BANCA D'ITALIA

EUROSISTEMA

Questioni di Economia e Finanza

Deflationary shocks and de-anchoring of inflation expectations

by Fabio Busetti, Giuseppe Ferrero, Andrea Gerali and Alberto Locarno

Questioni di Economia e Finanza

(Occasional papers)

Deflationary shocks and de-anchoring of inflation expectations

by Fabio Busetti, Giuseppe Ferrero, Andrea Gerali and Alberto Locarno

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of the Bank of Italy and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within the Bank of Italy, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

DEFLATIONARY SHOCKS AND DE-ANCHORING OF INFLATION EXPECTATIONS

by Fabio Busetti, Giuseppe Ferrero, Andrea Gerali and Alberto Locarno*

Abstract

A prolonged period of low inflation can heighten the risk of inflation expectations de-anchoring from the central bank's objective, particularly when monetary policy rates are near the zero lower bound. This paper investigates the effects of a sequence of deflationary shocks on expected/realized inflation and output. To do so we consider a simple New Keynesian model where agents have incomplete information about the working of the economy and form expectations through an adaptive learning process (in the sense that they behave like econometricians, using regressions to anticipate the future value of the variables of interest). The model is simulated with euro area data over the period 2014-16 under assumptions of both rational expectations and learning. The main findings are the following: (i) under learning, price dynamics in 2015-16 is 0.6 percentage points lower on average than in the case of fully rational agents, as inflation expectations are strongly affected by repeated deflationary shocks; (ii) the learning process implies a (data-driven) de-anchoring of inflation expectations from the central bank's target, which would be perceived by economic agents to fall to 0.8 per cent at the end of 2016; (iii) output expectations would also be lower in the case of learning, resulting in a slower recovery of economic activity.

Contents

---

* Bank of Italy, Directorate General for Economics, Statistics and Research.

1 Introduction

A prolonged period of low inflation, particularly in a situation of monetary policy rates near the zero lower bound, can heighten the risk of inflation expectations de-anchoring from the central bank objective. Indeed, over the recent months, market-based expectations of euro area inflation have progressively declined, even for those medium- to longer-term horizons that are typically relevant for monetary policy. The experiences of the past show how large the costs of policy errors in the wake of a sequence of inflationary or deflationary shocks can be: in the 1970s they gave rise to a stagflationary episode; in the 1990s they were responsible for Japan's lost decade; today they might be a prelude to a prolonged period of inflation in the euro area being well below the ECB definition of price stability.

Concerns about the outlook for inflation have been recently voiced by members of the Governing Council of the ECB. Speaking at the World Savings Day conference in Rome, Bank of Italy's Governor Visco stressed the link between actual and expected inflation and hinted that the de-anchoring of inflation expectations may easily turn into a harbinger of deflation.[1] President Draghi emphasised several times the risk that a too prolonged period of low inflation may become embedded in inflation expectations; he stated however that the ECB is willing to do what is needed in order to raise inflation and inflation expectations as fast as possible, as required by the price stability mandate.[2]

The purpose of this work is to assess the effects of a sequence of deflationary shocks, such as those that hit the euro area in the last one and half year, in an environment of loosely anchored expectations. To do so we depart from the assumption of rational expectations by considering a simple New Keynesian model where agents have limited information about the working of the economy and form expectations through an adaptive learning process (in the sense that they behave like econometricians, using regressions to anticipate the future value of the variables of interest). The model is first estimated using data up to 2012Q4; then, using a filtering technique, we recover the structural shocks that make the model dynamics consistent with the more recent data (from 2013Q1 to 2014Q3) and forecasts (from 2013Q4 to 2016Q4)[3] and then simulate the model under learning. In this way we measure the impact on economic outcomes of assuming a data-

---

[1]Specifically, Governor Visco said that "Inflation is lower than 1% in 15 out of 18 countries, and it's negative in 5 countries, including Italy, while price expectations for the next 5-10 years have dropped below 2%. We're not in deflation, but we cannot ignore the concrete risk of it ". See Visco (2014).

[2]See Draghi (2014)

[3]We have used the Eurosystem staff projections published in the September 2014 issue of the ECB Monthly Bulletin.

driven expectation mechanism that lacks the nominal anchor provided under rational expectations by the central bank's inflation objective.

The main findings of this paper are the followings.

(i) Under learning, price dynamics in 2015-16 are on average 0.6 percentage points lower than in the case of fully rational agents, as inflation expectations are strongly affected by repeated deflationary shocks.

(ii) The learning process implies a (data-driven) de-anchoring of inflation expectations from the central bank target, which would be perceived by economic agents to fall to 0.8% at the end of 2016.

(iii) Output expectations would also be lower in the case of learning, resulting in a slower recovery of economic activity (on average by 0.4 percent in 2015-16).

(iv) Long-run simulations, where all the model structural shocks are set to zero, indicate that under learning inflation tends to remain persistently low whereas it returns rapidly towards the central bank target if expectations are assumed to be rational.

The paper proceeds as follows. Section 2 provides a brief discussion of the hypothesis of learning against that of rational expectations, connecting the paper with the relevant literature. Section 3 contains an empirical motivation of our macroeconomic analysis, arguing that the recent behavior of longer-term inflation expectations appears heavily influenced by the current deflationary developments. The main results of the paper are in Section 4, where using a simple New Keynesian model we assess the macroeconomic impact of the assumption of imperfect information and learning of the agents in an environment of prolonged deflationary shocks. Section 5 concludes.

