DISCUSSION PAPER SERIES

DP14265

THE 3 E'S OF CENTRAL BANK COMMUNICATION WITH THE PUBLIC

Andrew Haldane, Alistair Macaulay and Michael McMahon

MONETARY ECONOMICS AND FLUCTUATIONS

ISSN 0265-8003

THE 3 E'S OF CENTRAL BANK COMMUNICATION WITH THE PUBLIC

Andrew Haldane, Alistair Macaulay and Michael McMahon

Discussion Paper DP14265
Published 01 January 2020
Submitted 08 November 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

THE 3 E'S OF CENTRAL BANK COMMUNICATION WITH THE PUBLIC

Abstract

In this paper we explore both theoretical and empirical evidence on communication with the general public. The model provides guidance for policy makers by high- lighting some potentially important risks in communicating simply with a broader audience. In particular, in a model where trust and engagement are low, there are benefits to engaging a wider audience. But doing so risks ultimately lowering welfare unless guided by the 3 E's of public communication: Explanation, Engagement and Education. Central banks have made great strides in all three, but numerous challenges remain.

JEL Classification: E52, E58

Keywords: monetary policy, communication, General Public

Andrew Haldane - andy.haldane@bankofengland.co.uk
Bank of England

Alistair Macaulay - alistair.macaulay@nuffield.ox.ac.uk
University of Oxford

Michael McMahon - michael.mcmahon@economics.ox.ac.uk
University of Oxford and CEPR

Acknowledgements
Paper prepared for the Central Bank of Chile's XXIII Annual Conference, to be held on July 22nd and 23rd, 2019 in Santiago, Chile. This paper is also part of the CEPR Research Policy Network on Central Bank Communication. We are particularly grateful to Kathryn Dominguez for her discussion of the first version of the paper. We have also benefited from comments, discussion and suggestions from Ben Broadbent, David Bradnum, Shiv Chowla, Michael Ehrmann, Al Firrell, Hope Gray, Refet Gürkaynak, Stephen Hansen, Andrew Hebden, David Nicholls, Chris Peacock and Emma Rockall, as well as numerous others for useful discussions about the issue in general at central bank communication conferences over the last 12 months especially the Central Bank of Chile's XXIII Annual Conference and the Annual Research Conference of the National Bank of Ukraine in May 2019. Nonetheless, the views expressed here do not necessarily represent the views of the Bank of England, the Monetary Policy Committee, Financial Policy Committee, the Irish Fiscal Advisory Council, or anyone else other than the authors. Any errors remain ours alone.

The 3 E's of Central Bank Communication with the Public*

Andrew Haldane† Alistair Macaulay‡ Michael McMahon§

December 2019

Abstract

In this paper we explore both theoretical and empirical evidence on communication with the general public. The model provides guidance for policy makers by highlighting some potentially important risks in communicating simply with a broader audience. In particular, in a model where trust and engagement are low, there are benefits to engaging a wider audience. But doing so risks ultimately lowering welfare unless guided by the 3 E's of public communication: Explanation, Engagement and Education. Central banks have made great strides in all three, but numerous challenges remain.

Keywords: Monetary Policy, Communication, General Public

JEL Codes: E52, E58

*Paper prepared for the Central Bank of Chile's XXIII Annual Conference, to be held on July 22nd and 23rd, 2019 in Santiago, Chile. This paper is also part of the CEPR Research Policy Network (RPN) on Central Bank Communication. We are particularly grateful to Kathryn Dominguez for her discussion of the first version of the paper. We have also benefited from comments, discussion and suggestions from Ben Broadbent, David Bradnum, Shiv Chowla, Michael Ehrmann, Al Firrell, Hope Gray, Refet Gürkaynak, Stephen Hansen, Andrew Hebden, David Nicholls, Chris Peacock and Emma Rockall, as well as numerous others for useful discussions about the issue in general at central bank communication conferences over the last 12 months especially the Central Bank of Chile's XXIII Annual Conference and the Annual Research Conference of the National Bank of Ukraine in May 2019. Nonetheless, the views expressed here do not necessarily represent the views of the Bank of England, the Monetary Policy Committee, Financial Policy Committee, the Irish Fiscal Advisory Council, or anyone else other than the authors. Any errors remain ours alone.

†Haldane: Bank of England, Threadneedle Street, London, EC2R 8AH, UK, andy.haldane@bankofengland.co.uk.

‡Macaulay: University of Oxford, Manor Road Building, Oxford, OX1 3UQ, UK, alistair.macaulay@nuffield.ox.ac.uk.

§McMahon: University of Oxford, Manor Road Building, Oxford, OX1 3UQ, UK, michael.mcmahon@economics.ox.ac.uk.

1 Communications Revolution

Central banks used to ask "Do we communicate this?" Now, as a rule, they ask "Why wouldn't we communicate this?" (Skingsley 2019). This first wave of the revolution in central bank communication is giving rise to a second wave; the question increasingly is "How should we communicate this in a way that engages a broader cross-section of society?" This addresses the challenge laid out by Blinder et al. (2008) that "It may be time to pay some attention to communication with the general public."

These two waves of the communication revolution followed naturally from the growing understanding of the central role that management of expectations plays in economic management (Blinder 2009, Woodford 2001) and the potentially important role that central bank communication has on expectations. But much remains to be understood, especially concerning the second wave.[1] In Haldane and McMahon (2018), we addressed issues of feasibility and desirabilty of communication with the general public. This paper explores, using a simple theoretical framework and supporting empirical analysis, some of the concerns that have been raised about such broader communication and especially the potential that these efforts may do more harm than good. We set out a simple framework of complementary activities – the 3 E's of Explanation, Engagement and Education – that may help the central bank to avoid the potential pitfalls.

As we show in section 2, the evidence suggests that many households may never engage with central bank communication because it is written in a way that they cannot understand. This contributes to a lack of trust in the central bank as an independent institution. These twin deficits (of understanding and trust) impinge on the efficacy of monetary policy and, potentially, limit the ability of operationally-independent central banks to meet the terms of their social contract to serve the whole population as well as possible. It is this realisation that has sparked the second stage of the revolution: shifting from the traditional audience for central bank communication (financial market participants and journalists) communicated via complex, carefully-crafted reports, speeches and statements, toward directly communicating with a broader audience of the general public.

While acknowledging the evidence on twin deficits of understanding and trust, and that broader engagement is important for democratic and political economy reasons, some economists have expressed concerns about this new focus for communication. In particular, there is concern that the economy is complex and, as a result, monetary policy is not simple but if communication is too simplified then people may develop a false sense of certainty about the central banks views about the economy. Delivering simple messages

[1] Important papers in this literature on communicating with the wider public include Kryvtsov and Petersen (2013), Binder (2017), Braun (2018), Bholat et al. (2018) and Coibion et al. (2019).

could ultimately lead the public to be disappointed if the central bank doesn't deliver on its communicated forecasts.

In order to explore this idea, we develop a framework in section 3 in the spirit of the rational inattention literature but include three important dimensions: (1) A second form of communication that is easier to read but that comes with the cost that the household misses the uncertainty around those forecasts. This means that when the world doesn't turn out exactly as the central bank predicted, households are surprised. (2) We change the structure of costs for different households reading the central bank communication. The household will no longer add idiosyncratic noise to the signal but rather some will, because of too high costs, simply choose to be uninformed. (3) We introduce a reduced-form concept of trust into the model. We assume that this trust evolves dynamically in the model, rising when the central bank engages the public but falling when the public are surprised by the outcomes in the economy. The cost of reading the central bank material is linked to the household's level of trust.

While clearly designed to emphasise the potential concerns about simplified communication, the model has both an optimistic and a cautionary message. Simplified communication can increase the proportion of the population paying attention to central bank messages which also builds trust and, as a result, increase welfare. However, this is a transitory state without further intervention. Trust ultimately falls when the household observes that reality did not exactly match the communicated signal; the net effect overall is that, in expectation, the trust of each household who pays attention to the simplified content for at least one period will be lower in the new steady state than in period 0, before the introduction of the new communication. Without intervention, welfare would be lowered in the kind of environment people have been concerned about.

While simplified communication alone is not enough, in such an environment central banks can take action to influence the speed of transition to the lower welfare steady state and thereby can extend the time during which welfare is boosted. We explore such complementary activities in section 4 under a framework of 3 E's of public communication: Explanation, Engagement and Education. These three pillars are clearly linked – more education increases the chances of engagement and makes explanation easier. Central banks have made great strides in all three in addressing the twin deficits. Our work suggests that these related endeavours may not simply be 'nice-to-haves'. Rather they may be 'need-to-haves' if central banks are to reach the people currently by-passed by central bank communication, maintain this reach and build durable levels of trust.

Central bank communication on monetary policy addresses both high frequency issues (such as current economic conditions and monetary policy decisions), as well as low frequency ones (such as the framework for monetary policy). The main focus of this

paper is adapting high frequency to be suitable for a wider audience is the most novel part of the recent push to communicate with the general public. Communication with a broader audience is at the heart of the inflation targeting framework; the target is itself a low frequency communication medium. But in a world of low trust, even low freqeuncy communication may need to be significantly enhanced. We conclude the paper with a brief discussion of the overlap of the 3 E's regarding low frequency communication as well as highlight some of the current efforts and challenges around them (section 5).

Will existing efforts to central banks' outreach, engagement and education be successful? Will the new approaches deliver significant penetration into previously unengaged parts of the population? The jury is still out; Blinder (2018) is pessimistic and believes that central banks are likely to continue to fail to land their messages with the general public. But given that this second wave of the central banking communication revolution is unlikely to disappear anytime soon, further research into this issue is a must. This should include continued assessment of the outcomes of new approaches, as well as providing suggestions to improve results with novel approaches or refinements to existing attempts.

2 Central Bank Communication and Trust

Does it make economic sense to have high-frequency communication with the general public? It does to the extent that expectations are important for economic dynamics (as in the New Keynesian model in Galí (2015)) and communication can aid expectations management (Blinder 2009, Woodford 2001). In standard economic environments, therefore, central banks wishing to control inflation can benefit from using communication to share any private information and influence inflation expectations.

In Haldane et al. (2019), we explore this question in a model in which agents do not become fully informed but rather are rationally inattentive. The main finding is that central banks should provide as much detail as possible even though some households will optimally pay little attention ('skim read') to the signals.[2] This finding is, essentially, the approach taken by most inflation-targeting central banks. They regularly release a large amount of highly-detailed information. Statements of policy decisions, Inflation Reports, minutes of meetings, speeches, forecast information, information on the models used, etc... are typically all available on central bank websites for anyone to read.

[2]Since households in this framework choose optimally how much attention to pay to signals about the shocks, and the central bank can vary the precision of its signal (more precise signals are more costly to process), the central bank optimally chooses how precisely to communicate in order minimise welfare losses. Making the signals easier to read involves making them less precise, but any such public noise is common to all households and so households co-ordinate on it leading to inefficient fluctuations in consumption. And so more central bank precision is optimal.

But the issue is not that households skim read the material. Rather most households do not read it at all. For many, in fact, they can not read it. As discussed in Haldane (2017), Coenen et al. (2017) and Haldane and McMahon (2018), the main central bank publications in many advanced economies including the UK and US had a reading grade level of between 14-18 according to the Flesch-Kincaid reading grade score. This is roughly equivalent to college-level and is, on the basis of the population distribution of literacy across the population, inaccessible to about 90% of the general public. The majority of people presumably do not even attempt to engage with the material in, for example, an Inflation Report. (Speeches by politicians, by contrast, are much simpler – around grade 8 level – and so accessbile to up to half the population.)

In such an environment, is it any surprise that many households have little understanding of monetary policy or the institutions that set policy? But it is not just a deficit of understanding that has concerned central banks recently. It is the fact that this deficit of understanding typically goes hand in hand with a deficit of trust in the institution (Haldane 2017).[3] This twin deficit is evident in responses to the Bank of England's Inflation Attitudes Survey; this is a survey of around 2,000 individuals conducted since 2001 (see also Jost (2017) and Rockall (2018)). To construct an index of monetary policy knowledge among the general public (hereafter called the "knowledge index"), we use responses to three questions about the institutional structure of monetary policy from the survey:

- Q11: Which group of people set Britain's basic interest rate level?
- Q12: Which of these groups do you think sets the interest rates?
- Q13: Which of these do you think best describes the Monetary Policy Committee?

For each question, respondents getting the correct answer adds +2 to the knowledge score, admitting you don't know yields +1 and getting it wrong yields 0. This index runs from a score of six ("perfect knowledge") through three ("admitted no knowledge") to zero ("Gets every answer wrong").

The top panel of Table 1 show the mean overall knowledge score in the UK survey over the past 17 years. At best, this has flat-lined despite the increase in communication by the Bank of England over the period suggesting that the public's understanding of monetary policy structures appears to have been largely immune to central banks' communication revolution. But the aggregate evolution masks significant stratification in knowledge scores by age, education and social class (as well as by income), with the young, less well-educated and poor being materially less knowledgeable. For example, those in social class AB (upper middle and middle class) have an index score 36 percentage points higher than those in grade DE (working and non-working class). This suggests that

[3]See also Braun (2016) who also discusses the issue of trust in communication with the general public.

central banks' current communications initiatives are by-passing large cohorts of society. The communications revolution has been selective (Haldane and McMahon 2018).

Using the survey answer to Q14, which asks "Overall, how satisfied or dissatisfied are you with the way the Bank of England is doing its job to set interest rates in order to control inflation?", we construct a measure of satisfaction with central banks' actions. This serves as a proxy for trust and runs from 5 (most satisfied / highest trust) to 1 (unsatisfied / lowest trust). The lower panel in Table 1 shows the mean of satisfaction/trust proxy score. As with other trust measures from other surveys, this declined during and following the financial crisis and has yet to fully recover. This pattern in satisfaction/trust scores in central banks' actions has been broadly-based across demographic groups and across countries.

Of course, one concern is that the measure satisfaction is not a good proxy for trust. We check this using the survey for 2017 when there was also a question about credibility; the first part of Q27 asks respondents to what extent they agree that the Bank of England is credible. In 2017 when we have both concepts measured, there is statistically significant positive correlation (0.46) between the credibility score and the trust proxy. Column (1) of table 2 shows that this correlation survives the inclusion of numerous demographic controls. Column (2) adds the institutional knowledge and economic knowledge scores too; the former is also an important correlate. Columns (3)-(5) instead focus on the correlates driving the trust proxy with (3) showing the reversed regression from (2), (4) shows that even excluding the Credibility measure, institutional knowledge is a signficant correlate in 2017, and (5) shows that this relationship in 2017 is very similar to the relationship acorss the whole sample (for which the credibility score is not available).

As argued in Haldane and McMahon (2018), one of the reasons that a central bank may want to communicate more directly with the general public is to try to build public understanding as a means of establishing trust and credibility about central banks and their policies. But why, apart from professional pride, should a central bank care whether people trust it? Shouldn't it simply get on with its job of setting the best interest rate which will, sometimes, involve difficult decisions? First, this is important for reasons of political accountability, ensuring operationally independent central banks are meeting the terms of their social contract with wider society.

Another reason to try to build trust is that trust helps manage expectations. The data in the UK is consistent with trust being an important driver of expected inflation (see analysis below which expands on the analysis in Haldane and McMahon (2018) and Rockall (2018)). There is growing evidence that inflation expectations affect economic choices made by households. This evidence includes effects on major purchase decisions and financial choices. Bachmann et al. (2015) show that higher expected inflation

Table 1: Knowlege of and Satisfaction with the Central Bank

	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Knowledge Score																			
Mean	4.2	4.2	4.2	4.2	4.3	4.1	4.2	4.2	4.1	4.2	4.2	4.1	4.0	4.0	4.0	3.9	4.1	4.0	4.0
AB Class	5.0	4.9	5.0	4.9	5.0	4.9	4.9	4.8	4.8	4.8	4.9	4.7	4.7	4.7	4.7	4.6	4.7	4.7	4.6
DE Class	3.6	3.6	3.6	3.8	3.7	3.5	3.6	3.6	3.6	3.6	3.6	3.6	3.4	3.4	3.3	3.3	3.5	3.5	3.4
University Degree	4.8	4.8	4.8	4.8	4.9	4.7	4.6	4.7	4.7	4.8	4.8	4.5	4.6	4.5	4.5	4.4	4.6	4.5	4.4
No Formal Education	4.0	4.1	4.0	4.2	4.1	4.0	4.1	3.7	3.7	3.8	3.7	3.7	3.5	3.4	3.5	3.4	3.6	3.5	3.6
Age: 15-24	3.2	3.4	3.4	3.3	3.3	3.2	3.2	3.4	3.3	3.2	3.1	3.1	3.1	3.1	3.1	3.1	3.2	3.1	3.3
Age: 55-64	4.6	4.5	4.6	4.7	4.7	4.7	4.8	4.6	4.5	4.7	4.6	4.6	4.4	4.5	4.3	4.4	4.5	4.6	4.6
Trust Proxy																			
Mean	3.5	3.7	3.6	3.6	3.6	3.6	3.5	3.4	3.0	3.3	3.1	3.2	3.1	3.3	3.4	3.4	3.4	3.4	3.3
AB Class	3.8	3.9	3.7	3.7	3.8	3.8	3.7	3.6	3.2	3.5	3.4	3.5	3.4	3.5	3.6	3.5	3.6	3.6	3.5
DE Class	3.4	3.5	3.4	3.4	3.5	3.5	3.3	3.2	2.9	3.1	2.9	3.0	2.9	3.1	3.3	3.2	3.2	3.2	3.1
University Degree	3.8	3.9	3.7	3.7	3.7	3.7	3.6	3.6	3.2	3.5	3.4	3.4	3.3	3.5	3.6	3.4	3.5	3.5	3.4
No Formal Education	3.5	3.7	3.6	3.5	3.6	3.6	3.4	3.3	2.8	3.2	3.0	3.0	2.9	3.2	3.4	3.3	3.3	3.3	3.2
Age: 15-24	3.3	3.4	3.4	3.2	3.4	3.4	3.2	3.2	3.0	3.1	3.0	3.0	3.1	3.2	3.3	3.2	3.3	3.2	3.2
Age: 55-64	3.7	3.8	3.7	3.7	3.7	3.7	3.6	3.5	2.9	3.4	3.2	3.2	3.2	3.4	3.5	3.5	3.5	3.5	3.4

Notes: Bank of England Attitudes to Inflation Survey. The knowledge score (upper panel) is between 0 and 6 where 6 indicates perfect knowledge of the institutions of monetary policy. The trust proxy score (lower panel) is between 5 (most satisfied / highest trust) and 1 (unsatisfied / lowest trust).

