling of forecast data (see e.g. Bonham and Cohen, 2001). However, there are some well-known difficulties with surveys. In markets the expectations that matter are those of the active market participants and surveys may not identify these people. This means that survey respondents may be inattentive and poorly informed. Other potential problems include truthfulness and accurate recall.

Some well-known surveys are the Livingston Series in the US and the Confederation of British Industry (CBI) survey in the UK, both of which survey professional economists. Early tests using these surveys concluded that expectations are not strictly rational because expectational errors are not serially uncorrelated over time, although some studies were not able to reject the assumption of unbiasedness conclusively (see Holden et al. 1985, Ch. 3). More recently, Thomas (1999) shows that inflation expectations in the Livingstone Survey clearly reject the unbiasedness test of the rational expectations hypothesis based on sub-samples from 1960 to 1980 and 1980 to 1997. Interestingly, however, he also finds that an unbiasedness test does not reject rational expectations on the full sample from 1960 to 1997, though Mankiw et al. (2003) are able to overturn this result using a slighlty longer sample period. Because the surveys used in these studies question professional economists, they should not be vulnerable to the criticism that participants are inactive or poorly informed. These studies would appear to provide convincing evidence that inflation expectations are not rational.

However, in a recent contribution, Andolfatto et al. (2008) provide a possible explanation for these rejections of rational expectations. They set up a New

Keynesian model where agents do not have full information about the state of the economy. In particular, monetary policy follows a Taylor rule with transitory monetary policy shocks and occasional persistent switches in the long run inflation target, but agents cannot observe these two shocks separately and therefore solve a signal extraction problem using the Kalman filter. Andolfatto et al. show that running standard tests of unbiasedness on simulated inflation expectations will incorrectly reject the rational expectations hypothesis. The key to this result is that since agents cannot disentangle the effects of transitory and persistent shocks, rational expectations will imply unbiased expectations in population, but not necessarily in small samples. This finding suggests that the results in Thomas (1999) and Mankiw et al. (2003) should not be interpreted as conclusive statistical evidence against rational expectations.

Support for rationality within models with less than full information has also been provided by Coibion and Gorodnichenko (2012). Using survey expectations data from several different sources, they show that forecast errors consistently move in the same direction as the variable being forecast in response to a variety of macro shocks. While this feature is predicted by both sticky-information and 'noisy information' models, there is more support for the predictions of the latter in the survey data.[23] It should be emphasised, however, that both Andolfatto et al. (2008) and Coibion and Gorodnichnko (2012) consider a weaker form of rational expectations that is typically assumed in the literature since the agents in these models do not have full information on the state of the economy. Given the potential problems with obtaining conclusive survey evidence on rational expectations, it is important to bring other sources of evidence to bear on whether expectations are rational.

4.2   Experimental evidence on rational expectations

An alternative approach to assessing expectations that has gained popularity in recent years is the use of lab experiments. Like surveys, experiments are flexible and can be targeted at individuals. The main advantage of the experimental approach is that the behaviour and expectations of participants can be studied within a particular 'economic model' constructed by the researcher in a carefully controlled setting. This approach has been used widely in the micro literature but more sparingly in macroeconomics.

An early contribution to the macro literature was Smith et al. (1988). They studied spot asset trading in an experimental environment and concluded that traders' expectations converge upon rational expectations as they acquire more experience through market interaction. This occured after several runs of the experiment (with the same subjects) that were characterised by speculative behaviour, suggesting an initial period of learning before convergence on the rational expectations equilibrium. However, Lei et al. (2001) found using a refined version of that experiment that at least some agents display irrational

---

[23] The term 'noisy information' is used by Coibion and Gorodnichenko (2012). For examples of these models, see Lucas (1972) and Sims (2003).

behaviour during the learning phase, while in a different experiment Bloomfield and Hales (2002) found that experimental subjects who were told that the data-generating process was a random walk did not have static expectations, the rational expectations solution in this case. Their beliefs were instead consistent with switches between a trending regime and a mean-reverting one.

The subsequent literature also contains mixed results, but several recent studies suggest that rational expectations might be a much better approximation than previously thought. For instance, the experimental data obtained by Heemejer et al. (2009) suggest that rational expectations might be a reasonable first approximation: the large majority of participants in the study had mean expectation errors which were not significantly different from zero, and expectation errors were not significantly autocorrelated.[24]. More recently, Bao et al. (2013) set out to understand how agents would learn a rational expectations equilibrium, with a focus on comparing the performance of individuals versus two-person teams. They set up an experiment where participants had to forecast future prices and make production decisions within a cobweb model. In both cases there is convergence on the rational expectations equilibrium, but the speed of convergence is somewhat faster for teams than individuals (10.7 periods versus 42.3 periods). These results suggest that teamwork may help participants to find the rational expectations equilibrium. This is an interesting observation given that many economic decisions are decided at the household level and not individually.

In relation to the IT versus PT debate, Amano et al. (2011) investigated whether experimental participants were able to accurately forecast inflation in a simulated PT regime. They found that forecasts did adjust somewhat under a PT regime, but not to the extent that the implications of a price-level target were fully reflected in expectations. It is important to note, however, that these findings were obtained under conditions relatively unfavourable to PT: participants were given minimal information about the model economy and had the shift in regime from IT to PT explained to them only once. It is also worth noting that inflation was forecast at the individual level, yet many economic decisions are taken at the level of the household in practice. Given the results of Bao et al. (2013), this may understate the true level of understanding of PT because it seems that 'two heads are better than one'.

In another experiment, Kryvtsov and Petersen (2013) set out to determine whether the stabilisation benefits of monetary policy through the expectations channel in New Keynesian models are likely to be borne out in practice. Around three-quarters of the stabilisation benefits obtainable under full rational expectations were realised in the experiment, suggesting that rational expectations are a much better approximation than strict adaptive expectations. At the same time, however, agents' forecasts appear to be partly backward-looking, with significant weights on last period's inflation and output gap realisations.

One weakness of experiments is that they are subject to some of the same

---

[24]For a detailed discussion of the inferences that can be drawn from this experiment, see Wagener (forthcoming). That paper also provides a useful survey of the experimental evidence on rational expectations.

criticisms as surveys: subjects may alter their behaviour under experimental conditions and active market participants may not be identified. It is also well-known that experimental outcomes can be sensitive to the 'rules of the game', as the results of Lei et al. (2001) and Bao et al. (2013) demonstrate. Lastly, the usefulness of experiments in addressing macro issues is likely to be limited by the need to rely on simple economic models and the fact that it is difficult to recreate the kind of economic environment that agents face in the real world in a laboratory setting. Given the potential problems with surveys and experiments, we focus on the applied macro literature in the remainder of the survey.

