DISCUSSION PAPER SERIES

DP13619
(v. 2)

INCENTIVE-DRIVEN INATTENTION

Raffaella Giacomini, Wagner Gaglianone, Joao Issler
and Vasiliki Skreta

INDUSTRIAL ORGANIZATION, MACROECONOMICS AND GROWTH AND MONETARY ECONOMICS AND FLUCTUATIONS

ISSN 0265-8003

INCENTIVE-DRIVEN INATTENTION

Raffaella Giacomini, Wagner Gaglianone, Joao Issler and Vasiliki Skreta

Discussion Paper DP13619
First Published 25 March 2019
This Revision 25 March 2019

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

INCENTIVE-DRIVEN INATTENTION

Abstract

"Rational inattention" is becoming increasingly prominent in economic modelling, but there is little empirical evidence for its central premise–that the choice of attention results from a cost-benefit optimization. Observational data typically do not allow researchers to infer attention choices from observables. We fill this gap in the literature by exploiting a unique dataset of professional forecasters who update their inflation forecasts at days of their choice. In the data we observe how many forecasters update (extensive margin of updating), the magnitude of the update (intensive margin), and the objective of optimiza- tion (forecast accuracy). There are also "shifters" in incentives: A contest that increases the benefit of accurate forecasting, and the release of official data that reduces the cost of information acquisition. These features allow us to link observables to attention and incentive parameters. We structurally estimate a model where the decision to update and the magnitude of the update are endogenous and the latter is the outcome of a rational inattention optimization. The model fits the data and gives realistic predictions. We find that shifts in incentives affect both extensive and intensive margins, but the shift in benefits from the contest has the largest aggregate effect. Counterfactuals reveal that accuracy is maximized if the contest coincides with the release of information, aligning higher benefits with lower costs of attention.

JEL Classification: E27, E37, D80, D83

Keywords: rational inattention, Contest, incentives, structural estimation, Survey Design

Raffaella Giacomini - r.giacomini@ucl.ac.uk
UCL and CEPR

Wagner Gaglianone - wagner.gaglianone@bcb.gov.br
Research Department, Banco Central do Brasi

Joao Issler - joao.issler@fgv.br
Graduate School of Economics and Finance – FGV EPGE

Vasiliki Skreta - vskreta@gmail.com
UCL and CEPR

Acknowledgements
We are grateful to Javier Turen and Francisco Lima for excellent research assistance and to the Investor Relations and Special Studies Department (Gerin) of the Banco Central do Brasil (BCB) for kindly providing the data used in this paper. We thank Oli Coibion, Philipp Kircher, Imran Rasul, Colin Stewart, Nico Voigtla¨nder, and participants at various workshops and seminars for useful comments and suggestions. We also acknowledge the support from CeMMAP, CNPq, FAPERJ, INCT and FGV on different grants. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Banco Central do Brasil.

Incentive-driven Inattention*

Wagner Piazza Gaglianone † Raffaella Giacomini ‡
João Victor Issler § Vasiliki Skreta ¶

February 11, 2019

Abstract

"Rational inattention" is becoming increasingly prominent in economic modelling, but there is little empirical evidence for its central premise–that the choice of attention results from a cost-benefit optimization. Observational data typically do not allow researchers to infer attention choices from observables. We fill this gap in the literature by exploiting a unique dataset of professional forecasters who update their inflation forecasts at days of their choice. In the data we observe how many forecasters update (extensive margin of updating), the magnitude of the update (intensive margin), and the objective of optimization (forecast accuracy). There are also "shifters" in incentives: A contest that increases the benefit of accurate forecasting, and the release of official data that reduces the cost of information acquisition. These features allow us to link observables to attention and incentive parameters. We structurally estimate a model where the decision to update and the magnitude of the update are endogenous and the latter is the outcome of a rational inattention optimization. The model fits the data and gives realistic predictions. We find that shifts in incentives affect both extensive and intensive margins, but the shift in benefits from the contest has the largest aggregate effect. Counterfactuals reveal that accuracy is maximized if the contest coincides with the release of information, aligning higher benefits with lower costs of attention.

Keywords: Rational Inattention; Contest; Incentives; Structural Estimation; Survey Design.

JEL Classification: E27, E37, D80, D83.

---

*We are grateful to Javier Turen and Francisco Lima for excellent research assistance and to the Investor Relations and Special Studies Department (Gerin) of the Banco Central do Brasil (BCB) for kindly providing the data used in this paper. We thank Oli Coibion, Philipp Kircher, Imran Rasul, Colin Stewart, Nico Voigtländer, and participants at various workshops and seminars for useful comments and suggestions. We also acknowledge the support from CeMMAP, CNPq, FAPERJ, INCT and FGV on different grants. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Banco Central do Brasil.

†Research Department, Banco Central do Brasil. Email: wagner.gaglianone@bcb.gov.br

‡University College London, Department of Economics. Email: r.giacomini@ucl.ac.uk

§Graduate School of Economics and Finance – FGV EPGE, Getulio Vargas Foundation. Email: joao.issler@fgv.br

¶UT Austin, University College London & CEPR. Email: vskreta@gmail.com

1 Introduction

Rational inattention (Sims, 2003) is the leading theoretical framework for endogenizing information frictions in economic models. Its central premise is that agents allocate their limited attention budget optimally when making decisions. A fast-growing theoretical literature models attention as an optimal cost-benefit decision in contexts ranging from price-setting (Mackowiak and Wiederholt, 2009) and investment choice (Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016), to dynamic stochastic general equilibrium models (Maćkowiak and Wiederholt, 2015). Yet, there is little empirical evidence that the choice of attention results from a cost-benefit optimization. One reason for the lack of evidence is the difficulty of observing attention choices, as well as the costs and benefits (i.e., the incentives) in observational data. This paper fills the gap in the literature by exploiting a unique dataset that enables us to identify the incentive parameters from observable decisions. We document new stylized facts and structurally estimate a rational inattention model that can explain the patterns in the data. Our results lend empirical credence to the emerging consensus on modelling inattention as an incentive-driven optimal decision.

The panel dataset we study is the little-known *Focus Survey* of professional forecasters maintained by the Central Bank of Brazil. We observe the decisions of participants to update their forecasts of current month's inflation (nowcasts) in Brazil's consumer price index, the IPCA. The data is unique for several reasons. First, forecasters can update any time they choose, so we can study both the decision to update and the magnitude of the update as endogenous decisions.[1] This allows us to study variations in both the extensive margin of updating (how many people update) and in the intensive margin (by how much they update), which matters for policy and survey design. Second, the objective function of the optimizing updating decision (the forecast accuracy) is observable in the data. This allows us to link the observables to the underlying incentive parameters, through optimal attention choices. Third, in addition to the informal incentives that it shares with other surveys analyzed in the literature, the survey has two "shifters" in the benefits and costs of processing information. The benefit

[1]This is in contrast to other common surveys, where forecasters are sampled at exogenously determined and infrequent times–e.g. monthly in the Consensus Forecasters or Blue Chip Analysts or quarterly in the Survey of Professional Forecasters.

Figure 1. Contest versus Information Release

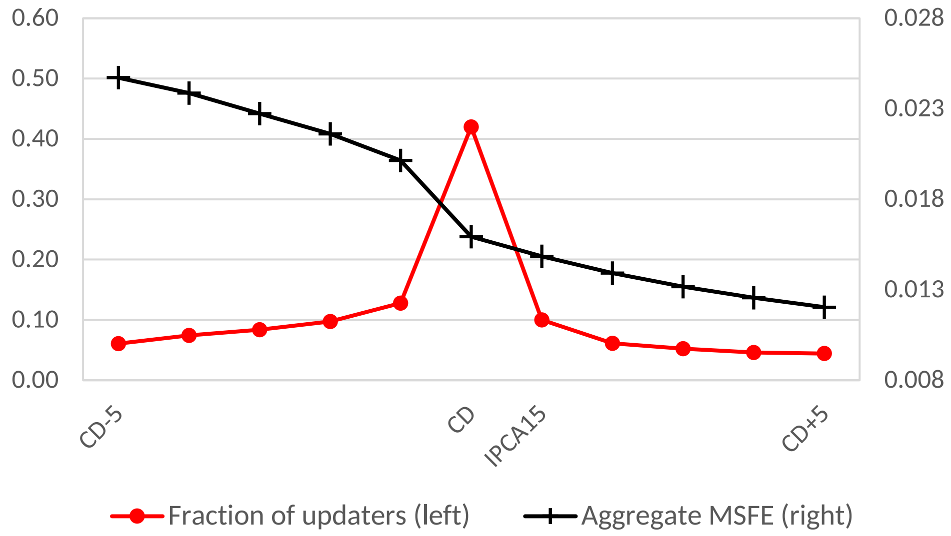

Notes: The figure shows the fraction of forecasters in the *Focus Survey* who update their nowcast of inflation on a five-day window around the contest (CD) and the information release (IPCA15) days, averaged over all months in the dataset. It also shows the aggregate MSFE, which is the average across forecasters of the individual Mean Squared Forecast Errors. The individual MSFE is the squared difference between the nowcast associated with each forecaster on that day and the realization of inflation for that month, averaged over all months. Accuracy is the negative of the MSFE.

shifter is a monthly contest that ranks participants based on the accuracy of their (most recent) forecast on a specific day. The cost shifter is the release of official information about inflation, the IPCA15 inflation, which occurs the day after the contest and which largely overlaps with the variable that agents seek to forecast.[2] These incentive shifters allow us to identify the cost and benefit parameters from shifts in the observables.

Figure 1 illustrates a striking empirical pattern in the raw data. On the contest day we see a large increase in both the fraction of updaters (from about 10% to 42%) and in the aggregate accuracy improvements (a sharp fall in the aggregate mean-squared forecast error, henceforth MSFE). In contrast, the information release on the day after the contest appears to have no aggregate effect: the fraction of updaters and the aggregate accuracy improvements on the IPCA15 day are similar to those on any other non-contest day. Panel regressions confirm that the contest is the strongest driver of the decision to update, and that it also improves updaters'

[2]The IPCA15 measures inflation from the 15th of the previous month to the 15th of the current month, whereas the IPCA measures inflation between the first and the last day of the current month.

accuracy. This suggests that the gain in aggregate accuracy on the contest in Figure 1 is due to changes in both extensive and intensive margins of updating. The structural model we develop allows us to quantify their relative contribution. A possible reason why the information release has no visible effect in Figure 1 is that almost no agents who update on the contest update again the day of information release, so they are not exploiting the additional accuracy improvement that would be available on that day. This raises the question whether the contest mis-aligns updates that would have otherwise occurred the day of information release. Our structural model can be used to investigate this and other questions related to survey design.