Table 2: Regression analysis of inflation attitudes survey

	(1) Credibility	(2) Credibility	(3) Trust Proxy	(4) Trust Proxy	(5) Trust Proxy
Trust Proxy	0.39***	0.37***			
	[0.00]	[0.00]			
Knowledge		0.052***	0.062***	0.10***	0.12***
		[0.00]	[0.00]	[0.00]	[0.00]
Econ Knowledge		-0.025	0.020	0.011	0.0070
		[0.18]	[0.33]	[0.62]	[0.15]
Credibility			0.49***		
			[0.00]		
Constant	2.19***	2.10***	1.20***	2.72***	2.77***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Observations	3,382	3,382	3,382	3,597	65,905
R-squared	0.272	0.280	0.266	0.102	0.087
Estimation	OLS	OLS	OLS	OLS	OLS
Demographic Controls	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	No	No
Sample	2017	2017	2017	2017	2001-2019

Notes: *Trust Proxy* measures respondent satisfaction with how the Bank is carrying out monetary policy to control inflation, *Knowledge* is their score in terms of understanding the institutions setting monetary policy, and *Econ Knowledge* is their score in terms of understanding of how monetary policy affects the economy. P-values constructed using robust standard errors are reported in brackets below the coefficient estimates. Demographic controls for gender, age, income, class, working status, housing tenure, education, and region are included.

slightly increases US consumers' readiness to spend on durables in normal times while Duca et al. (2018) find a similar effect in the euro-area consumers but the increase in likelihood of making a make major purchase is particularly strong at the effective lower bound. Malmendier and Nagel (2016) show household measures of inflation expectations explain household financial decisions such as whether to have a fixed or floating rate mortgage. Armantier et al. (2015) show consumer inflation expectations are correlated with their experiment-based investment choices, but also that those participants whose behaviour is not consistent with economic theory have lower education and economic literacy. Vellekoop and Wiederholt (2018) show that higher inflation expectations lead households to accumulate less net worth driven by both lower asset holdings (such as savings account, bonds, and stocks) and also lower liabilities.

Table 3 shows the relationship between our trust proxy and absolute values of deviations of household inflation expectations from the inflation target. There are two columns each for 1-year-ahead (columns (1)-(2)), 2-year-ahead (columns (3)-(4)) and 5-year-ahead inflation expectations (columns (5)-(6)). In these regressions, we control for the measures of both institutional knowledge and knowledge of the transmission mechanism, as well as time fixed-effects and various demographic factors (gender, age, income, class, working status, housing tenure, education, and region). Lower trust is associated with inflation expectations that are further from the inflation target. Moreover, including quadratic terms suggests that these deviations grow as trust falls. This suggests that the gains to building trust, as measured by the degree of anchoring of inflation expectation, will be largest if the central bank targets those with the lowest starting levels of trust.

Table 3: Effect of Trust on Inflation Expectations

	(1) $\lvert\mathbb{E}_t[\pi_{t+1}] - \pi*\rvert$	(2) $\lvert\mathbb{E}_t[\pi_{t+1}] - \pi*\rvert$	(3) $\lvert\mathbb{E}_t[\pi_{t+2}] - \pi*\rvert$	(4) $\lvert\mathbb{E}_t[\pi_{t+2}] - \pi*\rvert$	(5) $\lvert\mathbb{E}_t[\pi_{t+5}] - \pi*\rvert$	(6) $\lvert\mathbb{E}_t[\pi_{t+5}] - \pi*\rvert$
Trust Proxy	-0.17***	-0.79***	-0.23***	-0.95***	-0.23***	-0.98***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Knowledge	-0.034***	-0.041*	-0.065***	-0.097***	-0.064***	-0.055
	[0.00]	[0.08]	[0.00]	[0.01]	[0.00]	[0.19]
Econ Knowledge		0.011		0.012		0.054
		[0.67]		[0.80]		[0.32]
Trust Proxy^2		0.10***		0.12***		0.13***
		[0.00]		[0.00]		[0.00]
Knowledge^2		0.00076		0.0042		-0.0014
		[0.80]		[0.36]		[0.79]
Econ Knowledge^2		-0.0084		-0.013		-0.019
		[0.14]		[0.19]		[0.11]
Constant	1.97***	2.80***	2.97***	3.88***	3.26***	4.13***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Observations	58,150	58,150	29,139	29,139	25,870	25,870
R-squared	0.093	0.098	0.064	0.070	0.041	0.046
Estimation	OLS	OLS	OLS	OLS	OLS	OLS
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample	2001-2019	2001-2019	2009-2019	2009-2019	2009-2019	2009-2019

Notes: *Trust Proxy* measures respondent satisfaction with how the Bank is carrying out monetary policy to control inflation, *Knowledge* is their score in terms of understanding the institutions setting monetary policy, *Econ Knowledge* is their score in terms of understanding of how monetary policy affects the economy, π_t Perception measures their perception of current inflation, and $\mathbb{E}_t[\pi_{t+h}]$ is the respondent's expectation for h-years ahead inflation where $h = 1, 2\&5$. P-values constructed using robust standard errors are reported in brackets below the coefficient estimates. Demographic controls for gender, age, income, class, working status, housing tenure, education, and region are included.

3 A Model of Simple Communication and Trust

The previous section might make the use of simplified communication by central banks seem obvious. By communicating in a way that a broader cross section of society can understand, the central might be able to better anchor inflation expectations as well as, potentially, build trust.

But some economists have expressed concern that simplified communication might be too simple. The worry is that rather than boost trust and engagement, households will fail to understand the complex and stochastic nature of the environment which policymakers operate in. Over time, central banks might simply disappoint these newly-engaged households when they miss a target or have a large forecast error. The effect on household trust and engagement may not be so desirable over time.

To explore this, we build on the empirical evidence of the last section and try to develop a framework which incorporates a role for engagement and trust. But we also want to take seriously the concerns that have been raised. The framework, which assumes that some of the concerns will play out, enables us to ask to what extent, or under which circumstances, introducing more accessible communication could help address the twin deficits.

Of course, there are reasons that the concerns may not be correct. For instance, we will assume that simplicity of the message reduces its communicated uncertainty. It is not clear whether this is true or not. The model will also only allow engagement with central bank communication to increase trust whereas in reality there may be many different ways of building trust. And households may get their information on the central bank from other sources. We return to some of these issues in section 4.

3.1 Our Model Environment

The basic model environment is the simple, three-equation New Keynesian model. In order to have a role for communication with the public, we alter the informational assumptions. Specifically, we assume that, as in the textbook model, firms observe current shock realisations but, unlike the textbook model, households observe shocks only after a one-period lag. This can be thought of as the firms being 'close to the ground' and so seeing shocks to technology and costs first-hand, but households having to hear about the shocks after they have hit. Households can, however, learn about contemporaneous shocks from reading central bank communications.[4] In the equations below, $\mathbb{E}_t^F x$ is the

[4]Different households will receive idiosyncratic signals in our model. While the link from heterogeneous information to heterogeneous wealth is potentially interesting, it is beyond the scope of this paper. We therefore simplify by assuming that all households belong to a large family, which redistributes wealth among households at the end of each period.

expectation of x held by a fully-informed agent (who observes current shock realisations) in period t, and $\mathbb{E}_t^H x$ is the expectation of x held by households in period t.

Define c_t^* as the consumption that would be chosen by a household who observed the current realisations of all exogenous shocks. The Euler equation of a fully-informed household is:

$$c_t^* = \mathbb{E}_t^F c_{t+1}^* - \frac{1}{\sigma}(i_t - \mathbb{E}_t^F \pi_{t+1}) \tag{1}$$

Uninformed households maximise their expected utility by setting consumption at the level they expect a fully-informed household would choose, $c_t = \mathbb{E}_t^H c_t^*$. If, as in the standard NK model, households do observe the current realisations of exogenous shocks, then $c_t = c_t^*$ and this model collapses to the textbook three-equation NK model, but without real-time observation of shocks this is no longer the case.

So that our model is comparable to much of the existing literature, we want to ensure that our model admits the same New Keynesian Phillips Curve as in Galí (2015). The derivation of this requires that households are always on their product demand and labour supply curves, and so we have to assume that households directly observe relative prices and the real wage when making their consumption and labour decisions. This comes at the cost of assuming that agents cannot back out the shocks from these observations. To this end, we simplify by assuming that households observe wages and relative prices in the current period but they only observe the nominal interest rate with a lag, and they are unable to infer from wages and relative prices what the shocks and interest rate must be. This simplification keeps the analytic model tractable as it allows us to focus on i.i.d. shocks.[5]

$$\pi_t = \mathbb{E}_t^F \pi_{t+1} + \kappa \tilde{y}_t + v_t \tag{2}$$

The central bank follows a Taylor Rule:

$$i_t = \phi_\pi \pi_t \tag{3}$$

To complete the model there is a market clearing condition relating the output gap to aggregate consumption c_t.

$$\tilde{y}_t = c_t - y_t^n = \mathbb{E}_t^H c_t^* - \frac{1+\varphi}{\sigma+\varphi} a_t \tag{4}$$

There are two exogenous shocks: a technology shock a_t, and a cost-push shock v_t.

[5]With i.i.d. shocks, the nominal interest rate is the only way shocks can affect consumption in the household Euler equation. If households could observe the interest rate, they would therefore have no need of further information about the shocks and central bank communication would be irrelevant.

Both are assumed to be drawn from i.i.d. normal distributions:

$$a_t \sim N(0, \sigma_a^2), \quad v_t \sim N(0, \sigma_v^2) \tag{5}$$

3.2 Expectations, Central Bank Signals and Attention

As in Sims (2003), and the rest of the Rational Inattention literature, households form their expectations about the consumption they should be choosing $\mathbb{E}_t^H c_t^*$ by paying attention to signals about shocks. In this paper, we examine the role of central bank communication as the source of these signals. From the signals that they extract from the communication, the households will form expectations about current shock realisations, which they will then map to expected fully-informed consumption c_t^*.

As described above, Haldane et al. (2019), using an information environment that is similar to the typical RI environment, we show that when the central bank provides independent signals about the two shocks, welfare losses from the volatility of inflation and the output gap are minimised when the central bank communicates as much information as possible (which means that their signals contain as little noise as possible). This is because, being common to all households, any noise introduced by the central bank communication causes households to co-ordinate on this central bank noise leading to inefficient fluctuations in consumption. When households choose to pay less attention to signals, it introduces noise into expectations which is household-specific and so, unlike the central bank noise, it cancels out in aggregate.

In this paper, to allow us to consider the effects of the twin deficits of trust and understanding, as well as being more explicit in analysing the effects of introducing an alternative medium of communication, we make three changes to the communication and information setup:

1. Households face a household-specific fixed cost of processing central bank communication (μ_h).
2. There is a medium-specific cost of reading information which reflects the complexity of the medium (F_{medium}). This allows us to explore the effects of having a second form of central bank communication that is easier to read (lower processing cost).
3. We introduce, in a very reduced-form manner, the concept of trust, $\mathcal{T}_{ht}$, into the model and link this to the cost of reading the central bank material.

The overall household- and medium-specific cost of reading a communication combines the three aspects introduced:

$$C_{IR,h,t} = F_{IR}\frac{\mu_h}{\mathcal{T}_{ht}} \tag{6}$$

We will now discuss each element in turn.

3.2.1 Household-specific processing cost

The household specific cost of processing the central bank communication, which could be thought of as their ability to process information, or the particular importance of information to them compared with other households, is modelled as follows. A fraction $\underline{\lambda}$ of households are endowed with a cost of information $\mu_h = 0$, while the remaining $1-\underline{\lambda}$ have μ_h which is drawn (before period 0) from an exponential distribution $\mu_h \sim \exp(\psi)$.[6] This means that the proportion of households with zero cost of processing the Inflation Report is the $\underline{\lambda}$ and also those drawn from the exponential distribution to have zero cost ($f(\mu_h = 0) = \psi$). Thus the population probability that the household has no cost, $p(\mu_h = 0)$, is given by the combination of these two: $p(\mu_h = 0) = \underline{\lambda} + (1 - \underline{\lambda})\psi$. The $\mu_h = 0$ distribution, $p(\mu_h = 0)$, is depicted in figure 1.

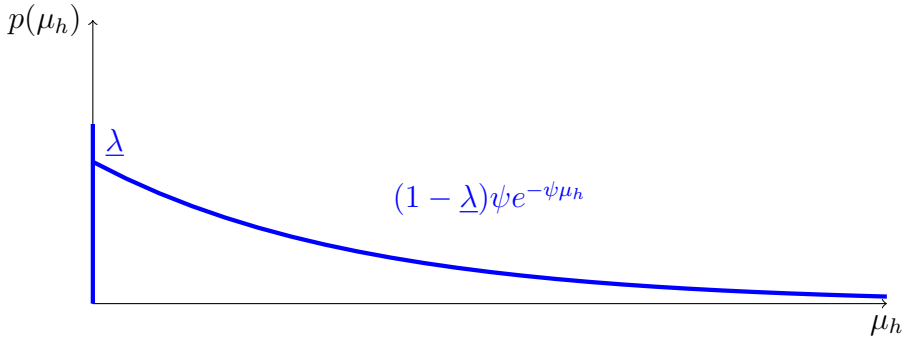

Figure 1: Distribution of Household Processing Costs μ_h

Notes: A fraction $\underline{\lambda}$ of households are endowed with a cost of information $\mu_h = 0$. The remaining $1 - \underline{\lambda}$ are drawn from an exponential distribution $\mu_h \sim \exp(\psi)$.

3.2.2 Traditional and simplified central bank communication

We will compare the effects of the central bank introducing a second form of communication that is easier to read. We will call these communications "Inflation Report" and "layered content" for consistency with the recent innovations with the Bank of England (discussed in Haldane and McMahon (2018) and below). We will assume the central bank always provides the Inflation Report and the decision is whether to introduce the second medium of communication.

Both forms of communication are costly to read. The layered content is easier to read because it contains less detail, which we model by setting $F_L < F_{IR}$. In line with the concerns discussed above, we impose that the layered content does not to communicate

[6] The probability density function (pdf) of an exponential distribution is defined over non-negative support as $f(x;\psi) = \psi e^{-\psi x}$.

fully the complex stochastic nature of the outlook. In other words, the easier to read content communicates the mean of the shocks at lower cost to the household but at the cost that the household underestimates the uncertainty around those forecasts.

Specifically, the layered content gives households the same expectations of all shocks as the full Inflation Report[7], but it does not say anything about the uncertainty around those expectations. Households misinterpret this to mean that there is no uncertainty.[8] This is clearly an extreme assumption. Household utility is unaffected by the uncertainty in the Inflation Report because this is a linearised model. However, the perceived certainty will lead households to be surprised by realisations that differ from their perceptions. These surprises, described formally below, reduce households' trust in the central bank.

3.2.3 Trust

We define a variable $\mathcal{T}_{ht} \in [0,1]$ to be the degree of trust household h has in the central bank. When a household trusts the central bank more, they will be more likely to pay attention to its communications which we model by including trust in the overall cost for a household when processing central bank signals. Trust evolves depending on the experiences of the household. We assume that trust in the central bank increases when the central bank communicates with a household. If, however, the communication leads the household to be surprised by the outcome, then trust will decline.

All households begin with $\mathcal{T}_h = 0.5$ and trust then evolves according to:

$$\mathcal{T}_{ht} = \begin{cases} 0 & \text{if } \hat{\mathcal{T}}_{ht} \leq 0 \\ \hat{\mathcal{T}}_{ht} & \text{if } \hat{\mathcal{T}}_{ht} \in (0,1) \\ 1 & \text{if } \hat{\mathcal{T}}_{ht} \geq 1 \end{cases} \tag{7}$$

Where:

$$\hat{\mathcal{T}}_{ht} = \mathcal{T}_{ht-1} + \delta_c \mathbb{1}_{\text{engage}} + \delta_s \mathbb{1}_{\text{surprise}} S(a_{t-1}, v_{t-1}, \epsilon^a_{t-1}, \epsilon^v_{t-1}), \quad \delta_c > 0, \delta_s < 0 \tag{8}$$

The indicator $\mathbb{1}_{\text{engage}}$ equals 1 when the household has processed some information

[7] As in Haldane et al. (2019), we assume that the Inflation report contains signals about each shock x_t given by $s^x_t = x_t + \epsilon^x_t$, where ϵ^x_t is an i.i.d. public noise shock. Households choosing to read the Inflation Report observe these signals and update their expectations about each fundamental shock x_t, and each noise shock ϵ^x_t, accordingly. Details of the resulting expectations can be found in Haldane et al. (2019).

[8] Strictly speaking, we are departing from Rational Inattention in the style of Sims (2003) here. Sims' information cost is proportional to the uncertainty reduction from processing the signal (measured by the expected entropy reduction between prior and posterior beliefs). Our simple signal reduces uncertainty to zero, and so would carry an infinite cost if we used this measure, which would not capture the intuitive notion that a point expectation is easier to communicate than the uncertainty around that expectation. This is why we specify the cost of processing signals in terms of the reduced-form constants F_{IR} and F_L.

from the central bank in period t. δ_c measures the responsiveness of trust to engagement. Some households will, in equilibrium, choose optimally to not read *any* communication.

In period t, the household observes the realisations of the shocks from period $t-1$. The indicator $\mathbb{1}_{\text{surprise}}$ equals 1 in period t if the realised shocks were outside that support of the household's expectations. In standard rational-inattention models, communication induces beliefs with an infinite support so these surprises never happen. This will, however, occur when we move to the simpler, 'layered content' for the reasons described above. The function $S(\cdot)$ measures how surprised the household is – how far realised shocks deviate from the edge of their beliefs – and it is defined formally below. δ_s measures the responsiveness of trust to surprises.