5 Testing the models used in evaluating price-level targeting

In this section, we provide a discussion of the applied macro literature on rational expectations. We survey recent developments – including estimated dynamic general equilibrium models and refinements and alternatives to rational expectations – before turning to a recent literature that has tested the restrictions implied by New Keynesian models with rational expectations using formal statistical tests.

5.1 The empirical performance of macro models

Macro models with rational expectations became commonplace after rational expectations revolution of the 1970s, yet few formal attempts were made to test these models against alternatives with different expectational assumptions. Fair (1993) was an early exception; we briefly discuss his results in order to bring out some important issues in testing and how the literature has responded to these. Fair's model was set-up to nest both rational and adaptive expectations, thus enabling a direct test of the rational expectations hypothesis. The equations of the model come from a medium-scale macro model with current and lagged values of macro variables and future expected values based on current information. The 'test' of rational expectations amounts to testing the joint significance of the estimated coefficients on the forward-looking variables in each individual equation of the model. In total, this test was conducted for 16 separate equations, including 7 describing the behaviour of the household sector, 5 describing the behaviour of firms, and additional equations for investment, employment, interest rates, and asset prices. Each equation was estimated using quarterly data from the postwar period.

The results show some support for the rational expectations hypothesis. In particular, half of the estimated equations have significant lead coefficients at the 1% level for at least one of the tests, with the strongest support coming from the 7 household equations, of which 5 are statistically significant. In the other 8 cases, the null hypothesis of adaptive expectations was not rejected. The testing procedure is not without problems, however, as recent research has highlighted. One important problem is that the model solution for the future expectations

of endogenous variables will typically include lags of the same variable; hence it is unclear whether these tests have much power or indeed that the relevant equations are identified. To give a pertinent example, it is widely recognised that equations with lagged inflation may be equivalent to the solution of those with future expectations.[25]

A second problem is that the test is a partial information test: each equation is tested on a individual basis, so there is no unambiguous accept/reject decision for the model as a whole. More importantly, this means that the cross-equation restrictions implied by rational expectations are left untested. The recent literature that we discuss below has overcome these limitations by using full-information methods. A final weakness of the test conducted by Fair (1993) is that it is vulnerable to the Lucas critique of policy evaluation (Lucas, 1976): the model is not derived from first-order conditions, so it does not contain 'deep parameters' that can be expected to be invariant to structural change.

In order to address these concerns, macroeconomic models are now built, as standard, from models of the economy with market clearing and optimising households and firms with rational expectations. These micro-founded dynamic stochastic general equilibrium models were popularised by the pioneering contribution of Kydland and Prescott (1982), who showed that, in such a model, technological shocks alone could account for a surprisingly large fraction of US output volatility. Since Kydland and Prescott's paper, researchers have augmented real business cycle models with New Keynesian nominal rigidities in order to provide a plausible stabilisation role for monetary policy. These models do a good job of matching several key features of aggregate data when hit with real and nominal shocks, as we discuss below. It is this combination of appealing theoretical foundations and empirical performance that has made New Keynesian models dominant in the macroeconomic stabilisation literature discussed in Section 3.

An important contribution to the modern applied macro literature was Christiano et al. (2005). They set up a medium-scale New Keynesian model with staggered wage contracts, variable capacity utilisation, investment adjustment costs, and indexation of prices and wages to past inflation. The model was estimated on quarterly US data over the postwar period using Bayesian methods. The estimated model was able to replicate the impulse response functions of several important macroeconomic aggregates in response to monetary policy shocks, as estimated using a vector autoregression (VAR). In particular, the model does a good job of accounting quantitatively for the impulse responses of inflation, output and real wages, as well as the lagged, hump-shaped responses of consumption, investment, profits and labour productivity. Therefore, augmenting the New Keynesian model with real and nominal frictions enables it to mimic the dynamic responses of several key variables to nominal shocks.

Smets and Wouters (2007) go even further. They show that a New Keyne-

---

[25] For example, the solution for inflation under optimal commitment in the baseline New Keynesian model is $\pi_t = \delta \pi_{t-1} + \phi \Delta u_t$, where ϕ depends on the parameters of the model (see Clarida et al. 1999, p. 1704). Hence $E_t \pi_{t+1} = \delta \pi_t + \phi E_t \Delta u_{t+1}$ is a function of π_{t-1}. Similar arguments apply to the output gap solution.

sian model with several sources of real and nominal rigidities and 7 orthogonal shocks (both real and nominal) can match several additional dynamic features of US data, and that the model as a whole performs well in out-of-sample forecast tests. In particular, the estimated model does as well as a Bayesian VAR with 4 lags based on both marginal likelihood and a more traditional out-of-sample forecasting exercise. Notably, in the out-of-sample forecasting exercise, the model performs comparably to the BVAR(4) over short horizons and does considerably better at long horizons such as 2 or 3 years.

Some recent papers have considered refinements of rational expectations in medium-scale New Keynesian models in the spirit of Christiano et al. (2005) and Smets and Wouters (2007). For instance, Schmitt-Grohé and Uribe (2012) show how one can model 'news', that is, the arrival of information on future macro shocks prior to the time when those shocks are actually realised. Fujiwara et al. (2011) and Khan and Tsoukalas (2011) both estimate medium-scale models with news shocks and sticky prices and wages. Fujiwara et al. (2011) find that the inclusion of news shocks significantly improves the fit of the New Keynesian model in the case of the US economy, but not for Japan. Khan and Tsoukalas (2011) show that the importance of news shocks is sensitive to model structure, shocks and the data used in estimation, making it hard to draw any general conclusions about the contribution of news shocks to business cycles. Other refinements to rational expectations that have had some success empirically include 'sticky information' (Mankiw and Reis, 2002; Mankiw and Reis, 2007), 'rational inattention' (Mackowiak and Wiederholt, 2009; Mackowiak and Wiederholt, 2011), and imperfect information (Levine et al., 2012).[26] The main appeal of these approaches is that they help macro models to match inertia in the data in the absence of ad hoc features such as indexation and habit formation which have little empirical support.[27] While these approaches have not yet been shown to convincingly outperform standard models with full rational expectations, they are clearly a promising avenue for future research.