2   Rational expectations vs learning

The idea of rational expectations (RE) rests on two pillars: individual rationality and mutual consistency of perceptions about the environment. The true stochastic process of the economy is assumed known, with unpredictable random shocks as the only source of uncertainty. RE are regarded by most researchers as the most appropriate hypothesis for economic analysis, since one necessary condition for optimization is that individuals eliminate any systematically erroneous component of their behavior, including in the formation of expectations; further, from a policy perspective, this assumption rules out policies designed to exploit patterns of suboptimal expectations. In other respects, however, the RE hypothesis is unappealing, since it clashes with the principle of cognitive consistency, as it implies that agents within the model are much smarter and have much

more knowledge than economists/econometricians, faced with problems of estimation and inference.[4] Besides, in most situations, there is no sufficient incentive to upgrade from incomplete to complete information, since the costs may be prohibitively large, while the benefits tend to be small. Finally, the RE assumption begs the crucial question of how agents arrive at having rational expectations in a changing environment.[5]

For these reasons several macroeconomic theorists have gradually moved away from the strict RE framework, toward models with imperfect information and learning. Different approaches have been developed to model learning, as reviewed for example in the monograph by Evans and Honkapohja (2001). The one considered here, *adaptive learning*, assumes that agents do not know the law of motion of the economy but instead use simple (potentially misspecified) forecasting equations, called perceived law of motion (PLM), to form expectations. This departure from the perfect/complete information assumption means that agents behave as econometricians, re-estimating the parameters of their PLM as information becomes available.[6] The law of motion of the variables becomes non-stationary and does not settle down, unless agents' subjective beliefs converge to the objective distribution of the variables. The conditions under which this happens are summarized in the Appendix for the interested reader, but are mainly of two types: first, and foremost, the environment must be stable, meaning that the structural parameters and the law of motion of the stochastic shocks must not change over time; second the learning scheme must satisfy some regularity conditions related to the speed with which it discounts past observations.

2.1 A short review of the literature

Academic interest in learning was originally prompted by the idea that it might justify the RE hypothesis. An equilibrium cannot be produced out of thin air; there must be forces at work that propel the economy toward it. How is an RE equilibrium achieved? How can agents eventually become fully rational, when at first they are not?

The early answers to these questions were hardly encouraging: Frydman (1982) for

---

[4]According to the definition in Evans and Honkapohja (2008), the principle of cognitive consistency is the requirement that private agents and policymakers in the economy behave like applied economists and econometricians.

[5]In the presence of structural or policy changes, agents might not know the exact nature of the equilibrium in which they find themselves; instead, RE insists that there is "mutual consistency of perceptions about the environment" but is completely silent about how this prerequisite is achieved.

[6]In their efforts to improve their understanding of the economy, agents introduce a self-referential element in the dynamics of the model. The dependency between outcomes and beliefs becomes *bidirectional*, as the expectations encoded in the PLM change as new observations become available, but the law of motion of the observables is in turn affected by agents' beliefs.

instance proposed a proof of the impossibility of rational learning when individuals cannot determine the average of other agents' forecasts. Later on the tide turned. First, an influential paper by Bray (1982) showed how to prove stability under learning of the RE equilibrium in an asset market model; then Evans (1985) introduced the notional time concept of expectational stability. Finally, Marcet and Sargent (1989a,b) proposed to adopt, as a plausible learning concept, recursive least squares, and showed how stochastic approximation theory could be applied to prove the learnability of the RE equilibria.

In addition to assessing the plausibility of RE equilibrium, learning can serve as a selection criterion when models have multiple solutions. As an example, Evans and Honkapohja (2001) considered the non-stochastic Cagan model with government spending financed by seignorage; the model has two steady-state solutions, one with low and one with high inflation, and only the first equilibrium is learnable. In more general models, learning does not necessarily select a unique RE equilibrium, but the set of plausible solutions is usually significantly smaller than the set of all solutions.

Monetary policy has been the privileged subject of a large stream of literature on adaptive learning. Initially, research efforts were devoted to choose the specific Taylor-type interest-rate rules ensuring that E-stability was achieved. Bullard and Mitra (2007) showed that for the simple 3-equation New Keynesian model the RE equilibrium is determinate and stable under learning if and only if the "Taylor principle" is satisfied, i.e. if the nominal interest rate responds more than one for one with inflation.[7]

Later studies considered other aspects of the learning process, like for instance the transitional dynamics. As stressed by Bullard (2006), under RE once a determinate equilibrium is shown to exist nothing else really matters, whereas under learning anything that affects how fast the private sector learns the RE equilibrium affects also social welfare. Among the few analytical studies of the transition to the RE equilibrium is that of Benveniste et al. (1990), who related the speed of learning to the eigenvalues of the Jacobian of the associated ODE and derived the conditions for root-t convergence of the parameters of the forecasting equation.[8] Marcet and Sargent (1995) subsequently suggested a simple numerical procedure, based on model simulations, to estimate the rate at which the PLM approaches the actual law of motion (ALM). This issue was studied by Ferrero (2007), who showed that in a New-Keynesian model the speed of learning depends upon the Taylor rule parameters.

---

[7]In models allowing for inertia in either the IS or Phillips curve, analytical results are not available and numeric methods must be used; see Evans and McGough (2010).

[8]Root-t is the speed at which, in classical econometrics, the mean of the distribution of the least square estimator approaches the asymptotic value; under root-t convergence, the effects of initial conditions die out at an exponential rate.