3.3 Welfare and Information Processing

The key to the model is the fact that information is costly to process but less-informed households suffer a welfare loss by making sub-optimal decisions. In this case, the household will decide whether or not to read central bank communication and, in doing so, will become somewhat informed. Households who optimally choose not to be informed will have no signals.

We follow a guess-and-verify approach:

1. We start with a guess for how shocks influence inflation and the consumption of a fully informed household; we assume that each is a linear functions of current shocks and public noise and refer to these as the policy functions.

   The policy function for fully informed consumption is $c_t^* = \beta_0 a_t + \beta_1 v_t + \beta_2 \epsilon_t^a + \beta_3 \epsilon_t^v$.

2. Given these relationships, we then find the consumption of inattentive households and the output gap implied by these linear rules. These choices of the inattentive households feed back into the model equations and determine the coefficients of the policy functions.
3. All of these policy function coefficients are dependent on the amount of attention households pay to central bank communication. The expected utility loss from being less than fully informed about shocks, to a quadratic approximation of the utility function[9], is proportional to the variance of $(c_t^* - c_t)$ – the gap between the consumption of a fully-informed household and actual consumption. We then solve for the household's decision to pay attention which depends on its time- and household-specific processing cost.
4. Since only a fraction of households process the central bank communication in

[9]We prove this finding in appendix A.

period t. Define this fraction Λ_t which is given by the following expression:

$$\Lambda_t = \underline{\lambda} + (1 - \underline{\lambda})\lambda_t \tag{9}$$

where λ_t is the fraction of households with positive information costs who process the communication. This variable will feed back into the behaviour of the economy.

5. Once we have the optimal household choices and the implied behaviour of the economy, we can explore the effects of introducing an alternative form of communication.

This guess-and-verify approach is necessary because of the role of higher-order beliefs in equilibrium. For a household to translate their expectations of each shock into a consumption choice, they must form a belief about how the interest rate responds to shocks. To do this they need to form beliefs about how other households respond to the shocks, and so each household must form beliefs about the (average) shock expectations of other households, and also about the beliefs of those other households about all other households, and so on. This is not an issue in full information models, where all households have the same expectations and those expectations are common knowledge. The guess-and-verify approach finds an equilibrium for the higher-order belief problem, and is common in the rational inattention literature (see e.g. Maćkowiak and Wiederholt (2009)).

3.4 To Read, Or Not To Read the Inflation Report?

We start by considering an environment in which there is only the Inflation Report, as in the baseline model. The key result from this section can be summarised as:

Result 1

When there is only the Inflation Report from the central bank, the equilibrium will be a steady state in which all households with zero idiosyncratic processing costs ($\mu_h = 0$), and some households with positive processing costs ($\mu_h > 0$), will read it.

Trust is constant in the steady state with all readers of the communication having full trust. Those who don't read anything remain with trust at its starting value because they never engage, but never form precise expectations so are also never surprised.

The decision is whether a household will read the Inflation Report or not read anything. Households processing the Inflation Report are not fully informed: they observe noisy signals s_t and set $c_{IR,t} = \mathbb{E}_t^H(c_t^*|s_t)$. Households not processing any communication get no information and so set $c_{N,t} = 0$.

If more households pay attention (i.e. if Λ_t rises), inflation is less volatile conditional on fundamental shocks a_t and v_t, because aggregate consumption is more responsive to

these shocks. Conversely, aggregate consumption is also more responsive to noise in the Inflation Report when Λ_t rises, which increases the volatility of inflation. The overall effect of an increase in the proportion of households who are attentive is that the variance of inflation falls, which means that the consumption of a fully-informed household is less volatile - which reduces the incentive for other households to pay attention.[10]

As noted above, the utility loss from lack of information is a constant multiple of the variance of the gap between actual consumption and the optimal consumption of a fully-informed household.[11] The utility loss from choosing no information, rather than reading the Inflation Report, is therefore a constant multiplied by the difference between these two variances. This simplifies to:

$$\text{Var}(c_t^* - c_{N,t}) - \text{Var}(c_t^* - c_{IR,t}) = \left(\frac{\kappa\phi_\pi(1+\varphi)}{(\sigma+\kappa\phi_\pi\Lambda_t)(\sigma+\varphi)}\right)^2 \tau_a\sigma_a^2 + \left(\frac{\phi_\pi}{\sigma+\kappa\phi_\pi\Lambda_t}\right)^2 \tau_v\sigma_v^2 \quad (10)$$

where τ_a and τ_v are the signal to noise ratios of the Inflation Report signals about the technology shock and the cost-push shock respectively.

We normalise the constant in front of this variance in the utility loss to 1 without loss of generality, as it just requires a rescaling of the complexity of information parameter F_{IR}. Households therefore choose to pay attention to the Inflation Report if:

$$\left(\frac{\kappa\phi_\pi(1+\varphi)}{(\sigma+\kappa\phi_\pi\Lambda_t)(\sigma+\varphi)}\right)^2 \tau_a\sigma_a^2 + \left(\frac{\phi_\pi}{\sigma+\kappa\phi_\pi\Lambda_t}\right)^2 \tau_v\sigma_v^2 > F_{IR}\frac{\mu_h}{\mathcal{T}_{ht}} \quad (11)$$

In the initial period when all households have trust equal to 0.5, all $\underline{\lambda}$ households with $\mu_h = 0$ pay attention to the Inflation Report. In addition, a fraction λ_0 of households with $\mu_h > 0$ pay attention, that is all households with a $\mu_h < \mu^*(\lambda_0)$, where:

$$\left(\frac{\kappa\phi_\pi(1+\varphi)}{(\sigma+\kappa\phi_\pi(\underline{\lambda}+(1-\underline{\lambda})\lambda_0))(\sigma+\varphi)}\right)^2 \tau_a\sigma_a^2 + \left(\frac{\phi_\pi}{\sigma+\kappa\phi_\pi(\underline{\lambda}+(1-\underline{\lambda})\lambda_0)}\right)^2 \tau_v\sigma_v^2 = F_{IR}\frac{\mu^*(\lambda_0)}{0.5+\delta_c} \quad (12)$$

[10] This is why we model a continuum of household information costs: with two types of households (low μ and high μ) there will not necessarily be an equilibrium where households play pure strategies of either paying attention or not.

[11] In order for this to be the relevant loss function here we assume that households do not take into account how the parameters in the optimal decision rule will change over time. That is, they assume that the current share of households processing information Λ_t will persist forever, though in fact with layered content it won't.

The exponential distribution of μ_h means that λ_0 is given by:[12]

$$-\frac{F_{IR}\ln(1-\lambda_0)(\sigma+\kappa\phi_\pi(\underline{\lambda}+(1-\underline{\lambda})\lambda_0))^2}{(0.5+\delta_c)\psi\phi_\pi^2}=\left(\frac{\kappa(1+\varphi)}{\sigma+\varphi}\right)^2\tau_a\sigma_a^2+\tau_v\sigma_v^2 \tag{13}$$

From this, $\frac{d\lambda_0}{dF_{IR}}<0$. That is, the more difficult the Inflation Report is to process, the fewer households process it.

After the initial period, all households with $\mu_h > \mu^*(\lambda_0)$ pay attention to the Inflation Report and see their trust rise until it reaches the maximum trust of 1. The steady state with the Inflation Report as the only possible communication from the central bank therefore has a share $\Lambda_0 = \underline{\lambda} + (1-\underline{\lambda})\lambda_0$ of households processing any information about shocks, and an average trust of:

$$\bar{\mathcal{T}}_0 = \Lambda_0 + \frac{1-\Lambda_0}{2} = \frac{1+\Lambda_0}{2} \tag{14}$$

3.5 Introducing Simplified Communication

Now, instead, imagine that in period 1, the central bank introduces the new form of easier to process communication. The key result from this section can be summarised as:

Result 2
Simplified Communication initially increases trust as more households engage with the central bank.
But when a large shock arrives households are surprised, lose trust, and stop engaging. Not reading the simplified communication is an absorbing state, as there is no way for trust to increase once a household has stopped reading the communication.
If trust starts out lower before the introduction of simplified communication, the initial gain in trust is larger, but the decay in engagement occurs more quickly.

In terms of the decision when there is a choice of media, we can distiguish the reaction of different types of household; those who were reading the Inflation Report and stay reading it; those who switch to reading the layered content; those who didn't read the inflation report but start reading layered content; and those who never engage with either medium.

Since household utility is unaffected by the uncertainty in the Inflation Report, households with $\mu_h = 0$ are indifferent between the Inflation Report and the simplified com-

[12]The quantile function of the exponential distribution is, conveniently, given by:

$$\mu^*(\lambda) = -\frac{\ln(1-\lambda)}{\psi}$$

munication. We assume that all households with $\mu_h = 0$ continue to read the Inflation Report.[13]

Households with $\mu_h > 0$, however, strictly prefer the simplified communication: it gives the same expected utility loss and is cheaper to process. All households with $\mu_h \in (0, \mu^*(\lambda_0)]$ therefore switch from processing the Inflation Report to paying attention to the simplified communication.

In addition, many households who were previously processing no information from the central bank will now read the simplified communication. This is true for households with $\mu_h \in (\mu^*(\lambda_0), \mu^*(\lambda_1)]$, where:

$$-\frac{F_L \ln(1-\lambda_1)(\sigma + \kappa\phi_\pi(\underline{\lambda} + (1-\underline{\lambda})\lambda_1))^2}{(0.5+\delta_c)\psi\phi_\pi^2} = \left(\frac{\kappa(1+\varphi)}{\sigma+\varphi}\right)^2 \tau_a\sigma_a^2 + \tau_v\sigma_v^2 \tag{15}$$

Note that $F_L < F_{IR}$ implies that $\lambda_1 > \lambda_0$, so we can be sure that some households switch from processing no information to paying attention to simplified communication.

These forces have opposing effects on trust. Processing the simplified communication increases the trust of these households in the central bank. In periods after switching, however, the households who read the simplified communication are subject to being surprised, which reduces trust.[14]

The degree to which their trust falls is determined by δ_s as well as how far the shocks are from the expectations given by the simplified communication which is determined by the $S(\cdot)$ function:

$$S(a_t, v_t, \epsilon_t^a, \epsilon_t^v) = (a_t - \mathbb{E}_t^L a_t)^2 + (v_t - \mathbb{E}_t^L v_t)^2 + (\epsilon_t^a - \mathbb{E}_t^L \epsilon_t^a)^2 + (\epsilon_t^v - \mathbb{E}_t^L \epsilon_t^v)^2 \tag{16}$$

Here $\mathbb{E}_t^L$ is the expectation induced by the simplified communication. By assumption, the simplified communication implies the same expectations of each shock as the Inflation Report, so $\mathbb{E}_t^L a_t = \tau_a(a_t + \epsilon_t^a)$, $\mathbb{E}_t^L v_t = \tau_v(v_t + \epsilon_t^v)$, $\mathbb{E}_t^L \epsilon_t^a = (1-\tau_a)(a_t + \epsilon_t^a)$, $\mathbb{E}_t^L \epsilon_t^v = (1-\tau_v)(v_t + \epsilon_t^v)$. Substituting this in to the definition of S we have:

$$S(a_t, v_t, \epsilon_t^a, \epsilon_t^v) = 2(a_t(1-\tau_a) - \tau_a\epsilon_t^a)^2 + 2(v_t(1-\tau_v) - \tau_v\epsilon_t^v)^2 \tag{17}$$

Note that the extent of surprise expected by the policymaker before shocks are realised

[13] This is necessary because all households who switch to simplified communication will eventually lose trust and switch to not processing any information. If all households did this, aggregate consumption would be completely unresponsive to the interest rate and the model solution would be indeterminate.

[14] They observe the true realisations of the previous period fundamental shocks a_{t-1} and v_{t-1}, and the noise shocks ϵ_{t-1}^a and ϵ_{t-1}^v. The values communicated in the simplified communication were combinations of fundamental and noise shocks, so the probability that these shocks exactly equal the values communicated in the simplified communication is zero. The shock realisations are therefore outside of the range households reading the simplified communication thought was possible, and so they lose trust.

is therefore:

$$\mathbb{E}_{t-1}S(a_t, v_t, \epsilon_t^a, \epsilon_t^v) = 2(1-\tau_a)\sigma_a^2 + 2(1-\tau_v)\sigma_v^2 \tag{18}$$

There is also a dynamic effect from the evolution of trust. If there are a few periods with small shocks, then the surprises S_t will be low and trust will rise. Eventually, however, there will be large enough shocks that trust falls, and when this happens households stop reading the simplified communication because the cost of processing it rises with falling trust. In the model, not reading the simplified communication is an absorbing state, as there is no way for trust to increase once a household has stopped reading the communication. This is an extreme assumption. This means that, eventually, there will be a series of sufficiently large shocks that the share of households processing simplified communication hits zero.[15] At that point, only the $\underline{\lambda}$ households with no information cost remain processing any information. We therefore eventually reach a new steady state with $\Lambda_t = \underline{\lambda}$. This is lower than the share of households processing information in period 0, before the introduction of the simplified communication.

The expected time path for λ_t, the share of households with positive information costs μ_h who process any information at all, is plotted in figure 2a for a quarterly calibration (discussed in appendix B). In this calibration, before the introduction of simplified communication a fraction $\lambda_0 = 0.1$ of households with positive information costs read the Inflation Report. In period 1, all of these switch to reading the simplified communication, and a further 20% of the households with $\mu_h > 0$ switch from not processing any information to reading the simplified communication. The new communication therefore initially has the effect that more households pay attention to the communication. Over time, however, the trust of households processing the simplified communication is eroded, and so households start to switch to no information processing.

The average trust households have in the central bank is expected to evolve according to the path plotted in figure 2b. Initially, trust rises when the simplified communication is introduced, because many households who were not paying any attention to central bank communication now read the simplified communication, and that contact with the central bank increases their trust. However, over time this boost is outweighed by the losses in trust when households see past realisations of shocks and realise that they were outside of the support of their beliefs, which they were given by the central bank through the simplified communication. Trust then falls. The rate at which it falls is decreasing over time (the time path is convex) because a household's trust only continues to fall for as long as they pay attention to the simplified communication. As time passes, fewer

[15] Interestingly, the fact that not reading any information is an absorbing state means that even if $\delta_c + \delta_s \mathbb{E}S \geq 0$, i.e. if trust would rise over time if surprises were of their expected magnitude, the model eventually ends up at the low-trust steady state.

and fewer households are still paying attention to that communication, and so the rate of decrease of average trust slows down. Eventually, no households are left paying attention to the simplified communication and average trust reaches a new lower steady state.

With this calibration, average trust is above its initial (pre-simplified communication) level for 11 quarters on average, and the share of households engaging with simplified communication remains above its initial level for 50 quarters. This continues to be higher than the initial level long after trust is below its initial value because the simplified communication has a lower processing cost than the Inflation Report. Trust and engagement reach their new lower steady state after approximately 250 quarters.

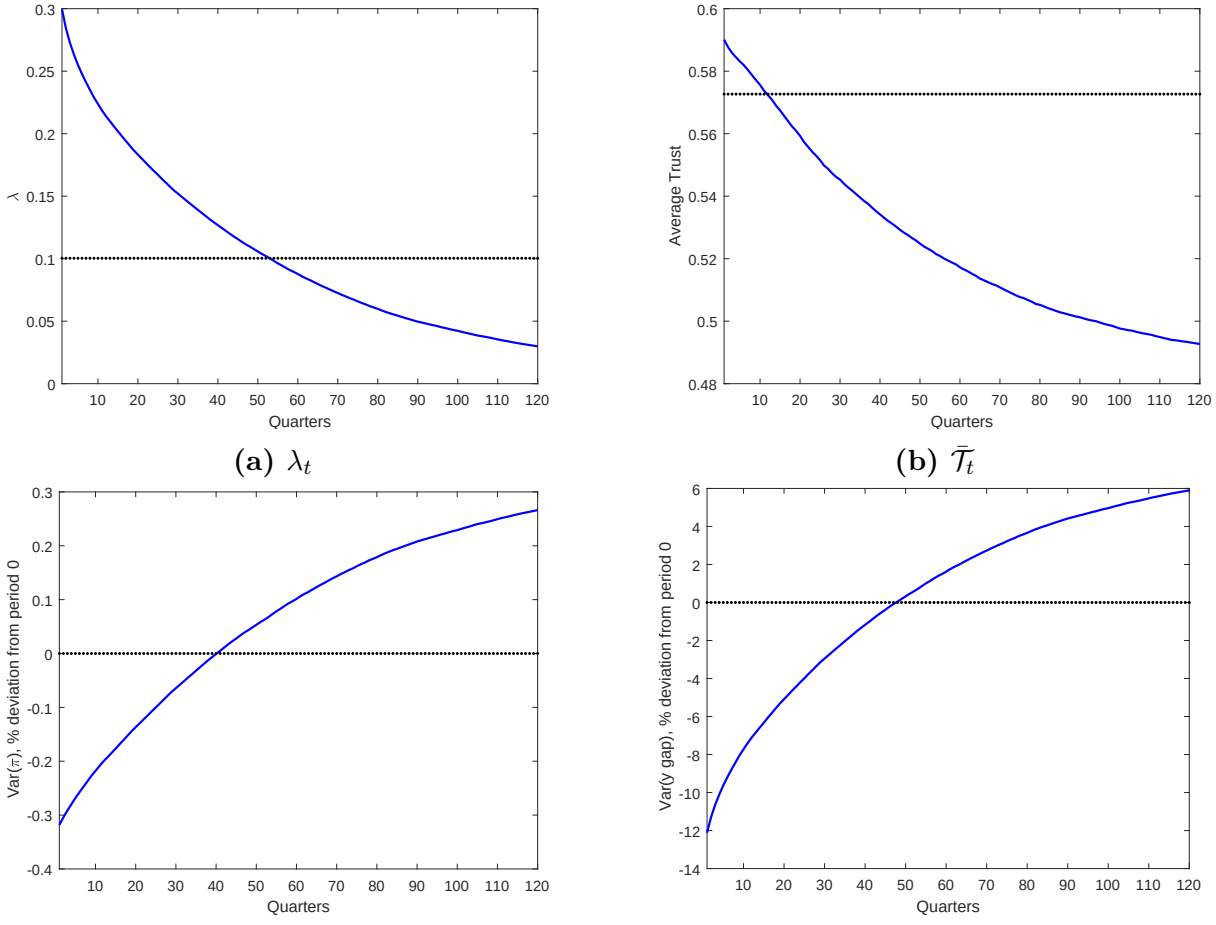

(a) λ_t **(b)** $\bar{\mathcal{T}}_t$

(c) Varπ **(d)** Var$\tilde{y}_t$

Figure 2: Time path of λ_t, $\bar{\mathcal{T}}_t$, Varπ and Var$\tilde{y}_t$ after the introduction of simplified communication

Notes: The blue line is the expected path of either share of processing households or average trust. The black dotted line is the steady state where no simplified communication has ever existed.