A more substantial deviation from rational expectations that has gained traction recently is adaptive learning (Evans and Honkapohja, 2001). Under this approach, agents in the model behave like econometricians, running simple regressions in order to form forecasts, and updating these forecasts over time as new data arrives. The rational expectations equilibrium is often 'learnable' in these models – i.e. adaptive learning may converge on the rational expectations equilibrium – but this is not guaranteed. Estimated New Keynedian models under adaptive learning have become common in recent years. For instance, Milani (2007) estimated a New Keynesian model and found that the estimated coeffi-

---

[26]Levine et al. (2012) consider a rational expectations model with imperfect information and also a model with adaptive expectations, which we discuss below. In the model with imperfect information, agents do not have full information about state variables in the system and shocks hitting the economy, so they solve real-time filtering problems to make inferences about the values of these variables. Milani (2012) provides a useful survey of alternative approaches to modelling expectations in general equilibrium models.

[27]For instance, the empirical evidence on price-setting cited in Section 2.2 indicates that nominal prices remain fixed for substantial periods. This is strongly at odds with the assumption that firms who cannot re-optimise index their prices to past inflation.

cients on habit formation and indexation of prices to past inflation fall to almost zero under adaptive learning because it generates substantial endogenous persistence in inflation and output. More recently, Slobodyan and Wouters (2012) estimated the Smets and Wouters (2007) model using an updated dataset and found that adaptive learning improves upon the fit of the rational expectations version of the model.

In Section 5.4, we consider a test of rational expectations against behavioural expectations (i.e. heuristics), as used by Brock and Hommes (1997), Brazier et al. (2006), De Grauwe and Grimaldi (2006) and De Grauwe (2010), amongst others. Behavioural expectations bear some resemblance to adaptive learning in the sense that agents use statistical methods in order to make their forecasts (in this case by choosing the weights on different forecasting rules based on their past performance), but these expectations are probably closest to those in the composite 'behavioural model' of Levine et al. (2012) in which some agents have adaptive expectations.[28] They estimate, using Bayesian methods, a New Keynesian model in which a fraction of agents have full rational expectations and the remaining fraction form expectations adaptively. The composite model clearly outperforms a model with only rational expectations based on the likelihood race methodology.[29] In contrast to these findings and those of the adaptive learning literature, the formal statistical test that we discuss in Section 5.4 strongly favours a New Keynesian model with rational expectations over one with non-rational expectations. We emphasise that this test is not vulnerable to criticisms of the assumption that prices are indexed to past inflation because it is based on a model where price indexation is absent.

5.1.1   Summary

New Keynesian models have been successful matching the dynamics of aggregate macro data. The literature has recently considered refinements of rational expectations that might improve the fit of these models further. The logical next step is to test these models against the data using formal statistical tests that accept or reject the model and variants of it. Such tests should narrow down the set of models that need to be considered in future research by rejecting those that cannot match the data. As we discuss below, this challenge has been taken on by a recent strand of applied macro literature. We discuss this literature in the next section before turning to a direct test of rational expectations versus behavioural expectations in a New Keynesian model.

---

[28]Levine et al. (2012) note in passing (see p. 1288) that the expectations in their model are similar in spirit to both adaptive learning and behavioural expectations.

[29]As Levine et al. (2012, p. 1298) note, a limitation of the likelihood race methodology is that the assessment of model fit is only relative to rival models with different restrictions. Hence, the outperforming model may still be poor at capturing important dynamics in the data, making it difficult to draw any firm conclusions from these results.

5.2 Testing macro models: indirect inference

Indirect inference is a simulation-based method used for estimating or evaluating economic models (Smith, 1993; Gourieoux and Monfort, 1996). Its distinguishing feature is the use of an 'auxiliary model' – which need not be correctly specified – to represent the time series properties in the data. Estimation and evaluation are based upon the auxiliary model, which acts as a criterion function that selects important features in the data; this is the sense in which inference is 'indirect'. Indirect inference has most commonly been used to estimate structural economic models . An important advantage of this approach is that estimation is possible when the likelihood function is difficult to evaluate or analytically cumbersome (Canova, 2007). Indeed, indirect inference can be used to estimate almost any economic model from which data can be simulated. Indirect inference basically chooses the parameters of the macro model so that, from the point of the view of the auxiliary model, the actual and simulated data look similar. In this respect, researchers can use indirect inference to focus on matching those aspects in the data which they view as most important. The extension of indirect inference to the case of evaluating macro models is discussed in detail below, as this our main focus in this section.

Indirect inference can be used to *test* whether a macro model can simulate behaviour that is 'like' the behaviour in the data, where the data behaviour is summarised by the reduced form representation of some unknown true model, or an approximation to it. Typically, the reduced form approximation is a VAR or a VECM.[30] Then the question is whether the structural macro model could have been the generating mechanism of these coefficients. We can answer this question by simulating the model over the same data period with repeated samples of its own errors – these samples will give us 'pseudo-histories' that the model was capable of generating, and we can then ask whether the actual data (as captured by its estimated reduced form coefficients) could have been one of these histories.

To formalise this intuition, we explain the indirect inference testing procedure. The method works as follows. Suppose we have a macro model with a fixed vector of structural parameters θ which can be taken as given, having been reached either by estimation or calibration. Let the vector of auxiliary VAR parameters associated with simulated data of length T periods from the model be denoted $\alpha_T(\theta)$. The corresponding parameter vector from the VAR on the actual data of length T periods is denoted a_T. Under the null hypothesis that the model is correct, the Wald statistic $W(\theta)$ for a test of the model against the data is based on the difference between the VAR parameters estimated from the data a_T, and the mean VAR parameter vector $\overline{\alpha_{T,N}(\theta)}$ estimated from N

---

[30] Wickens (2011, pp. 506-8) has shown that DSGE models will have a VARMA or VARIMA representation. Indirect inference can proceed with approximations to this because the researcher runs a VAR (or VECM) using the data from the simulated model; hence the model's simulations of the approximation give rise to the distribution against which the approximating VAR or VECM on the data is tested for whether it comes from this distribution.

bootstrapped samples from the macro model:[31]

$$W(\theta) = d'\Omega(\theta)^{-1}d \qquad (9)$$

where $d \equiv a_T - \overline{\alpha_{T,N}(\theta)}$ and $\Omega(\theta)^{-1}$ is the inverse variance-covariance matrix of the distribution of d.

Notice that the Wald test statistic can be interpreted as a quadratic loss function in the deviations of the VAR parameters estimated on the data from those implied (on average) by the bootstrapped macro model. Essentially, the Wald statistic is assessing whether the VAR parameters for the data fall outside the joint confidence limit of the model-simulated parameter distribution. High values of the Wald statistic will reject the null hypothesis that the model is the data-generating process because they tell us that it is unlikely that the VAR parameters in the data could have been generated by the simulated model.