Motivated by the empirical findings,[3] we develop a decision-theoretic model of rational inattention where agents have limited resources to allocate to produce an accuracy-maximizing forecast, while facing varying costs and benefits. The model assumes that each month, agents use a realistic statistical model (an Autoregressive Moving Average–ARMA) to produce an initial forecast. On each subsequent day, they make two decisions: 1) whether to update their forecast and 2) how much attention to allocate to process information in order to update. The first decision is driven by the *opportunity cost* of time devoted to forecast updating. An agent updates if the benefit of forecasting is greater than the benefit of time devoted to other activities. The second decision is the result of a rational inattention optimization problem that depends on the information available to the agent and on the marginal cost and benefit of attention. The two decisions are endogenous, as they are driven by the same incentive parameters. Parameters can vary across agents, which implies that not everyone updates and accuracy is heterogeneous.

There are two dynamic dimensions in our setting: the month-to-month problem of forecasting inflation at the beginning of the month, and the within-month problem of updating the forecast. We use results from aggregation of ARMA processes to make the statistical model for the initial forecast compatible with the statistical model used for the update. The statistical model allows us to decompose the forecast accuracy into a component that depends on the resolution of uncertainty as the forecast horizon decreases and a component that depends on attention. We can then obtain optimal attention as a function of the forecast horizon and the incentive parameters.

[3]In addition, we rule out strategic behaviour and self-selection based on ability on the contest.

The model yields simple and intuitive analytical expressions for the theoretical counterparts of the observables in the data–the fraction of updaters and the accuracy of the update–as functions of the model's parameters. The parameters can be divided into ARMA parameters and incentive parameters. The estimates of the ARMA parameters can be used to assess the external validity of the model. The variation in incentive parameters on the contest and the IPCA15 days is our main interest and it is identified by the observed shifts in extensive and intensive margins of updating on these days. The incentive parameters enter the model as ratios, so we need to impose normalization restrictions if we want to identify them beyond their patterns of time variation. We report results for two sets of such restrictions. The conclusions about the variation in incentive parameters are robust to these and alternative normalizations, which also don't affect the estimates of the ARMA parameters.

We structurally estimate the model by Simulated Method of Moments, in order to match the *joint* dynamics of updating frequencies and aggregate accuracy reported in Figure 1. Both set of restrictions fit the data well and pass the specification test. The robust conclusion of the estimation is that the incentives follow a realistic pattern: There is a constant incentive for participation to the survey on "normal" days. The contest provides an additional benefit that is strong on the contest day but is also felt before the contest (since even "old" forecasts count for the contest). The cost of attention is lower when information is released. Remarkably, the estimates of the ARMA parameters are almost identical to those implied by Brazil's inflation data. These data are not used in the estimation, so this finding reinforces the external validity of the model. Overall, the estimation results provide empirical support for modelling updating and inattention as incentive-driven decisions.

We use the estimated structural model to quantify how changes in aggregate accuracy are affected by changes in the extensive and intensive margins. We find that 70% of the accuracy improvement on the contest that is visible in Figure 1 is due to more agents updating (the extensive margin) and 30% to agents paying more attention (the intensive margin).

We further perform counterfactual exercises to investigate alternative survey designs. First, we find that holding the contest on any given day of the month would result in an accuracy improvement from the previous day that is 3 to 4 times larger than it would be without the contest. Second, we show that the optimal contest day is the IPCA15 day. On this day the increase in benefits from the contest is amplified by the availability of low-cost information. Finally, we investigate the extent to which the contest mis-aligns updates from the more "natural" IPCA15 day. We find that without the contest average accuracy is worse, even though most updates happen on the IPCA15 day. This underscores that the coordinated updates that occur because of the contest are crucial for the survey's aggregate accuracy.

For the sake of tractability, the model relies on two main simplifying assumptions. The first is that the two decisions of whether to update and how much attention to allocate to updating are not the result of a joint optimization. Rather, the forecaster decides sequentially whether to update and, if so, how much attention to optimally devote to processing information. We model only the second decision as a problem of rational inattention, while the first decision boils down to comparing the opportunity cost of time to a fixed threshold.[4] Both decisions are nonetheless endogeneous, as they depend on the same incentive parameters. The second assumption concerns the information available to agents who update. We assume that updaters have access to *past* public signals that are more accurate than any past private signals. This reduces the dynamic rational inattention problem to a sequence of static problems: An agent only needs to process information for the current day, which depends on the current benefit and cost parameters.[5]

The paper is organized as follows. Section 2 relates the paper to the literature; Section 3 discusses the data; Section 4 presents the empirical findings. Section 5 discusses the model; Section 6 presents the structural estimation results, and Section 7 the counterfactual analysis. Section 8 concludes.

[4]We believe this does not diminish the contribution of the paper in terms of providing empirical credence to models of rational inattention, as the same conclusion would have emerged by focusing only on the second decision. As discussed by Woodford (2009), applying rational inattention to a timing decision like the first decision here would present additional challenges, so attempting to solve a joint rational inattention optimization for the two decisions would substantially complicate the analysis without necessarily adding new insights. See Section 5.3 for further discussion.

[5]In the likely presence of both public and private information, we argue that this assumption is more realistic than the opposite extreme assumption that updaters only rely on their past, less accurate, private signals. Relaxing the assumption would imply that the accuracy of updaters depends on past updating decisions. We show that this prediction is not supported by the data. See Section 5.3 for further discussion.

2 Related Literature and Contribution

This paper makes several contributions to the literature.

We follow the recent theoretical literature on rational inattention (Maćkowiak and Wiederholt, 2015) in endogenizing the attention choice, as opposed to assuming that the bound is exogenous as in, e.g., Mackowiak and Wiederholt (2009), Caplin and Dean (2015), and Steiner, Stewart, and Matějka (2017). By exploiting a unique dataset to show that the choice of attention is driven by cost-benefit considerations, we provide an empirical foundation for this increasingly prominent theoretical mechanism in models of rational inattention.

This paper brings a rational inattention model to the data, contributing to a literature still in its infancy. The typical approach to validation of rational inattention models involves calibrating and deriving testable implications of the model. We go beyond the existing literature by structurally estimating the model. We can do this because our data allows us to overcome some of the challenges typically encountered in validating these models using observational data.[6] One such key challenge is the difficulty in separately identifying the unobservable attention from the (usually also unobservable) prior uncertainty. Caplin, Leahy, and Matejka (2016) and Csaba (2018) make important steps towards overcoming this challenge in the context of discrete choice analysis. In our data we can separately identify the two components.

We consider expectations data, similarly to the literature studying the role of information frictions in explaining expectation formation and the dynamics of economic variables (see the survey of Woodford, 2013). Coibion and Gorodnichenko (2012) provide evidence of such frictions in expectations data, and Coibion and Gorodnichenko (2015) find empirical support for some predictions of models of exogenous information frictions. This paper takes a step forward by endogenizing the information frictions that are assumed exogenous in this literature. This sheds light into what affects the quality of professional forecasts, which are a key input in business and governmental decisions. It also allows us to study the implications of survey design for forecast accuracy.

[6]There are important experimental studies, however, including Cheremukhin, Popova, Tutino et al. (2011), Caplin and Dean (2013), Dean and Neligh (2017), Martin (2016) and Cavallo, Cruces, and Perez-Truglia (2017).

Our findings show that a forecasting contest improves accuracy. Marinovic, Ottaviani, and Sørensen (2013) study theoretically the effect of a forecasting contest in a strategic model without information frictions. They find that the effect of the contest on forecast accuracy can be ambiguous. Forecasters in our dataset seem to ignore the strategic component and focus on overcoming the information barriers in order to provide accurate forecasts.[7]

The link that we document between attention and incentives speaks to the broader question of what are the productivity drivers in economics (see the survey of Syverson, 2011)). Lazear (2000) studies the effect of monetary incentives on output, while Shearer (2004) shows how the structure of compensation contracts affects productivity using data from a field experiment. In the psychology literature, Reeve, Olson, and Cole (1985) consider the role of incentives and competition in motivation and performance. In the same spirit, Glaeser, Hillis, Kominers, and Luca (2016) argue that tournaments can be a cost-effective tool to outsource public services. Viewed from this broader perspective, our study contributes to establish a clear link between incentives and performance.

3 Data

Our panel data are from the Central Bank of Brazil's (BCB) survey of professional forecasters, the *Focus Survey*. We study forecasts–nowcasts–of current month's inflation in the consumer price index (IPCA), which is the official inflation measure and the target of monetary policy at the BCB.[8] The panel includes all forecasters who provide forecasts that are confirmed or updated within 30 days from the first forecast considered. The panel is unbalanced since not all forecasters participate each month and the number of participants is generally increasing over time. It consists of forecasts for a given month that each participant can produce every working day of the month, starting from January 8th, 2004 to January 8th, 2015, amounting to

[7]This was also confirmed in personal interviews with some participants. One reason why participants seem not to act strategically may be that the survey is confidential and anonymous.

[8]The inflation index measures the change in prices of a fixed set of goods and services. The price research is done daily, covering thirteen cities in Brazil. Inflation is closely monitored by economic agents in Brazil for tracking monetary policy and also because it is used to index treasury bonds, in wage negotiations, and as an adjustment for certain contractually regulated prices.

a total of 2,751 daily forecasts for 132 months, with an average of 85.3 forecasters.[9] We treat months and the forecasts associated with each of them as events, that repeat one after another until the end of the sample. Each event entails a decreasing-horizon forecasting problem. The events are connected by inflation, which is a continuous process over the whole sample.

The BCB provides forecasters with a software (the *Market Expectation System*) that they can access any time to provide forecasts for a number of economic variables.

Forecast Updates: Any time a forecaster logs in the system, she can change a forecast or confirm it. For forecasters who do not log in, the system copies the previous forecasts. We say that a forecast is *updated* if the forecaster changed the forecast. This means that we do not consider the forecasters who confirmed the previous forecast. The reason is that there are very few examples of this in the data and it would unnecessarily complicate the analysis to try and account for the distinction.

