ISSN 1518-3548

BANCO CENTRAL DO BRASIL

Incentive-driven Inattention

Wagner Piazza Gaglianone, Raffaella Giacomini, João Victor Issler and Vasiliki Skreta

October 2018

Working Papers

485

ISSN 1518-3548
CGC 00.038.166/0001-05

Working Paper Series	Brasília	no. 485	October	2018	p. 1-43

Working Paper Series

Edited by the Research Department (Depep) – E-mail: workingpaper@bcb.gov.br

Editor: Francisco Marcos Rodrigues Figueiredo – E-mail: francisco-marcos.figueiredo@bcb.gov.br

Co-editor: José Valentim Machado Vicente – E-mail: jose.valentim@bcb.gov.br

Head of the Research Department: André Minella – E-mail: andre.minella@bcb.gov.br

The Banco Central do Brasil Working Papers are all evaluated in double-blind refereeing process.

Authorized by Carlos Viana de Carvalho, Deputy Governor for Economic Policy.

General Control of Publications

Banco Central do Brasil
Comun/Divip
SBS – Quadra 3 – Bloco B – Edifício-Sede – 2º subsolo
Caixa Postal 8.670
70074-900 Brasília – DF – Brazil
Phones: +55 (61) 3414-3710 and 3414-3565
Fax: +55 (61) 3414-1898
E-mail: identidadevisual.comun@bcb.gov.br

Citizen Service Division

Banco Central do Brasil
Deati/Diate
SBS – Quadra 3 – Bloco B – Edifício-Sede – 2º subsolo
70074-900 Brasília – DF – Brazil
Toll Free: 0800 9792345
Fax: +55 (61) 3414-2553
Internet: http//www.bcb.gov.br/?CONTACTUS

Non-technical Summary

The findings in this paper connect two central ideas in economics: that attention is limited and incentives matter. Information frictions play a key role in realistically capturing decision-making and have been at the forefront of modern macroeconomics and finance. Limited attention is an important source of such frictions and has been the subject of extensive study in psychology and economics starting with the influential work of Kahneman (1973).

The rational inattention paradigm (Sims, 2003) is a convenient and tractable theoretical framework for capturing how agents optimally allocate their limited attention resources when making decisions, such as choosing prices or investment portfolios. Nevertheless, what influences the amount of attention that forward-looking agents avail at each point in time has been less studied theoretically, let alone empirically.

This paper builds a model of rational inattention, where agents have a budget of attention, or cognitive resources, to devote to updating expectations. Both the amount of attention and the expectations' accuracy are endogenous and linked to the cost and benefit of updating, which can vary across agents and over time. We structurally estimate the model using data from a unique panel of professional forecasters – the Central Bank of Brazil's Focus survey – in which agents choose when to update and there is a recurring contest (Top5) that ranks agents based on their accuracy.

The empirical facts that we document are novel and serve as a guide for building the model: The incentives linked to the contest are the primary drivers of updates and accuracy gains. The model fits the data well and allows us to perform counterfactual analysis to understand the value of the contest and to investigate alternative survey designs.

Sumário Não Técnico

Este artigo conecta duas ideias centrais em economia: que atenção é limitada e incentivos são importantes. Fricções informacionais desempenham um papel fundamental nas decisões de agentes econômicos e têm estado na fronteira da macroeconomia e das finanças modernas. A atenção limitada é uma fonte importante de tais fricções e tem sido objeto de extenso estudo em psicologia e economia desde o influente trabalho de Kahneman (1973).

O paradigma da inatenção racional (Sims, 2003) é um arcabouço teórico conveniente e tratável para analisar como os agentes alocam de maneira ótima seus recursos de atenção limitada ao tomar decisões, tais como escolher preços ou carteiras de investimento. No entanto, o que influencia a quantidade de atenção que os agentes usam em cada instante no tempo tem sido tema menos estudado teoricamente, e muito menos empiricamente.

Este artigo desenvolve um modelo de inatenção racional, onde os agentes têm uma dotação orçamentária de atenção, ou de recursos cognitivos, para dedicar à atualização de suas expectativas. Tanto o grau de atenção quanto a precisão das expectativas são endógenos e vinculados ao custo e benefício de atualização, que pode variar entre agentes e ao longo do tempo. O modelo é estimado estruturalmente usando dados de painel de analistas profissionais – a pesquisa Focus do Banco Central do Brasil – no qual os agentes escolhem quando atualizar e há um concurso recorrente (*Top5*) que classifica os agentes com base na acurácia de suas expectativas.

As evidências empíricas documentadas no artigo são novas e servem de guia para a construção do modelo: os incentivos ligados ao concurso são os principais fatores responsáveis pela atualização das expectativas e ganhos de precisão. O modelo apresenta uma boa aderência aos dados e permite realizar análises contrafactuais para entender o valor do concurso e investigar desenhos alternativos de *survey*.

Incentive-driven Inattention*

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

Abstract

This paper establishes a link between incentives and limited attention in economic decision-making. We explore how agents' attention reacts to incentives versus the arrival of information and how each factor affects the quality of decisions. We analyze a unique survey dataset where professional forecasters decide when to update a forecast and there is a formal incentive in the form of a contest rewarding forecast accuracy. There is also a major piece of information arriving right after the contest. We empirically establish that the contest is the primary driver of updating decisions and accuracy improvements. Then, we develop and structurally estimate a rational inattention model where agents choose how much attention to allocate to updating. The estimated model fits the data and allows us to perform counterfactuals to quantify the value of the contest and how it affects updates and accuracy, as well as to establish the optimal timing of the contest.

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, Carlos Viana de Carvalho, Philipp Kircher, André Minella, 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.

†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

The findings in this paper connect two central ideas in economics: that attention is limited and that incentives matter. Well-designed incentives can improve productivity by inducing workers to work harder (see, for example Lazear (2000)), CEOs to make better decisions (Murphy (1999)) and firms to innovate more (Manso (2011)). Limited attention is an important source of information frictions and has been the subject of extensive study in psychology and economics, starting with the influential work of Kahneman (1973). The rational inattention paradigm (Sims, 2003) is the leading theoretical framework used to study how agents optimally allocate their limited attention budget when making decisions such as choosing prices (Mackowiak and Wiederholt, 2009) or investment portfolios (Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016).

What affects the amount of attention that agents avail at each point in time has been less studied theoretically, let alone empirically. In other words, we know little about what affects the "attention budget" of economic agents. This question is important since the amount of mental resources devoted to a decision could impact its quality. In this paper we highlight the role of formal incentives by endogenizing the attention budget, and by establishing a strong link between the quality of the decision and the incentives. This link provides a microfoundation for the mental capacity a decision-maker avails to a decision.