The critical trust level at which a household with information cost μ_h stops processing

the simplified communication is given by:

$$\mathcal{T}_h^* = \frac{F_L \mu_h (\sigma + \kappa\phi_\pi(\underline{\lambda} + (1-\underline{\lambda})(1-e^{-\psi\mu_h})))^2}{\phi_\pi^2 \left(\left(\frac{\kappa(1+\varphi)}{\sigma+\varphi} \right)^2 \tau_a \sigma_a^2 + \tau_v \sigma_v^2 \right)} \tag{19}$$

This critical trust is increasing in μ_h, so households who face higher information costs stop processing simplified communication earlier, when their trust has fallen only a small amount. Once a household has stopped processing the simplified communication, their trust from the next period onwards is $\mathcal{T}_h^* + \delta_s S(a_{t^*}, v_{t^*}, \epsilon^a_{t^*}, \epsilon^v_{t^*})$, where t^* is the last period in which they processed the simplified communication.[16] This model has the implication, consistent with the UK data, that the households with the highest trust are also those with most engagement and understanding; the $\underline{\lambda}$ households.

The effect of this on welfare is clear. When the fraction of households processing information about shocks increases, the unconditional variance of inflation and the output gap decrease, boosting welfare. This is because attentive households respond appropriately to changes in the interest rate, where inattentive households do not. A greater share of responsive households therefore has the same effect in the model as increasing the Taylor Rule coefficient ϕ_π. However, inflation and the output gap are more volatile in the new steady state, because fewer households ultimately process information about shocks. The time-path for the volatility of inflation and the output gap is plotted in figure 2c and 2d.

This means that even if the policymaker does not care about trust for its own sake, introducing simplified communication can have negative long-run welfare effects. This is because it causes some households who were previously paying attention to the Inflation Report to switch to simplified communication and therefore lose trust in the central bank, which means that the long-run share of households processing information from the central bank falls, and that increases the volatility of inflation.

3.6 Factors Affecting the Balance Between the Two Effects

As described at the beginning of this section, this model is engineered to give a central role to the concern that simple communication doesn't communicate uncertainty appropriately and this can lead households to become surprised. Given the result above, why would a central bank in our model environment adopt the simple communication strategy? In this subsection we describe the key model parameters that alter the magnitude of, and speed of moving between, the positive and negative welfare effects. The frame-

[16] The extra $\delta_s S$ comes from the surprise they receive in the period after they stop processing simplified communication, when they realise that the shocks in period t^* were not within the support of their expectations.

work also allows us to begin to explore the extent to which central banks introducing simplified communication may wish to also engage in other outreach activities to try to prevent this disengagement and reduction in welfare over time. In the next section, we relate these model parameters to more practical concepts in the real world and emphasise the 3 E's.

3.6.1 Myopia

In assessing the decision to introduce the simple forms of communication, a central banker needs to weigh near-term welfare gains with longer term losses. Since the potential costs come only over time, a more myopic central banker will be more likely to want to switch as the future welfare losses will be discounted toward zero.

3.6.2 Luck

If there are periods of shocks with smaller than average magnitudes then average trust will rise, and no households will switch away from reading the simplified communication. As soon as larger shocks come along, though, trust will fall. To see this, figures 3a and 3b plot the paths of λ_t and $\bar{\mathcal{T}}_t$ for two simulations of the model.

The effect of these different time paths of trust and engagement is reflected in markedly different welfare effects. Figure 3c (3d) shows that inflation (output gap) volatilty decreases (decreases) when simplified communication is introduced, and if shocks are benign it stays low, as in the first 55 periods of the green simulation. This makes the adoption of simplified communication much more beneficial in the green (solid) simulation than in the red (dashed), where large shocks early on after the introduction of simplified communication cause large falls in trust and engagement.

3.6.3 Less sensitivity to surprises (or greater sensitivity to communication)

Another obviously important aspect of the model, and one that may potentially be influenced be the central bank, is the speed at which the central bank gains or loses trust ($\delta_c > 0$ and $\delta_s < 0$). Interestingly, we reach the new lower-welfare steady state even in the case where the trust loss from the expected surprise is smaller than the trust gain from communication (i.e. $\delta_c + \delta_s \mathbb{E}S > 0$). This is because not reading any communication is an absorbing state: once trust has fallen below the critical level for a household, it is assumed they stop reading any communication and there is no way for trust to rise again. (In reality, the central bank will have to adopt alternative engagement techniques to reestablish trust.) After many periods, there will eventually be enough large shocks to ensure that trust falls to the level needed to reach the new steady state. This is helped

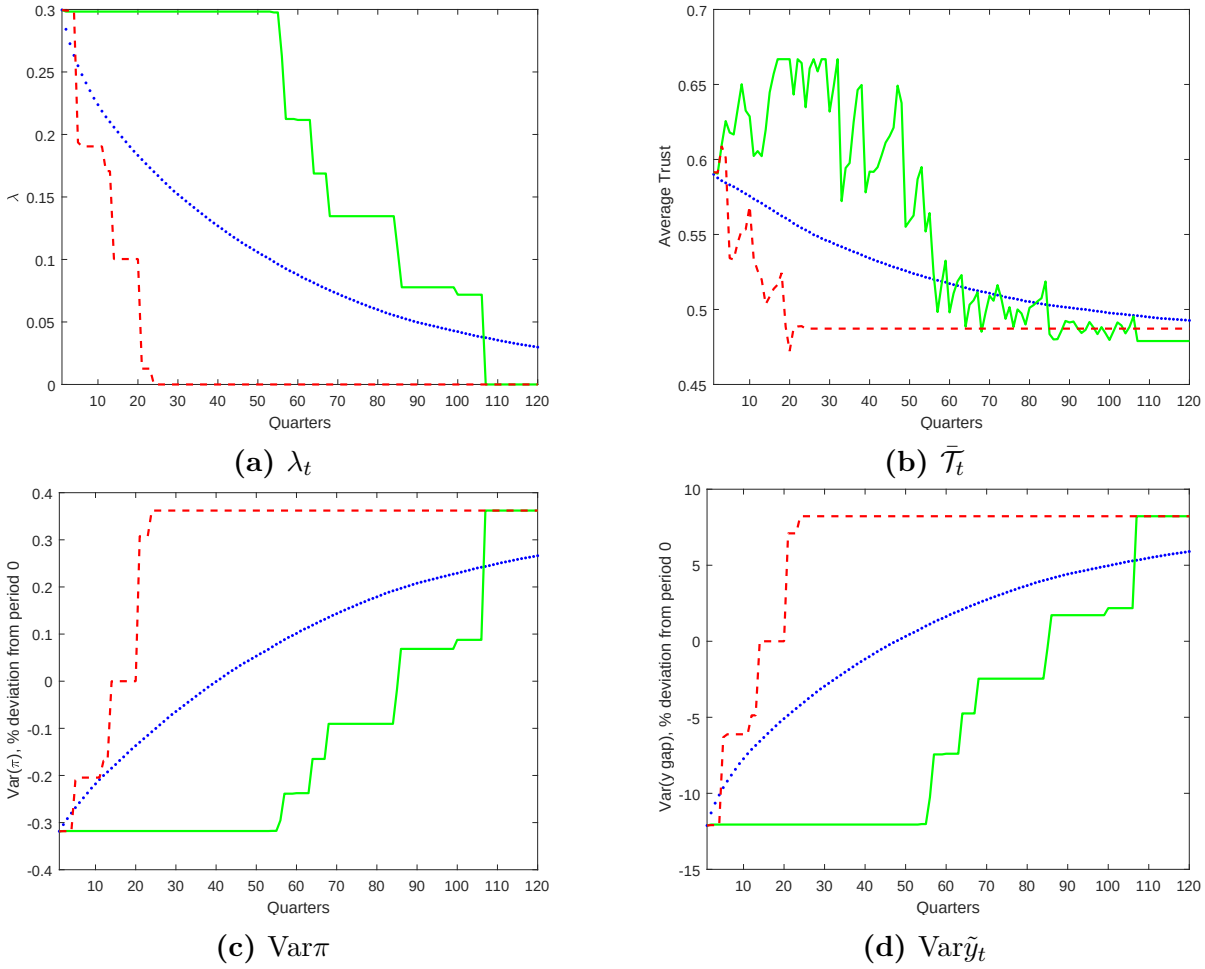

(a) λ_t **(b)** $\bar{\mathcal{T}}_t$

(c) Varπ **(d)** Var$\tilde{y}_t$

Figure 3: Time path of λ_t, $\bar{\mathcal{T}}_t$, Varπ and Var$\tilde{y}_t$ after the introduction of simplified communication: the effect of bengin or volatile times

The blue dotted line is the expected path of share of processing households, average trust, the variance of inflation or the variance of the output gap. The green (solid) and red (dashed) lines are these variables for two simulations of the model.

by the fact that trust is bounded above by 1, so many periods of reasonably accurate communication does not imply trust continually improving. Figure 4 plots the same expected time paths of model variables considered in figures 2, comparing the baseline results with the corresponding paths if penalty from surprises (δ_s) has been reduced so that the effect of communication on trust after an average-sized surprise is positive.

It takes much longer for households to stop reading the simplified communication in this setting, and so there are many more periods before engagement with central bank communication is expected to fall below its initial level. The economy does still arrive at the new steady state in which no household with $\mu_h > 0$ reads any central bank communication eventually, however. In this calibration that is expected to occur after approximately 100 years (400 quarters).

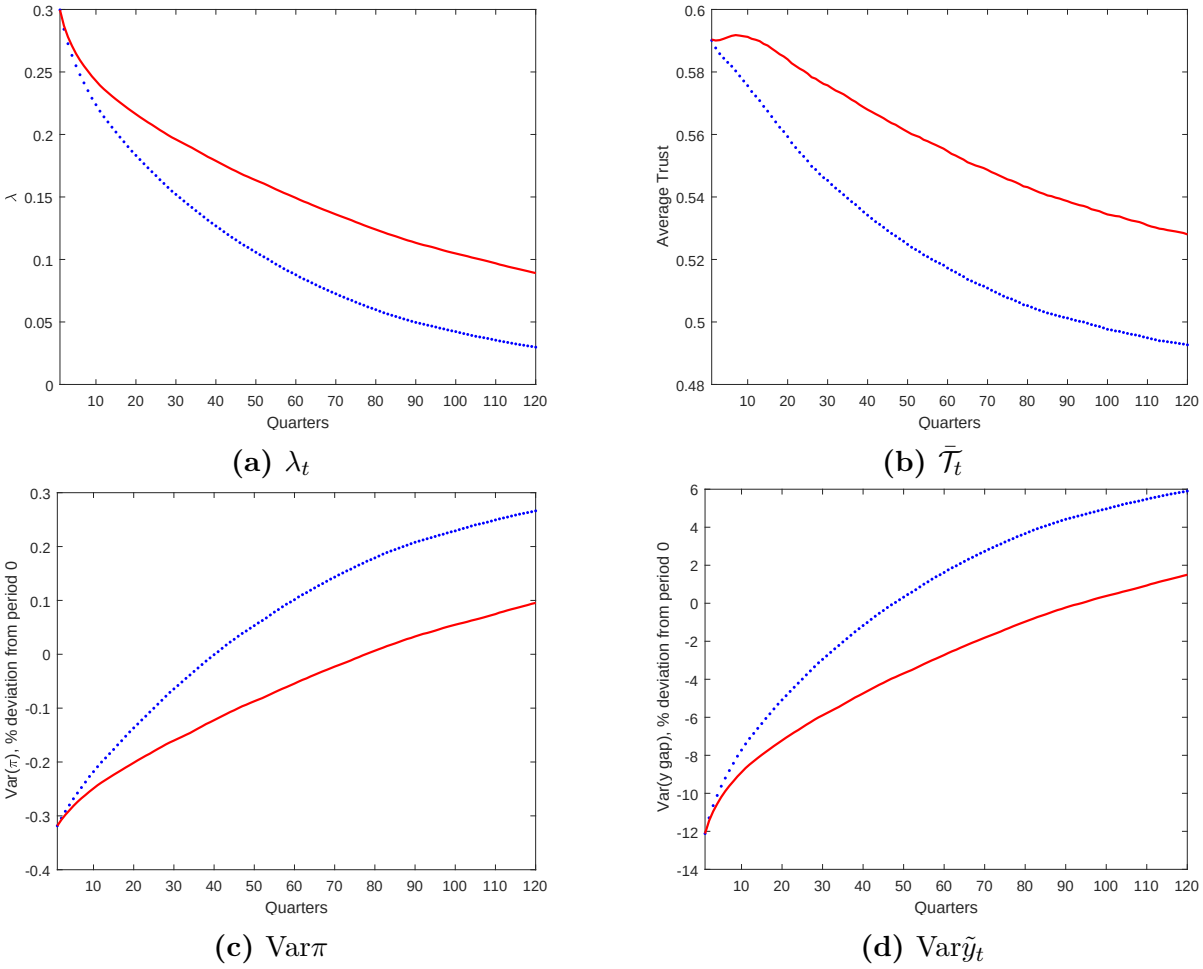

(a) λ_t **(b)** $\bar{\mathcal{T}}_t$

(c) Varπ **(d)** Var$\tilde{y}_t$

Figure 4: Time path of λ_t, $\bar{\mathcal{T}}_t$, Varπ and Var$\tilde{y}_t$ after the introduction of simplified communication: the effect of less sensitivity via $\delta_c > 0$ and $\delta_s < 0$

Notes: The blue dotted line is the expected path share of processing households, average trust, the variance of inflation or the output gap in the baseline model. The red solid line is the same expected paths for a 10% smaller (less negative) value of δ_s.

3.6.4 Starting levels of trust matter

The introduction of simplified communication at the Bank of England did not take place in a vacuum. It was, in part, a response to a general fall in the trust households placed in the institution (and in public institutions in general) after the Great Recession (as highlighted above). Here we show the effects of introducing simplified communication in this model differ depending on whether it is done in an era of high trust (i.e. pre-crisis) or after an external shock has reduced the trust of all households (post-crisis).

Figure 5 plots the expected paths of the share of households with positive information processing costs engaging with central bank communications λ_t, average trust, and the volatility of inflation and the output gap, after the introduction of simplified communi-

cation for two starting points. In the first (drawn in blue), trust in period zero before the introduction of simplified communication is high for all households, whereas in the second (in red) initial trust is low for all households, even those who have been reading the Inflation Report for many periods[17]. In both cases, the expected paths of all variables are plotted as percentage deviations from the respective values of these variables in the period before the introduction of simplified communication.

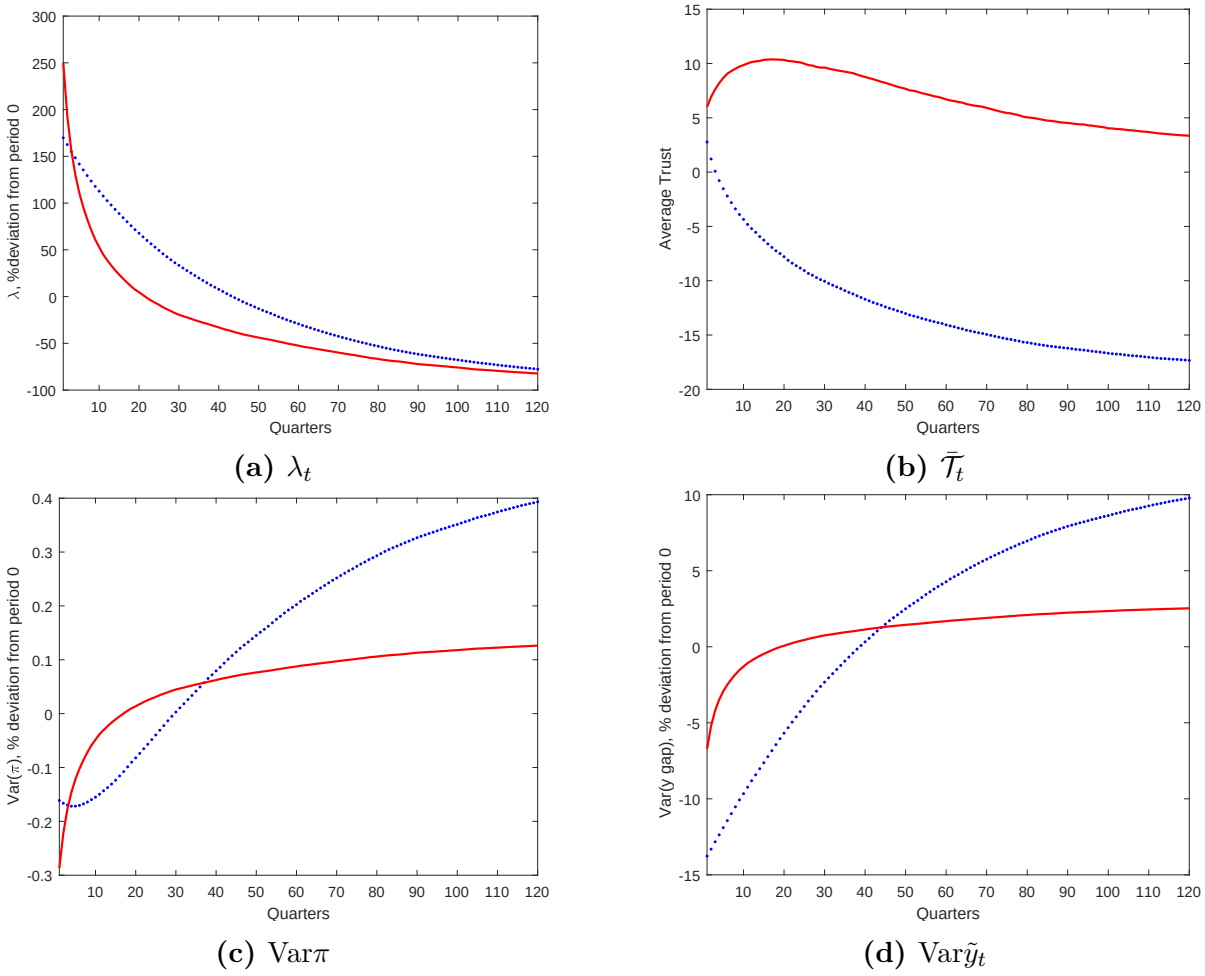

(a) λ_t **(b)** $\bar{\mathcal{T}}_t$

(c) Varπ **(d)** Var$\tilde{y}_t$

Figure 5: Time path of λ_t, $\bar{\mathcal{T}}_t$, Varπ and Var$\tilde{y}_t$ after the introduction of simplified communication: the effect of different starting levels of trust.