Informal Incentives: Similarly to the other surveys of professional forecasters that have been analyzed in the literature, (the "Blue Chip," the "Consensus" or the Fed's "Survey of Professional Forecasters") the *Focus Survey* has several informal incentives for updating and accuracy. First, every Monday the BCB publishes the highly visible in the media "Focus-Market Readout."[10] The readout only considers forecasts that were updated during the previous thirty days. Second, forecasters who are inactive for more than thirty days are removed from the system. Those who remain inactive for six months are blocked from the system, and need to request a renewal of their login and password. Third, some of the active participants are invited to BCB meetings to provide opinions about the economic outlook.

The Contest: The survey's main formal incentive is a contest. The monthly 'contest dates' are announced by the bank before the beginning of each calendar year. Every month, upon the release of the realization of the variable, the forecasters are ranked based on the accuracy

[9]We start the sample in 2004 because there were too few participants prior to this year.

[10]The readout reports key aggregate statistics from the *Focus Survey* based on data collected at 5 PM of the previous Friday. See Marques (2013) for further details.

of the forecast that was on the Market Expectation System on the pre-announced day of the previous month, the *contest day*. The names of the five most accurate forecasters (institutions) according to the absolute forecast error are then published on the BCB website. The contest is highly valued by the survey participants and the top-five forecasting institutions usually publicize their contest accomplishments on their websites or advertising material. Figure 2 shows as an example the outcome of the monthly contest for February 2017.[11]

Figure 2. Example of Contest Outcome

BANCO CENTRAL DO BRASIL

Gerin

Top 5 Forecasting Institutions - February 2017

March 10, 2017

The Investor Relations and Special Studies Department (Gerin) has announced the Top 5 forecasting institutions for February 2017.

Table 1

Top 5 Forecasting Institutions - Short-Run

February 2017

	IPCA	Deviation
1	Flag Gestora de Recursos	0.0717
1	Petros Fundação de Seguridade Social -	0.0717
3	Quantitas Asset Management	0.0833
4	Banco Bradesco S.A.	0.0852
5	ICAP Brasil	0.0867

	IGP-DI	Deviation
1	Banco Itaú S.A.	0.0883
2	BBM Investimentos	0.1217
2	SPX Capital	0.1217
4	Haitong Banco de Investimento do Brasil	0.1317
5	J. Safra Asset Management	0.1350
5	Verde Asset Management	0.1350

	IGP-M	Deviation
1	LCA Consultores S/C Ltda.	0.0683
2	Haitong Banco de Investimento do Brasil	0.0733
3	Icatu Vanguarda Administração de Recursos	0.0833
4	Banco Itaú S.A.	0.0883
5	Banco Fibra S.A.	0.0983

	Exchange Rate	Deviation
1	Telefônica / Vivo	0.0691
2	Rosenberg & Associados S/C Ltda.	0.0693
3	BB DTVM S.A.	0.0741
4	Tendências Consultoria Integrada	0.0746
5	Banco do Brasil S.A.	0.0751

	Over Selic	Deviation
1	Barclays Capital	0.0417
1	Bozano Gestão de Recursos	0.0417
1	CSHG Gauss	0.0417
1	M. Safra	0.0417
5	Banco do Brasil S.A.	0.0625
5	Banco Itaú S.A.	0.0625
5	Banco Original do Agronegócio	0.0625
5	Brasilprev Seguros e Previdência S.A.	0.0625
5	BW Gestão de Investimentos Ltda.	0.0625
5	Caixa Asset	0.0625
5	Daiwa Asset Management	0.0625
5	Deutsche Bank - Banco Alemão S.A.	0.0625
5	Fapes - BNDES	0.0625
5	Flag Gestora de Recursos	0.0625
5	Ibiuna Investimentos Ltda.	0.0625
5	Icatu Vanguarda Administração de Recursos	0.0625
5	Kondor Admin. e Gest. de Rec. Financ. Ltda.	0.0625
5	MCM Consultores	0.0625
5	PREVI Caixa Previd Funci Banco Brasil	0.0625
5	Quantitas Asset Management	0.0625
5	Quest Investimentos Ltda.	0.0625
5	Santander Asset Management	0.0625
5	Sul America Investimentos	0.0625
5	Vintage Investimentos	0.0625

Information for unrestricted disclosure. It is not intended to bind Banco Central do Brasil in its monetary or foreign exchange policy actions. Questions and comments to gerin@bcb.gov.br

[11]See http://www4.bcb.gov.br/pec/gci/ingl/focus/top5.asp for further details about the contest.

Figure 3. Example of Forecast Timeline

forecast begins	contest day	IPCA15 released	IPCA released and used to evaluate forecasts
8 Feb	21 Feb	22 Feb	8 March

Information Releases: The main information release is the monthly release of IPCA15 inflation, which measures inflation between the 15th of the current month and the 15th of the previous month. The date of release of the IPCA15 changes from month to month, but it is always the day after the contest. In the panel regressions, we further consider the release of the minutes of the meeting of the BCB Monetary Policy Committee (MPC), which occurs less frequently and at irregular times.

Forecast Timeline: Forecasters know the dates of the contest and data releases in advance. The number of workdays in the month (i.e., the duration of the forecasting period) and the timing of the contest can vary across months. The chronology of relevant events within a representative month is depicted in Figure 3: the first forecast for February's inflation can be given on the day of release of the IPCA for January, which occurs most often on the 8th of February. The contest most often takes place on the 21st of February which is always the day *before* the release of IPCA15 inflation (measuring inflation between the 15th of February and 15th of January). Forecasters can provide a new forecast on each working day between the 8th of February and the day of the release of IPCA for February–(most often) the 8th of March.

Survey Participants: Participants include non-financial institutions, commercial banks, asset-management firms and consulting firms.

Confidentiality: The data are proprietary and the identity of the forecasters is not known to us nor is it revealed to the public, except for the winners of the contest (cf. Figure 2).

4 Stylized Facts

In this section we document new stylized facts about the drivers of forecast updates and accuracy improvements. We first focus on the *individual* level, by means of panel regressions. We then visualize the *aggregate* dynamic behavior of updates and accuracy around the contest day and around days of information releases.

Drivers of the Decision to Update: We first consider the extensive margin of updating (how many forecasters update) by estimating a panel logit model for forecast updates:

$$\Pr\left(z_{it} = 1 \left| x_{it} \right.\right) = G\left(\alpha_i + x'_{it}\beta\right), \tag{1}$$

where G is the logistic function and

$$z_{it} = \begin{cases} 1 & \text{if forecaster } i \text{ updates on day } t \\ 0 & \text{otherwise.} \end{cases}$$

The regressors x_{it} include dummy variables for the day of the contest (d_t^{CD}), the day of release of the IPCA15 (d_t^{IPCA15}), the day before or after these (d_t^{CD-1} and $d_t^{IPCA15+1}$) and the day when the MPC minutes are released (d_t^{MPC}). Other regressors are dummy variables for Mondays and Fridays and the $EMBI_{t-1}$, the Emerging Markets Bond Index Plus for Brazil (EMBI+BR)–a measure of uncertainty on the previous day.

Table 1 reports the coefficient estimates and the average marginal effects (in square brackets). The table shows that the contest is not the only driver of updates, as forecasters update at other times as well. The contest has, however, the largest effect on the extensive margin of updating (how many people update): The probability of updating goes up by 38.9 percentage points (p.p.) on the contest.[12] There is also a "contest anticipation" effect with a 18.8 p.p. increase in the probability of updating one day *before* the contest. The release of information has a smaller effect on the probability of updating (the IPCA15 is associated with a 12.1 p.p.

[12]We focus the discussion on marginal effects.

Table 1. Drivers of the Decision to Update

Regressors	Logit Fixed Effect Coefficients	
	(1)	(2)
d_t^{CD-1}	0.911***	0.888***
	(0.031)	(0.031)
	[0.203]	[0.188]
d_t^{CD}	2.609***	2.714***
	(0.023)	(0.024)
	[0.417]	[0.389]
d_t^{IPCA15}	0.539***	0.546***
	(0.034)	(0.035)
	[0.125]	[0.121]
$d_t^{IPCA15+1}$	0.023	-0.015
	(0.042)	(0.042)
	[0.005]	[-0.004]
d_t^{MPC}	0.103**	0.261***
	(0.042)	(0.043)
	[0.024]	[0.059]
$EMBI_{t-1}$	-	0.024**
		(0.01)
		[0.006]
d_t^{MON}	-	0.383***
		(0.021)
		[0.087]
d_t^{FRI}	-	0.494***
		(0.022)
		[0.112]
Log likelihood	-55258.8	-52989.0

Notes: Model for the probability that a forecaster updates on day t. Sample from January 8th, 2004 to January 8th, 2015. Number of observations (model 1) = 228,157. Robust standard errors in parentheses. ***, ** and * indicate, respectively, significance at the 1%, 5% and 10% level. Average marginal effects are in square brackets.

increase and the MPC with a 5.9 p.p. increase). The Friday dummy is also significant, which may reflect the importance of, in this case more informal, incentives in the survey, as summary statistics about the forecasts collected on Fridays are released on the following Monday as part of the Focus-Market Readout. The table also reveals that forecasters are more likely to update when there is higher uncertainty, as indicated by the coefficient for $EMBI_{t-1}$. This finding is consistent with one of the main predictions of rational inattention models.

Drivers of Accuracy Improvements for Updaters: We next analyze the drivers of accuracy improvements, conditional on updating. While we expect to find that information releases improve accuracy, it is less clear a priori whether the contest would affect the accuracy of updaters. An affirmative answer would support the hypothesis that the contest not only induces more forecasters to update, but also makes them exert more "effort" into producing accurate forecasts. This would link the intensive margin of updating to the quality of the decision and suggest that forecast accuracy is the objective of the update.

Analyzing accuracy improvements is complicated by the fact that there are confounding factors that cause time variation in forecast accuracy and in the effects we want to investigate. First, the decreasing-horizon setting means that the forecast accuracy is expected to improve during the month and that the accuracy improvement is not necessarily constant. Second, the contest and IPCA15 days fall on different dates each month so they are also associated with different horizons from month to month. Finally, the accuracy improvement by construction depends on the time between updates. To partly control for these factors, we add the forecast horizon and the "duration" (the days since the previous update) both as regressors and as interaction terms when investigating the effects of the contest and information releases.

We consider only observations for which agent i updated on day t and estimate the following panel regression:

$$ln(e_{it-1}^2) - ln(e_{it}^2) = \alpha_i + x_{it}'\beta + u_{it}, \tag{2}$$

where e_{it} denotes the forecast error for forecaster i on day t and e_{it-1} the forecast error on the day that forecaster i previously updated. The regressors x_{it} considered by the various models are reported in the first column of Table 2.