Some authors focused on the assumption of 'constant gain learning', which implies discounting the past observations in the expectation formation mechanism. The idea is that the most recent data convey more accurate information on the economy's law of motion, making the mechanism more robust to (gradual) structural changes; see the appendix for further details.

Orphanides and Williams (2005) analyzed the impact of constant-gain learning on the effectiveness of central bank's strategies, showing that policy that are optimal under RE can perform poorly when adopted in an environment where agents have imperfect knowledge and learn adaptively. In particular they showed that: (i) learning leads to a bias towards more "hawkish" policies; (ii) persistent deviations of inflation expectations from target can arise following a sequence of unfavorable shocks.

3 Longer-term inflation expectations and economic conditions

As discussed in the previous section, when agents have imperfect information and learn adaptively long term inflation expectations are not necessarily firmly anchored to the central bank target, as they react sluggishly to the current economic conditions. A commonly adopted market-based measure of inflation expectations are the *inflation swap rates*, obtained from derivative contracts that involve an exchange of a fixed payment for realized inflation over a predetermined horizon. These are available daily over a wide range of horizons.[9] This measure of inflation expectation implicitly contains some risk premium component, that is however relatively small.[10]

Figure 1 shows the behavior of annual inflation expectations in the euro area for 2, 5 and 10 years ahead, over the January 2005-September 2014 decade. Regardless of the horizon, before the financial crisis inflation was anticipated to fluctuate around 2 per cent, i.e. broadly in line with the ECB target. Clearly, expectations at longer horizons display smaller fluctuations. Starting from 2008 the volatility of inflation expectations increased distinctively for all horizons, partly in connection with large movements of commodity prices and exchange rates. Since the second half of 2013, there has been a progressive decline of inflation expectations, which departed from the ECB target even at the 10 years ahead horizon. This reduction mirrors the effect of a prolonged period

---

[9]See for instance ECB (2011).
[10]See among others Garcia and Werner (2010) and Pericoli (2012).

Figure 1: Longer term inflation expectations, measured with inflation swap rates.

of low inflation and weak economic activity, where inflation data have most of the time surprised on the downside.

Figure 2 depicts the series of 'inflation surprises', defined as the difference between the euro area realized HICP inflation (the monthly series of the year-on-year inflation rate) and its average forecast, defined as the median of the analysts surveyed by Bloomberg just before the data are released. Since 2013 surprises have been mostly negative, like in the international financial crisis of 2008-09.

Recent empirical work has showed that longer term market-based inflation expectations respond (with a positive coefficient) to current surprises, taking also into account cyclical conditions and commodity prices. In particular, a 10 basis points negative surprise reduces two-year ahead expected inflation by about 3 basis points.[11] [12]

This evidence clearly contradicts the rational expectation hypothesis, but is consistent with adaptive learning behavior of the agents.

---

[11]See Di Cesare et al. (2014). In particular, they estimate the following linear regression on monthly data,

$$\pi_{t+1}^E = \beta_0 + \beta_1 \pi_t^E + \beta_2 surprise + \delta' z_t,$$

where π_{t+1}^E are market-based longer term inflation expectations and z_t is a proxy of current economic conditions, including the unemployment rate and the change in commodity prices.

[12]Along similar lines, ECB (2013) explores the link between the long-term inflation expectations and short-term movements in actual inflation.
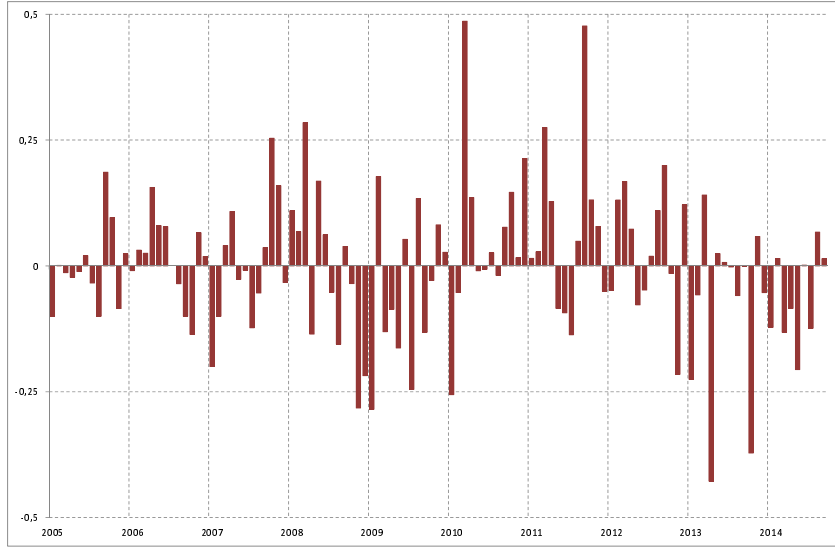
Figure 2: Inflation surprises (difference between monthly realized inflation and its forecast immediately before the data release).

4 The macroeconomic impact of deflationary shocks under non-anchored expectations

This section provides a model-based assessment of the potential impact of the persistent sequence of deflationary shocks that hit the euro area starting in 2013Q1, once we drop the assumption of RE and allow for the possibility of loosely anchored expectations.