Notes: The blue dotted line is the expected path of either share of processing households, average trust, the variance of inflation or of the output gap relative to initial values with high initial trust. The red solid line is the expected path relative to period 0 of the same variables in the case where initial trust is low.

The share of households engaging with central bank communication λ_t increases when simplified communication is introduced for both initial levels of trust, but this increase is

[17]The initial trust before simplified communication of those not reading any communications and those reading the Inflation Report is 0.9 and 1 respectively in the high trust case, and 0.1 and 0.2 respectively in the low trust case.

substantially larger when initial trust is low. However, low initial trust also leads to a more rapid decline in λ_t. This is because the total cost to a household of processing central bank communications is the complexity of that information F multiplied by $\frac{\mu_h}{\mathcal{T}_h}$. The difference between the cost of processing the simplified communication and the Inflation Report is therefore higher when trust is low:

$$C_{L,h,t} - C_{IR,h,t} = \frac{(F_L - F_{IR})\mu_h}{\mathcal{T}_h} \tag{20}$$

When trust is low, introducing simplified communication therefore makes a greater difference to the costs of processing central bank communication, and so the initial rise in λ_t when simplified communication is introduced is greater when trust is low. The rate at which processing cost falls as trust rises is also greater when trust is low. This is why λ_t falls more quickly over time in the low initial trust case:

$$\frac{dC_{L,h,t}}{d\mathcal{T}_h} = -\frac{F_L\mu_h}{\mathcal{T}_h^2} \tag{21}$$

These paths for λ_t imply that the fall in the volatility of inflation on the introduction of simplified communication is greater when initial trust is low, but that inflation volatility also rises more quickly in this case. The low-trust steady state that is reached after many periods of simplified communication and household surprises is the same irrespective of the initial levels of trust. As the variance of inflation before simplified communication is higher when trust is lower, the increase in inflation volatility from pre-simplified communication to the new steady state is smaller for lower initial trust[18].

Interestingly, average trust may actually be higher in the new steady state after the introduction of simplified communication than it was with just the Inflation Report. This is because there is a large number of households who engage with the simplified communication and so see their trust rise. They stop engaging when they receive a surprise and their trust falls, but it is still above the level when the Inflation Report was the only way to engage with the central bank. This is not the case in our baseline with medium initial trust, or with high initial trust.

4 The 3 E's of Public Communication

We now turn to consider the practical steps a central bank can take in conjunction with adopting simplified communication. If the concerns built in to the model are correct, these

[18]Lower initial trust implies higher initial inflation volatility because with lower trust, fewer households are engaging with the Inflation Report, and so fewer households are informed about shocks.

are necessary steps to ensure the longest possible benefits in terms of welfare and trust. But even if not, these are likely to be desirable as part of central banks' commitment to be accountable to the whole economy.

While most of the central bank communication literature focuses on management of expectations, we adopt a focus on 3 E's that play an important role in the management of expectations:[19]

- Explanation
- Engagement
- Education

We shall discuss each in turn and relate the ideas our model's predictions. We also discuss the Bank of England's activities under each heading.

4.1 Explanation

This is the core of communication in the effort to manage expectations. Explanation is about ensuring the people form their expectations with the best possible information. In the model, it is the sending of signals. In reality it is much harder. The economy is not summarised by two independent shocks but is, rather, a high dimensional and extremely complex system.

In the model, we embedded the complexity of the explanation in the common cost of the communication F_{IR} or F_L. The idea of the model is that clearer explanations that are easier to read (related to the earlier material on readability measures), build trust but ultimately may lead to the household being unduly confident about the future outcome such they are surprised by actual developments.

Haldane and McMahon (2018) undertook an experiment using the communication from the Bank of England's November 2017 release of more-easily-understood communication alongside the traditional quarterly Inflation Report (IR) and Monetary Policy Summary.[20] The new, broader-interest version of the IR became known as its layered content; different layers spoke to less-specialist audiences. In that paper, we presented the results of these experiments conducted immediately after the November 2017 Inflation Report launch. There were two groups surveyed: a survey of 285 members of the UK general public ("Public sample") and a sample of first-year graduate students in the department of economics at the University of Oxford ("MPhil sample"). Here we relate those results to the analysis in the paper, as well as update the discussion for more recent

[19]There are numerous other "3 E's" in different fields such as the 3 E's of sustainability (Environmental, Economic, and Ethical) as in Goodland (1995).

[20]Experiments in macroeconomics are more common than often considered to be the case. For example, see Petersen et al. (2014) for a discussion.

analyses of the issue.

The layered content achieved its aim of being easier to read. It had a Flesch-Kincaid Grade Level of 7.8 (eighth grade level), which compares with the Monetary Policy Summary, which was released at the same time, with a Flesch-Kincaid Grade Level of 13.4.

We randomly assigned participants to read the new content or the traditional content and measured their expectations for the UK economy at the time using equation (22). The dummy variable, *D(Layers)*, indicates those participants that read the new style communication. We use a series of demographic controls, X_i, in the public sample, though these are not available in the MPhil sample.[21] As a proxy for knowledge, we use whether or not the person has studied economincs (D(Econ)) One of the questions asks "To what extent do you have confidence in the Bank of England as a public institution to implement macroeconomic policy?"; we use this, Trust, as our proxy measure for existing trust in the Bank of England.

$$Y_i = \gamma_0 + \gamma_1 D(\text{Layers}) + \gamma_2 D(\text{Econ}) + \gamma_3 \text{Trust} + \Gamma' X_i + \zeta_i \tag{22}$$

Here we replicate and expand on that earlier analysis to show how the responses depend on both knowledge of economics and the proxy for pre-existing trust in the institution. In order to emphasise the latter point, we also run a regression, equation (23) that includes an interaction between existing trust and exposure to the new content:

$$Y_i = \gamma_0 + \gamma_1 D(\text{Layers}) + \gamma_2 D(\text{Econ}) + \gamma_3 \text{Trust} + \gamma_4 \text{Prior Trust} \times D(\text{Layers}) + \Gamma' X_i + \zeta_i \tag{23}$$

We assessed the effect of the new style on responses to three questions:

1. "To what extent are you able to understand the content and messages of the material you just read?". Participants selected from a five-point scale from which we created a numeric variable *Understand* which ranges from 1 ("None or nearly none of it") to 5 ("All or nearly all of it").
2. "How has reading the excerpt from the Inflation Report summary changed your views or expectations on the outlook for the UK economy, if at all?" From this question, along with knowledge of how participants differed from the IR forecasts, we define a dummy variable *D(Adjust)* which is 1 if the participant appropriately adjusts their expectations, and 0 otherwise.[22]
3. "Learning that this is typical of the type of communication in the Bank of England's

[21] Excluding controls does not significantly affect the results for the public sample.

[22] Participants provided their two-years expectations for CPI inflation, unemployment and interest rates on a five-point scale from "Fall significantly (-2)" through "Broadly unchanged (0)" to "Increase significantly (2)". The November 2017 IR projections were also mapped to this scale. This allowed us to work out whether converging on the IR projections meant that the participant *should* become more pessimistic (higher inflation, unemployment and/or interest rates) or optimistic.

quarterly Inflation Report, how has the Inflation Report summary affected your perceptions of the Bank of England, if at all?" The five-point numeric scale, measuring $\Delta Perception$, runs from "Worsened significantly (1)" to through "Broadly unchanged (3)" to "Improved significantly (5)".

Table 4 presents the results of regressions of *D(Layers)* on participant understanding from the two different samples. Columns (1)-(3) present the results for the public sample and (4)-(5) for the MPhil survey. The main result is that, for both samples, the new layered content is easier to read and understand, even for technically-advanced MPhil students. This improvement in understanding was statistically significant for both samples, at the 1% level, and averaged 0.68 points across the two. To contextualise these benefits, the effect of the layered content on understanding is larger than the effect on understanding of studying economics as part of a university degree. The MPhil sample results suggests that even the traditional, technically-trained audiences may benefit from clarifying and simplifying communication.

Table 4: Regression analysis of communication experiment on Understanding

	(1) Understand	(2) Understand	(3) Understand	(4) Understand	(5) Understand
D(Layers)	0.71***	0.83***	0.73***	0.63***	0.46
	[0.00]	[0.00]	[0.00]	[0.00]	[0.14]
Trust x D(Layers)			0.21		0.26
			[0.11]		[0.47]
D(Economics)	0.54***		0.54***		
	[0.00]		[0.00]		
Trust	0.10	0.15**	0.016	0.16	-0.015
	[0.10]	[0.03]	[0.81]	[0.29]	[0.96]
Constant	2.68***	2.49***	2.68***	3.63***	3.74***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Observations	285	235	285	68	68
R-squared	0.226	0.247	0.235	0.140	0.150
Estimation	OLS	OLS	OLS	OLS	OLS
Demographic Controls	Yes	Yes	Yes	No	No
Sample	Public	Non-Econ	Public	MPhil	MPhil

Notes: *D(Layers)* is 1 if the participant was randomly assigned the new, layered content in the experiment. D(Economics) is a dummy variable which is 1 if the participant has studied economics as part of a university degree course. BoE Confidence is a numeric variable rating the participant's confidence in the Bank. Demographic controls, available only for the public survey, are separate dummy variables equal to 1 indicating the respondent is Female, English-speaking, British nationality, Student or Fulltime Employed. P-values constructed using robust standard errors are reported in brackets below the coefficient estimates.

Columns (3) and (5) report the estimates of (23). The results of different levels of prior trust on the effect of *D(Layers)* is presented graphically for the two samples in figure

6. The sample estimates are very close across the two samples. In particular, those who have highest existing trust find the new content to be an even bigger improvement.

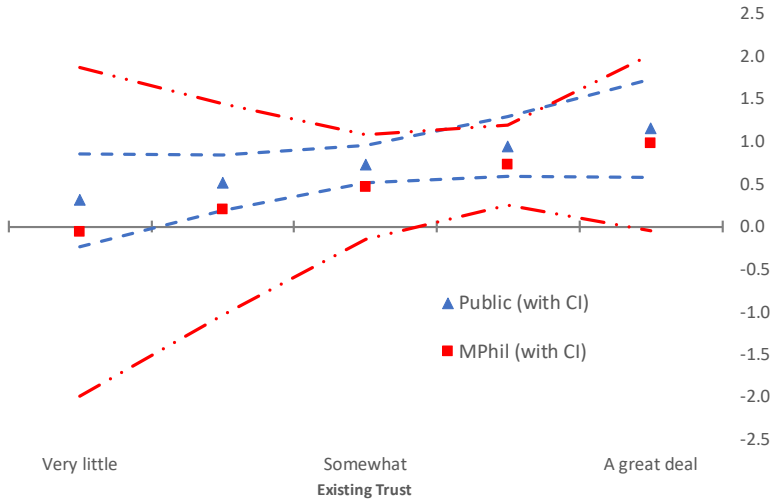

Figure 6: Marginal Effect of $D(Layers)$ on $D(Understand)$ by Trust

Notes: The blue triangles (red squares) show the estimated coefficient from the Public (MPhil) sample of the effect of reading the Layered Content ($D(Layers)$) on reported understanding by different levels of Trust. The lines around the point estimates indicate the 95% confidence intervals.

We now repeat the analysis using the $D(Adjust)$ dummy variable to see if participants brought their expectations into line with the Bank of England forecast. As the dependent variable is a dummy variable, we use a probit model for equations (22) and (23). Table 5 and figure 6 present the results as before. The effect of the more readable communication on expectations differs between the two samples. In the case of the general public survey, layered communication boosts the chance that the participant update her/his beliefs to become more closely aligned with the Bank's forecasts. This effect is more significant for the less trusting (which is the bulk of the public sample). For MPhil students, the average coefficients are possitive, but the results are not statistically significantly.

Finally, table 6 and figure 8 examine whether participants reading the new content tended to develop an improved perception of (trust in) the institution. While the mean effect is not statistically significant in the public survey, it is highly significant in the MPhil sample. The inclusion of the interaction term, as with the regressions on understanding the content, shows the two samples are quite similar. The interaction term highlights that the layered content tends to significantly increase perceptions of those with existing high levels of trust. The different mean estimates seems to reflect the fact that the existing levels of trust are, on average, higher in the MPhil sample. There is, in addtion, a difference whereby the technically-trained MPhil respondents seem to appreciate efforts to "talk to the layperson" more. The takeaway from this is that on-going efforts may be

Table 5: Regression analysis of communication experiment

	(1) D(Adjust)	(2) D(Adjust)	(3) D(Adjust)	(4) D(Adjust)	(5) D(Adjust)
D(Layers)	0.35**	0.43**	0.33*	0.090	-0.16
	[0.04]	[0.02]	[0.06]	[0.78]	[0.76]
Trust x D(Layers)			-0.089		0.37
			[0.64]		[0.50]
D(Economics)	-0.24		-0.24		
	[0.32]		[0.33]		
Trust	-0.11	-0.070	-0.065	0.28	0.044
	[0.28]	[0.51]	[0.63]	[0.26]	[0.91]
Constant	-0.21	0.036	-0.21	-0.81***	-0.66*
	[0.52]	[0.92]	[0.51]	[0.01]	[0.07]
Observations	285	235	285	68	68
Estimation	Probit	Probit	Probit	Probit	Probit
Demographic Controls	Yes	Yes	Yes	No	No
Sample	Public	Non-Econ	Public	MPhil	MPhil

Notes: *D(Layers)* is 1 if the participant was randomly assigned the new, layered content in the experiment. D(Economics) is a dummy variable which is 1 if the participant has studied economics as part of a university degree course. BoE Confidence is a numeric variable rating the participant's confidence in the Bank. Demographic controls, available only for the public survey, are separate dummy variables equal to 1 indicating the respondent is Female, English-speaking, British nationality, Student or Fulltime Employed. P-values constructed using robust standard errors are reported in brackets below the coefficient estimates.

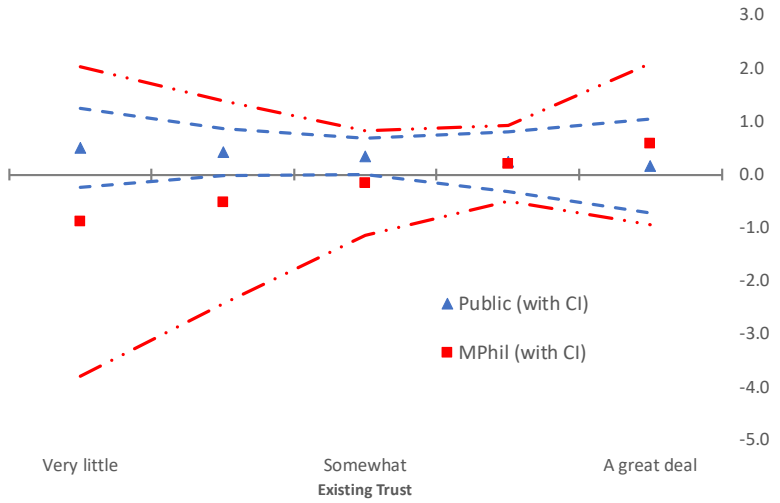

Figure 7: Marginal Effect of $D(Layers)$ on $D(Adjust)$ by Trust

Notes: The blue triangles (red squares) show the estimated coefficient from the Public (MPhil) sample of the effect of reading the Layered Content ($D(Layers)$), by different levels of Trust, on whether the respondent adjusted their expectations in the direction of the Bank of England forecast. The lines around the point estimates indicate the 95% confidence intervals.

Table 6: Regression analysis of communication experiment

	(1) Δ Perception	(2) Δ Perception	(3) Δ Perception	(4) Δ Perception	(5) Δ Perception
D(Layers)	0.083	0.086	0.098	0.35**	0.16
	[0.33]	[0.36]	[0.23]	[0.01]	[0.49]
Trust x D(Layers)			0.12		0.30
			[0.19]		[0.18]
D(Economics)	-0.032		-0.033		
	[0.76]		[0.75]		
Trust	0.15***	0.16***	0.10*	-0.14	-0.31*
	[0.00]	[0.00]	[0.10]	[0.19]	[0.07]
Constant	3.19***	3.13***	3.20***	3.12***	3.23***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Observations	285	235	285	66	66
R-squared	0.055	0.065	0.062	0.111	0.138
Estimation	OLS	OLS	OLS	OLS	OLS
Demographic Controls	Yes	Yes	Yes	No	No
Sample	Public	Non-Econ	Public	MPhil	MPhil

Notes: *D(Layers)* is 1 if the participant was randomly assigned the new, layered content in the experiment. D(Economics) is a dummy variable which is 1 if the participant has studied economics as part of a university degree course. BoE Confidence is a numeric variable rating the participant's confidence in the Bank. Demographic controls, available only for the public survey, are separate dummy variables equal to 1 indicating the respondent is Female, English-speaking, British nationality, Student or Fulltime Employed. P-values constructed using robust standard errors are reported in brackets below the coefficient estimates.

needed to reach and convince those parts of the public most mistrustful of central banks to begin with. This speaks to improved communication alongside improved economics education for this less specialist audience (see below).