Table 2. Drivers of Accuracy Improvements for Updaters

Regressors	**Panel Fixed Effect Coefficients**			
	(1)	(2)	(3)	(4)
d_t^{CD-1}	-3.694	-3.616	-3.445	-3.398
	(3.227)	(3.289)	(3.282)	(3.281)
d_t^{CD}	7.132**	6.334**	−50.101***	−50.443***
	(3.017)	(3.219)	(17.794)	(17.646)
d_t^{IPCA15}	33.656***	35.299***	63.632*	35.542***
	(4.839)	(5.002)	(33.84)	(4.995)
$d_t^{IPCA15+1}$	35.272***	33.549***	33.581***	33.478***
	(5.647)	(5.766)	(5.752)	(5.772)
d_t^{MPC}	20.660***	22.847***	22.369***	22.421***
	(5.585)	(5.536)	(5.558)	(5.563)
$duration_t$	2.423***	2.381***	2.443***	2.369***
	(0.197)	(0.198)	(0.248)	(0.197)
$horizon_t$	−0.490***	−0.490***	−0.560***	−0.569***
	(0.169)	(0.171)	(0.175)	(0.176)
$EMBI_{t-1}$	-	-0.995	-0.832	-0.840
		(0.835)	(0.84)	(0.842)
d_t^{MON}	-	4.797*	4.703*	4.761*
		(2.693)	(2.693)	(2.692)
d_t^{FRI}	-	3.968	4.041	3.985
		(2.679)	(2.668)	(2.677)
$duration_t \times d_t^{CD}$	-	-	-0.038	-
			(0.48)	
$horizon_t \times d_t^{CD}$	-	-	4.589***	4.594***
			(1.433)	(1.432)
$duration_t \times d_t^{IPCA15}$	-	-	-1.384	-
			(1.01)	
$horizon_t \times d_t^{IPCA15}$	-	-	-1.897	-
			(2.908)	
$constant$	26.015***	24.150***	24.496***	25.087***
	(2.375)	(2.613)	(2.749)	(2.651)

Notes: Dependent variable is minus the change in the log of the squared forecast error for an *updater* on day t, relative to the previous update. Sample from January 8th, 2004 to January 8th, 2015. Number of observations (model 1) = 26,911. Standard errors in parentheses. ***, ** and * indicate, respectively, significance at the 1%, 5% and 10% level.

Table 2 reports the estimation results. Column (1) confirms the expected finding that information releases are associated with accuracy improvements, as the coefficients for IPCA15 and MPC are large and significant. Perhaps more surprisingly, it shows that the contest also has an effect, as accuracy improvements go up by 7.1 p.p. on the contest.[13] In contrast, column (2) shows that other variables that according to Table 1 are associated with an increased probability of updating–Mondays, Fridays and the EMBI–are not associated with an increase in accuracy improvements (except for the coefficient for the Monday dummy, which is however only significant at the 10% level).

Aggregate Dynamics of Updates and Accuracy: We now focus on the aggregate dynamics of updates and accuracy around the contest and around days associated with information releases (the MPC meetings and the IPCA15). Since these dates change across months, we consider a window of five days around these days. The left panel of Figure 4 is the same as Figure 1, and the right panel considers a window of five days around the MPC day. The figure shows that the main driver of updates and accuracy improvements at the aggregate level is the contest, as there are no visible similar changes on the days associated with information releases. The figure also confirms the finding from the panel regressions that forecasters update outside the contest and information release days (about 10% of forecasters update on each non-contest day), leading to the conclusion that informal incentives also matter in the survey.

The conclusion from the left panel of Figure 4 is that, although there is a small asymmetry in updating behavior before and after the contest, the fraction of updaters is approximately constant on non-contest days, but it rises substantially on the contest. The $MSFE$ declines as the forecast horizon shrinks, which is an expected consequence of the natural resolution of uncertainty leading up to the revelation of the forecasted variable. The effect of the contest is to induce a sizable level shift downwards in the $MSFE$ curve, resulting in a much larger improvement in accuracy on the contest day (and consequently for the rest of the month), relative to the (approximately constant) improvement we see on any other day.

The documented jump in aggregate accuracy on the contest could be caused by both changes

[13] Using regression in column (4) one can compute that the average effect of the contest is 6.4 p.p., which is comparable to the estimates from models in columns (1) and (2).

Figure 4. Dynamics of Updates and Aggregate MSFE

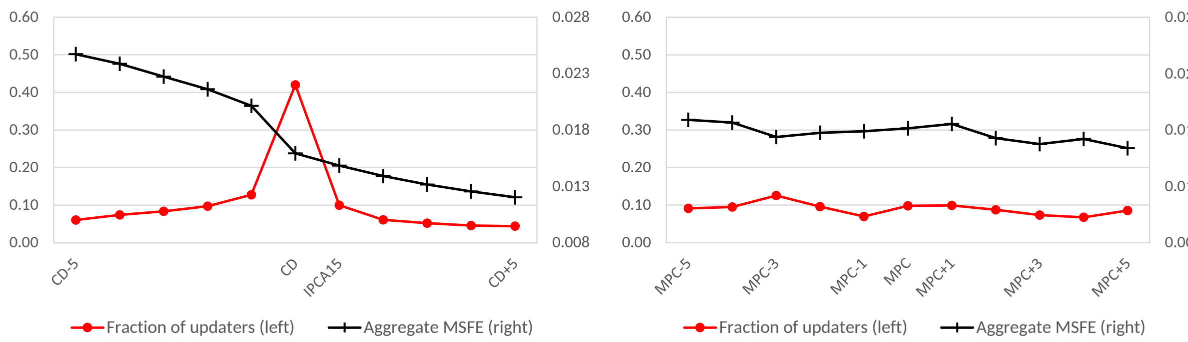

Notes: : Daily evolution of the fraction of updaters and aggregate $MSFE$ around the contest and IPCA15 day (left graph) and around the MPC day (right graph).

in the extensive margin (if more forecasters update, their average accuracy is higher) and in the intensive margin (each forecaster may be putting more effort into obtaining an accurate forecast). The estimated structural model allows us to decompose the aggregate accuracy improvement on the contest into the contribution of changes along both margins.

Our findings suggest that the contest may be crowding-out updates on other days. Since the contest is the day before, few forecasters update on the IPCA15 day, even though doing so improves accuracy. In fact, the cases when a forecaster updates on both the contest and the IPCA15 day constitute only 0.65% of the sample. Another counterfactual exercise allows us to shed light onto a potential crowding-out effect of the contest on aggregate accuracy.

Forecaster Heterogeneity: The observable dimensions of heterogeneity in our data are the updating behavior and the forecast accuracy across forecasters and over time. To investigate whether they are driven by time-invariant fixed effects (e.g., forecasters who are always frequent updaters and/or the most accurate) we compute measures of mobility in the cross-sectional distributions of both updating probability and accuracy. Specifically, we consider forecasters who participated in the last two years of the sample and compute the normalized trace measure of Shorrocks (1978), by dividing the cross-sectional distribution of average MSFE or total number of updates during each year into 5, 10 or 20 quantiles and computing transition probabilities among the different quantiles.

Table 3. Mobility Index of Shorrocks (1978)

Number of Quantiles	Frequency of Updates	$MSFE$
5	0.691	0.749
10	0.760	0.793
20	0.872	0.886

Notes: Index based on the last two years of the sample and on dividing the cross-sectional distribution of the yearly frequency of updates and the average MSFE over the year into different quantiles. Immobility=0 and perfect mobility=1.

Table 3 shows that the indexes are close to 1 (which corresponds to perfect mobility). This suggests that there is significant mobility in the distribution of both accuracy and updating probability across forecasters. This result supports the conclusion that our findings are not primarily driven by time-invariant heterogeneity, and it leads us to rule out the potential presence of *positive selection* on the contest that may induce more able forecasters to update on that day. The general conclusion also motivates our modelling of heterogeneity in the theory that we develop below: We assume that forecasters face heterogeneous incentive parameters that result in both heterogeneous updating histories and heterogeneous attention choices.

Figure 5. Forecast Disagreement (Standard Deviation of Forecasts)

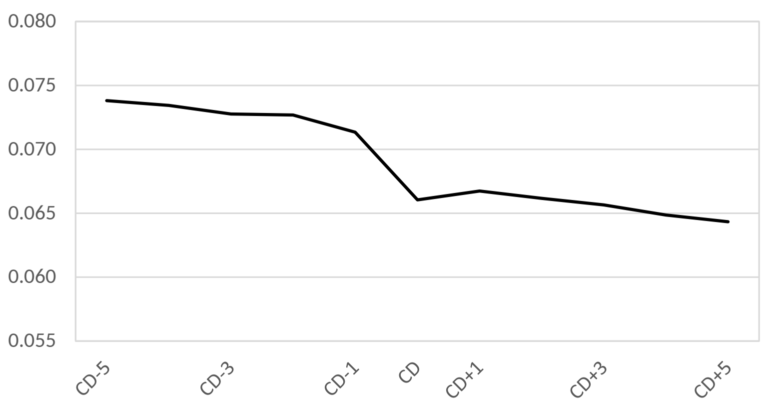

Forecaster's Objective and Strategic Behaviour: The empirical facts document that the contest affects both the updating decision and the accuracy of forecasts. So a reasonable objective function is that forecasters seek to win the contest and maybe have another objective on other days. In the simpler contest with one prize studied in Marinovic et al. (2013), the probability of winning increases the more accurate is a forecast but, also, the more it differs from the competing forecasts: If, say, N forecasters submit the same forecast, the probability that one of them wins is 1/N. Marinovic et al. (2013) argue that this encourages strategic forecasters to put more weight on private signals, compared to what would have been optimal in a decision-theoretic setup, leading to an increase in disagreement among forecasters on the contest day. This is however inconsistent with the data (see Figure 5), as the disagreement *decreases* on the contest day. We take this as evidence that forecasters in our data do not behave strategically (this was confirmed in conversation by a few participants). Based on these observations, we find it more suitable (as well as tractable) to employ a decision-theoretic setup, rather than a game-theoretic one, and assume that a forecaster's objective is to maximize accuracy.