4.1 The model and the expectations formation mechanism

We consider a simple version of a New Keynesian DSGE, featuring nominal price rigidities and intrinsic inertia in inflation and output, as reviewed inter alia in Clarida, Gali and Gertler (1999). The linearized reduced-form of the model forms the following three equation system:

$$x_t = \lambda E_t x_{t+1} + (1-\lambda)x_{t-1} - \varphi(i_t - E_t \pi_{t+1} - \bar{r}) + \varepsilon_{x,t} \tag{1}$$

$$\pi_t = \psi E_t \pi_{t+1} + (1-\psi)\pi_{t-1} + \kappa x_t + \varepsilon_{\pi,t} \tag{2}$$

$$i_t = \rho i_{t-1} + (1-\rho)\bar{i} + \alpha_\pi(\pi_t - \bar{\pi}) + \alpha_x x_t + \varepsilon_{i,t} \tag{3}$$

11

where x_t is the output gap, π_t the inflation rate, i_t the monetary policy interest rate, $\bar{\pi}$ the central bank inflation target, $\bar{i}$ ($\bar{r}$) the equilibrium nominal (real) interest rate; E_t indicates expectation formed at time t.

Equation (1) is the IS curve, while (2) represents the Phillips curve and (3) the interest rate rule; $\varepsilon_{x,t}$, $\varepsilon_{\pi,t}$ and $\varepsilon_{i,t}$ are the corresponding demand, supply and monetary policy shocks, which are assumed to be AR(1) processes to account for the persistence in the data. All the parameters of the model and the shock processes have been estimated with Bayesian techniques using data for the euro area over the period 1999Q1-2012Q4.[13]

Under rational expectations, economic agents know the model, the parameters and the shocks up to time t and use this knowledge to form expectations of inflation and output. Using the *minimum state variable* representation for the solution of the model, the expectations can be written as

$$E_t z_{t+1} = \alpha w_t$$

where z_t collects the variables whose expectations have to computed (in our case $z_t = (x_t', \pi_t')'$); α is a matrix of coefficients and w_t is a vector containing the (minimum number of) state variables for the model at hand. The vector w_t typically includes both observables variables and unobservable shocks.

Once the rational expectations assumption is dropped, the learning mechanism through which economic agents form expectations about next-period inflation and output gap must be specified (the expectations under learning will be denoted as $\hat{E}_t$). In the experiment described in the next section we assume that agents use regression equations whose specification has the following form:

$$\hat{E}_t z_{t+1} = \alpha_t \tilde{w}_t$$

where $\tilde{w}_t$ contains only the subset of states within w_t that are observables (in our case: lagged inflation, lagged estimate of the output gap and lagged interest rate, plus a constant) and α_t is a matrix of coefficients updated every period according to the following recursive algorithm:

$$
\begin{aligned}
\alpha_t &= \alpha_{t-1} + \gamma R_{t-1}^{-1} \tilde{w}_{t-1}(z_t - \alpha_{t-1}\tilde{w}_{t-1})' \\
R_t &= R_{t-1} + \gamma(\tilde{w}_{t-1}\tilde{w}_{t-1}' - R_{t-1})
\end{aligned}
\tag{4}
$$

---

[13]For inflation we used the year on year change in the HICP; the output gap estimates were taken from the ECB but in our graph we showed the OECD series, due to confidentiality issues with the ECB estimates; the policy rate was proxied by the 3-month Euribor rate.

γ is a so called "gain parameter", that, in this learning scheme, assigns higher weights to the most recent observations, as opposed to a standard least-square learning where each observations, even from the distant past, is equally weighted (see e.g. Evans and Honkapohja 2001).

When the agents use a learning scheme of this type to forecast inflation, the perceived central bank's inflation objective is mainly reflected in the coefficient attached to the constant, while the inertia in expectation formation is encoded in the coefficient attached to past inflation. Intuitively, a dis-anchoring of inflation expectations can happen if a sequence of negative "surprises"(i.e. a persistent deviation of actual from expected inflation) induces a downward revision of the perceived central bank's inflation objective (i.e. a fall in the constant coefficient) and an increase in the weight assigned to past inflation. The larger and more persistent the shocks, the stronger is their macroeconomic impact through the change in the parameters of the expectations equations.

4.2   Design of the experiment and results

The objective of the experiment is to assess whether a sequence of deflationary shocks, like those that hit the euro-area economy since 2013Q1, can trigger a de-anchoring of inflation expectations. To do so, we compare the outcomes obtained when simulating the model for the period 2013Q1-2016Q4 under two alternative expectations formation mechanisms: rational expectations and adaptive learning.

As a preliminary step, we use a Kalman filter to filter out the values of the structural shocks $\varepsilon_{y,t}$, $\varepsilon_{\pi,t}$ and $\varepsilon_{i,t}$ necessary to align the dynamics of the model to the more recent data (for the period 2013Q1 - 2014Q3) and forecasts (for the period 2014Q4 - 2016Q4) available.[14] The values of the model parameters are kept fixed at the posterior mean throughout the exercise.