The decisions we study in this paper are the sequential decisions of forecasters to update their forecast of a future event during the period before the event occurs. We build a dynamic model of rational inattention that links both the amount of attention dedicated to updating a forecast and its accuracy to the cost of processing information and the benefit derived from being accurate. The benefit can vary depending on the formal and informal incentives a forecaster faces; the cost can vary depending on the arrival of information. The model provides analytical expressions for the dynamic evolution of both the decision to update and the quality of the decision (the forecast accuracy). This enables us to structurally estimate it using panel data from a unique survey of professional forecasters where updating decisions are observable. The key feature of the survey is a formal incentive in the form of a monthly contest. This feature allows us to link attention and incentives. In addition, the survey has clear times when crucial information arrives. This allows us to empirically assess the relative effects of information arrival versus the incentive (the contest) on the probability to update and on accuracy improvements. Finally, our structural estimates allow us to counterfactually evaluate the effect of the contest on aggregate accuracy and to determine the optimal contest day in a number of "survey design" exercises we perform.

The survey of professional forecasters we consider is the little-known *Focus Survey* maintained by the Central Bank of Brazil. We study forecasts of current month's inflation (nowcasts). We exploit two unique features of the data. First, forecasters can update an existing forecast any time they choose–even daily,[1] so we can learn what affects this decision. Second, in addition to the informal incentives that it shares with other surveys of professional forecasters analyzed in the literature, this survey has a formal incentive in the form of a monthly contest that ranks participants based on the accuracy of their forecast on a specific day. For the case of inflation (the IPCA–Broad Consumer Price Index), the contest always occurs the day *before* the release of a key piece of public information about inflation–the IPCA15.[2]

The empirical facts that we document are novel and of separate interest. The first is the striking pattern visible in Figure 1, showing how incentives and information affect *aggregate* behavior: 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 sharper fall in MSFE). In contrast, the information release on the day after the contest appears to have a lesser effect on aggregate behavior, as the fraction of updaters and the aggregate accuracy improvements on the IPCA15 day appear similar to those on any other non-contest day. This is surprising, as one would expect to see more forecasters updating and larger accuracy improvements on a day when it is presumably less costly to process information.

We then investigate the updating behavior and forecast accuracy at the *individual level*, by means of panel regressions. We find evidence that the contest induces more forecasters to update (the *extensive margin*) and each forecaster to generate a more accurate forecast

[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.*

[2] IPCA15 measures inflation from the 15th of one month to the 15th of the next, whereas IPCA measures inflation from the first day of the month to the last.

Figure 1. Aggregate effects of incentives versus information: Daily evolution of the fraction of updaters and aggregate $MSFE$

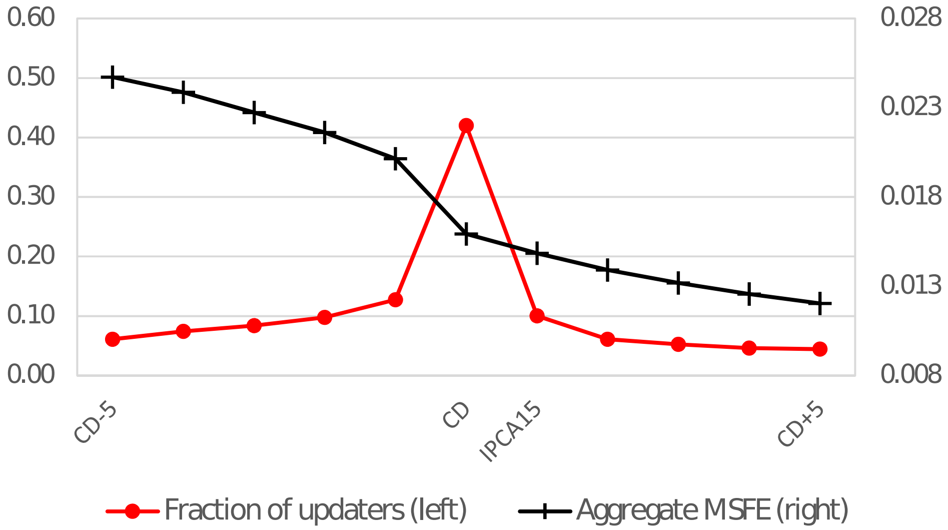

Notes: The figure shows the fraction of forecasters in the Focus Survey who update their nowcast of inflation on a five-day window around the contest (CD) and the information release (IPCA15) days, averaged over all months in the dataset. It also shows the aggregate $MSFE$, which is the average across 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 (which could be a non-updated nowcast) and the realization of inflation for that month, averaged over all months. Accuracy is the negative of the $MSFE$.

(the *intensive margin* of effort). Additional findings provide support for modelling choices: Increases in updating probabilities and in aggregate accuracy on the contest are more likely due to increased attention than to self-selection based on ability; there is heterogeneity both across forecasters and over time in updating behavior and accuracy; forecasters do not appear to act strategically as they provide less dispersed forecasts on the contest day.[3]

Motivated by the empirical findings, we develop a decision-theoretic model where rationally inattentive agents choose how much attention to devote to processing daily information in order to improve the accuracy of their forecast.[4] We assume that each month, agents use a realistic statistical model (an Autoregressive-Moving Average–ARMA model) to produce an initial forecast. Agents then sequentially decide at the beginning of each working day whether to update their forecast and how much attention to allocate. These decisions depend on both the information available up to that time and on the costs and benefits of providing a forecast update. In the model, the information is in the form of both public and private signals. The benefits can vary across agents and over time. The model yields simple and intuitive analytical expressions for the theoretical counterparts of the observables in the data: the fraction of updaters and agents' individual forecast accuracy on each day. The agents' optimal choice of attention implies that each day only a (potentially time-varying) fraction of agents update their forecast and that the accuracy of the updated forecast depends on the amount of attention. In particular, the accuracy on a given day is the sum of a common component due to past public signals and an idiosyncratic component due to agents' attention to process private signals on that day.

We structurally estimate the model by Simulated Method of Moments, in order to match the dynamics of updates and aggregate accuracy reported in Figure 1. Focusing on a time window around the contest allows us to uncover how the parameters characterizing the costs and benefits of updating change over time. The model is parsimoniously parameterized and fits the patterns in Figure 1 well. The estimates indicate that the benefits of updating are higher before the contest than after the contest and spike up substantially on the contest. The cost of processing information is lower on the IPCA15 day. Remarkably, the estimates of the ARMA

[3]As we later elaborate, this phenomenon seems in contrast with the prediction of the strategic model of Marinovic, Ottaviani, and Sørensen (2013).

[4]We later explain in detail why this objective function is the most compelling one.

model parameters are almost identical to those implied by Brazil's inflation data (data that is not used in the estimation). The estimated model allows us to quantify how the aggregate accuracy improvements on the contest are affected by changes in both extensive and intensive margins. We find that 74% of the aggregate accuracy improvements are due to more agents updating (extensive margin) and 26% to agents exerting more effort (intensive margin).