Since our original analysis, others have conducted similar work. Also focusing on the Bank of England's introduction of layered content, Bholat et al. (2018) tested four different ways of communicating the February 2018 Inflation Report: (1) the traditional Monetary Policy Summary, (2) the layered content, (3) a reduced text summary, and (4) a relatable summary. The latter two were designed by the joint BIT/Bank of England team. The relatable summary aimed to make the material more relatable to the lives of the participants, as well as expressing costs in absolute rather than relative or growth terms. This relatable summary was found to be most effective at increasing comprehension scores (+42% compared to the traditional Monetary Policy Summary) and it was also the most effective for applicable understanding. For example, readers of it were best able to predict what a basket of groceries costing £100 *should* cost next year based on the information.

In the US context, Coibion et al. (2019) conduct a large (20,000 participant) RCT examining eight different communication types about inflation. They find that reading

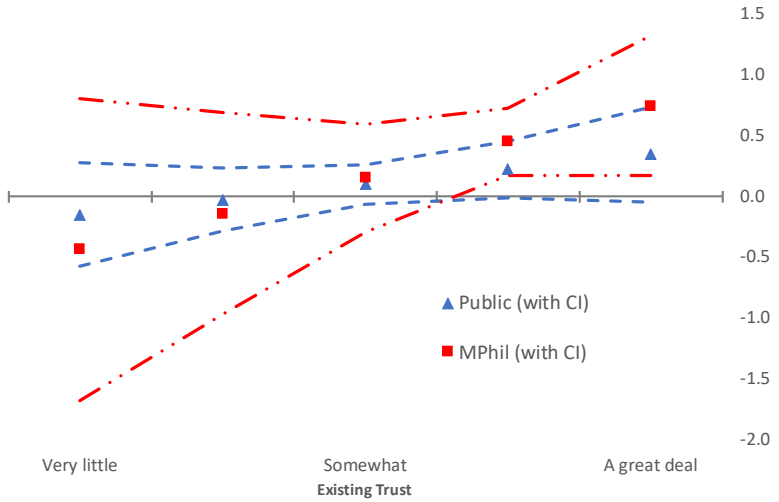

Figure 8: Marginal Effect of $D(Layers)$ on ΔPerceptions by Trust

Notes: The blue triangles (red squares) show the estimated coefficient from the Public (MPhil) sample of the effect of reading the Layered Content ($D(Layers)$) on reported perceptions about the institution by different levels of Trust. The lines around the point estimates indicate the 95% confidence intervals.

the FOMC statement changes inflation expectations by the same as the latest inflation data. The effect is economically significant – households' average inflation forecast is reduced, from a high level, by around 1.2 percentage points. They also found that relying on news intermediaries, such as the media, gives rise to effects that are smaller and less persistent. This is particularly the case for some lower-income, lower-education participants when reading "USA Today". As well as pointing to a need for further research on the role of the media in expectation formation, this also suggests a potentially potent role for direct communication rather than relying on message intermediaries.

Binder and Rodrigue (2018), also in the US context, find that households' long-run inflation forecasts react to communication about the prevailing or recent inflation rate, or the inflation target. This suggests that, for some households, even a very simple message such as the inflation target could be very powerful in acnhoring expectations, but only if those households can be reached.

4.2 Engagement

Clearly-explained communication may count for nothing if people don't engage with these communication in the first place. The effects that we find in our experiment come about *after* participants were incentivised to engage. But how likely is it that people do engage? To give a sense of the challenge facing central banks if they stick with their traditional medium of explanation, we asked the sample participants in our survey of the public in

November 2017 about their familiarity with the IR. Most participants (66%) claimed to have heard of the IR, although less than 6% had ever read it (and with only 1 participant who claimed to read it regularly). The remaining 34% had never heard of it.

For these reasons, engagement is, like explanation, core to the objectives of the central bank when improving its policy effectiveness. And as with explanation, engagement in theory is easier than in practice. Moreover, a key message of our trust model is that simple communication on its own might not be enough. To build and maintain trust, it might require extra action. Engagement in itself might contribute to building and, in particular, maintaining that trust. Or, put differently, trust is less likely to depreciate (or evaporate) the greater the degree of engagement.

One aspect of our model to consider is what happens to trust (and therefore the potential for future engagement) if a household doesn't read any of the information? In the model we make the assumption that unengaged households are never surprised and so their trust doesn't change. It is also assumed that, once engaged but surprised, trust is lost and gone forever. Both are likely to be far too strong. There may be a risk that if the central bank is not communicating with individuals then their trust might fall anyway. This is especially true in an era of social media engagements targeted at previously unengaged areas of the population.

An alternative formulation would be to acknowledge that, if the central bank isn't talking to people, someone else will fill the void with possibly noisier messages. A way to model this would be to follow the application of Bernoulli's model of infectious diseases to social dynamics as in Burnside et al. (2016).[23] While they apply it to the housing market, the idea in terms of central bank communication is that of being the narrative entrepreneur who can help people to make sense of the economy, and form reasonable expectations. The central bank, by engaging and educating people, can help households to form better expectations. If a household has no engagement with the central bank, then they do not receive the best guidance and are more susceptible to believing other (nosier) opinions about the outlook and the institution. This would admit a stronger role for engagement because, absent engagement, the baseline could be progressively less well- informed opinions on central banks.

4.3 Media and the Narrative Channel

Shiller (2017) stresses the important role that "popular narratives" can play in determining behaviour in the macroeconomy. One advantage could be that simplified content enables greater coverage and penetration of the policy narrative. And this better under-

[23]Bailey et al. (2017) also discuss the role of social contagion in driving housing market behaviour.

standing of the factors driving the decisions could help to reduce the incidence of such self-reinforcing expectational swings in sentiment and behaviour.

Communications may need to be simple, relevant and story-based to become convincing and credible to a wider audience. Traditional central bank communications tend to fail on all three fronts. Therefore, to be more engaging, central banks need to create a context. They need to create stories. The availability of simplified central bank messaging may also help traditional information intermediaries, such as the mainstream media, which further facilitates the process of message transmission to a wider audience.

A risk, related to concerns in Morris and Shin (2002), is that such simple messages create an incentive for people to stop investing in their own information collection. This is a problem because, if it is common to all households, then any noise in it leads to inefficient variation in consumption.

4.3.1 Social media: Opportunity and challenges

New media channels, especially but not exclusively social media, provide new opportunities and new challenges.[24] The obvious benefit is that it is likely easier to target the uninformed because many of the uninformed view large amounts of news material on Facebook, Twitter, Instagram, Youtube and other social media every day. The challenge is that, in a saturated market for news and stories, how can the central bank compete with cat videos?

Most central banks are now on social media platforms. McMahon et al. (2018) reports followers data from a number of major central banks. While some have large followers, none have more than 0.5% of their national population. To put this in context, the US Federal Reserve has around 0.5m followers, while US President Donald Trump has over 61m followers or nearly 20% of the US population.[25] The most followed accounts include Katy Perry (108m followers), Barack Obama (107m followers), Justin Bieber (106m followers), Rihanna (92m followers) and Taylor Swift (84m followers).

Nonetheless, easier to understand communication should improve the reach of the central bank's communication. To examine this, we compare reach for the November 2017 Inflation Report (which had the layered content but also the UK's first rate rise for a decade) with two counter-factual events:

1. August 2017 (previous) IR - this is without layered content but also without any major monetary news
2. August 2016 IR - this also had no layered content but is associated with significant monetary news (a 25bps rate reduction and an additional QE package).

[24] See also Binder (2017) on the implications of new media.

[25] Of course, in both cases some followers will be international.

Table 7: Analysis of IR Reach

	August 2016 IR	August 2017 IR	November 2017 IR
Website hits	16,600	12,460	30,900
o/w Layer 2	n/a	n/a	16,200
Tweets	1,745	320	1,566
o/w Layer 1	n/a	n/a	845

Notes: Bank of England data. The Tweets numbers represent the total number of retweets of, quotes of and replies to all BoE tweets relating to the Inflation Report and Bank rate announcement in the time period up to 24 hours after each period's announcements. Tweets about the Inflation Report from Twitter accounts other than the Bank's which are not retweets of, quotes of or replies to BoE tweets are not included. Layer 1 refers to just a tweet of the basic announcement that Bank Rate went up by 25bps. Layer 2 is the Inflation Report Visual Summary webpage content on www.inflationreport.co.uk.

Table 7 summarises the website and Twitter activity associated with the three events, over the course of the subsequent 24 hours. There was a large increase in direct website traffic associated with the November 2017 IR. Even relative to August 2016, website hits almost doubled. Moreover, almost all of this increase was associated with hits on the new, simplified content, with hits on the existing technical material largely unaffected. This is consistent with the new communications having achieved a somewhat broader reach with a somewhat different audience.[26]

An analysis of social media engagement, measured by Twitter traffic, suggests a more nuanced picture. Numbers of tweets and retweets associated with the IR were highest in August 2016. Nonetheless, Twitter traffic was 4.9 times higher in November 2017 than in August, and the Bank itself issued more than twice as many tweets in August 2016 than in November 2017.

An alternative window on social window engagement is provided by looking at the Twitter networks associated with the monetary policy and IR events. Figure 9 shows these in the subsequent 24-hour period, with nodes sized according to the number of followers and colour-coded according to different types of institution.[27]

The August 2016 and November 2017 Twitter networks are similar in their reach and penetration. By contrast, the August 2016 network involves significantly fewer tweets in total and the network was simpler and more sparse. There is also evidence of far less media engagement. Overall, this preliminary analysis is a nuanced good news message.

[26]Our data does not allow us to show that the extra hits on the website hosting the new layers (www.inflationreport.co.uk) were unique. However, the majority of hits to the new microsite came via paid search, which is unlikely to be relevant for the usual IR readers. Moreover, we can measure the click-through from the main IR page to the new microsite (and vice versa) and it is a very small percentage of the total hits on each; this suggests the users are different.

[27]Red indicates Bank of England tweets. Blue and green indicate, respectively, official UK and global media organisations while turquoise represents the accounts of specific journalists. Purple are politicians' accounts, orange are economists' accounts and yellow are celebrity' accounts. Other accounts are grey.

It is clear, however, that monetary policy news itself, rather than the means by which it is communicated, is the largest single factor determining the reach of central bank communications. This makes detecting the marginal impact of changes to communications strategy problematic using traffic data alone.

Looking at the time-series data on both website hits, two points stand out. First, hits on the visual summary have remained about constant each IR (November 2018 is an exception). This is very positive given the additional marketing effort that accompanied the first Visual Summary. Second, the data on Twitter retweets and the hits to the Monetary Policy Summary website make clear that it is interest rate changes that lead to greatest engagement. Even the May 2018 surprise decision to not increase rates did not lead to the same interest.

4.3.2 Direct engagement: Business contacts and citizens panels

Central banks can also engage people in a more direct way. Central banks regularly engage business contacts through established networks across the country. For example, the Bank of England has a network of 12 Regional Agencies across the UK, with regular contact with almost 10,000 companies. These hundreds of engagements each month allow for a two-way flow of information. The information gathered is fed into the policy process and senior policy makers often join the agents on visits.

Policymakers now participate in a range of bespoke engagements, designed and delivered in partnership with organisations such as charities, social enterprises and faith groups. These groups often represent some of the hardest-to-reach groups in society including, for example, those living in significant poverty, facing severe debt issues, refugees, the homeless and even prisoners.

And the Bank has started to setting up citizens' panels. The idea is to assemble, via a publicity campaign in local print and social media, a group of around 20-24 people in each of the 12 agency regions and to hold two meetings a year. The people, who are selected to be broadly representative of the local population, will have a regular chance to explain their worries and concerns, as well as to discuss current policy issues.

Other central banks are using social media for such attempts to generate direct engagement. For example, Stefan Ingves, Governor of the Riksbank, takes part in regular online Q&As, as does Minneapolis Fed President Neel Kashkari on Twitter with his "#AskNeel" sessions. The recent "Fed Listens conference" is another example.

Monetary policy decisions are largely an exercise in information aggregation (Hansen et al. 2014) and policy makers who bring a broader coverage of information likely become more influential (as in Hansen et al. (2017)). Is there any evidence that listening to a wider audience leads to a change in policy? Perhaps not directly, but such information

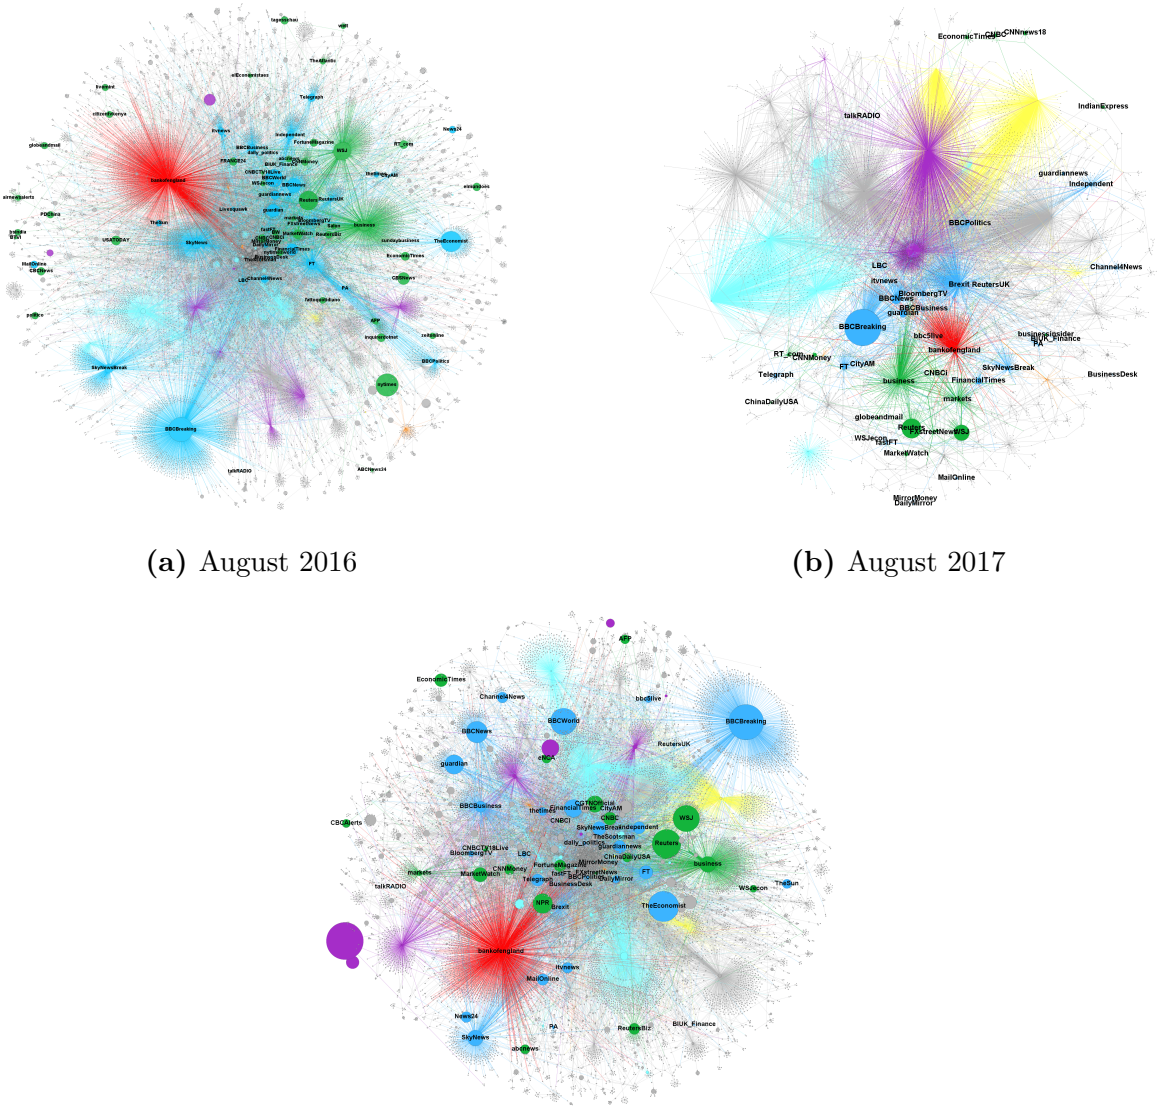

(a) August 2016

(b) August 2017

(c) November 2017

Figure 9: Twitter network in 24 hours after the IR release

Notes: Bank of England data. Account nodes are sized according to the number of followers they have. Red indicates Bank of England tweets. Blue and green indicate, respectively, official UK and global media organisations while turquoise represents the accounts of specific journalists. Purple are politicians' accounts, orange are economists' accounts and yellow are celebrity' accounts. Other accounts are grey.

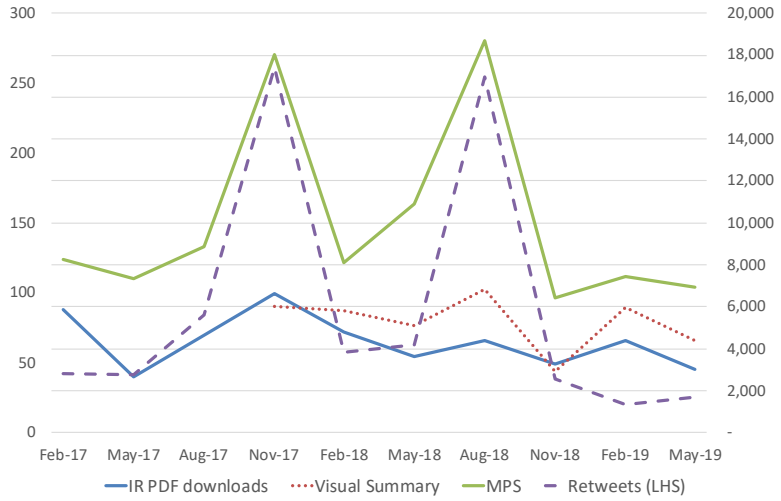

Figure 10: Website Hits and Twitter Retweets around the IR launch

Notes: Bank of England data.

can help to contextualise the more traditional data and highlight potential solutions to data puzzles. It may also help policy makers to ensure their communications are conveyed in a way that addresses peoples' concerns.