Then, starting from 2012Q4, we use these shocks to simulate the model under two different assumptions about the rule agents use to forecast future values for inflation and output gap. In the first simulation, agents have rational expectations; in the second, they use an adaptive rule based on all the observables in the model (inflation, the interest

---

[14]We use the Eurosystem staff projections published in the September 2014 issue of the ECB Monthly Bulletin. Due to confidentiality reasons involving the Eurosystem estimates of the output gap, figure 3 reports - for the case of rational expectations - the output gap series taken from the OECD (the output gap under learning showed in the figure is likewise obtained applying the distance from the case of rational expectations obtained in our estimates).
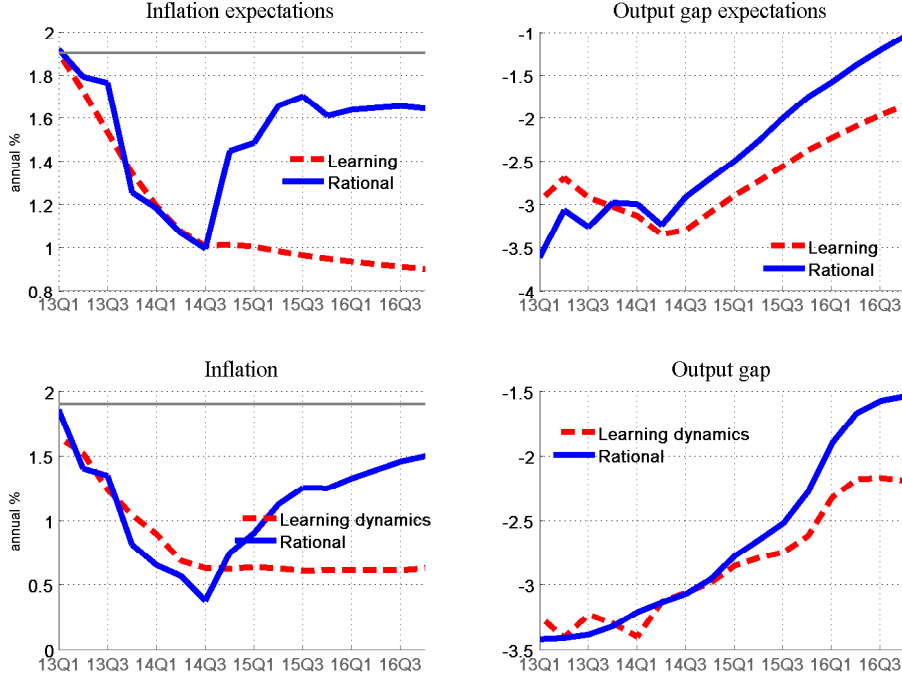
Figure 3: Simulation results under different expectation mechanisms

rate and the gap):

$$\hat{E}_t \left[\begin{array}{c} \pi_{t+1} \\ x_{t+1} \end{array} \right] = \begin{array}{c} \alpha^\pi_{0,t} + \alpha^\pi_{1,t}\pi_{t-1} + \alpha^\pi_{2,t}x_{t-1} + \alpha^\pi_{3,t}i_{t-1} \\ \alpha^x_{0,t} + \alpha^x_{1,t}\pi_{t-1} + \alpha^x_{2,t}x_{t-1} + \alpha^x_{3,t}i_{t-1} \end{array}$$

In this second case, the matrix of coefficients $\alpha_t = [\alpha^\pi_t, \alpha^x_t]$ is initialized in 2012Q4 to values as close as possible to the corresponding entries in the RE vector, while delivering at the same time the same expectations for 2013Q1 as under RE. The gain parameter γ in the learning rule is set to $\frac{1}{40}$, implying that agents use a rolling window of 40 quarters of the most recent data in their adaptive rules.

Figure 3 shows expectations and outcomes in the two cases. Inflation expectations start out in 2013q1 at around 1.9 per cent. As inflation keeps surprising on the downside, under learning agents progressively revise their expectations downward more persistently than in the case of RE: in the average of 2015, expectations under learning are about 0.6 percentage points lower. Once anticipated inflation starts to drift downward, the mechanism becomes self-reinforcing, as actual and expected inflation push each other
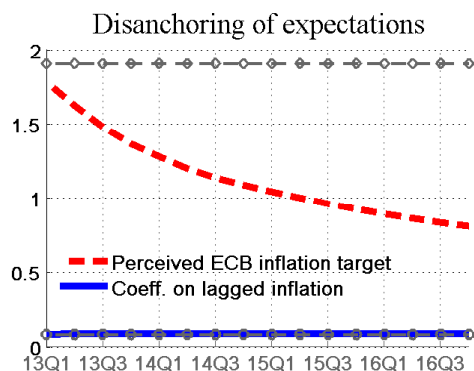
Figure 4: Perceived long-run inflation target and coefficient on inflation inertia under learning.

down, creating a downward price spiral. The simulations under RE and learning diverge markedly: in particular, the latter does not feature the sudden rebound in inflation, both expected and actual, that starts in 2014Q4 in the former, even though the underlying sequence of structural shocks is exactly the same.

With respect to the benchmark case of rational expectations, under learning price dynamics would be lower by about 0.5 percentage points in 2015 and by 0.8 in 2016. Output gap expectations deteriorate as well, although not as much as the ones about inflation. The actual gap under learning is roughly in line with the RE benchmark in 2014, but starts to widen afterwards: at the end of the simulation horizon it is almost a full percentage point below than it would be if agents were fully informed and rational.

The main reason explaining this striking difference in outcomes is that under learning a de-anchoring of inflation expectations from the ECB target occurs. The agents' perceived "long-run equilibrium level" of inflation can be recovered from the coefficients of the learning rule, and is shown in Figure 4. Long-run equilibrium inflation gradually falls to 0.8 per cent by the end of 2016, whereas the ECB target of 1.9 per cent is maintained under the assumptions of rational expectations. The coefficient measuring the inertia in the forecasting rule instead does not change significantly (the blue line in Figure 4 is almost flat), as the way the forecast rule adjusts to track price dynamics is by decreasing the perceived target, which is incorporated in the intercept, as opposed to increasing the weight assigned to lagged inflation.