We further perform counterfactual exercises. First, we compute the value of the contest and 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 effect of the formal incentive is amplified by the availability of reliable and low cost information. Finally, we investigate the extent to which the contest mis-aligns updates from the more "natural" IPCA15 day. Assuming the total number of updates in a month remains fixed, we leverage the model to back out how the benefits would be redistributed in the absence of the contest. We then compute how the frequency of updates and the aggregate accuracy would change in this counterfactual scenario. We find that without the contest average accuracy is worse, even though most updates happen on the IPCA15. This underscores that the coordinated updates that occur because of the contest are crucial for the survey's aggregate accuracy.

This paper contributes to the literature on inattention and the literature on how incentives affect the decisions of economic agents.

Following Sims (2003), we model limited attention as a bound on the reduction of uncertainty (measured by entropy) that a rational decision maker can achieve by processing new information. In Mackowiak and Wiederholt (2009), Caplin and Dean (2015), Mackowiak, Matejka, and Wiederholt (2016) and Steiner, Stewart, and Matějka (2017) the agent's attention bound is *exogenous*, while in our model the agent chooses it at each point in time based on a cost-benefit calculation as in Wiederholt (2010). In addition, the agent in our model has access to a statistical model that she uses to understand how the information that she has processed reduces future uncertainty. In Steiner et al. (2017), the agent chooses ex-ante the information structure (the data generating process of signals) as a function of the past history of actions and signals.

The forecasting decisions analyzed in this paper also provide a new and ideal test bed for bringing rational inattention models to the data, thus contributing to a literature that is still in its infancy: Rational inattention models are just starting to be evaluated empirically. This is partly due to the challenges of bringing these models to the data.[5] One key challenge is the difficulty in separately identifying the unobservable attention from the (usually also unobservable) prior uncertainty. Caplin, Leahy, and Matejka (2016) make an important step towards overcoming this challenge in the context of discrete choice analysis. We focus on dynamics and structurally estimate a model of rational inattention that fits the data, which, to our knowledge, is new to the literature. In the data we observe the benefits of attention since they correspond to higher forecast accuracy. Our model has closed-form solutions for these observables, which allows us to disentangle how their dynamics are driven by reduction in uncertainty versus attention.

A large literature in macroeconomics has highlighted the role of information frictions in explaining expectation formation and the dynamic behavior 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 both the sticky information model of Mankiw and Reis (2002) and the noisy information model of Woodford (2001).[6] This paper takes a step forward by endogenizing the information frictions that are assumed exogeneous in this literature. This, in turn, 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 quantify how survey design affects the quality of forecasts: Incentives matter, as they induce forecasters to pay 'more attention', resulting in more forecasters updating and in better forecasts. The timing of the incentives also matters, and can be exploited to maximize the aggregate accuracy of the survey.

Marinovic et al. (2013) study theoretically the effect of a forecasting contest in a strategic model without information frictions. They find that the contest's effect on forecast accuracy can be ambiguous. Forecasters in our dataset seem to ignore the strategic component and focus on overcoming the information barriers aiming to provide the best forecast they can.[7]

[5]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).

[6]See also Mankiw, Reis, and Wolfers (2004) and Andrade and Le Bihan (2013), among others.

[7]This was also confirmed in personal interviews with some participants. One reason why participants seem

The strong link we document between attention and accuracy on one hand, and the contest on the other, speaks to the broader question of what are the productivity drivers in economics (see the survey of Syverson (2011)). Lazaer (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, recently, 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 formal incentives and performance.

The paper is organized as follows. Section 2 discusses the data, and section 3 the stylized facts about updating behavior and accuracy. The theoretical model is introduced in section 4. Section 5 presents the structural estimation results, and section 6 the counterfactual analysis. Section 7 concludes.

2 Data

Our 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 target. 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 a total of 2,751 daily forecasts for 132 months, with an average of 85.3 forecasters.[8] We treat months and the forecasts associated with each of them as events, that repeat one after another until the end of the sample. As events, they entail fixed-horizon planning problems. However, all of these events are connected by inflation, which is a continuous process over the whole sample.

not to act strategically may be that the survey is confidential and anonymous.

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

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 and for different forecast horizons.

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 logged in and either confirmed or changed the forecast.

Informal Incentives: Similarly to the other surveys of professional forecasters that have been analyzed in the literature, (the "Blue Chips," 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."[9] 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 at 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 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) are then published each month on the BCB website. The competition 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.[10]

[9] 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.

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

Figure 2. Example of contest outcome

Top 5 Forecasting Institutions - February 2017

Gerin

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

Information Releases: There are two major information releases. The first 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. The second is the release of the minutes of the meeting of the BCB Monetary Policy Committee (MPC), which occurs less frequently and at irregular times.

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

Forecast Timeline: Forecasters know the dates of the data releases as well as the contest days at the beginning of the year. The number of workdays in the month (i.e., the duration of the forecasting period), the forecast horizon and the timing of the contest can vary across months. The chronology of relevant events within each month is depicted in Figure 3, which shows that the first forecast for (say) 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).

3 Stylized Facts

In this section we document new stylized facts about the drivers of forecast updates and accuracy improvements at the individual level. We also analyze the aggregate dynamic behavior of updates and accuracy around the contest and around information releases.

Drivers of Individual Forecast Updates: We first investigate how information and incentives affect " the extensive margin" (how many forecasters update). We do so 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 marginal effects (in square brackets). The table shows that the contest is not the only driver of updates, as forecasters update at other times outside the contest. The contest has, however, the largest effect on the extensive margin: the probability of updating goes up by 38.9 percentage points (p.p.) on the contest.[11] There is also a "contest anticipation" effect with a 18.8 p.p. increase in the probability of updating one day *before* the contest. Forecasters also respond to the release of information, but this has a smaller effect on the probability of updating (the IPCA15 is associated with a 12.1 p.p. 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 on updating behavior, 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 positive and significant coefficient for $EMBI_{t-1}$. This finding is consistent with one of the main predictions of rational

[11] We focus the discussion on marginal effects.

Table 1. Drivers of forecast updates

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: 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.

inattention models.

Drivers of Individual Accuracy Gains: We next analyze the drivers of accuracy improvements, conditional on updating (computed as minus the change in the log of the squared forecast error relative to the previous update). While we expect to find that information releases improve accuracy, it is less clear a priori whether the contest would affect accuracy. An affirmative answer would support the hypothesis that the contest not only induces more forecasters to update (the extensive margin documented before), but also makes them exert more effort into producing accurate forecasts, the *intensive margin*.

The regression is complicated by the fact that the accuracy on different days is associated with different forecast horizons, and accuracy improvements are not necessarily constant during the month. Moreover the accuracy improvements depend on the time elapsed between subsequent updates. To partly control for these factors we add as regressors the forecast horizon and the duration between updates. In addition, the regression coefficients for the contest and information releases could be non-constant, e.g., because the dates when these occur are associated with different forecast horizons each month. We thus consider two variants where these coefficients can depend on the forecast horizon as well as the duration between updates. The results are reported in Table 2 in columns 3 and 4.