It should be noted that the Wald statistic has typically been compared to its simulated bootstrap distribution and not the asymptotic distribution implied by theory under the null. It is therefore the percentiles of the bootstrap distribution that provide the basis for accepting or rejecting a particular model, with a Wald statistic higher than the 95% critical value being the standard rejection criterion. The Wald test is thus an example of what Canova (2007) calls a 'size test' of an economic model. In a simple univariate size test the researcher might compare a particular correlation in the data with the distribution of correlations implied by many simulations of an economic model, rejecting the model if the real-world correlation does not fall within a particular confidence interval.[32] The Wald test is based on the same logic, but it is more general because it enables the researcher to conduct an overall test of a model based on its *joint* parameter distribution.

Moreover, it is important to note that although the Wald statistic in Equation (9) considers only a direct test of the VAR parameters, the test can easily be augmented to include the variances of key macro variables or the impulse response functions implied by the VAR coefficients. The papers we discuss below take this more general approach, so we provide some technical details in a footnote for interested readers.[33]

5.3 Testing the degree of nominal rigidity in the US economy using indirect inference

Le et al. (2011) use indirect inference to investigate the degree of nominal rigidity in the US economy in the postwar period. This paper provides a useful

---

[31]Bootstrapping N samples avoids the weaknesses associated with a test based on a single random sample of simulated data.

[32]For example, Lim and McNelis (2008, Ch. 9, pp. 171–72) simulate the distribution of the real-wage employment correlation in an open economy New Keynesian model and discuss how a size test could be conducted using this distribution.

[33]If we let $g(.)$ denote a vector valued function that includes the different aspects of the data that are to be tested, the Wald statistic will be amended to $W(\theta) = d'\Omega(\theta)^{-1}d$, where now $d \equiv g(a_T) - \overline{g(\alpha_{T,N}(\theta))}$.

example of indirect inference in action and is relevant to our discussion here because it sheds light on the performance of New Keynesian models in formal statistical tests. Le et al. consider a weighted version of the Smets and Wouters (2007) model in which a New Classical sector with flexible wages and prices and a one-quarter information lag are introduced. The coefficients of the model – including the relative weight on the New Classical sector – are chosen to minimise a Wald statistic based on the VAR parameters and the variances of key macro variables. Several notable results emerge. The most striking results come from the post-1984 Great Moderation period: the Wald-minimising model has a weight of almost 1 on the New Keynesian wage and price sectors and passes the test comfortably with a Wald percentile of 83.8%. These results show that New Keynesian models with rational expectations can successfully mimic some key features of the data *and* pass a stringent statistical test. Interestingly, when the model is estimated on the entire postwar period, the Wald-minimising weights on the New Keynesian part of the model are noticeably lower (suggesting that nominal rigidity has increased over time) and the minimum Wald of 98.7% exceeds the standard rejection percentile of 95%, so that the model is rejected. Overall, the results of Le et al. (2011) provide support for New Keynesian models while suggesting, at the same time, that there is considerable scope for improvement. In the next section we ask whether a behavioural expectations version of the New Keynesian model can improve on the baseline model.

We focus on behavioural expectations for two reasons. First, these expectations have no forward-looking component, so they imply a model in which IT will clearly dominate PT (see Section 3); therefore, a finding in favour of behavioural expectations would imply that adopting PT could be costly and thus lend fresh support to the Bank of Canada's justifications for sticking with IT. Second, models with behavioural expectations have had some success in accounting for puzzles in the finance literature. For instance, Bernatzi and Thaler (1995) show how 'loss aversion' could account for the equity premium puzzle, while De Grauwe and Grimaldi (2006) show that a behavioral finance model of the foreign exchange market can produce excess kurtosis and fat-tails in exchange rate returns.

These models have also had some success in formal statistical tests against the data. As shown by ap Gwilym (2010), for example, a simple behavioural finance model cannot be rejected as the data-generating process of the FTSE. One problem, however, is that empirical tests of single series suffer from low power to reject the null that the model is the true data-generating process; indeed, Meenagh et al. (2007) show that an efficient markets model with rational expectations and regime-switching can also match the time series properties of the FTSE. This difficulty points to need for higher-power multivariate tests to distinguish between rational and behavioural expectations. We consider such a test in the next section.

5.4 A test of rational versus behavioural expectations in a macro model of the US economy

In this section we discuss a test of behavioural versus rational expectations based on indirect inference. As noted above, this method tests whether a macro model can match the *joint* statistical distributions of moments in the data – a powerful test. In considering such tests, we emphasise that these are tests of a joint hypothesis about expectations and New Keynesian models. It may well be the case that alternative macro models or expectational assumptions could pass these statistical tests, but as far as we know there is no evidence bearing on this question available at this time. Clearly, if the model with rational expectations is rejected we cannot not know for sure whether this is down to the model itself or the rational expectations assumption; either way PT must come into question. However, if the model is not rejected, then it also follows that we cannot reject the arguments for PT that come from these models. Thus there is an important asymmetry in these tests: rejection leads to doubt and rethinking, including of PT, whereas non-rejection leads to survival of the model and benefit of the doubt for PT. This is what we turn to now in examining such a test.

Liu and Minford (2012) test a rational expectations version of the baseline New Keynesian model against a version based on behavioural expectations; we go into it in some detail because it is the only available test of these rival models of which we are aware. The behavioural model is a standard New Keynesian-type model similar to that in De Grauwe (2010). It consists of an IS curve, a Phillips curve, and a Taylor rule:

$$\widetilde{x}_t = \tilde{E}_t \widetilde{x}_{t+1} - a_1 (R_t - \tilde{E}_t \pi_{t+1}) + \varepsilon_{1t} \tag{10}$$

$$\pi_t = b_1 \widetilde{x}_t + \beta \tilde{E}_t \pi_{t+1} + k \varepsilon_{2t} \tag{11}$$

$$R_t = (1 - c_1)(c_2 \pi_t + c_3 \widetilde{x}_t) + c_1 R_{t-1} + \varepsilon_{3t} \tag{12}$$

where $\widetilde{x}_t$ is the output gap, π_t is the rate of inflation, R_t is the nominal interest rate, and ε_{1t}, ε_{2t}, and ε_{3t} are the demand error, supply error and policy error respectively.