Overall, the results suggest that the macroeconomic effects of a sequence of negative
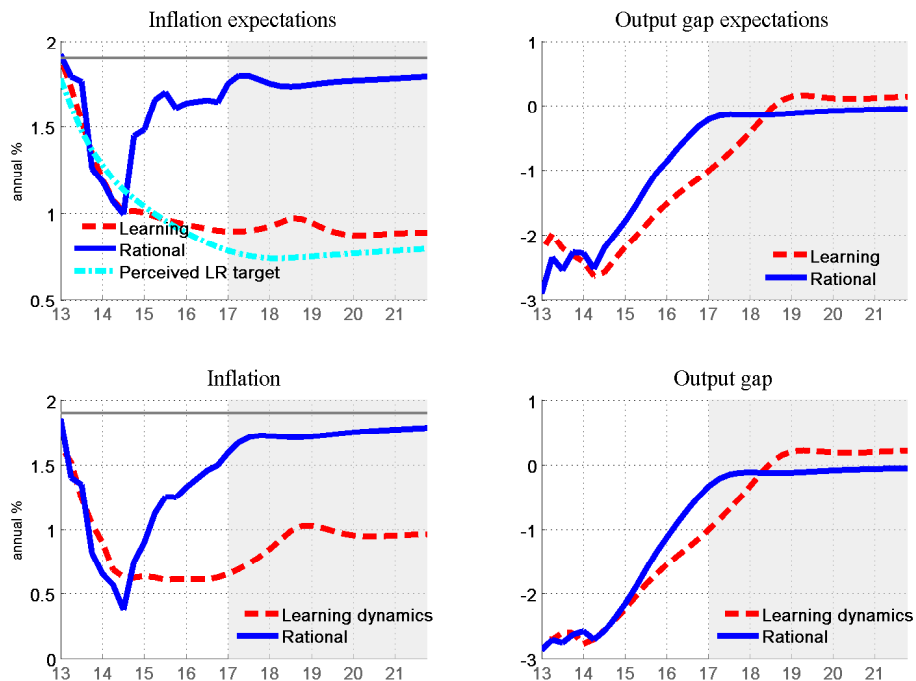
Figure 5: Longer term effects of the dis-anchoring.

inflation surprises can be much higher if the credibility of the long-run inflation target of the central bank is linked to performance. Effects seem particularly persistent since, once this process is set in motion, the lack of a nominal anchor represented by the central bank target means that inflation expectations may go adrift.

To have a feeling of how persistent these effects can be, we run a counter-factual experiment in which we extend the simulations up to 2021, setting all structural shocks in the model to zero from 2017q1 onward. The purpose is to investigate how long it takes for expectations to revert to target absent new deflationary shocks. The results are shown in Figure 5.

Inflation developments are severely affected by the de-anchoring. At the end of 2021, both actual and expected inflation are still below 1.0 per cent; by contrast, under RE price dynamics has completely recovered from the disinflationary episode and appears back in line with the price stability objective of the ECB. Output developments, on the other end, are alike: under both scenarios, output gap expectations rise steadily and level off in positive territory by the end of 2019; the actual gap recovers faster under RE, reaching -0.2 in 2017, while under learning it approaches that level only one year

later, though in the end settles on a higher level.

4.3 Robustness

One important dimension along which these results ought to be checked is robustness to alternative speeds of adjustment of the learning rule. This is encoded in the gain parameter γ. For any given current observation, a lower value for γ means that the learning rule dictates a smaller adjustment in the coefficients of the expectation functions; conversely, a high γ is associated with learning rules that adapt more quickly to a changing environment, as they give relatively more weight to the current observation. The inverse of γ, in fact, represents the size of the rolling window of past observations that are kept in the sample (together with the current observation) when adjusting the parameters in (4). A limiting case is the standard *least-squares* learning algorithm, in which the constant γ is replaced by the factor $\frac{1}{t}$ and no past observations are discarded from the sample.

In order to check whether the results reported in the main text are consistent with a wide variety of possible values of γ, Figure 6 reports actual inflation under different learning simulations, when the size of the rolling window is allowed to vary from a minimum of 16 (four years of data are kept in the sample) to a maximum of 90 observations (roughly the sample length used to estimate the model's structural parameters); the least-squares case is indicated with 'all' in the figure. The dis-anchoring of inflation clearly appears to be a robust feature of any simulation under learning, including the least-squares case; in particular, our benchmark case (the grey dotted line corresponding to $\gamma = \frac{1}{40}$) can be considered a middle case scenario in this respect.

5 Conclusions

The paper offers a model-based assessment of the macroeconomic effects of a sequence of deflationary shocks not confronted with proper policy actions. Besides its theoretical interest, the issue is of practical relevance, as it mirrors the situation the euro area economy might become trapped in. We use a simple New Keynesian model, in which agents do not possess the amount of information and knowledge needed to form expectations rationally but rather behave like econometricians, using regression equations to predict the future value of the variables of interest. The RE solution is used as a benchmark, which helps showing that the likelihood that the economy converges to a low growth-low inflation equilibrium increases sizeably when the assumption of full rationality is
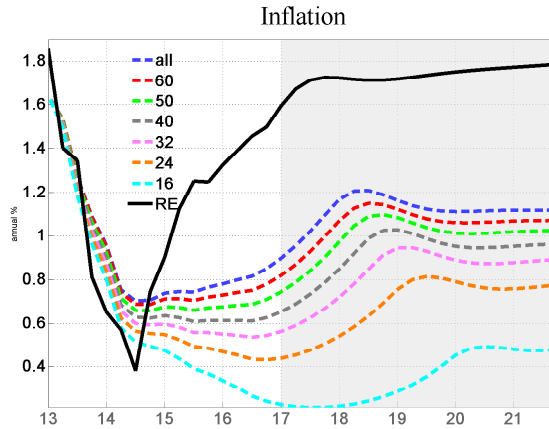
Figure 6: Effect of changing the speed of adjustment in the learning rule.

dropped. The main findings are the followings: (i) actual and expected inflation diverge much more easily under learning than under RE; (ii) provided that the de-anchoring of expectations occurs, inflation may remain low for a protracted period of time; (iii) the cost in terms of GDP of the loss of credibility of the central bank target may be sizable.