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 reports the estimation results. Column 1 confirms the expected result that information releases are associated with accuracy improvements, as the coefficients for IPCA15 and MPC are large and significant. It also shows evidence supporting the hypothesis that the contest affects the intensive margin of forecasting efforts, as accuracy improvements go up by

Table 2. Drivers of accuracy gains conditional on updating

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: 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.

7.1 p.p. on the contest. In contrast, column 2 shows that other variables that according to Table 1 were 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). Columns 3 and 4 show only mild evidence that the coefficients could be non-constant, as the effect of the contest depends on the forecast horizon.[12]

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 shows the fraction of updaters and the average $MSFE$ across forecasters on 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), so informal incentives also matter for aggregate behavior.

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 that occurs during the month as the forecast horizon diminishes. 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 improvement we see on any other day.

The documented jump in aggregate accuracy on the contest could be caused by both changes in the extensive margin (if more forecasters update, their average accuracy is higher) and in

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

Figure 4. Dynamics of updates and accuracy: daily evolution of the fraction of updaters and aggregate $MSFE$ around the contest and IPCA15 (left graph) and around the MPC day (right graph)

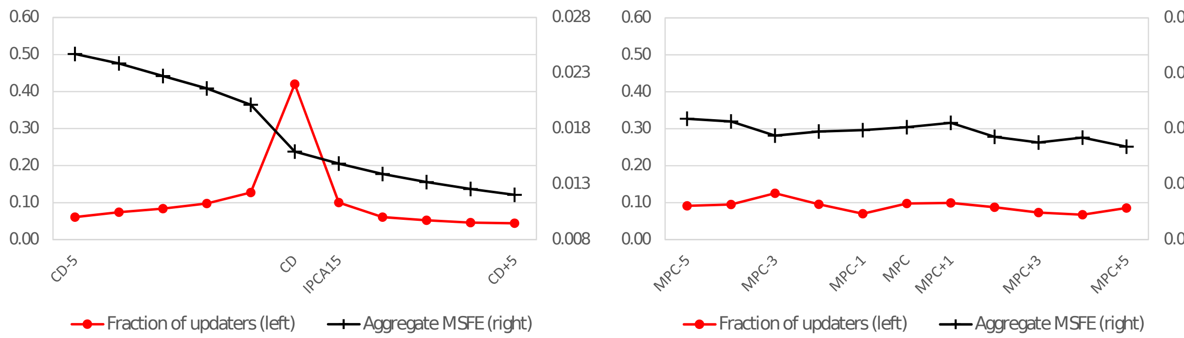

the intensive margin (each forecaster may be putting more effort into obtaining an accurate forecast). In one of the counterfactual exercises, we leverage our theoretical model to decompose the aggregate accuracy improvement on the contest into the contribution of changes along both margins.

Our findings also suggest that the contest may be crowding-out updates on other days. In particular, since the contest is the day before the release of information, few forecasters update on the IPCA15, even though doing so improves accuracy. In fact, we computed that the instances in which a forecaster updates on both the contest and the IPCA15 constitute only 0.65% of the sample. Another counterfactual exercise allows us to shed light onto this potential crowding-out effect 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 such heterogeneity is driven by time-invariant fixed effects (e.g., the presence of forecasters who are always frequent updaters and/or the most accurate forecasters) 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), based on dividing the cross-sectional distribution of both average $MSFE$ and total number of updates during each year

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.

into 5, 10 or 20 quantiles and computing transition probabilities among the different quantiles.

Table 1 shows that the indexes are close to 1 (where 1 indicates perfect mobility), which 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 are ex-ante identical, but behave differently ex-post depending on a random draw of the heterogeneous opportunity costs of updating that they face.

Figure 5. Forecast Disagreement: daily evolution of the standard deviation of individual forecasts around the contest day

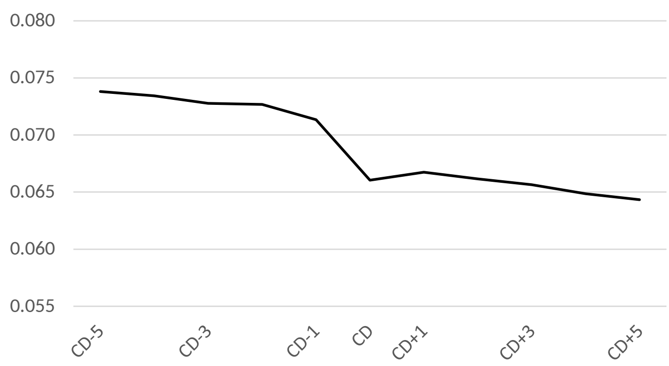

Forecasters' Objective: The empirical facts clearly document that the contest affects both the updating probability and the accuracy improvement among updaters. So a reasonable objective function is that forecasters seek to win the contest when submitting a forecast on the contest and maybe another objective on other days. Let's consider a simpler contest with one prize as that studied in Marinovic et al. (2013). In that setting the probability of winning the contest increases the more accurate is a forecast but, also, the more it differs from the competing forecast: If say all N forecasters submit the same forecast (even it is fully accurate), the probability that one of them wins is 1/N. Marinovic et al. (2013) argue that this force to differ induced by the contest 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 what we see in the data.

Figure 5 reports the evolution of disagreement around the contest–the standard deviation of individual forecasts–which displays a sizeable reduction on the contest day. We take this as evidence that forecasters are just simply trying to win the contest by maximizing accuracy and do not take into account the "strategic" effect that would induce them to differentiate their forecast. This is also supported by informal conversations we had with survey participants, who asserted that their objective is to provide accurate forecasts. Based on these observations, together with the fact that the survey is anonymous and confidential, we find it more suitable (as well as tractable) to assume that a forecaster's objective is to maximize accuracy, which we express as a desire to minimize mean squared forecast error. Thus we employ a decision-theoretic setup, rather than a game-theoretic one.

4 Theory

We build on the theory of rational inattention (Sims, 2003) to link the observable dynamics of both forecast accuracy and updating behavior to the available information and to the unobservable preference parameters characterizing each agent's individual decision problem. Our theory is inspired by the way forecasters behave in reality: They employ state-of-the art statistical models in which they input publicly available information but also 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 time-varying patterns of updates and accuracy improvements suggest that forecasters might face time-varying costs and benefits of producing a forecast. Moreover, there are no systematic patterns in terms of persistent differences across forecasters as suggested by the mobility indexes in Table 3. In the model, forecasters choose the amount of resources to employ in order to formulate a forecast that is accuracy-maximizing given those resources. This optimal amount depends on time-varying costs and benefits of processing information which is what we ultimately estimate. In the theory the costs and benefits are treated as known parameters to each forecaster. The rational inattention model allows for the endogenous precision of information that forecasters feed into their model, which translates into analytical expressions for the probability of updating and for the accuracy of forecasts.