These errors are assumed to be autoregressive processes. The errors are extracted from the model and the data; thus the model implies the errors, conditional on the data. Equation (10) is the IS curve, where $\tilde{E}$ refers to expectations that are not formed rationally. The IS curve is standard, including the expectation of output gap in the next period and the real interest rate. Equation (11) is a New Keynesian Phillips curve augmented with behavioural expectations, while Equation (12) is a Taylor rule with interest rate smoothing.

The difference between the behavioural and rational expectations models lies in expectations formation. The expectation term in the behavioural model, $\tilde{E}$ is the weighted average of two kinds of forecasting rule. One is the fundamental

forecasting rule, by which agents forecast the output gap or inflation at their steady state values. The other one is the extrapolative rule, by which individuals extrapolate the most recent value into the future. Thus:

$$\tilde{E}_t^f \tilde{x}_{t+1} = 0 \tag{13}$$

$$\tilde{E}_t^e \tilde{x}_{t+1} = \tilde{x}_{t-1} \tag{14}$$

$$\tilde{E}_t^{tar} \pi_{t+1} = 0 \tag{15}$$

$$\tilde{E}_t^{ext} \pi_{t+1} = \pi_{t-1} \tag{16}$$

Equations (13) and (14) are the forecasting rules for the output gap, while Equations (15) and (16) are the forecasting rules for inflation. The steady state output gap is zero, while the inflation target in the Taylor Rule is the steady state inflation rate, which is set at zero since the data is linearly detrended and demeaned.

In De Grauwe (2010), it is assumed that the market forecast is the weighted average of the fundamentalist and extrapolative rules. Equation (17) is the market forecast for the output gap, while Equation (18) is the market forecast for inflation:

$$\tilde{E}_t \tilde{x}_{t+1} = \alpha_{f,t} * 0 + \alpha_{e,t} \tilde{x}_{t-1} = \alpha_{e,t} \tilde{x}_{t-1} \tag{17}$$

$$\tilde{E}_t \pi_{t+1} = \beta_{tar,t} * 0 + \beta_{ext,t} \pi_{t-1} = \beta_{ext,t} \pi_{t-1} \tag{18}$$

where $\alpha_{f,t}$ and $\alpha_{e,t}$ are the probabilities that agents will use a fundamentalist and extrapolative rule for forecasting the output gap, $\beta_{tar,t}$ and $\beta_{ext,t}$ are the equivalents for inflation. These probabilities sum to one and are determined by past forecast success of the two rules in an intuitive way.

The solution method to the behavioural model is obtained by substituting the expectation formation of Equations (17) and (18) into Equations (10) and (11). The model therefore becomes

$$\tilde{x}_t = \alpha_{e,t} \tilde{x}_{t-1} - a_1(R_t - \beta_{ext,t} \pi_{t-1}) + \varepsilon_{1t} \tag{19}$$

$$\pi_t = b_1 \tilde{x}_t + \beta(\beta_{ext,t} \pi_{t-1}) + k\varepsilon_{2t} \tag{20}$$

$$R_t = (1 - c_1)(c_2 \pi_t + c_3 \tilde{x}_t) + c_1 R_{t-1} + \varepsilon_{3t} \tag{21}$$

The rational expectations model is defined as Equation (10)-(12), except that the expectations are formed rationally. In other words, it is a version of the baseline New Keynesian model. Notice that since the only specification difference

between the two models is in the nature of these expectations, the comparison tests the different expectational assumptions while allowing each model the benefit of reestimation of parameter values. The rational expectations version of the model can be solved in the standard way; Dynare (Adjemian, 2011) is used for this. As Equations (19)-(21) make clear, the behavioural expectations model is purely backward-looking. Therefore, the test here is between the baseline New Keynesian model in which PT is clearly dominant (Vestin, 2006) and a backward-looking version of the same model in which IT clearly dominates PT.

The results are shown in Table 1, which includes the Wald percentiles for both models, broken down by variances alone ('volatility'), VAR coefficients alone ('dynamics'), and the full vector of descriptors. In all aspects the behavioural model is strongly rejected, while the rational expectations model comfortably passes the Wald test. It should be emphasised that this is after allowing each model to explore all possible values for all the model's parameters to find the set of 'best' parameter values that gets closest to the data behaviour.

Wald percentile (%)	Behavioural	Rational
Dynamics (VAR coeffs)	100	90.0
Volatility (variances)	96.0	24.2
Overall (all α)	100	79.8

Table 1: Comparison of Behavioural and Rational Expectations Models, indirectly re-estimated Parameters

It may seem counter-intuitive that a theory of expectations so apparently unrealistic as rational expectations could replicate macro behaviour so much better than behavioural expectations. However, as Muth (1961, p. 318) points out, in an economy with informed and misinformed agents, the well-informed agents could profit by selling their information to the misinformed agents. This mechanism may explain how superior information initially held by a small number of agents eventually makes its way into the hands of economically active agents. For instance, Minford and Peel (2002) argue that rational expectations rests on the ability of competitive markets in information to process it efficiently: industries grow up to make these markets as efficient as possible – such as analysts, portfolio advisers, forecasters, hedge funds, and investment banks. The 'ordinary person' may not have literal rational expectations but is enabled to access sources with superior information. Under this argument, realistic but non-rational expectations will be driven out by these sources, so that models based on them struggle to fit the data behaviour.[34]

In short, we do not find it plausible that relevant information will be neglected in favour of heuristics as under behavioural expectations. And while we

---

[34]This argument relies on the assumption that less-informed agents are not deliberately misled by those who have an informational advantage. While examples of such behaviour are observed in the real world, these actions are unlikely to be a systematic feature because they will lead to punishment by withdrawal of custom or enforcement of the law.

find approaches that involve learning more plausible, we believe that rational expectations may be a reasonable approximation for studying many policy issues – e.g. the numerous (stochastic) steady-state analyses of IT and PT reviewed in this survey. Furthermore, in cases such as the transition to PT where some learning may be required, we would expect agents to learn about the new regime quite quickly as long as it is communicated in an effective manner. Consistent with this position, the empirical evidence that we have reviewed does not enable us to reject the rational expectations hypothesis. This finding has relevance for policymakers such as the Bank of Canada who remain sceptical about PT due to its reliance on rational expectations.