5 Theory

We build on the theory of rational inattention (Sims, 2003) to link the observable dynamics of updating decisions to the available information and to the unobservable cost and benefit parameters that are driving agents' decisions. The theory is inspired by the way forecasters behave in reality: They use state-of-the art statistical models in which they input both publicly available information and privately collected and processed signals.

The theory is also motivated by the empirical regularities: The fact that forecasters do not update every day suggests that they have limited resources to do so. The patterns of updates and accuracy improvements suggest that forecasters might face time-varying costs and benefits of producing a forecast. In the model, forecasters choose the amount of mental capacity–henceforth "attention"–to employ in order to formulate a forecast that is accuracy-maximizing given those resources. This optimal amount depends on the time-varying costs and benefits of processing information, which is what we ultimately estimate.

5.1 Agents' Decision Problem

At the beginning of each month all agents produce an initial forecast. On each subsequent day t, agent i makes two decisions: (i) whether or not to update and (ii) how much attention, k_{it}, to devote to collecting/processing information in order to update the forecast. Thus, attention is *endogenous*.[14] The following two subsections discuss the two decisions separately. In Subsection 5.2 we discuss how we link them when we bring the model to the data.

5.1.1 Decision of Whether to Update

Agents face not only explicit costs of processing information, but also opportunity costs of time or mental effort—e.g. consultants need to travel, employees at financial institutions have meetings or other inflexible work obligations. Let w_{it}^{o} denote the marginal benefit of time devoted to activities other than forecasting and let $\tilde{w}_{it}$ be the marginal benefit of time devoted to updating a forecast for agent i at day t. Then, the opportunity cost of time is:

$$C_{it} \equiv \frac{w_{it}^{o}}{\tilde{w}_{it}} \tag{3}$$

and it captures a *"fixed" cost* of updating–that is a cost that has to be incurred regardless of the level of attention choice. If agent i has an opportunity cost below 1, so

$$\tilde{w}_{it} > w_{it}^{o}, \tag{4}$$

she finds it worthwhile to update on day t.

[14]In contrast, in leading dynamic rational inattention models (e.g., Steiner et al., 2017) the decision-maker faces some exogenously *fixed* capacity and, given this, chooses and commits at $t = 0$ to a full contingent plan of which signals to observe in each period.

5.1.2 Decision of How Much Attention to Allocate

Agents update in order to improve accuracy, or, equivalently, reduce Mean Squared Forecast Error (MSFE). The day t MSFE is defined as:

$$MSFE_{it} = E[(y_m - f_{it})^2], \tag{5}$$

where f_{it} is agent i's forecast of monthly inflation y_m on day t.

At the beginning of day t an agent chooses how much attention k_{it} to allocate in order to minimize the sum of future MSFEs, conditional on updating. The MSFE on day t can in principle depend on all past attention choices as well as the current choice. Let $k_i^t = \{k_{i1}, \ldots, k_{it}\}$ denote this sequence of choices. The optimization problem for agent i on day t can thus be stated as follows:[15]

$$\min_{k_{i\tau}} \sum_{\tau=t}^{T} \left(\frac{w_{i\tau}}{2\ln 2} MSFE_{i\tau}(k_i^\tau) + c_{i\tau} k_{i\tau} \right) \mathbb{1}_{\tilde{w}_{i\tau} > w_{i\tau}^o}, \tag{6}$$

where T is the number of working days in the month. The parameter $c_{i\tau}$ is the marginal cost of attention and $w_{i\tau}$ is the marginal disutility of MSFE.

In what follows, we describe how to obtain an analytical expression for $MSFE_{i\tau}$ in equation (6) in terms of attention and then solve the optimization problem. We proceed in five steps: Step 1 specifies agents' statistical model; Step 2 specifies agents' information set–the information they need to process to feed into their statistical model; Step 3 derives an analytical expression of $MSFE_{i\tau}$ which depends on the statistical model and the precision of information; Step 4 expresses precision and ultimately $MSFE_{i\tau}$ as a function of attention and Step 5 derives an expression for optimal attention.

Step 1. Specifying a Statistical Model: Monthly inflation y_m –the difference between the log of the price index at the end of the current month and that at the end of the previous

[15]The division by $2ln2$ in the minimization problem is a redefinition of w_{it} which makes equations easier to read.

month–can be written as the sum of daily inflation x_t (the difference between the logs of the prices on days t and $t-1$):

$$y_m = \sum_{t=1}^{T} x_t. \tag{7}$$

We assume that agents model monthly inflation as an ARMA, which implies (using results from temporal aggregation of ARMA models, e.g., Amemiya and Wu, 1972) that daily inflation is also an ARMA, and the orders and parameters of the two models can be related analytically. In the case of Brazil, the ARMA that best fits monthly inflation according to the BIC is an ARMA(1,1):

$$y_m = a + \psi y_{m-1} + v_m + \theta v_{m-1}, v_m \sim i.i.d.\mathcal{N}(0, \sigma_v^2), \tag{8}$$

which implies that daily inflation is an AR(1):

$$x_t = b + \phi x_{t-1} + \varepsilon_t, \varepsilon_t \sim i.i.d.\mathcal{N}(0, \sigma_\varepsilon^2), \tag{9}$$

with $b = \frac{a(1-\psi^{1/T})}{T(1-\psi)}$, $\phi = \psi^{1/T}$ and $\sigma_v^2 = (1 + (1+\phi)^2 + ... + (1 + \phi + \phi^2 + ... + \phi^{T-1})^2)\sigma_\varepsilon^2$.

Note that there are two dynamic dimensions in our setting: the month-to-month problem of forecasting current-month inflation at the beginning of the month, and the within-month problem of updating the forecast. We make the two problems coherent by assuming that the initial forecast is based on the ARMA(1,1) in (8) and the updates are based on the AR(1) in (9). Our main focus here is on the within-month dynamic problem.

Step 2: Specifying Agents' Information Set: During the forecasting period, agents can not only collect and process private information, but also have access to public information. We make the following assumptions.

Assumption 1 (Public Signals) On day t, the public signal contains *past* values of daily inflation:[16]

$$s_p^t = \{x_{t-1}, x_{t-2}, ...\}. \tag{10}$$

[16] This assumes perfect observability of past inflation realizations, but measurement error could be easily accommodated. This would, on the one hand, add a parameter that would give more flexibility in the estimation, but, on the other, would require additional assumptions about the relative accuracy of public and private signals.

Assumption 2 (Private Signals) On day t, *current* daily inflation x_t is not observed but agent i can obtain a noisy signal s_{it} about it.

Assumptions 1 and 2 imply that the agent's information set on day t is:

$$s_i^t \equiv (s_{it}, s_p^t) = \{s_{it}, x_{t-1}, x_{t-2}, ...\} . \tag{11}$$

The precision of the current private signal is endogenous and depends on how much attention the agent decides to allocate to information gathering and processing.

Assumptions 1 and 2 are crucial in making the optimization problem tractable. The next two steps show that these assumptions imply that $MSFE_{i\tau}(k_i^\tau)$ in (6) is only a function of current attention $k_{i\tau}$ and not of the entire sequence of past choices. This implies that the agent's problem is equivalent to a myopic choice and the dynamic problem turns into a sequence of static problems. We discuss and motivate these assumptions in Section 5.3.

Step 3. Forecast Updates and MSFE: The optimal forecast is the conditional expectation of y_m, based on the information set available to the agent. The initial forecast is based on the ARMA(1,1) model in equation (8) and is given by $E\left[y_m|y_{m-1}, y_{m-2}, ...\right]$. This corresponds to an initial MSFE that is constant across agents and months:

$$MSFE_0 = \sigma_v^2 = [1 + (1+\phi)^2 + ... + (1 + \phi + \phi^2 + ... + \phi^{T-1})^2]\sigma_\epsilon^2. \tag{12}$$

On a given day $1 \leq t \leq T$ of the month, the forecast update is the conditional expectation of y_m based on each agent's information set s_i^t, $E\left[y_m|s_i^t\right]$. Combining (11), (7) and (9) implies:

$$E\left[y_m|s_i^t\right] = \sum_{j=1}^{t-1} x_j + \sum_{j=t}^{T}\left(\frac{b(1-\phi^{j-t})}{1-\phi} + \phi^{j-t}E\left[x_t|s_i^t\right]\right), \tag{13}$$

so that $MSFE_{it} = E\left[(y_m - E\left[y_m|s_i^t\right])^2\right] = \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \left(\sum_{j=t}^T \phi^{j-t}\right)^2 E\left[(x_t - E[x_t|s_i^t])^2\right]$, or

$$MSFE_{it} = \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \left(\sum_{j=t}^{T} \phi^{j-t}\right)^2 E\left[\sigma^2_{x_t|s_i^t}\right], \quad (14)$$

where $\sigma^2_{x_t|s_i^t}$ denotes the conditional variance of x_t.

Note that Assumptions 1 and 2 imply that the MSFE depends only on the conditional variance of the current signal and not on past private signals.

Step 4. MSFE as a Function of Attention: We now express the expectation $E\left[\sigma^2_{x_t|s_i^t}\right]$ in equation (14) in terms of attention. Following the rational inattention literature, the *additional* information content of the signal s_{it} is captured by the relative conditional entropy based on the information sets with and without the signal (respectively (s_{it}, s_p^t) and s_p^t):

$$I(x_t; s_{it}|s_p^t) = H(x_t|s_p^t) - E_{s_{it}}[H(x_t|s_{it}, s_p^t)|s_p^t] \leq k_{it}. \quad (15)$$

In our Gaussian-quadratic objective framework it is well-known that the optimal distribution of signals is normal. Then, under the assumption that x_t and s_i^t have a joint normal distribution, the conditional entropy of $x_t|s_i^t$ is: $H(x_t|s_i^t) = \frac{1}{2}\log_2(2\pi e \sigma^2_{x_t|s_i^t})$. The inequality (15) holds with equality because the agent exhausts all capacity. The agent chooses the distribution of s_{it} so that $\sigma^2_{x_t|s_p^t, s_{it}}$ (which is $\sigma^2_{x_t|s_i^t}$) is the same for each signal realization s_{it} and, hence, $E\left[\sigma^2_{x_t|s_i^t}\right] = \sigma^2_{x_t|s_i^t}$. This implies that $\frac{\sigma^2_{x_t|s_p^t}}{\sigma^2_{x_t|s_p^t, s_{it}}} = 2^{2k_{it}}$ or

$$E\left[\sigma^2_{x_t|s_i^t}\right] = \sigma^2_{x_t|s_i^t} = \sigma^2_{x_t|s_p^t}(2^{2k_{it}})^{-1}. \quad (16)$$

By substituting (16) into (14) and using the fact the AR(1) model implies $\sigma^2_{x_t|s_p^t} = \sigma_\varepsilon^2$, we obtain:

$$MSFE_{it} = \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \left(\sum_{j=t}^{T}\phi^{j-t}\right)^2 \sigma_\varepsilon^2 (2^{2k_{it}})^{-1}. \quad (17)$$

Equation (17) shows that the $MSFE_{it}$ depends only on current attention k_{it} and not on past attention choices and it has two components: The first is common across agents and captures the resolution of uncertainty due to the public signal. The second depends on how much attention each agent allocates to obtaining a better signal for current-day inflation (i.e., on the choice of k_{it}), and on how this feeds into the monthly forecast.