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 month[13]–can be written as the sum of daily inflation x_t (the difference between the logs of the prices on days t and $t-1$):

$$y_m = \sum_{t=1}^{T} x_t, \tag{3}$$

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

$$y_m = a + \psi y_{m-1} + v_m + \theta v_{m-1}, v_m \sim i.i.d.\mathcal{N}(0, \sigma_v^2), \tag{4}$$

[13]This is a reasonable approximation for actual single-digit inflation, which is the monthly increase in the average price level.

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

$$x_t = c + \phi x_{t-1} + \varepsilon_t, \varepsilon_t \sim i.i.d.\mathcal{N}(0, \sigma_\varepsilon^2), \tag{5}$$

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

There are two dynamic dimensions in our setting: the month-to-month problem of forming a forecast of current-month inflation at the beginning of the month and the within-month problem of updating the initial forecast. The two dynamic problems are coherent since the initial monthly forecast for all agents is based on the ARMA(1,1) model in (4) and the updates are based on the daily AR(1) model in (5). Our main focus here is, however, on the within-month dynamic problem and how it is linked to the cost and benefit of updates.

Agents' Objective and Initial $MSFE$: We assume that the agent's objective is to maximize forecast accuracy or, equivalently, to minimize the $MSFE$,

$$MSFE = E[(y - f)^2], \tag{6}$$

where f indicates the forecast of monthly inflation y.[14] This objective function implies that it is optimal to set the forecast equal to the conditional expectation of y, based on the information set available to the agent. As a consequence, every month all agents form the initial forecast optimally using the ARMA(1,1) model, which means that the initial $MSFE$ is the same for all agents and for all months and it equals

$$MSFE_0 = \sigma_v^2 = [1 + (1 + \phi)^2 + ... + (1 + \phi + \phi^2 + ... + \phi^{T-1})^2]\sigma_\epsilon^2. \tag{7}$$

Agents' Decision Problem: Agents choose each day how much capacity to devote to processing information in order to obtain a more accurate forecast. This decision depends on the costs and benefits of updating, which are parameters that can vary across agents and over time. The key difference from many leading dynamic rational inattention models is that we

[14]For notational simplicity, we henceforth drop the subscript "m" from the monthly variable but retain the subscript "t" for the daily variable.

endogeneize the capacity choice whereas in those papers (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 what signals to observe in each period.

Sources of Information: Reductions in $MSFE$ during the month are not only due to agents collecting/processing private signals, but also to the revelation of public information. We make the following assumption about the types of signals an agent can collect/process.

Assumption 1 (Public signals) On day t, the public signal contains past values of daily inflation for the current month:

$$s_p^t = \{x_{t-1}, x_{t-2}, ..., x_1\} . \tag{8}$$

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. The signal's precision is endogenous and depends on how much 'capacity' or mental effort the agent decides to allocate to information gathering and processing.

The assumption that agents have access to past public signals that are more accurate than past private signals is a realistic one in an environment where information is released at higher frequencies than the variable being forecasted. For example, consider the problem of updating yearly inflation forecasts at a monthly frequency, in which case forecasters have access to monthly inflation data releases. These public releases are plausibly more accurate than any past private signals the forecasters may have collected, and forecasters may not have collected private signals for some past months and thus have an information set with missing data. As a result, it wouldn't be optimal for the agent to ignore the public past information and instead base a forecast on an incomplete sequence of past noisy signals. Even for daily inflation, the assumption is plausible since agents have access to public information about past daily inflation, for example in the form of the daily releases of gasoline prices.

Information Costs: Following the rational inattention literature, the cost of the private signal is proportional to the reduction in uncertainty that is associated with it as measured by the mutual information–the difference in the conditional entropy with and without this signal. The bound on the mutual information captures capacity or mental effort. Such effort is costly, as agents in our model face not only explicit effort 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. The *marginal* cost of mental capacity or effort of agent i at date t is denoted by c_{it}. The benefit of this effort comes in the form of reduction in forecast error. We let $w_{it} > 0$ denote agent i's marginal benefit of forecasting (i.e, of reduction in $MSFE_{it}$) at date t.

There is a *fixed* cost of updating C_{it} that captures the opportunity cost of forecasters' time. Let w^o be the marginal benefit of time devoted to other activities than forecasting. Given that we focus on the dynamics of behavior relating to participation in the survey, we assume this benefit to be fixed across agents and time. Then, the opportunity cost of updating is

$$C_{it} \equiv \frac{w^o}{w_{it}}.$$

If agent i has an opportunity cost below 1, so $w_{it} > w^o$ then she finds it worthwhile to update and in this case she observes the free public signals. As we discuss later, we assume that an agent who decides to update finds it worthwhile to also observe a private signal, resulting in the information set:

$$s_i^t = \{s_{it}, x_{t-1}, x_{t-2}, ..., x_1\}. \tag{9}$$

The information sets s_i^t differ across agents and evolve over time in a way that induces heterogeneity across agents, depending on the individual time-patterns of opportunity cost realizations. Figure 2 shows an example of the evolution of the information set over the first three days of the month:

Discussion: The assumption that forecast updates are driven by agents allocating mental capacity ("attention") to acquire information about the variable of interest is a realistic assumption in the context of nowcasting. This is supported by our private communications with some survey participants who claimed that at their institutions inflation nowcasting relies on

Figure 6. Evolution of information sets

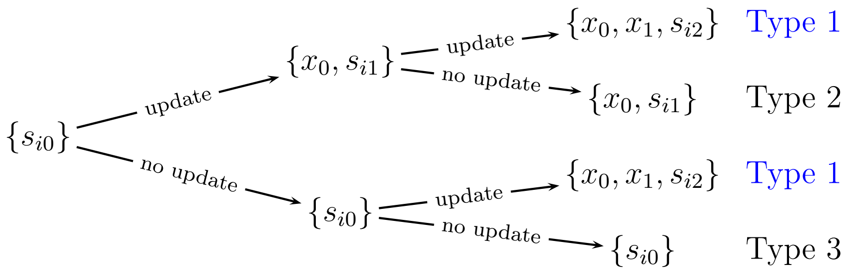

direct collection of price data.

Forecast updates and *MSFEs*: On day $1 \leq t \leq T$ of the month, the forecast update is the conditional expectation of y based on each agent's information set $s_i^t \equiv (s_p^t, s_{it})$. Combining (3), (5) and (9) implies that the conditional expectation equals:

$$E\left[y|s_i^t\right] = \sum_{j=1}^{t-1} x_j + \sum_{j=t}^{T}\left(\frac{c(1-\phi^{j-t})}{1-\phi} + \phi^{j-t}E\left[x_t|s_i^t\right]\right), \tag{10}$$

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

$$MSFE_{it} = \left[1 + (1+\phi)^2 + ... + \left(1+\phi+...+\phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \left(\sum_{j=t}^{T}\phi^{j-t}\right)^2 E\left[\sigma^2_{x_t|s_i^t}\right], \tag{11}$$

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

Note that Assumptions 1 and 2 imply that the forecast error at day t depends only on the conditional variance of the current signal and not on past privately-collected signals. This is not only realistic–since it essentially says that a "fresh" public signal is more accurate than past ones, but we later see that it also simplifies the forecaster's dynamic choice problem into a sequence of static, rather than sequential dynamic rational inattention problems.