6 Conclusion

We have surveyed recent literature comparing inflation targeting (IT) and price-level targeting (PT) as macroeconomic stabilisation policies, focusing in particular on New Keynesian models and areas that have seen significant developments since Ambler's (2009) survey: optimal monetary policy under commitment and discretion; the zero lower bound; financial frictions; and the costs of transition from IT to PT. The main conclusion reached by Ambler was that the ability of PT to improve social welfare in New Keynesian models rests with the assumption of rational expectations. The recent literature suggests that things are more complicated than this. In particular, we highlighted an important split in the literature: PT consistently outperforms IT when policymakers commit to simple Taylor-type rules, but results favouring PT which are derived from models where central banks minimise loss functions can be overturned with small deviations from the baseline New Keynesian model. We thus emerge with a more precise version of Ambler's original conclusion that we view as relevant for real-world central banks: PT tends to improve stabilisation in New Keynesian models when policymakers commit to simple Taylor-type rules.

Several of the developments we surveyed are initial attempts to address issues raised by the financial crisis. An important finding that emerges is that PT is potentially very attractive in the context of the zero lower bound because, in conditions of deflationary recession, a price-level target induces expectations of higher than usual inflation. These expectations in turn induce negative real interest rates which stimulate economic activity out of the recession, putting an end to deflation and ending lower bound episodes rapidly. This mechanism raises the possibility of safely lowering trend inflation under a PT regime. The recent literature suggests that this dual mandate – a price-level target with a lower trend inflation rate – could bring substantial welfare gains, though the assumption that PT is highly credible is crucial for this result. There is also evidence that augmenting New Keynesian models with financial frictions leaves intact the potential welfare gains from PT, albeit that this literature is at an early stage. Finally, the recent literature re-affirms the importance of rational expectations in the sense that the transition from IT to PT could be costly if agents need time to learn about the new regime and have doubts about PT.

Since the beneficial effects of PT appear to hang on the joint assumption that economic agents are rational and the economy New Keynesian, we devoted the second half of the survey to empirical evidence on rational expectations and the applied macro literature on the performance of New Keynesian models. In addition, we surveyed a more recent strand of applied literature that has formally tested New Keynesian models and alternative types of expectation formation. Overall, the evidence on rational expectations and New Keynesian models is not clear-cut, but we note that New Keynesian models can pass formal statistical tests against macro data and that models with rational expectations outperform those with behavioural expectations (i.e. heuristics) in direct statistical tests. These findings suggest that the benefits of PT might be higher than envisaged by many policymakers. We therefore conclude that policymakers should continue to pay attention to PT in the future.

Acknowledgements

References

Adam, K. and Billi, R.M. 2007. Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics* 54(3), 728–52.

Adam, K. and Billi, R.M. 2006. Optimal monetary policy under commitment with a zero bound on interest rates. *Journal of Money, Credit and Banking* 38(7), 1877–1905.

Adjemian, S., Bastani, H., Juillard, M., Karamé, F. Mihoubi, F., Perendia, G., Pfeifer, J., Ratto, M., and Villemot, S., 2011. Dynare: Reference Manual, Version 4. Dynare Working Paper No. 1, CEPREMAP.

Amano, R., Ambler, S., and Shukayev, M. 2012. Optimal price-level drift under commitment in the canonical New Keynesian model. *Canadian Journal of Economics* 45(3), 1023-36.

Amano, R., Engle-Warnick, J. and Shukayev, M. 2011. Price-level targeting and inflation expectations: experimental evidence. Bank of Canada Working Paper 2011-18.

Amano, R., Ambler, S. and Rebei, N. 2007. The macroeconomic effects of nonzero trend inflation. *Journal of Money, Credit and Banking* 39(7), 1821–1838.

Ambler, S. 2009. Price-level targeting and stabilisation policy: a survey. *Journal of Economic Surveys* 23(5), 974–97.

Ambler, S. 2007. The costs of inflation in New Keynesian models. *Bank of Canada Review* (Winter), 5–14.

Andolfatto, D., Hendry, S. and Moran, K. 2008. Are inflation expectations rational? *Journal of Monetary Economics* 55(2), 406–22.

Aoki, K. 2001. Optimal monetary policy reponses to relative price changes. *Journal of Monetary Economics* 48(1), 55–80.

ap Gwilym, R. 2010. Can behavioral finance models account for historical asset prices? *Economics Letters* 108(2), 187–89.

Ascari, G., 2004. Staggered prices and trend inflation: some nuisances. *Review of Economic Dynamics* 7(3), 642–67.

Bailey, M.J. 1956. The welfare cost of inflationary finance. *Journal of Political Economy* 64(2), 93–110.

Bailliu, J., Meh, C. and Zhang, Y. 2012. Macroprudential rules and monetary policy when financial frictions matter. Bank of Canada Working Paper 2012-6.

Bank of Canada 2011. Renewal of the inflation-control target. Ottawa. Available at: http://www.bankofcanada.ca/wp-content/uploads/2011/11/background_nov11.pdf.

Bank of Canada 2006. Inflation target renewal agreement background document. Ottawa.

Bao, T., Duffy, J., and Hommes, C. 2013. Learning, forecasting and optimizing: an experimental study. *European Economic Review* 61, 186–204.

Barro, R. J. and Gordon, D.B. 1983. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91(4), 589–610.

Batini, N. and Yates, T. 2003. Hybrid inflation and price-level targeting. *Journal of Money, Credit and Banking* 35 (3), pp. 283-300.

Berg, C., Jonung, L. 1999. Pioneering price-level targeting: the Swedish experience 1931-1937. *Journal of Monetary Economics* 43(3), 525–51.

Bernanke, B., Gertler, M. and Gilchrist, S. 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J.B. and Woodford, M. eds. Handbook of Macroeconomics. Amsterdam: North-Holland.

Bernatzi, S. and Thaler, R.H. 1995. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110(1), 73–92.

Bils, M. and Klenow, P. 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–85.

Blake, A.P. and Kirsanova, T. 2012. Discretionary policy and multiple equilibria in LQ RE models. *Review of Economic Studies* 79 (4), 1309–39.

Blake, A.P., Kirsanova, T. and Yates, A. 2011. The gains from delegation revisited: price-level targeting, speed limit and interest rate smoothing policies. Bank of England Working Paper 415.

Bloomfield, R. and Hales, J. 2002. Predicting the next step of a random walk: experimental evidence of regime-shifting beliefs. *Journal of Financial Economics,* 65(3), 397–414.

Bonham, C.S. and Cohen, R.H. 2001. To aggregate, pool or neither: testing the rational expectations hypothesis using survey data. J*ournal of Business and Economic Statistics* 19 (3), 278–291.

Brazier, A., Harrison, R., King, M. and Yates, T. 2006. The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model. *International Journal of Central Banking* 4(2), 219–54.

Brock, W.A., and Hommes, C.H. 1997. A rational route to randomness. *Econometrica* 65 (5),1059–95.