Step 5. Optimal Attention: As usual, we solve the problem of sequential decisions backwards. Consider agent i's problem at the last period $t = T$:

$$\min_{k_{iT}} \left(\frac{w_{iT}}{2\ln 2} MSFE_{iT} + c_{iT}k_{iT}\right)\mathbb{1}_{\tilde{w}_{iT} > w^o_{iT}} \text{ subject to } k_{iT} \geq 0, \text{ (16) and (17).} \quad (18)$$

The agent can only control the part of the $MSFE$ in (17) that depends on collecting information about the current daily signal, so optimal attention solves:

$$\min_{k_{iT}} \left(\frac{w_{iT}}{2\ln 2}\sigma_\varepsilon^2 \left(2^{2k_{iT}}\right)^{-1} + c_{iT}k_{iT}\right)\mathbb{1}_{\tilde{w}_{iT} > w^o_{iT}} \text{ s.t. } k_{iT} \geq 0.$$

Differentiating with respect to k_{iT} and rearranging gives optimal attention as:[17]

$$k^*_{iT} = \begin{cases} \frac{1}{2}\log_2\left(\frac{w_{iT}}{c_{iT}}\sigma_\varepsilon^2\right) & \text{if } \frac{w_{iT}}{c_{iT}}\sigma_\varepsilon^2 > 1 \text{ and } \tilde{w}_{iT} > w^o_{iT} \\ 0 & \text{otherwise} \end{cases}$$

As anticipated, our assumption that the public signal is the realization of the variable implies that it is by construction weakly more precise than any past private signal which would require infinite attention to be as precise. The agent knows that t's choice of attention does not affect $t+1$'s decision, because the prior uncertainty that she reduces the following day is not based

[17]Note that if $\tilde{w}_{iT} < w^o_{iT}$ the objective is constant and equal to zero so any choice of k is optimal, and we choose $k^*_{iT} = 0$.

on past private signals, but on the more precise public signal. Hence, the agent's dynamic problem turns into a sequence of static problems. In particular, at the beginning of day t an agent solves:

$$\min_{k_{it}} \left(\frac{w_{it}}{2\ln 2} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 \left(2^{2k_{it}}\right)^{-1} + c_{it} k_{it} \right) \mathbb{1}_{\tilde{w}_{it} > w_{it}^o} \quad \text{s.t. } k_{it} \geq 0, \tag{19}$$

which gives

$$k_{it}^* = \begin{cases} \frac{1}{2}\log_2 \left(\frac{w_{it}}{c_{it}} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 \right) & \text{if } \frac{w_{it}}{c_{it}} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 > 1 \text{ and } \tilde{w}_{it} > w_{it}^o \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

The formula implies that attention is higher the larger the current benefit-cost ratio, the earlier the day is in the month, and the larger the prior variance of the signal (measured by σ_ε^2). Optimal attention varies over time because of two reasons. The first has to do with the resolution of uncertainty due to revelation of the public signal, which is common across agents. The second is due to agents possibly facing different cost/benefit ratios on different days.

5.2 Bringing the Model to the Data

We now link the two decisions described in the previous sections and impose normalization restrictions that enable us to identify the structural parameters from the observables in the data: the fraction of updaters and their MSFE.

Since, due to data limitations, we equate updating with forecast changes (which require positive attention), we cannot separately identify the parameters driving the decision to update and those affecting the choice of attention. We therefore link the two decisions, by assuming that the same marginal benefit parameter drives *both* decisions:

Assumption 3 $\tilde{w}_{it} = w_{it}$.

The parameters of the model can be divided into ARMA parameters and incentive parameters. The estimates of the ARMA parameters can be used to assess the external validity of the model. The variation in incentive parameters on the contest and the IPCA15 days is identified by the shifts in observables on these days. Further normalizations are needed if one wants to identify the incentive parameters, because the decisions to update and the choice of attention respectively depend on the ratios $\frac{w_{it}}{w_{it}^o}$ and $\frac{w_{it}}{c_{it}}$. We report results for two sets of such restrictions: a "shifting benefits" and a "shifting costs and benefits" model.

5.2.1 Shifting Benefits

Here we assume that the only variation in parameters is in the benefits of updating:

Assumption 4 $w_{it}^o = w^o$ and $c_{it} \equiv c$ are constant across agents and time.

To eliminate w^o as a free parameter, we make the following normalization assumption.

Assumption 5 The fixed cost of updating ($C_{it} = \frac{w^o}{w_{it}}$) is small relative to the marginal cost of attention, so an agent who pays the fixed cost (i.e., $w_{it} > w^o$) also finds it worthwhile to pay the marginal cost and choose positive attention (i.e., $k_{it}^* > 0$).

Any choice of w^o such that $w^o \geq \max_{1 \leq t \leq T} \frac{c}{\left(\sum_{j=t}^{T} \phi^{j-t}\right)^2 \sigma_\varepsilon^2} = \frac{c}{\sigma_\varepsilon^2}$ satisfies Assumption 5, because in this case if $w_{it} > w^o$, the first condition for a positive k_{it}^* in (20) is also satisfied. We thus impose the following normalization:

$$w^o = \frac{c}{\sigma_\varepsilon^2}. \tag{21}$$

This restriction is not essential to the conclusions of the analysis. In unreported results, we impose the alternative normalization $w^o = 1$ and estimate both this and the model below. The conclusions regarding the time variation in incentive parameters and the estimates of the ARMA parameters are essentially unaffected.

The parameters of this model are: $\theta = (\mu_t, \sigma_w^2, c, \phi, \sigma_\epsilon^2)$ and the theoretical counterparts of the observables in the data are as follows.

Fraction of updaters: We model the heterogeneity in the benefits by assuming that the cross-sectional distribution for w_{it} is a truncated normal $TN(\mu_t, \sigma_w^2)$. The particular distributional choice is not crucial because we focus on how the cross-sectional mean changes on the days around the contest. The fraction of updaters on day t then equals the probability that $w_{it} > \frac{c}{\sigma_\varepsilon^2}$ implied by the truncated normal:

$$\lambda_t = P\left(w_{it} > \frac{c}{\sigma_\varepsilon^2}\right). \tag{22}$$

MSFE: Substituting (20) into (17), together with (19) and the fact that agents who don't update maintain their previous forecast, gives the optimal MSFE for agent i on day t as:

$$MSFE_{it}^* = \begin{cases} \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \frac{c}{w_{it}} & \text{if } w_{it} > \frac{c}{\sigma_\varepsilon^2} \\ MSFE_{it-1}^* \text{ otherwise} & \end{cases}. \tag{23}$$

5.2.2 Shifting Costs and Benefits

Here we allow the marginal cost c to be different on the IPCA15 day but, for identification purposes, restrict the time variation in μ_t by modelling it as a step function. The parameters are $\theta = (\mu_t, \sigma_w^2, c_t, \phi, \sigma_\epsilon^2)$, where[18]:

$$\mu_t = \begin{cases} \mu_B & \text{if } t < CD - 1 \\ \mu_{CD-1} & \text{if } t = CD - 1 \\ \mu_{CD} & \text{if } t = CD \\ \mu_A & \text{if } t > CD \end{cases} \tag{24}$$

$$c_t = \begin{cases} c_{IPCA15} & \text{if } t = IPCA15 \\ c & \text{otherwise.} \end{cases}$$

The moments are (22) and (23) with c substituted by c_t. Note that shifts in benefits and/or costs here affect both the opportunity cost of updating and the amount of attention.

[18]The additional restrictions $\mu_B = \mu_A$, $\mu_{CD-1} = \mu_{CD}$, $c_{IPCA15} = c$ are rejected by the data.

5.3 Discussion of Assumptions

For the sake of tractability, we make three main simplifying assumptions.

The first simplification is that we consider a decision-theoretic rather than a strategic setting. This assumption buys tractability and is well-grounded, because the empirical regularities are at odds with the key predictions of strategic models.

The second simplification is the assumption that the decision of whether or not to update boils down to comparing the opportunity cost of time to a fixed threshold. Similar decision rules are considered in the literature on price setting with menu costs (e.g., Midrigan, 2011; Gertler and Leahy, 2008), where it is shown that a rule of the form "update price if the menu cost is below some fixed threshold" is a good approximation to the optimal decision rule. We believe this simplification does not diminish the contribution of the paper in terms of providing empirical credence to models of rational inattention, as the same conclusions would have emerged by focusing only on the intensive margin of updating and validating this part of the model using the patterns of accuracy in the data. As discussed by Woodford (2009) in a different context, applying rational inattention to a timing decision presents additional challenges, so attempting to solve a joint rational inattention optimization for the two decisions would substantially complicate the analysis without necessarily adding new insights. We consider both decisions in this paper for quantitative realism and because both extensive and intensive margins matter for policy and survey design. We can investigate these implications because the two decisions are endogenous in our model, since they are driven by the same incentive parameters.

The third simplification is that updaters rely on past public information rather than past private signals (Assumptions 1 and 2). Arguably, Assumptions 1 and 2 are more realistic and plausible than the opposite extreme of assuming that updaters only use past private signals: If we want the theory to be generally applicable, we should note that our survey is an exception in its high frequency–typical surveys are collected monthly or quarterly. When updating at these frequencies, agents surely have access to past official monthly or quarterly data releases. It is doubtful that agents would rather use a sequence of incomplete and less accurate private signals instead of the complete set of accurate public information. Even in our high-frequency case, there is plausibly public information about past daily inflation, for example daily releases of gasoline prices. These assumptions make the model tractable by turning the dynamic rational inattention problem into a sequence of static decisions. Otherwise, the choice of attention would depend on all past and current cost and benefit parameters. The prediction of our model that the accuracy of updaters on a given day only depends on the incentives on that day and not on past updating choices is supported by the data. Table 4 shows the result of a panel regression for the accuracy of updaters considering the same regressors as Table 2, column (1).