Attention-Driven Precision of Private Signals: Following the rational inattention literature, the *additional* information content of the signal is captured by the relative entropy based on the information sets with and without the signal:

$$I(x_t; s_{it}|s_p^t) = H(x_t|s_p^t) - E_{s_{it}}[H(x_t|s_p^t, s_{it})|s_p^t] \leq k_{it}. \tag{12}$$

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 constraint (12) binds at an optimum and thus the agent chooses the distribution of s_{it} so that $\sigma^2_{x_t|s_p^t, s_{it}}$ is constant and exhausts all capacity for each signal realization s_{it}, which implies that $\frac{\sigma^2_{x_t|s_p^t}}{\sigma^2_{x_t|s_p^t, s_{it}}} = 2^{2k_{it}}$ or

$$\sigma^2_{x_t|s_i^t} = \sigma^2_{x_t|s_p^t}(2^{2k_{it}})^{-1}. \tag{13}$$

By substituting (13) into (11) and using the fact that $\sigma^2_{x_t|s_i^t}$ is constant and that the AR(1) model implies $\sigma^2_{x_t|s_p^t} = \sigma^2_\varepsilon$, we obtain:

$$MSFE_{it} = \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma^2_\varepsilon + \left(\sum_{j=t}^{T}\phi^{j-t}\right)^2 \sigma^2_\varepsilon (2^{2k_{it}})^{-1}. \tag{14}$$

Equation (14) shows that the dynamics of individual $MSFEs$ depend on two components: The first is common to all updaters and captures the resolution of uncertainty due to the public signal. The second is agent-specific and depends on how much attention the 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.

Optimal Decision: At the beginning of each day τ agent i decides how much attention $k_{i\tau}$ to allocate to obtaining a forecast that minimizes the sum of future errors:

$$\min_{k_{i\tau}} \sum_{t=\tau}^{T}\left[\frac{w_{it}}{2\ln 2}MSFE_{it} + c_{it}k_{it}\right].$$

Note that this formulation assumes that i knows the current and future cost-benefit realizations. In a moment we explain why he only needs to know the current.

As usual, we solve this problem of sequential decisions backwards. Consider i's problem at the last period. For $t = T$ we have:

$$\min_{k_{iT}} \left[\frac{w_{iT}}{2 \ln 2} MSFE_{iT} + c_{iT} k_{iT} \right] \quad \text{subject to } k_{iT} \geq 0, \text{ (13) and (14).} \tag{15}$$

Agents can only control the part of the $MSFE$ in (14) that depends on collecting information about the current daily signal, so optimal attention solves:

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

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

$$k_{iT}^{*} = \begin{cases} \frac{1}{2} \log_2 \left(\frac{w_{iT}}{c_{iT}} \sigma_{\varepsilon}^{2} \right) & \text{if } \frac{w_{iT}}{c_{iT}} \sigma_{\varepsilon}^{2} > 1 \\ 0 & \text{otherwise} \end{cases}.$$

Note that this level of optimal attention does not depend on past choices, which implies that the forecast error at T does not depend on past choices: Our assumption that the public signal is the actual realization of the variable implies that it is by construction weakly more precise than any 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 on past private signals, but, rather, on the $t+1$ public signal which is more precise than the private signal. Hence, under our assumptions an agent's dynamic problem turns into a sequence of static problems. In particular, at the beginning of a working day t an agent solves:

$$\min_{k_{it}} \left[\frac{w_{it}}{2 \ln 2} MSFE_{it} + c_{it} k_{it} \right] \quad \text{subject to } k_{it} \geq 0, \text{ (13) and (14),} \tag{16}$$

so, as in the case of $t = T$, optimal attention solves:

$$\min_{k_{it}} \frac{w_{it}}{2\ln 2} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 \left(2^{2k_{it}} \right)^{-1} + c_{it} k_{it} \ \text{ s.t. } \ k_{it} > 0,$$

which gives

$$k_{it}^* = \begin{cases} \frac{1}{2}\log_2 \left(\frac{w_{it}}{c_{it}} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 \right) & \text{if } \frac{w_{it}}{c_{it}} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 > 1 \\ 0 & \text{otherwise} \end{cases}. \tag{17}$$

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 (which in this case is measured by σ_ε^2). Notice that optimal attention varies over time because of two different reasons. The first has to do with the resolution of uncertainty during the month due to the public signal, which is common to all agents. The second is due to agents possibly facing different cost/benefit ratios on different days.

Fraction of Updaters: The probability of updating on day t is given by:

$$\lambda_t = P\left(w_{it} > w^o \right). \tag{18}$$

We assume that the marginal cost of processing information is small enough relative to the fixed cost of updating, implying that if an agent finds it worthwhile to pay the fixed cost and decides to update (i.e., $w_{it} > w^o$), she also puts positive effort into collecting/processing private information (i.e., $\frac{w_{it}}{c_{it}} \left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2 > 1$). A sufficient condition for this assumption is:

$$w^o = \max_{1 \le t \le T} \frac{c_{it}}{\left(\sum_{j=t}^{T} \phi^{j-t} \right)^2 \sigma_\varepsilon^2} = \frac{c_{it}}{\sigma_\varepsilon^2}.$$

For now assume that $c_{it} \equiv c$, so the marginal cost is constant across forecasters and time (we relax this assumption in the counterfactuals section of the paper). These assumptions combined imply the following probability of updating:

$$\lambda_t = P\left(w_{it} > \frac{c}{\sigma_\varepsilon^2}\right). \tag{19}$$

We next derive an analytical expression for the evolution of individual $MSFE$s.

Dynamics of $MSFE$: Substituting (17) into (14), together with (19) and the fact that agents who don't update maintain their previous forecast, gives the evolution of the optimal $MSFE$ as

$$MSFE_{it}^* = \begin{cases} \left[1 + (1+\phi)^2 + ... + \left(1 + \phi + ... + \phi^{T-1-t}\right)^2\right]\sigma_\varepsilon^2 + \frac{c}{w_{it}} & \text{if } w_{it} > \frac{c}{\sigma_\varepsilon^2} \\ MSFE_{it-1}^* \text{ otherwise} \end{cases}. \tag{20}$$

Equations (20) and (19) are the basis of the structural estimation that we present next.