Calvo, G. 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), pp. 383–98.

Canova, F. 2007. *Methods for applied macroeconomic research.* NJ: Princeton University Press.

Cateau, G. and Dorich, J. 2012. Price-level targeting, the zero lower bound on the nominal interest rate and imperfect credibility. Manuscript, Bank of Canada.

Catueau, G., Kryvtsov, O., Sukayev, M. and Ueberfeldt, A. 2009. Adopting price-level targeting under imperfect credibility in ToTEM. Bank of Canada Working Paper No. 2009-17.

Christiano, L.J., Eichenbaum, M.S. and Evans, C.L. 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113(1), 1–45.

Clarida, R., Galí, J. and Gertler, M .1999. The science of monetary policy: a new Keynesian perspective. *Journal of Economic Literature* 37(4), 1661–1707.

Coibion, O., Gorodnichenko, Y. and Wieland, J. 2012. The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies* 79(4), 1371–1406.

Coibion, O., Gorodnichenko, Y. 2012. What can survey forecasts tell us about information rigidities? *Journal of Political Economy* 120(1), 116–159.

Coletti, D., Lalonde, R., Masson, P., Muir, D. and Snudden, S. 2012. Commodities and monetary policy: implications for inflation and price level targeting. Bank of Canada Working Paper 2012-16.

Coletti, D., Lalonde, R. and Muir, D. 2008. Inflation targeting and price-level-path targeting in the Global Economy Model: some open economy considerations. *IMF Staff Papers* 55(2), 326-338.

Cournède, B. and Moccero, D. 2009. Is there a case for price-level targeting? OECD Economics Department Working Papers No. 721, OECD Publishing.

Covas, F. and Zhang, Y. 2010. Price-level versus inflation targeting with financial market imperfections. *Canadian Journal of Economics* 43(4), 1302–32.

De Grauwe, P. 2010. Top-down versus bottom-up macroeconomics. *CESifo Economic Studies* 56(4), 465–97.

De Grauwe, P. and Grimaldi, M. 2006. *The exchange rate in a behavioral finance framework*. NJ: Princeton University Press.

Dhyne, E., Alvarez, L.J., Le Bihan, H., Veronese, G., Dias, D., Hoffmann, J., Jonker, N., Lunnemann, P., Rumler, F. and Vilmunen, J. 2006. Price setting in the Euro Area and the United States: some facts from individual consumer price data. *Journal of Economic Perspectives* 20(2), 171–192.

Dib, A., Mendicino, C. and Zhang, Y. 2013. Price-level targeting rules and financial shocks: the case of Canada. *Economic Modelling* 30(C), pp. 941–53.

Dib, A., Mendicino, C. and Zhang, Y. 2008. Price-level targeting in a small open economy with financial frictions: welfare analysis. Bank of Canada Working Paper No. 2008-40.

Dixon, H. and Le Bihan, H. 2012. Generalised Taylor and generalised Calvo wage and price setting: micro-evidence with macro implications. *Economic Journal* 122(560), 532–54.

Eggertsson, G.B. and Woodford, M. 2003. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 1:2003, 139–211.

Erceg, C., Henderson, D. and Levin, A.T. 2000. Optimal monetary policy with staggered wage and price contracts. Journal of Monetary Economics 46(2), 281–313.

Evans, G.W. and Honkapohja, S.M.S. 2006. Monetary policy, expectations and commitment. *Scandinavian Journal of Economics* 108(1), 15–38.

Evans, G.W. and Honkapohja, S.M.S. 2001. *Learning and expectations in macroeconomics*. NJ: Princeton University Press.

Fair, R.C. 1993. Testing the rational expectations hypothesis in macroeconometric models. *Oxford Economic Papers* 45(2), 169–90.

Fujiwara, I., Hirose, Y. and Shintani, M. 2011. Can news be a major source of aggregate fluctuations? *Journal of Money, Credit and Banking* 43, 1–29.

Galí, J. and Gertler, M. 1999. Inflation dynamics: a structural econometric analysis. *Journal of Monetary Economics* 44(2), 195–222.

Gaspar, V., Smets, F. and Vestin, D. 2007. Is the time ripe for price level path stability? European Central Bank Working Paper 818.

Gerberding, C., Gerke, R. and Hammermann, F. 2012. Price-level targeting when there is price-level drift. *Journal of Macroeconomics* 34(3), 757–68.

Giannoni, M. 2010. Optimal interest rate rules in a forward-looking model, and inflation stabilization versus price-level stabilization. NBER Working Paper 15986.

Golosov, M. and Lucas, R.E. Jr. 2007. Menu costs and Phillips curves. *Econometrica* 115(2), 171–99.

Gourieroux, C. and Monfort, A. 1996. Simulation-based econometric methods. Oxford: Oxford University Press.

Haldane, A.G. and Salmon, C.K. 1995. Three issues on inflation targets: some United Kingdom evidence. In A.G. Haldane (ed.), Targeting Inflation (pp. 170–201). London: Bank of England.

Heemejer, P., Hommes, C., Sonnemans, J. Tuinstra, J. 2009. Price stability and volatility in markets with positive and negative expectations feedback: an experimental investigation. *Journal of Economic Dynamics and Control* 33(5), 1052–72.

Holden, K., Peel, D.A., and Thompson, J.L. 1985. *Expectations – theory and evidence*. London: Macmillan.

Kahn, G.A. 2009. Beyond inflation targeting: should central banks target the price level? *Economic Review, Federal Reserve Bank of Kansas City* (2009:3), 35–64.

Karadi, P. and Reiff, A. 2012. Large shocks in menu cost models. European Central Bank Working Paper No.1453.

Khan, H.U. and Tsoukalas, J. 2012. The quantitative importance of news shocks in estimated DSGE models. *Journal of Money, Credit and Banking* 44(8), 1535–61.

Kryvtsov, O. and Petersen, L. 2013. Expectations and monetary policy: experimental evidence. Bank of Canada Working Paper 2013–44.

Kryvtsov, O., Shukayev, M. and Ueberfeldt, A. 2008. Adopting price-level targeting under imperfect credibility: an update. Bank of Canada Working Paper No. 2008-37.

Kydland, F.E. and Prescott, E.C. 1982. Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–70.

Kydland, F E. and Prescott, E.C. 1977. Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy* 85(3), 473–91.

Laxton, D. 2008. Getting to know the Global Economy Model and its philosophy. *IMF Staff Papers* 55(2), 213–42.