In addition to information, these engagements should help build trust. In the model, we show that a stronger positive reaction to engagements could help central banks improve welfare. Figure 11 shows two examples of this from the Bank of England's direct engagements. Figure 11a shows the results of a survey of their business contacts carried out by the Bank of England's Regional Agents immediately following release of the November 2017 IR. The survey asked specifically about the new layered content. Overall, more than 70% of respondents felt the new layered summary helped them to better understand the messages of the IR. Moreover, as figure 11a shows (with results broken down by company size) around 60% of respondents felt the new communication improved either 'somewhat' or 'a lot' their perception of the Bank.

Figure 11b shows aggregated results of surveys carried out following a few of the Bank's Citizens' Panels. Asked to rate how the session (a) increased your knowledge of the Bank's responsibilities, (b) increased your trust in the Bank, and (c) improved your understanding of the economy. The evidence is that the events have helped on all counts: 90% either 'somewhat agree' or 'strongly agree' that the event increased their knowledge of the Bank's responsibilities, and confidence; the proportion is 75% for increasing trust in the Bank; and 76% believe it improved their knowledge. Of course, such survey results should be interpreted carefully due to the possibility of self-selection by companies, and likely self-selection by citizens' panels participants.

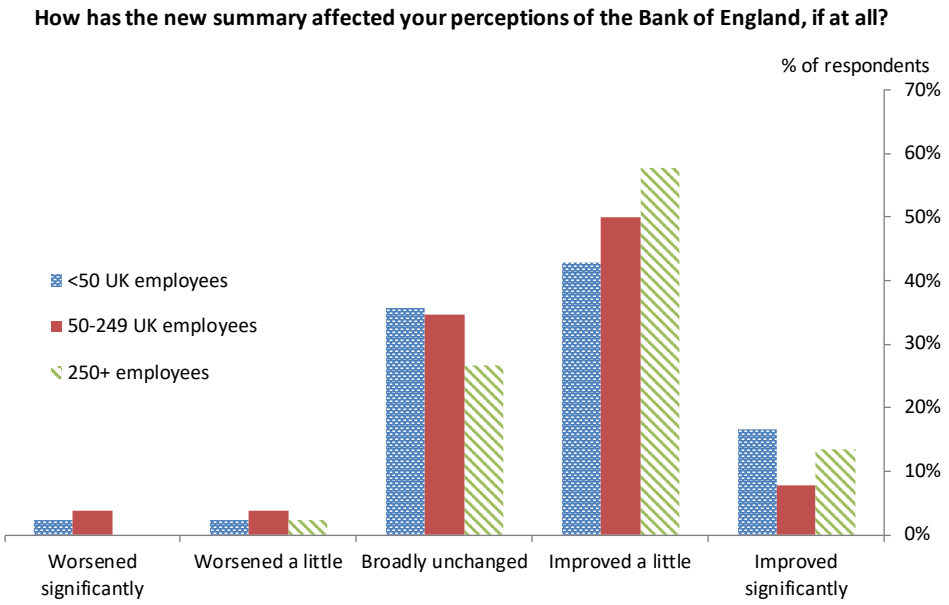

(a) Agents Survey of Regional Contacts

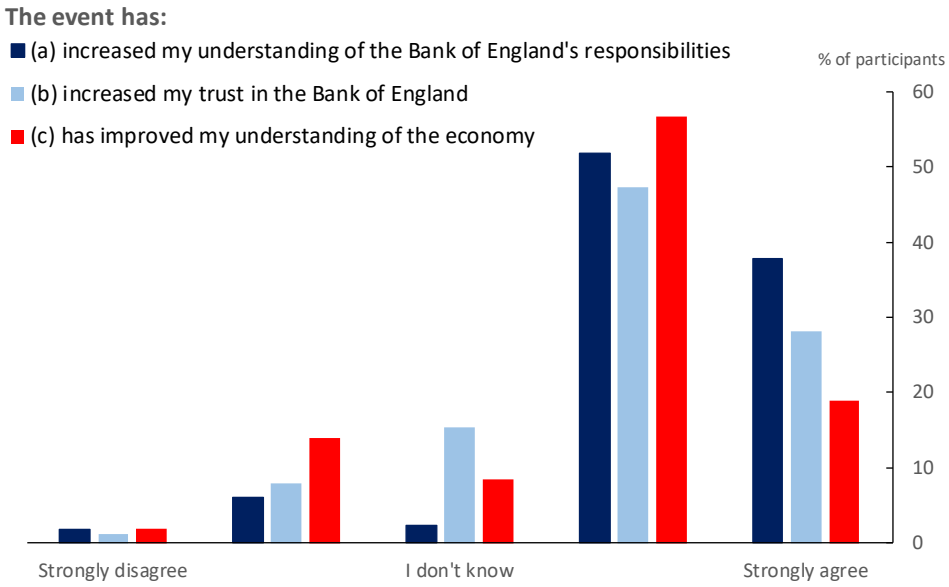

(b) Citizens' Panels

Figure 11: Effects of Direct Engagement and Simple Communication

Notes: Survey of regional corporate contacts carried out by the Bank of England's regional agents in November 2017 and Bank of England 2019 Citizens' Panels.

Finally, these engagements could actually have further reach than is easy to measure, e.g. the Bank of England uses local media to promote its activity in the regions. As in the example of social dynamics discussed above, central banks need torchbearers to carry the story and narrative forward. Direct engagement may help to provide such torchbearers in the local economy.

4.4 Education

In the model, a key challenge arises from the fact that the households that are newly-engaged by the simply communication fail to understand the complexity and stochastic nature of the economy. Better education may reduce the costs of engagement and reduce the reaction to surprises; it may slow the depreciation rate of trust (as in 3.6.3 above). Also the evidence has suggested that those with a better understanding have higher levels of trust which, in the model, would translate into better engagement and higher welfare.

The central bank has a primary role for educating the public on its framework, strategy, analysis and policy decisions. This entails education on both the high frequency and low frequency aspects that the central bank communicates on. Better informed agents may, at a higher frequency, form more appropriate expectations for inflation, output and interest rates. But, equally, high levels of trust and understanding may help to sustain democratic legitimacy as an independent institution and improve the resilience of trust. In this section, we discuss attempts to educate existing economic decision makers, leaving efforts to educate younger audiences to the next section.

Even engaged and technical audiences need regular educational briefing. This includes briefings with, notes for and videos aimed at businesses and major banks explaining new ideas on the economy. This is especially necessary when the central bank sees fit to deploy new tools, or to vary how it will operate the existing ones. Such decisions now always come with additional explanation and extra materials.

But there is a larger population of less engaged and less technical decision makers. One example of how education influences the high frequency nature of the central bank's communication strategy concerns understanding of economic concepts. Key words such as "inflation" and "GDP", which are central to policy discussions, are understood by only small minorities of the general public (Haldane 2017). Focus groups highlight, therefore, that the public rarely understand there may be a relationship between inflation and unemployment. This makes it clear that explanation is linked to the ability to engage which itself depends on the extent of successful past education.

One reaction by a number of central banks, as already discussed, is to adapt their communications strategies to improve their reach to the general public through more-accessible language and more direct engagement. The other is the increasing provision

of videos such as those explaining the decisions made, or simply videos explaining recent issues or research in layman's terms. Other resources, aimed at educating on the lower frequency dimensions of monetary policy, include guides to how the economy and monetary policy interact, and the mechanisms that are at play. Specific examples include the Federal Reserve Bank of St Louis' "In Plain English: Making Sense of the Federal Reserve" material, the Bank of England's "Knowledge Bank: The economy made simple" website and the ECB's "The ECB Explains".

Aimed at existing college students or graduates, the Fed also hosts videos of "Chairman Bernanke's College Lecture Series". These are four lectures delivered in March 2012 by Ben Bernanke (then Fed Chairman) about the Federal Reserve and the financial crisis that emerged in 2007.

As is the case with explanation, a big challenge in educating household and business decision makers is engagement. This is particularly tricky when there is a large population of people who don't understand how the aggregate economy and monetary work, but they think that they do. At least this shows that people want to understand. But how do we feed their interest? Where is the monetary policy equivalent of Sir David Attenborough (the nature documentary maker) to succeed in creating widespread wonder in how central banks work? The Bank of England has recently been the subject of a two-part, behind-the-scenes documentary on national TV in the UK. Below we also discuss the Bank Of Jamaica's attempts at engagement using reggae music videos.

5 Lower Frequency Monetary Communication

While the focus of this paper has been on the decision of central banks to communicate at a relatively high frequency, the last section made clear that educational efforts don't have as clear a distinction between high and low frequency. And central banks must also communicate at a lower frequency. They must explain their framework and, where appropriate, target, and they must engage and educate them to understand what they do and why. Here we briefly examine some of the way in which low frequency communication is also about the 3 E's and give some examples of the activities of central banks in each regard.

5.1 Explanation: Inflation Targets

The widespread adoption of inflation targeting since the Reserve Bank of New Zealand did it in 1990 can be viewed as a communication tool. The idea was that indirect targets such as monetary rules or exchange rate targets didn't provide the majority of people

with a sufficient nominal anchor. Inflation targets, it was hoped, would be easier to understand and this has largely turned out to be true. For example, Crowe (2010) provides cross-country evidence on the usefulness of an inflation target in anchoring inflation expectations. And in the case of the US, Binder (2017) shows that the Federal Reserve's adoption of a formal 2% inflation target contributed to better anchored households' inflation expectations. This work also relates to issues of the twin deficits as the analysis also shows that better-informed households' expectations were more affected (in terms of becoming better anchored) relative to less-informed households.

The importance of low frequency communication cannot be over-stated. Coibion et al. (2019), discussed above, find that communicating the Fed's inflation target has the same statistically-significant effect on households inflation expectations as communicating the FOMC's inflation forecast or the FOMC statement.

One important issue that affects communication on low frequency issues is how to communicate changes to existing frameworks. While the above analysis suggests that adopting an inflation target can aid the management of inflation expectations, it is less clear how easily established inflation targeting regimes could be changed. This has come to be discussed because in an era of low nominal rates, higher inflation targets are seen by some as low-hanging fruit to build a buffer away from hitting the Effective Lower Bound (ELB) again soon. This requires skeptical a credible shift such that expectations move and become reanchored at the new target.

One difficulty with this is that changing the regime may also signal that the regime *can* change. In the UK, for example, the current inflation targetting framework with an operationally-independent central bank is over 20 years old. In that time there has been one variation in the framework; in December 2003, the inflation index used to calculate the measure of inflation in the target was changed from RPIX to CPI. In line with methodological differences in the two indices, the target changed from 2.5% RPIX to 2% CPI. It was emphasised that this was a non-change.

Such care with credible and established regimes is warranted. The US Federal Reserve has recently announced a review of its monetary framework. However, Vice Chair of the Federal Reserve Board and FOMC member Richard Clarida suggested that it will be "more likely to produce evolution, not a revolution" (April 2019 speech).

5.2 Engagement: Recent Novel Approaches

As with higher frequency analysis, it is important for the impact of the communication that households and businesses engage with it. They need to read or see it and they need to take the message on board. Reis (2011) examines a rational inattention model in which a central bank must decide when to make public a low frequency announcement such as

change in the monetary framework. His analysis emphasises that economic agents trade-off being more informed about today (and responding better to today's environment) and being better informed about the future (and so preparing better for the change). The central bank also needs to balance the clarity of the message it can send (which grows over time) with the risk that the public will inefficiently coordinate on its announcement.

In practice, new technology has provided a mechanism for direct engagement on these lower frequency messages too. For example, the ECB has used popular YouTuber Simon Clark to explain what the central and specifically the ECB is.

The Bank of Jamaica's (BoJ) move from a focus on control of the exchange rate to "full-fledged inflation targeting' (FFIT)' has been widely discussed for the innovative ways in which the BoJ has communicated the move with the public.[28]

The BoJ faced a public that was more familiar with a policy focus on the exchange rate. In order to speak the language of the public, they have released a serious of videos including top reggae stars (such as Tarrus Riley) and inflation control to the baseline in reggae music. Through their 'Low, Stable, Predictable inflation' narrative, made available on TV, radio and on social media platforms such as YouTube and Twitter, they hope to establish both support for and understanding of their new framework.

While this is a great example of an innovative engagement effort with a wider audience, the benefits are more difficult to measure. Businesses' perception of the authorities' control of inflation, calculated as 100 plus the number of satisfied survey respondents minus the number of dissatisfied respondents, decreased in April 2019 although it has generally been increasing since the move toward FFIT. But this also coincides with the underlying state of the economy. Further analysis will be warranted to see if this campaign yields longer term benefits and trust in the FFIT framework.

5.3 Education: From Comics to Classrooms

As pointed out in section 4.4, the distinction between high and low frequency communication is less pertinent. Since most of the discussion above concerned both high and low frequency objectives toward people who need to learn now, here we discuss some of the efforts of central banks be involved in educating younger audiences in a more gradual fashion before they become economic decision makers. This can be justified by realising that children who understand the economy and the role of the central bank from an earlier age will be less susceptible to attempts to undermine central bank independence.

[28]The Bank of Jamaica (Amendment) Bill, currently under review by a Joint Select Committee of Parliament, will amend that The Bank of Jamaica Act to clarify its mandate as well as some other changes. This includes clarification that "The mandate of the Bank is the maintenance of price stability and financial system stability with the primary objective being the maintenance of price stability".

Also, today's youth are tomorrow's politicians and decision makers.

When we think of education of young people, it is not obvious that the central bank is the main entity with responsibility. Decisions such as how much to teach about interest rates in school rest, typically, with educational boards and the Ministry for Education.

But central banks have taken on the role of providing, in addition to the videos and other engagement mechanisms discussed above, free classroom materials. These range from resources about how the economy works, to to what the central bank does. Many central banks split the resources into material for different target age groups. The Federal Reserve Bank of New York has even developed a series of comic books to describe the Fed, monetary policy and how money works.

Many central banks also offer competitions for school and/or university students. These events raise awareness of the central bank and its objectives, as well as provide opportunities for personal development for the participants. The Bank of England has number of efforts in this direction. For example, 'EconoME' is a free education resource created jointly by educational experts and the Bank. It is designed to help young people aged 11 to 16 understand the economy better and provide them with the analytical skills to make informed decisions. The Bank also provides inflation and interest calculators to help households with financial planning. It previously ran a national monetary policy competition across UK schools (Target 2.0).

All of these are potentially useful exercises to engage, explain and educate. Of course, central banks are constrained by what resources they have available. Two activities will help focus the allocation of resources in the future. First, listening to a wide array of stakeholder is one way to learn where to target the educational efforts. Second, careful examination and appraisal of the successes and failures of different approaches should be undertaken.

6 Conclusion

The last decade has seen central banks respond to the challenges posed by the fallout of the financial crisis by engaging more and more with a broader audience about monetary policy. Providing clarity is likely important but this paper argues that explanation through simplified communication may, alone, be necessary but not sufficient. Central banks need complementary efforts in engagement and education.

There is much still to be done to understand the optimal design and use of communication with the general public. This includes further research, and further practical experimentation in terms of communication with the public. Such experiments should be scrutinised for the lessons of what worked, what didn't and why. The CEPR has recently

initiated a Research Policy Network, together with the European Central Bank and with membership of many central banks, academics, journalists and professional economists. The objective is to encourage such research efforts and the dissemination of findings to both researchers, those involved in communication in central banks, journalists as well as other interested stakeholders.

Moreover, most central banks now have remits that extend beyond monetary policy. The design of communication strategies are likely specific to each objective especially since the audiences are possibly different. For example, the communication about prudential policies may give rise to even tougher challenges. This is because the policies aims might be harder to communicate and the tools available are more varied both within and across countries.

Central banks must remain steadfast in their efforts to reach a broader audience. Given the necessary degree of trial and error, there will be mistakes. But success should not be measured by the ability to reach everyone, but rather by engaging even limited audiences beyond the current small core audience of technical specialists and information intermediaries.

References

Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., and Zafar, B. (2015). Inflation expectations and behavior: Do survey respondents act on their beliefs? *International Economic Review*, 56(2):505–536.

Bachmann, R., Berg, T. O., and Sims, E. R. (2015). Inflation expectations and readiness to spend: Cross-sectional evidence. *American Economic Journal: Economic Policy*, 7(1):1–35.

Bailey, M., Cao, R., Kuchler, T., and Ströbel, J. (2017). The Economic Effects Of Social Networks: Evidence From The Housing Market. *Journal of Political Economy*, Forthcoming.

Bholat, D., Broughton, N., Parker, A., Ter Meer, J., and Walczak, E. (2018). Enhancing central bank communications with behavioural insights. Bank of England working papers 750, Bank of England.

Binder, C. (2017). Fed speak on main street: Central bank communication and household expectations. *Journal of Macroeconomics*, 52:238 – 251.

Binder, C. and Rodrigue, A. (2018). Household Informedness and Long-Run Inflation Expectations: Experimental Evidence. *Southern Economic Journal*, 85(2):580–598.

Blinder, A. (2009). Talking about monetary policy: the virtues (and vice?) of central bank communication. BIS Working Papers 274, Bank for International Settlements.

Blinder, A. S. (2018). Through a crystal ball darkly: The future of monetary policy communication. *AEA Papers and Proceedings*, 108:567–71.

Blinder, A. S., Ehrmann, M., Fratzscher, M., Haan, J. D., and Jansen, D.-J. (2008). Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence. *Journal of Economic Literature, American Economic Association*, 46(4):910–45.

Braun, B. (2016). Speaking to the people? Money, trust, and central bank legitimacy in the age of quantitative easing. *Review of International Political Economy*, 23(6):1064–1092.