Table 4. Drivers of Accuracy for Updaters

Regressors	**Panel Fixed Effect Coefficients**
d_t^{CD-1}	0.00049
	(0.00048)
d_t^{CD}	0.00208***
	(0.00035)
d_t^{IPCA15}	0.00331***
	(0.00045)
$d_t^{IPCA15+1}$	0.00467***
	(0.00039)
d_t^{MPC}	-0.00069
	(0.00054)
$duration_t$	−0.00006*
	(0.00004)
$horizon_t$	−0.00087***
	(0.00003)
$constant$	−0.00200***
	(0.00035)

Notes: Dependent variable is minus the squared forecast error for agent i on day t. Sample from January 8th, 2004 to January 8th, 2015. Number of observations 31,319. Standard errors in parentheses. ***, ** and * indicate, respectively, significance at the 1%, 5% and 10% level.

Table 4 shows that the duration between updates (which is linked to past updating choices), has an effect that is only significant at the 10% level. In contrast, the contest and IPCA15 dummies have large and highly significant effects, as our model predicts (see equation (23), which predicts higher accuracy –lower MSFE–due to the higher benefits and lower costs respectively associated with these days).

6 Estimation

We estimate the two versions of the model presented in Section 5.2 by Simulated Method of Moments (SMM) (e.g., Gourieroux and Monfort, 1996; Duffie and Singleton, 1993; Ruge-Murcia, 2012), which involves matching empirical moments with their theoretical counterparts. We only discuss how to simulate theoretical moments for the Shifting Benefits case, as the modifications for the Shifting Benefits and Costs case will be obvious to the reader.

The moments we aim to jointly match are the fraction of updaters and the aggregate $MSFE$ in a window of days around the contest, as reported in Figure 1.[19] The simulation is based on τM months and N agents, where $M = 132$ and $N = 85$ as in the data, and τ is an arbitrary number of replications.[20]

Every month the initial $MSFE$ for all agents is given by (12). On every subsequent working day $t = 1, ..., T$ of the month, each agent receives a random draw of the benefit w_{it} from a $TN(\mu_t, \sigma_w^2)$. The benefit draws determine the fraction of updaters on that day according to (22). The $MSFE_{it}$ for agents who update is given by the first line of (23). Agents who don't update keep their previous $MSFE$, $MSFE_{it-1}$. The same simulation is repeated for all months, changing only the number of working days T and the date of the contest to match those in the corresponding month in the data. The moments we match are the average (over different months) fraction of updaters and the average (over different months) of the average $MSFE$ across agents computed each day within a five-day window around the contest.[21]

Estimation Results: Table 5 shows that the mean benefit of in the "shifting benefits" model is never zero, indicating that agents find it worthwhile to update even outside the contest. The benefit increases slightly on the lead-up to the contest and jumps up on the contest (to a level that is 15% higher than 5 days before). It then jumps down on the IPCA15 to a level comparable to before the contest and then decreases further and settles on an approximately

[19]Results are robust to considering different window lengths; however note that a much larger window would run the risk of going outside the current month, as for some months in the sample the contest day is early or late in the month.

[20]Following Duffie and Singleton (1993), the requirement is to have $\tau M \to \infty$ as $M \to \infty$. The reported results are for $\tau = 5$.

[21]Due to the high correlation among moments we use a diagonal weighting matrix in the SMM estimation.

Table 5. Shifting Benefits Model Estimation

Parameter	Estimate	Standard Error	Parameter	Estimate	Standard Error
μ_{CD-5}	0.470	2.4E-08	μ_{CD+3}	0.467	2.4E-08
μ_{CD-4}	0.475	2.4E-08	μ_{CD+4}	0.464	2.4E-08
μ_{CD-3}	0.479	3.2E-08	μ_{CD+5}	0.462	2.4E-08
μ_{CD-2}	0.483	4.9E-08	σ_w	0.051	2.2E-07
μ_{CD-1}	0.491	2.3E-08	c	1.37E-05	1.5E-08
μ_{CD}	0.539	2.1E-08	ϕ	0.922	2.9E-09
μ_{CD+1}	0.483	2.3E-08	σ_ε	0.005	5.4E-06
μ_{CD+2}	0.470	2.4E-08			

Note: p-value of the J-test = 0.99.

Table 6. Shifting Costs and Benefits Model Estimation

Parameter	Estimate	Standard Error	Parameter	Estimate	Standard Error
μ_B	0.476	4.1E-18	c_{IPCA15}	1.34E-05	1.4E-13
μ_{CD-1}	0.488	4.0E-18	c	1.38E-05	1.4E-13
μ_{CD}	0.542	3.6E-18	ϕ	0.925	2.1E-18
μ_A	0.463	4.2E-18	σ_ε	0.005	3.9E-16
σ_w	0.058	3.4E-17			

Note: p-value of the J-test = 0.49.

constant level that is lower than before the contest. The model fits the data well and it passes the J-test of overidentifying restrictions with a p-value close to 1. The most remarkable finding is that the estimates of the AR(1) parameters–which are not part of the moments matched by the estimation–are very close to the estimates of the same parameters in Brazilian inflation data:[22] The autoregressive coefficient equals .922 in the model and .963 in the data; the error standard deviation equals 5.0E-03 in the model and 3.36E-03 in the data.

The estimation results for the "shifting costs and benefits" model (estimated using a window of 3 days around the contest) are in Table 6. Figure 6 shows the fit of this model.

[22]We obtain the estimates of the AR(1) parameter for (unobservable) daily inflation by estimating an ARMA(1,1) on observable monthly inflation data in Brazil from January 2004 to December 2014 and assuming 21 working days in each month, then using the formulas after equation (9) to back out the AR(1) parameters.

Figure 6. Shifting Costs and Benefits Model Predictions versus Data. Fraction of Updaters (Left) and Aggregate MSFE (Right)

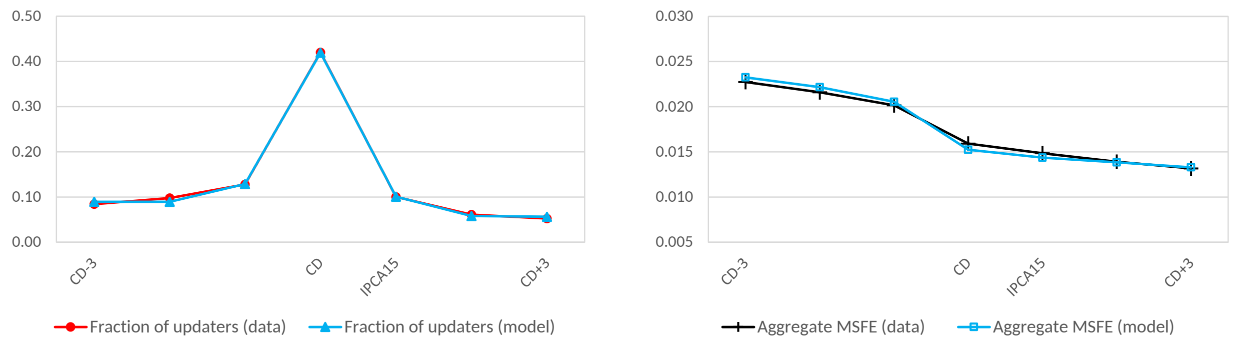

The model also passes the J-test with a p-value of 0.49 and gives almost identical estimates for the AR(1) parameters. The estimates confirm that there is a constant incentive for participation to the survey on "normal" days but the contest provides an additional benefit. The additional benefit is high on the contest day but is also present before the contest (since a forecast submitted before the contest but not updated still counts for the contest). The cost of attention is lower on the IPCA15 day.

7 Counterfactuals

In this section we perform a number of counterfactual analyses, using the estimates of the Shifting Costs and Benefits model reported in Table 6.

Aggregate Accuracy and Changes in Extensive versus Intensive Margins: The estimates from Table 6 imply that the average $MSFE$ across agents falls from 0.0205 the day before the contest to 0.0152 on the contest day, which is due to both an increase in the number of updaters and to a shift in the benefit distribution across agents (so agents who update put more effort). We then assume that the number of updaters remains the same as before the contest, but that they receive draws from the shifted distribution of benefits that characterizes the contest. This would make the $MSFE$ fall to 0.0189, which implies that 30% of the accuracy improvement on the contest is due to agents paying more attention (intensive margin) and 70% to more agents updating (extensive margin).

Quantifying the Value of the Contest: To assess the value of the contest, we let the contest fall on each possible day in a five-day window around the IPCA15 and generate counterfactual $MSFE$s as in (23), using the estimates from Table 6. Figure 7 reports the reduction in the average $MSFE$ across agents on each potential contest day (relative to the previous day) and compares it to the reduction in average $MSFE$ that one would observe between two consecutive days in the absence of the contest. We assume that the mean benefit without a contest would be the constant benefit we now only observe after the contest, i.e., $\mu_t = \mu_A$ for all t. Figure 7 shows that having the contest on any day has a very large effect on accuracy. The largest accuracy improvement is obtained by having the contest on the IPCA15 day. It amounts to a 347% accuracy improvement relative to not having a contest[23] and a 31% improvement relative to having the contest the day before, as it is now in the survey.

Figure 7. MSFE Daily Reduction With and Without the Contest on Different Days

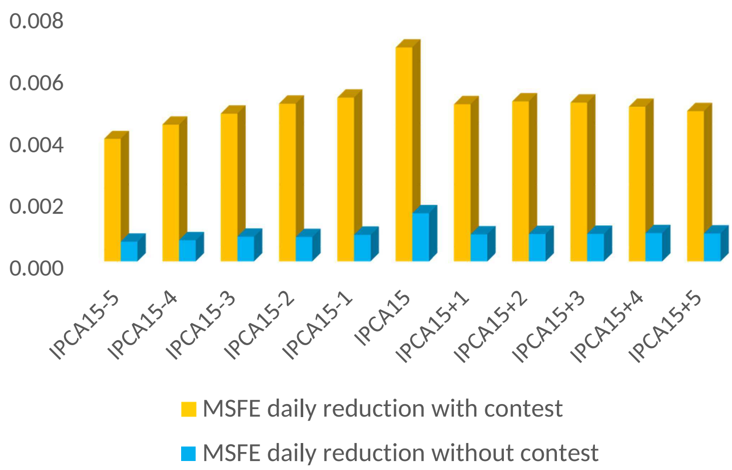

[23]Even without a contest the IPCA15 day would benefit from a larger MSFE reduction than other days, due to the lower cost.