5 Estimation

We estimate the model presented in Section 4 by Simulated Method of Moments (SMM) (see Gourieroux and Monfort (1996), Duffie and Singleton (1993), Ruge-Murcia (2012)), which involves matching empirical moments with their theoretical counterparts.

The goal is to jointly match the fraction of updaters and the aggregate accuracy in a window of days around the contest, as reported in Figure 1.[15] 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.[16]

In order to simulate theoretical moments, every month we start from an initial $MSFE$ for all agents given by (7). As previously discussed, this makes the within-month updating

[15]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.

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

problem coherent with the month-to-month forecasting problem. On every subsequent day t of the month, each agent receives a random draw of the benefit w_{it}. We specify the cross-sectional distribution for w_{it} to be a truncated normal $TN(\mu_t, \sigma_w^2)$ and let the variance σ_w^2 be constant over time. The particular distributional choice is not crucial because the main feature that we want to capture is how the cross-sectional mean changes on the days around the contest.

The benefit draws determine the fraction of updaters on that day according to (19) and the $MSFE_{it}$ for agents who update is given by the first line of (20). Agents who don't update on day t 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 in a five-day window around the contest.[17]

We consider two different parameterizations of the model. In the first, we keep μ_t unrestricted and assume a common and constant marginal cost parameter, denoted by c. In the second, we impose restrictions on the time variation in μ_t and allow the marginal cost to be different on the IPCA15. The parameters of the first parameterization are: $\theta = (\mu_t, \sigma_w^2, c, \phi, \sigma_\epsilon^2)$, for $t = CD - 5, \ldots, CD + 5$.

As we later see, the estimates from the first parameterization of the model suggest that the mean benefit μ_t is approximately constant outside the contest. We thus consider imposing various restrictions on the time variation in μ_t. The model that still passes the J-test restricts the mean benefit to take four possible different values, while also allowing for a different marginal cost on the IPCA15. We refer to it as the "parsimonious model." The parameters are $\theta = (\mu_t, \sigma_w^2, c_t, \phi, \sigma_\epsilon^2)$, where:[18]

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

[18]In unreported results, we estimate the model imposing combinations of the additional restrictions $\mu_B = \mu_A$, $\mu_{CD-1} = \mu_{CD}$, $c_{IPCA15} = c$. All these parameterizations are rejected by the J-test.

$$\mu_t = \begin{cases} \mu_B & \text{if } t < CD-1 \\ \mu_{CD-1} & \text{if } t = CD-1 \\ \mu_{CD} & \text{if } t = CD \\ \mu_A & \text{if } t > CD \end{cases} \tag{21}$$

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

Table 4. Model Estimation via Simulated Method of Moments (SMM)

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 5. Parsimonious Model Estimation via Simulated Method of Moments (SMM)

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.

Estimation Results: Table 4 shows that the mean benefit of updating increases slightly on the lead-up to the contest, jumps up on the contest (it is 15% higher on the contest than 5

Figure 7. Comparison between the parsimonious model predictions and data for the daily fraction of updaters (left graph) and the average $MSFE$ (right graph) around the contest day

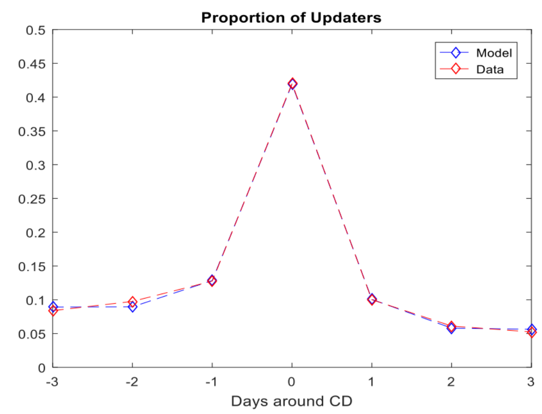

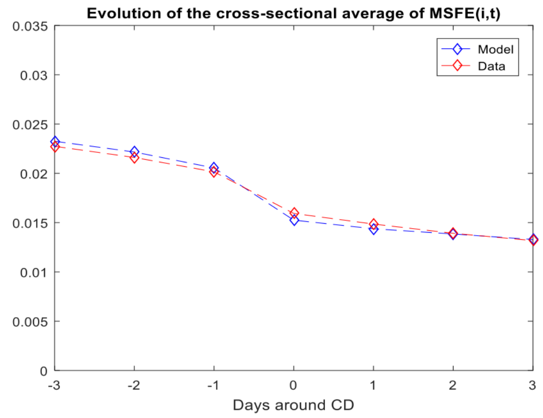

days before), jumps down on the IPCA15 to a level comparable to before the contest and then decreases afterwards, to a slightly lower level than before the contest. The model fits the data very well, as can be seen from Figure 7 and it passes the J-test of overidentifying restrictions with a p-value 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 estimates of the same parameters in Brazilian inflation data:[19] 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 parsimonious model in equation (21) (estimated using a window of 3 days around the contest) are in Table 5. Figure 7 shows the fit of this parsimonious model.[20]

The model 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 the mean benefit increases on the contest and the day before and that it is higher before the contest than after the contest. The cost is lower on the IPCA15 than on other days.

[19]We obtain the estimate 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 (5) to back out the AR(1) parameters.

[20]The fully parameterized model fits slightly better. Figures are available upon request.

6 Counterfactuals

In this section we perform a number of counterfactual analyses, using the estimates of the constrained model reported in Table 5.

Aggregate Accuracy and Changes in Extensive vs. Intensive Margins: The estimates from the model in Table 5 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 putting more effort (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 (20), using the estimates from Table 5.

Figure 8 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 with no contest would be $\mu_t = \mu_A$ for all t). The graph shows that the daily $MSFE$ reduction on each potential contest day is several times larger with the contest than without the contest (on average, the improvement is 472%). The largest reduction in daily $MSFE$ is obtained by putting the contest on the IPCA15, and it amounts to a 347% improvement relative to not having a contest (this is because even without a contest the IPCA15 would benefit from a larger MSFE reduction than other days, due to the lower cost), and a 31% improvement relative to having the contest the day before as it is now in the survey.