Some words of caution are due. Notwithstanding its successes, the theory of econometric learning presents a few shortcomings. The enhancements scored in terms of cognitive consistency are paralleled by the indeterminacy in choosing the specification of the PLMs or in selecting the value of the gain parameter. In both cases we make choices that either do not bias the results of the analysis or go against the thesis we are interested in. Indeed, the PLMs are chosen so as to be identical to the minimum-state-variable solution of the RE equilibrium, implying the least possible deviation from the benchmark, and the value of gamma is selected after a careful sensitivity analysis, to ensure robustness of the results.

References

[1] Benveniste, A., M. Metivier and P. Priouret (1990), Adaptive Algorithms and Stochastic Approximations, Springer-Verlag, Berlin Heidelberg.

[2] Bray, M.M. (1982), Learning, Estimation, and the Stability of Rational Expectations, Journal of Economic Theory.

[3] Bullard, J. (2006), The Learnability Criterion and Monetary Policy, Federal Reserve Bank of Saint Louis Review, May/June.

[4] Bullard, J. and K. Mitra (2007), Determinacy, Learnability, and Monetary Policy Inertia, Journal of Money, Credit and Banking.

[5] Clarida, R. Gali, J. and M. Gertler (1999), The Science of Monetary Policy: A New Keynesian Perspective, Journal of Economic Literature.

[6] Di Cesare, A., Miccoli, M., Neri, S. and M. Pericoli (2014), The impact of inflation surprises on survey and market-based expectations, mimeo.

[7] Draghi, M. (2014), Monetary policy in the euro area, Keynote speech at the Frankfurt European Banking Congress, 21 November 2014, available at: https://www.ecb.europa.eu/press/key/date/2014/html/sp141121.en.html.

[8] European Central Bank (2011), Monthly Bulletin, February 2011.

[9] European Central Bank (2013), Monthly Bulletin, October 2013.

[10] Evans, G.W. (1985), Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models, Quarterly Journal of Economics.

[11] Evans, G. W. and S. Honkapohja (2001), Learning and Expectations in Macroeconomics, Princeton University Press.

[12] Evans, G. W. and S. Honkapohja (2003), Adaptive Learning and Monetary Policy Design, Journal of Money, Credit and Banking.

[13] Evans, G. W. and S. Honkapohja (2008), Expectations, Learning and Monetary Policy: An Overview of Recent Research, Centre for Dynamic Macroeconomic Analysis, Working Paper Series no. 08/02.

[14] Evans, G.W. and B. McGough (2005), Monetary policy, indeterminacy, and learning, Journal of Economic Dynamics and Control.

[15] Ferrero, G. (2007), Monetary policy, learning and the speed of convergence, Journal of Economic Dynamics and Control.

[16] Frydman, R. (1982), Towards an Understanding of Market Processes: Individual Expectations, Learning and Convergence to Rational Expectations Equilibrium, American Economic Review.

[17] Garcia J.A. and T. Werner (2010), "Inflation risk and inflation risk premia", ECB Working Paper, No.1162, March.

[18] Marcet, A. and T.J. Sargent (1989a), Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information, Journal of Political Economy.

[19] Marcet, A. and T.J. Sargent (1989b), Convergence of Least Squares Learning Mechnisms in Self-Referential Linear Stochastic Models, Journal of Economic Theory.

[20] Marcet, A. and T.J. Sargent (1995), Speed of convergence of recursive least squares: learning with autoregressive moving-average perceptions, in Kirman, A. and M. Salmon (eds), Learning and Rationality in Economics.

[21] Orphanides, A. and J.C. Williams (2004), Imperfect Knowledge, Inflation Expectations, and Monetary Policy, in The Inflation-Targeting Debate, edited by B.S. Bernanke and M. Woodford, University of Chicago Press.

[22] Pericoli, M. (2012), Expected inflation and inflation risk premium in the euro area and in the United States, Bank of Italy Discussion Papers, No. 842.

APPENDIX:
Determinacy and stability under learning

A model with adaptive learning has two main ingredients: (i) an equation describing agents' beliefs on the dynamics of economic variables and (ii) a temporary equilibrium of the system generated by the interaction between expectations and the structure of the economy.