Figure 8. Cumulative MSFE (Left) and MSFE (Right) with Contest on Different Days

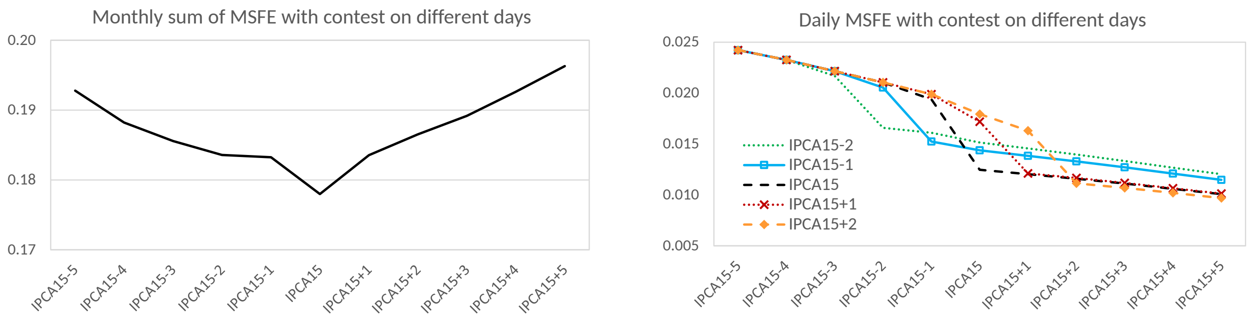

Optimal Timing of the Contest: We next determine the optimal timing of the contest. The right panel of Figure 8 shows the counterfactual *daily* average $MSFE$ when the contest is on different days within the five-day window. The daily $MSFE$ decreases almost linearly, and the contest induces a downward shift in the line. The left panel of Figure 8 plots the counterfactual *cumulative* $MSFE$ (the monthly sum of the average $MSFE$). The figure shows that the optimal timing in terms of both daily and cumulative improvements in accuracy is the IPCA15. The percentage improvement in cumulative $MSFE$ of having the contest on the IPCA15 instead of on the day before (as it is currently in the survey) would be 3%. Note that the cumulative $MSFE$ is a U-shaped curve around the IPCA15. Intuitively, this is because there is a tradeoff between holding the contest earlier in the month, when agents observe fewer past signals but there are more days left to lower the path of the $MSFE$, and holding it later, when more signals are observed but there are fewer remaining days to lower the $MSFE$.

Contest versus No-Contest: We finally investigate the extent to which the contest crowds-out updates–or more precisely mis-aligns updates from the more "natural" IPCA15 day when it is cheaper to process information. One challenge we face is that we do not know how the total number of updates within a month would change in the absence of the contest. In what follows we assume that this number stays constant. This stacks the cards in favour of the no-contest scenario, since one would expect fewer updates without the contest. Holding the total number of updates in a month fixed, we use the model to back out the parameter μ (constant on each day of the month but possibly varying from month to month) that would deliver the same

Figure 9. Fraction of Updaters and Accuracy With and Without the Contest.

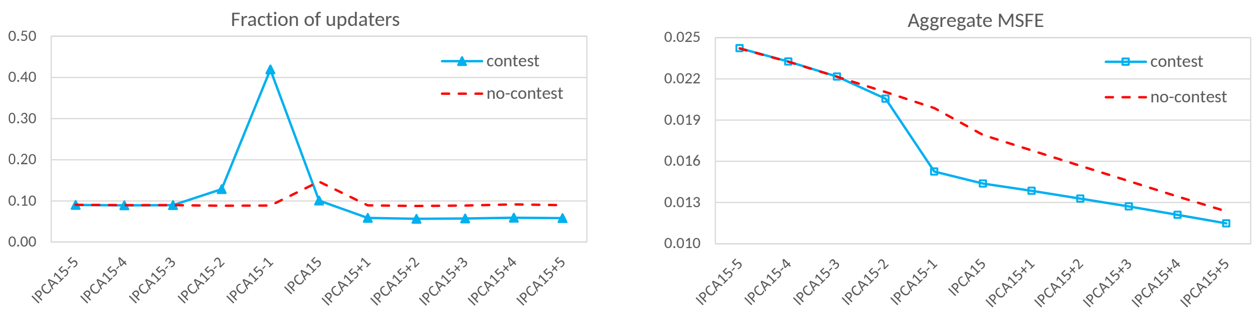

number of updates as in the estimated model. We maintain the same estimates for the other parameters as in Table 6. Then, the IPCA15 is the date with the smallest cost-to-benefit ratio.

Using this counterfactual parameter we then simulate the fraction of updaters and the average MSFE for each agent and compare them to those obtained in the presence of the contest. Figure 9 reports the results. The results are eye-opening: Accuracy is worse overall without the contest, even though in that case most updates happen on the IPCA15 and there are more updates in the days after the IPCA15 than in the presence of the contest. This underscores that the coordinated updates that occur because of the contest are important for the survey's average accuracy.

8 Conclusions

This paper connects two important ideas in economics: that attention is limited and that incentives matter. We analysed panel data from a unique survey of professional forecasters where the forecast updating decisions of participants are observable and incentive "shifters" are present. The empirical findings are consistent with a rational inattention model in which forecast updating in general, and inattention in particular, are incentive-driven decisions.

Kacperczyk et al. (2016) conclude "*While information choices have consequences for real outcomes that are poorly understood because they are difficult to measure.*" Our model has predictions for the observables in our data and thus ties information choices to outcomes.

The empirical patterns we document and the counterfactuals underscore the importance of a contest, and formal incentives more broadly, for accuracy. The role of competition among forecasters on the quality of forecasts is also underlined in the influential book Tetlock and Gardner (2016) within the framework of forecasting election outcomes. We show that a contest makes more forecasters participate and each forecaster put more effort, resulting in an increase in both individual and aggregate accuracy. Aligning the contest with information releases would lead to further increases in accuracy. Our findings can be of interest to central banks and private institutions that run surveys of professional forecasters. Such surveys are increasingly becoming a key input in economic and policy decisions by governments and firms. Despite that many policy institutions worldwide have been running surveys for years and private sector surveys are a thriving and growing industry,[24] virtually no attention has been paid in the literature to how survey design affects forecast quality.

More broadly, our results have implications for general settings where a collection of agents have limited resources to devote to processing information in order to make a decision. Examples are soliciting expert opinions, choices of employees' savings, retirement plans and investment choices. Our results suggest that people devote more attention when they compete. For example, a contest for best retirement portfolio returns among employees of an organization could encourage more attention and active participation.

[24]Interestingly, most private firms focus on nowcasts (short-term) forecasts, as we do in this paper.

References

Amemiya, T. and R. Y. Wu (1972): "The effect of aggregation on prediction in the autoregressive model," *Journal of the American Statistical Association*, 67, 628–632.

Caplin, A. and M. Dean (2013): "Behavioral implications of rational inattention with shannon entropy," Tech. rep., National Bureau of Economic Research.

——— (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.

Caplin, A., J. Leahy, and F. Matejka (2016): "Rational Inattention and Inference from Market Share Data," *Working Paper*.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017): "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments," *American Economic Journal: Macroeconomics*, 9, 1–35.

Cheremukhin, A., A. Popova, A. Tutino, et al. (2011): "Experimental evidence on rational inattention," *Federal Reserve Bank of Dallas Working Paper*, 1112.

Coibion, O. and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120, 116 – 159.

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644 – 2678.

Csaba, D. (2018): "Attentional Complements," Tech. rep., Working Paper, NYU.

Dean, M. and N. Neligh (2017): "Experimental tests of rational inattention," Tech. rep., Working Paper, Columbia University.

Duffie, D. and K. J. Singleton (1993): "Simulated Moment Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–952.

Gertler, M. and J. Leahy (2008): "A Phillips curve with an Ss foundation," *Journal of Political Economy*, 116, 533–572.

Glaeser, E. L., A. Hillis, S. D. Kominers, and M. Luca (2016): "Crowdsourcing city government: Using tournaments to improve inspection accuracy," *American Economic Review*, 106, 114–18.

Gourieroux, C. and A. Monfort (1996): *Simulation-Based Econometric Methods*, Oxford University Press.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp (2016): "A rational theory of mutual funds' attention allocation," *Econometrica*, 84, 571–626.

Lazear, E. P. (2000): "Performance Pay and Productivity," *The American Economic Review*, 90, 1346–1361.

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

Maćkowiak, B. and M. Wiederholt (2015): "Business cycle dynamics under rational inattention," *The Review of Economic Studies*, 82, 1502–1532.

Marinovic, I., M. Ottaviani, and P. N. Sørensen (2013): "Forecasters' objectives and strategies," *Handbook of economic forecasting*, 2, 691–720.

Marques, A. B. C. (2013): "Central Bank of Brazil's market expectations system: a tool for monetary policy," *IFC Bulletins - Bank for International Settlements*, 36, 304–324.

Martin, D. (2016): "Rational inattention in games: experimental evidence," Tech. rep.

Midrigan, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.

Reeve, J., B. C. Olson, and S. G. Cole (1985): "Motivation and performance: Two consequences of winning and losing in competition," *Motivation and Emotion*, 9, 291–298.

Ruge-Murcia, F. (2012): "Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles," *Journal of Economic Dynamics and Control*, 36, 914–938.

SHEARER, B. (2004): "Piece rates, fixed wages and incentives: Evidence from a field experiment," *The Review of Economic Studies*, 71, 513–534.

SHORROCKS, A. F. (1978): "The measurement of mobility," *Econometrica: Journal of the Econometric Society*, 1013–1024.

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

STEINER, J., C. STEWART, AND F. MATĚJKA (2017): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica*, 85, 521–553.

SYVERSON, C. (2011): "What determines productivity?" *Journal of Economic Literature*, 49, 326–365.

TETLOCK, P. E. AND D. GARDNER (2016): *Superforecasting: The art and science of prediction*, Random House.

WOODFORD, M. (2009): "Information-constrained state-dependent pricing," *Journal of Monetary Economics*, 56, S100–S124.

——— (2013): "Macroeconomic analysis without the rational expectations hypothesis," Tech. rep., National Bureau of Economic Research.