Le, V.P.M., Meenagh, D., Minford, P. and Wickens, M.R. 2011. How much nominal rigidity is there in the US economy? Testing a New Keynesian DSGE model using indirect inference. *Journal of Economic Dynamics and Control* 35(12), 2078–2104.

Lebow, D.E., Roberts, J.M. and Stockton, D.J. 1992. Economic performance under price stability. US Board of Governors of the Federal Reserve Board, Economic Activity Section Working Paper No. 125.

Lei, V., Noussair, C.N., and Plott, C.R. 2001. Nonspeculative bubbles in experimental asset markets: lack of common knowledge of rationality vs. actual irrationality. *Econometrica* 69(4), 831–59.

Levin, A., López-Salido, D., Nelson, E. and Yun, T., 2010. Limitations on the effectiveness of forward guidance at zero lower bound. *International Journal of Central Banking* 6, 143–89.

Levine , P., Pearlman, J., Perendia, G. and Yang, B. 2012. Endogenous persistence in an estimated DSGE model under imperfect information. *Economic Journal* 122, 1287–1312.

Lim, G.C. and McNelis, P.D. 2008. *Computational macroeconomics for the open economy*. MA: MIT Press.

Liu, C. and Minford, P. 2012. Comparing behavioural and rational expectations for the US post-war economy. CEPR Discussion Paper 9132.

Lucas, R.E. Jr. 1976. Econometric policy evaluation: a critique. *Carnegie-Rochester Series on Public Policy* 1(1), 19–46.

Lucas, R.E. Jr. 1972. Expectations and the neutrality of money. *Journal of Economic Theory* 4(2), 103–24.

Mackowiak, B. and Wiederholt, M. 2011. Rational inattention and macroeconomic dynamics. ECB Working Paper No. 1331.

Mackowiak, B. and Wiederholt, M. 2009. Optimal sticky prices under rational inattention. *American Economic Review* 99(3), 769–803.

Mankiw, N.G. and Reis, R., 2007. Sticky information in general equilibrium. *Journal of the European Economic Association* 5(2–3), 603–13.

Mankiw, N.G. and Reis, R., 2002. Sticky infomation versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4), 1295–1328.

Mankiw, N.G., Reis, R., Wolfers, J., 2003. Disagreement about inflation expectations. *NBER Macroeconomics Annual* 18 (1), 209–248.

Masson, P.R. and Shukayev, M.D. 2012. Are bygones not bygones? Modeling price-level targeting with an escape clause and lessons from the gold standard. *Journal of Macroeconomics* 33(2), 162-75.

Meenagh, D., Minford, P. and Peel, D.A. 2007. Simulating stock returns under switching regimes – a new test of market efficiency. *Economics Letters* 94(2), 235–39.

Midrigan, V. 2011. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica* 79(4), 1139–80.

Milani, F. 2012. The modeling of expectations in empirical DSGE models: a survey. *Advances in Econometrics* 28, 3–38.

Milani, F. 2007. Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics* 54(7), 2065–82.

Minford, P. and Peel, D.A. 2002. *Advanced macroeconomics: a primer*, London: Blackwell.

Murchison, S. 2010. Price-level targeting and relative price shocks. *Bank of Canada Review* (Summer), 11–21.

Muth, J.F. 1961. Rational expectations and the theory of price movements. *Econometrica* 29(3), 315–35.

Nakamura, E. and Steinsson, J. 2008. Five facts about prices: a reevaluation of menu cost models. *Quarterly Journal of Economics* 123(4), 1415–1464.

Nakov, A. 2008. Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking* 4(2), 73–127.

Nessén, M. and Vestin, D. 2005. Average inflation targeting. *Journal of Money, Credit and Banking* 37(5), 837–63.

Rogoff, K. 1985. The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics* 100(4), 1169–89.

Røisland, O. 2006. Inflation inertia and the optimal hybrid inflation/price-level target. *Journal of Money, Credit and Banking* 38(8), 2247–51.

Schmitt-Grohé, S. and Uribe, M., 2012. What's news in business cycles?, *Econometrica* 80(6), 2733-2764.

Sims, C.A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–90.

Slobodyan, S. and Wouters, R. 2012. Learning in a medium-scale DSGE model with expectations based on small forecasting models. *American Economic Journal: Macroeconomics* 4(2), 65–101.

Smets, F. and Wouters, R. 2007. Shocks and frictions in US business cycles: a Bayesian DSGE approach. *American Economic Review* 97(3), 586–606.

Smith, A.A. Jr. 1993. Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics* 8(S1), S63–S84.

Smith, V.L., Suchanek, G.L. and Williams, A.W. 1988. Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica* 56(5), 1119–51.

Steinsson, J. 2003. Optimal monetary policy in an economy with inflation persistence. J*ournal of Monetary Economics* 50(7), 1425–56.

Svensson, L.E.O. 1999. Price-level targeting versus inflation targeting: a free lunch? *Journal of Money, Credit and Banking* 31(3), 277–95.

Svensson, L.E.O. 1997. Optimal inflation targets, 'conservative' central banks, and linear inflation contracts. *American Economic Review* 87(1), pp. 98–114.

Taylor, J.B. 1999. The robustness and efficiency of monetary policy rules as guidelines for interest rate setting by the European Central Bank. *Journal of Monetary Economics* 43(3), 655–79.

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

Taylor, J.B. 1980. Aggregate dynamics and staggered contracts, *Journal of Political Economy* 88(1), 1-23.

Thomas, L.B. Jr. 1999. Survey measures of expected U.S. inflation. *Journal of Economic Perspectives* 13(4), 125–144.

Vestin, D. 2006. Price-level versus inflation targeting. *Journal of Monetary Economics* 53(7), 1361–76.

Wagener, F. (Forthcoming). Expectations in experiments. Forthcoming in *Annual Review of Economics*.

Walsh, C.E. 2010. Monetary theory and policy (3rd Ed.). MIT Press.

Walsh, C.E. 1995. Optimal contracts for central bankers. *American Economic Review* 85(1), 150–67.

Wickens, M.R. 2012. *Macroeconomic theory: a dynamic general equilibrium approach* (2nd ed.). NJ: Princeton University Press.

Wolman, A.L. 2005. Real implications of the zero bound on nominal interest rates. *Journal of Money, Credit and Banking* 37(2), 273–296.

Woodford, M. 2003. Interest and prices: Foundations of a theory of monetary policy. NJ: Princeton University Press.

Yates, A. 2004. Monetary policy and the zero bound to interest rates: a review. *Journal of Economic Surveys* 18(3), 427–81.