Braun, B. (2018). *Central bank planning? Unconventional monetary policy and the price of bending the yield curve*, pages 194 – 216. Oxford University Press.

Burnside, C., Eichenbaum, M., and Rebelo, S. (2016). Understanding Booms and Busts in Housing Markets. *Journal of Political Economy*, 124(4):1088–1147.

Coenen, G., Ehrmann, M., Gaballo, G., Hoffmann, P., Nakov, A., Nardelli, S., Persson, E., and Strasser, G. (2017). Communication of monetary policy in unconventional times. Working Paper Series 2080, European Central Bank.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2019). Monetary Policy Communications and their Effects on Household Inflation Expectations. NBER Working Papers 25482, National Bureau of Economic Research, Inc.

Crowe, C. (2010). Testing the transparency benefits of inflation targeting: Evidence from private sector forecasts. *Journal of Monetary Economics*, 57(2):226–232.

Duca, I. A., Kenny, G., and Reuter, A. (2018). Inflation Expectations, Consumption and the Lower Bound: Micro Evidence from a Large EU Survey.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press, Princeton, USA and Oxford, UK, 2 edition.

Goodland, R. (1995). The concept of environmental sustainability. *Annual Review of Ecology and Systematics*, 26:1–24.

Haldane, A., Macaulay, A., and McMahon, M. (2019). Central Bank Communication in a World of Rational Inattention. Technical report.

Haldane, A. and McMahon, M. (2018). Central bank communications and the general public. *AEA Papers and Proceedings*, 108:578–83.

Haldane, A. G. (2017). A Little More Conversation, A Little Less Action. Speech given at the Federal Reserve Bank of San Francisco Macroeconomics and Monetary Policy Conference, 31 March 2017.

Hansen, S., McMahon, M., and Prat, A. (2017). Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*. Forthcoming.

Hansen, S., McMahon, M., and Velasco, C. (2014). How Experts Decide: Preferences or Private Assessments on a Monetary Policy Committee? *Journal of Monetary Economics*.

Jost, A. (2017). Is Monetary Policy Too Complex for the Public? Evidence from the UK. Technical report.

Kryvtsov, O. and Petersen, L. (2013). Expectations and Monetary Policy: Experimental Evidence. Technical report.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

McMahon, M., Schipke, A., and Li, X. (2018). China's Monetary Policy Communication: Frameworks, Impact, and Recommendations. IMF Working Papers 18/244, International Monetary Fund.

Morris, S. and Shin, H. S. (2002). Social Value of Public Information. *American Economic Review*, 92(5):1521–1534.

Petersen, L., Amano, R., and Kryvtsov, O. (2014). Recent Developments in Experimental Macroeconomics. Discussion Papers dp14-08, Department of Economics, Simon Fraser University.

Reis, R. (2011). When Should Policymakers Make Announcements? Mimeograph.

Rockall, E. (2018). Could knowledge about Central banks impact households' expectations? `https://bankunderground.co.uk/2018/10/22/could-knowledge-about-central-banks-impact-households-expectations/`.

Shiller, R. J. (2017). Narrative Economics. *American Economic Review*, 107(4):967–1004.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Skingsley, C. (2019). How to open up a central bank – and why stay open. Speech given at the Annual Research Conference of the National Bank of Ukraine. 23 May, 2019.

Vellekoop, N. and Wiederholt, M. (2018). Inflation Expectations and Choices of Households.

Woodford, M. (2001). Monetary policy in the information economy. *Proceedings - Economic Policy Symposium - Jackson Hole*, pages 297–370.

A Appendix: Expected utility loss from inattention

This derivation follows the steps of appendix D in Maćkowiak and Wiederholt (2015). Define $\tilde{U}$ as the log-quadratic approximation of the discounted household utility function U, and let $\tilde{U}^*$ be the equivalent for the fully informed household. The approximation is taken about the steady state. Note that since in steady state all shocks equal the household prior beliefs of zero, inattention plays no role in the determination of the steady state.

It can be shown that the expected utility loss from inattention is:

$$\tilde{U}^* - \tilde{U} = -\mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2}(x_t - x_t^*) H_0 (x_t - x_t^*)' + (x_t - x_t^*) H_1 (x_{t+1} - x_{t+1}^*)' \right] \tag{24}$$

Where $x_t = [b_t \quad n_t]$, and:

$$H_0 = \begin{bmatrix} U_{bb} & U_{bn} \\ U_{bn} & U_{nn} \end{bmatrix} \tag{25}$$

$$H_1 = \begin{bmatrix} U_{bb_1} & U_{bn_1} \\ U_{nb_1} & U_{nn_1} \end{bmatrix} \tag{26}$$

U_{ij} is the second derivative of discounted utility U (before approximation) with respect to i and j, evaluated at the steady state. Note that lower case x_t is the log-deviation of x from steady state in period t. Furthermore, denote the steady state of x by the capital X, and let $\tilde{x}_t = x_t - x_t^*$. U_{ij_1} is the second derivative with respect to i_t j_{t+1}.

In this particular model, substituting in for U_{ij} and substituting out for bonds using the budget constraint, we have:

$$\begin{aligned} \tilde{U}^* - \tilde{U} = -\mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \Bigg[& -\frac{1}{2}\sigma B^2 C^{-\sigma-1} \left(1 + \frac{1}{\beta}\right) \tilde{b}_t^2 + \sigma W N B C^{-\sigma-1} \tilde{b}_t \tilde{n}_t - \frac{\varphi}{2} N^{1+\varphi} \tilde{n}_t^2 \\ & - \frac{\sigma}{2} W^2 N^2 C^{-\sigma-1} \tilde{n}_t^2 + \sigma B^2 C^{-\sigma-1} \tilde{b}_t \tilde{b}_{t+1} - \sigma B W N C^{-\sigma-1} \tilde{b}_t \tilde{n}_{t+1} \Bigg] \end{aligned} \tag{27}$$

Factorising:

$$\begin{aligned} \tilde{U}^* - \tilde{U} = \mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \Bigg[& \frac{\varphi}{2} N^{1+\varphi} \tilde{n}_t^2 - C^{-\sigma-1} \Bigg(-\frac{1}{2}\sigma B^2 \left(1 + \frac{1}{\beta}\right) \tilde{b}_t^2 + \sigma W N B \tilde{b}_t \tilde{n}_t \\ & - \frac{\sigma}{2} W^2 N^2 \tilde{n}_t^2 + \sigma B^2 \tilde{b}_t \tilde{b}_{t+1} - \sigma B W N \tilde{b}_t \tilde{n}_{t+1} \Bigg) \Bigg] \end{aligned} \tag{28}$$

Now define three new variables:

$$\Delta_t = B\tilde{b}_t \tag{29}$$

$$\Delta_{c,t} = \frac{1}{\beta}\Delta_{c,t-1} - C\tilde{c}_t \tag{30}$$

$$\Delta_{n,t} = \frac{1}{\beta}\Delta_{n,t-1} + WN\tilde{n}_t \tag{31}$$

The log-linearised budget constraint implies:

$$C\tilde{c}_t + B\tilde{b}_t = \frac{1}{\beta}B\tilde{b}_{t-1} + WN\tilde{n}_t \tag{32}$$

From this, we obtain:

$$\Delta_t = \Delta_{c,t} + \Delta_{n,t} \tag{33}$$

Taking the term in round brackets in equation 28 we substitute out for $\tilde{b}$ and $\tilde{n}$ using these new variables to obtain:

$$-\frac{1}{2}\sigma\left(1+\frac{1}{\beta}\right)(\Delta_{c,t}+\Delta_{n,t})^2 + \sigma(\Delta_{c,t}+\Delta_{n,t})\left(\Delta_{n,t} - \frac{1}{\beta}\Delta_{n,t-1}\right) - \frac{1}{2}\sigma\left(\Delta_{n,t} - \frac{1}{\beta}\Delta_{n,t-1}\right)^2 + \sigma(\Delta_{c,t}+\Delta_{n,t})(\Delta_{c,t+1}+\Delta_{n,t+1}) - \sigma(\Delta_{c,t}+\Delta_{n,t})\left(\Delta_{n,t+1} - \frac{1}{\beta}\Delta_{n,t}\right) \tag{34}$$

Expanding the brackets and cancelling terms we obtain:

$$-\frac{1}{2}\sigma\left(1+\frac{1}{\beta}\right)\Delta_{c,t}^2 + \sigma\Delta_{c,t}\Delta_{c,t+1} + \sigma\left(\Delta_{n,t}\Delta_{c,t+1} - \frac{1}{\beta}\Delta_{n,t-1}\Delta_{c,t}\right) + \frac{1}{2\beta}\sigma\left(\Delta_{n,t}^2 - \frac{1}{\beta}\Delta_{n,t-1}^2\right) \tag{35}$$

Now we take the first two terms of this expression and write them as:

$$-\frac{\sigma}{2}\Delta_{c,t}^2 - \frac{\sigma}{2\beta}\Delta_{c,t}^2 + \sigma\Delta_{c,t}\Delta_{c,t+1} \tag{36}$$

Substitute out for $\Delta_{c,t}$ in the first term of this, and for $\Delta_{c,t+1}$ in the third term, using equation 30, to obtain:

$$-\frac{\sigma}{2}\left(\frac{1}{\beta^2}\Delta_{c,t-1}^2 - \frac{2C}{\beta}\tilde{c}_t\Delta_{c,t-1} + C^2\tilde{c}_t^2\right) - \frac{\sigma}{2\beta}\Delta_{c,t}^2 + \frac{\sigma}{\beta}\Delta_{c,t}^2 - \sigma C\tilde{c}_{t+1}\Delta_{c,t} \tag{37}$$

Rearranging:

$$-\frac{\sigma C^2}{2}\tilde{c}_t^2 + \frac{\sigma}{2\beta}\left(\Delta_{c,t}^2 - \frac{1}{\beta}\Delta_{c,t-1}^2\right) - \sigma C\tilde{c}_{t+1}\Delta_{c,t} + \frac{\sigma C}{\beta}\tilde{c}_t\Delta_{c,t-1} \tag{38}$$

Using these expressions the utility loss from inattention becomes:

$$\tilde{U}^* - \tilde{U} = \mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \left[\frac{\varphi}{2} N^{1+\varphi} \tilde{n}_t^2 + \frac{\sigma C^{1-\sigma}}{2} \tilde{c}_t^2 \right.$$
$$- C^{-\sigma-1} \left(\frac{\sigma}{2\beta} \left(\Delta_{c,t}^2 - \frac{1}{\beta} \Delta_{c,t-1}^2 \right) - \sigma C \tilde{c}_{t+1} \Delta_{c,t} + \frac{\sigma C}{\beta} \tilde{c}_t \Delta_{c,t-1} + \right.$$
$$\left. \left. \sigma \left(\Delta_{n,t} \Delta_{c,t+1} - \frac{1}{\beta} \Delta_{n,t-1} \Delta_{c,t} \right) + \frac{1}{2\beta} \sigma \left(\Delta_{n,t}^2 - \frac{1}{\beta} \Delta_{n,t-1}^2 \right) \right) \right] \quad (39)$$

Notice that every term within the round brackets cancels with a corresponding term in another period. Using $\lim_{T\to\infty} \beta^T \mathbb{E}_0[\Delta_{c,T}^2] = \lim_{T\to\infty} \beta^T \mathbb{E}_0[\Delta_{n,T}^2] = \lim_{T\to\infty} \beta^T \mathbb{E}_0[\Delta_{n,T} \Delta_{c,T+1}] = \lim_{T\to\infty} \beta^T \mathbb{E}_0[\Delta_{c,T} \tilde{c}_{T+1}] = 0$, we therefore have:

$$\tilde{U}^* - \tilde{U} = \mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \left[\frac{\varphi}{2} N^{1+\varphi} \tilde{n}_t^2 + \frac{\sigma C^{1-\sigma}}{2} \tilde{c}_t^2 \right] \quad (40)$$

Finally, note that through the log-linearised labour supply condition, $\tilde{n}_t = -\frac{\sigma}{\varphi} \tilde{c}_t$, so:

$$\tilde{U}^* - \tilde{U} = \frac{\sigma}{2} \mathbb{E}_0^H \sum_{t=0}^{\infty} \beta^t \left[C^{1-\sigma} + \frac{\sigma}{\varphi} N^{1+\varphi} \right] \tilde{c}_t^2 \quad (41)$$

Since the model is stationary, the expected loss from inattention is therefore proportional to the variance of $\tilde{c}_t = (c_t^* - c_t)$.

B Appendix: Calibration

In section 3 we use a standard quarterly calibration, with values as in the table below.

Parameter	Name	Value
β	Discount factor	0.99
σ	Coefficient of risk aversion	1
φ	Disutility of labour	1
ϵ	Elasticity of substitution	9
κ	Phillips curve slope	0.34
ϕ_π	Taylor rule coefficient	1.5
σ_a^2	Variance of technology shocks	0.01
σ_v^2	Variance of cost-push shocks	0.01

Notes: Calibration used in the simulations of the model in section 3.

We set the parameters of the attention decision as:

Parameter	Name	Value
F_{IR}	Complexity of the Inflation Report	1
F_L	Complexity of layered content	0.25
τ_a	Signal to noise in technology signal	0.9
τ_v	Signal to noise in cost signal	0.9
$\underline{\lambda}$	Proportion with no processing cost	0.05
δ_c	Trust improvement from engagement	0.1
δ_s	Trust change from surprise	$-\frac{0.105}{\mathbb{E}S}$
ψ	Parameter in μ_h distribution	9

Notes: Calibration used in the simulations of the model in section 3.

These parameters imply that before layered content, 14.5% of all households read the Inflation Report. In the first period of the layered content, 5% read the Inflation Report and 28.5% read the layered content.

When we study changes in the initial level of trust, the high trust case has those reading the Inflation Report with full trust, and those not reading with trust 0.9. The low trust case has these households on trust 0.2 and 0.1 respectively. The higher δ_c we consider is 0.11, and the lower δ_s we consider is $-\frac{0.095}{\mathbb{E}S}$.

C Effect of Starting Level of Trust Relative to Medium Trust Baseline

Figure 12 plots the equivalent of Figure 4 in the main text but in these figures the deviations are relative to the baseline starting level in the medium trust case.

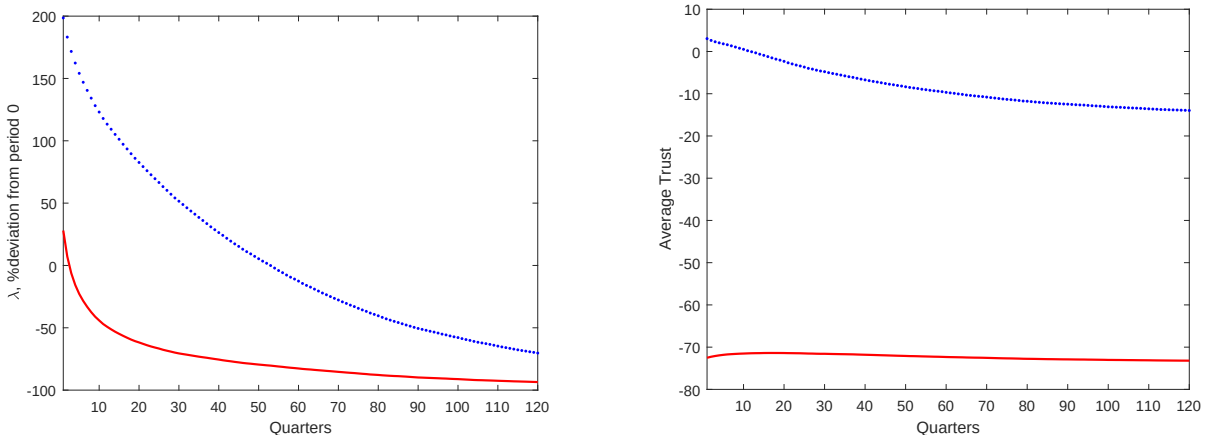

(a) λ_t

(b) $\bar{\mathcal{T}}_t$

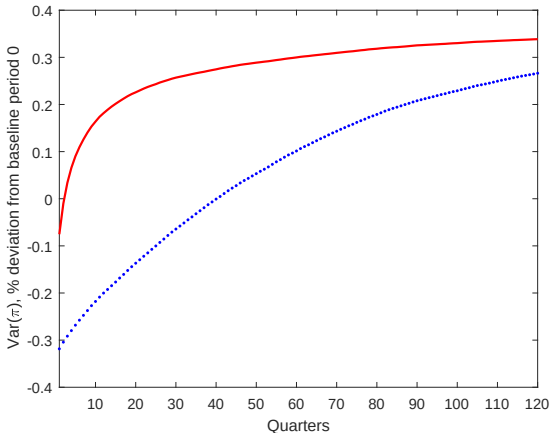

(c) Varπ

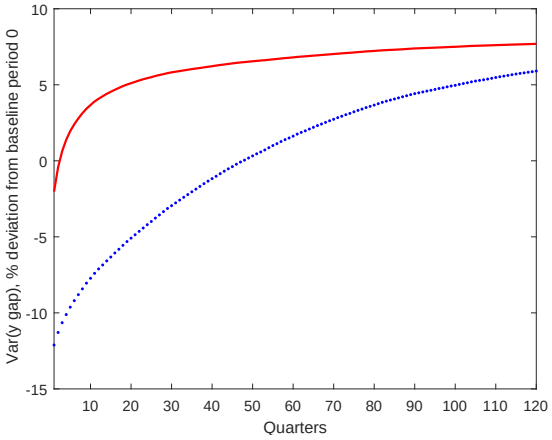

(d) Var$\tilde{y}_t$

Figure 12: Time path of λ_t, $\bar{\mathcal{T}}_t$, Varπ and Var$\tilde{y}_t$ after the introduction of simplified communication: the effect of starting with higher or lower trust

Notes: The blue dotted line is the expected path of either share of processing households, average trust, the variance of inflation or of the output gap relative to initial values in the baseline (medium trust) case. The red solid line is the expected path relative to the medium-trust period 0 baseline of the same variables in the case where initial trust is low.