Let us assume that the economy is described by the following linear system:

$$
\begin{aligned}
y_t &= A_1 y_{t-1} + A_2 E_{t-1} y_{t+1} + B x_{t-1} + u_t \\
x_t &= F x_{t-1} + v_t
\end{aligned}
\tag{5}
$$

where y_t and x_t are the vectors of endogenous and exogenous variables respectively. Agents' beliefs are described by a forecasting model, the so-called perceived law of motion (PLM), which usually has the same functional form as the (minimum state variable) solution of the RE equilibrium:

$$
\widehat{E}_{t-1} y_t = a_{1,t-1} y_{t-1} + b_{t-1} x_{t-1}
\tag{6}
$$

The operator $\widehat{E}$ refers to subjective beliefs, which may vary across individuals, and does not coincide with conditional expectations. The coefficients of the forecasting model are re-estimated in every period by recursive least squares (RLS).[15] The learning process is described by the following set of recursive equations:

$$
\begin{aligned}
\theta_t &= \theta_{t-1} + \gamma_t R_t^{-1} z_{t-1} \left(y_t - \theta_{t-1}^T z_{t-1} \right) \\
R_t &= R_{t-1} + \gamma_t \left(z_{t-1} z_{t-1}^T - R_{t-1} \right)
\end{aligned}
\tag{7}
$$

where the gain sequence $\{\gamma_t\}_{t=k}^{\infty}$ is equal to $\{t^{-1}\}_{t=k}^{\infty}$, and $\theta_t = (a_{1,t}, b_t)^T$ and $z_t = \left(y_t^T, x_t^T \right)^T$. Given the forecasts, the economy attains a temporary equilibrium, the so-called actual law of motion (ALM), which is equal to:

$$
\begin{aligned}
y_t &= A_1 y_{t-1} + A_2 \left(a_{1,t-1} \widehat{E}_{t-1} y_{t+1} + b_{t-1} \widehat{E}_{t-1} x_t \right) + B x_{t-1} + u_t \\
&= \left(A_1 + A_2 a_{1,t-1}^2 \right) y_{t-1} + \left(B + A_2 a_{1,t-1} b_{t-1} + A_2 b_{t-1} F \right) x_{t-1} + u_t \\
&= T \left(\theta_{t-1} \right)^T z_{t-1} + u_t
\end{aligned}
\tag{8}
$$

---

[15]Econometric learning can be alternatively modeled using a (generalized) stochastic gradient (SG) updating rule. Equation (7) modifies to $\theta_t = \theta_{t-1} + \frac{1}{t}\Gamma z_{t-1}\left(y_t - \theta_{t-1}^T z_{t-1}\right)$, where the time-invariant matrix Γ is usually set equal to either the identity matrix or $\mathrm{E}\left(z_t z_t^T\right)$.

where $T(\theta_{t-1})^T = T(a_{1,t-1}, b_{t-1}) = \left(A_1 + A_2 a_{1,t-1}^2, B + A_2 a_{1,t-1} b_{t-1} + A_2 b_{t-1} F\right)$ is the mapping that describes the evolution of the RLS estimator θ_t. Once the ALM is substituted for y_t in (7), the dynamics of the system is fully described by the recursive least squares equations:

$$
\begin{aligned}
\theta_t &= \theta_{t-1} + \tfrac{1}{t} R_t^{-1} z_{t-1} \left(T(\theta_{t-1}) + u_t - \theta_{t-1}\right) \\
R_t &= R_{t-1} + \tfrac{1}{t} \left(z_{t-1} z_{t-1}^T - R_{t-1}\right)
\end{aligned}
\tag{9}
$$

With the shift $S_{t-1} = R_t$, (9) becomes a stochastic recursive algorithm (SRA), whose behavior is well approximated by an ordinary differential equation (ODE)

$$
\tfrac{d\phi}{d\tau} = h(\phi) = T(\phi) - \phi
\tag{10}
$$

where $\phi_t \equiv vec(\theta_t, S_t)$ and $h(\phi)$ is obtained by computing the asymptotic limit of the expectation of the 2^{nd} term (the updating function) on the right-hand side of (9): the zeros of the ODE represent the only possible limit points of the SRA and the corresponding equilibria are stable if the (real part of the) eigenvalues of the Jacobian of $h(\phi)$ are negative. When $(a_{1,t}, b_t) \rightarrow (A_1, B)$, i.e. when the PLM comes to coincide with the ALM, an RE equilibrium is attained and agents have learned the rational expectations equilibrium.

In addition to a few regularity conditions on the exogenous processes and on the updating function, asymptotic convergence of the learning algorithm depends on the (positive and non-stochastic) gain sequence $\{\gamma_t\}_{t=k}^{\infty}$ being such that $\sum\limits_{t=k}^{\infty} \gamma_t = \infty$ and $\sum\limits_{t=k}^{\infty} \gamma_t^2 < \infty$. These assumptions are necessary to avoid convergence of ϕ_t to a non-equilibrium point and to ensure the asymptotic elimination of all residual fluctuations.

Some authors have departed from this framework, studying the implications for the equilibrium outcomes of a constant gain sequence, i.e. of setting $\gamma_t = \gamma$ for all values of t. Constant gain learning precludes the convergence to the RE equilibrium: as long as the solution is stable, agents' expectations are correct on average but keeps on fluctuating around rather than at the equilibrium. This happens because observations are not assigned equal weight: those far in the past are discounted at an exponential rate, so that information does not accumulate fast enough to completely remove the randomness in the data.[16] A constant γ is justified when agents suspect that the economy is under-

---

[16]Orphanides and Williams refer to constant gain learning as *perpetual learning*, to stress the fact that full information about the structure of the economy is never achieved.

going structural changes. Although in principle they might attempt to model structural change, this would call for an amount of knowledge comparable to that needed for RE; a reasonable alternative is to recognize, in adjusting the parameter estimates, that the more recent observations convey more accurate and less distorted information on the economy's laws of motion, which is the idea captured by constant gain learning.