Figure 8. Daily $MSFE$ reduction with and without the contest on different days

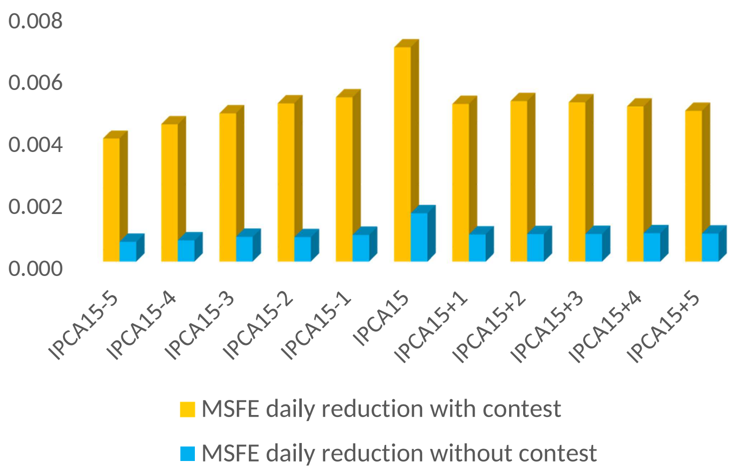

Optimal Timing of the Contest: We next determine the optimal timing of the contest within the five-day window around the IPCA15, in terms of obtaining the smallest *monthly sum* of average $MSFE$ (cumulative $MSFE$) across agents. The right panel of Figure 9 shows the evolution of the counterfactual *daily* average $MSFE$ when the contest is put on different days. Note that the daily $MSFE$ decreases essentially linearly, and the effect of the contest is a downward shift of the line. The left panel of Figure 9 plots the counterfactual cumulative $MSFE$. It shows that the optimal timing in terms of cumulative accuracy is the IPCA15 (which from Figure 8 is also the optimal timing in terms of the daily accuracy improvement). 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%. A noteworthy observation is 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 have more days afterwards to lower the path of the daily $MSFE$, as opposed to later in the month, when more signals are observed but there are fewer days afterwards to lower the daily $MSFE$. An implication of the figure is that the current timing of the contest on the day before the IPCA15 is the second best choice.

Figure 9. Cumulative $MSFE$ with contest on different days (left graph) and daily average $MSFE$ with contest on different days (right graph)

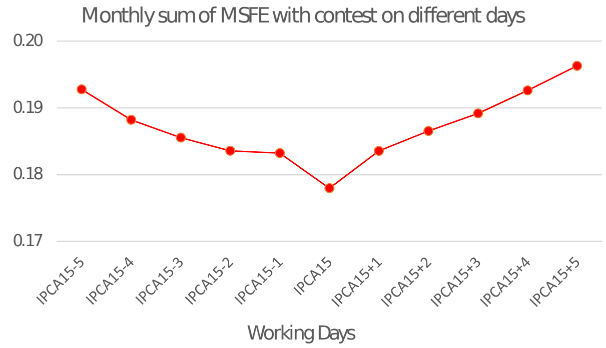

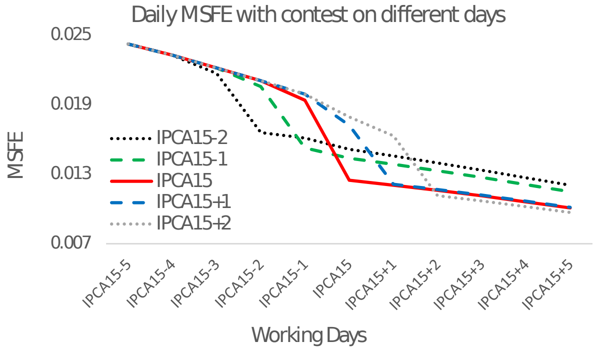

Figure 10. Counterfactual Updates and Accuracy Evolution without Contest (red). Blue curves simulated model with contest.

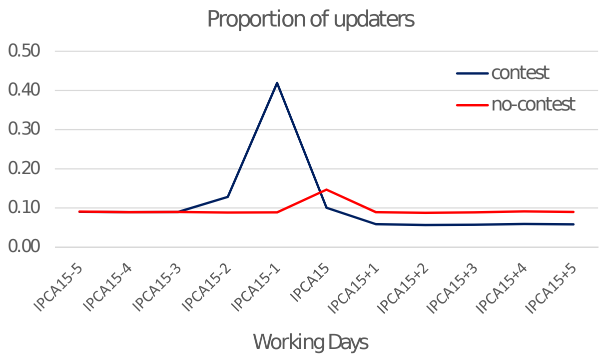

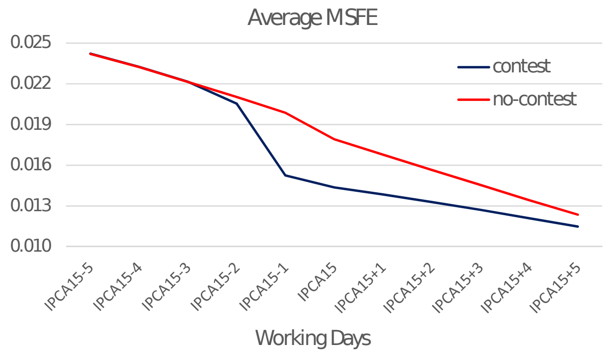

Contest versus No-Contest: We finally investigate the extent to which the contest crowds-out updates–or more precisely mis-aligns updates from more "natural days" such as the IPCA15 where our estimates show (as expected) that it is cheaper to process information. What would happen to the evolution of accuracy if there is no contest?

One challenge in investigating what would happen is that we do not know how the total number of updates within a month would change if there is no contest. In what follows we assume that this number stays constant even without the contest. 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 can leverage 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 number of updates as in the estimated model. We maintain the same estimates for the other parameters as in Table 5. 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 10 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 very important for the survey's average accuracy. This insight is quite surprising.

7 Conclusions

We analysed panel data from a unique and proprietary survey of professional forecasters where the forecast updating decisions of participants are observable and both incentives and information releases are present. The empirical findings are consistent with a model in which agents respond not only to information, but also to incentives to collect private signals. We build a rational inattention model where agents have a budget of attention or cognitive resources to devote to forecast updating. Updates respond to available information and to preference parameters corresponding to the costs and benefits associated with being more attentive, and so does the accuracy of forecasts. The model has predictions for the individual dynamics of forecast accuracy and the fraction of agents who update, which we observe in the data. We structurally estimate the model and use the estimates to perform counterfactuals about alternative survey designs and to assess the value of the contest.

The empirical patterns we document and the counterfactual exercises underscore the importance of a contest, and formal incentives more broadly, for accuracy. A contest makes more forecasters participate and each forecaster put more effort, resulting in an increase in both individual and aggregate accuracy. This 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 of professional forecasters are a thriving and growing industry,[21] virtually no attention has been paid in the literature to how survey design affects forecast quality.

Kacperczyk et al. (2016) conclude "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 (the evolution of forecast updating choices and forecast accuracy) and thus ties information choices to outcomes (accuracy). The unique dataset we use allows us to structurally estimate a model of endogenous information choice and to highlight the link between effort and incentives. Viewed from this perspective, this paper connects two important ideas in economics, namely that attention is limited and that incentives (captured here by the contest) matter.[22]

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, which may lead to a better fund allocation across assets.

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.

Andrade, P. and H. Le Bihan (2013): "Inattentive professional forecasters," *Journal of Monetary Economics*, 60, 967–982.

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

[22] 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.

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.

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.

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.

Kahneman, D. (1973): *Attention and effort*, vol. 1063, Prentice-Hall Englewood Cliffs, NJ.

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

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

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

Mackowiak, B. A., F. Matejka, and M. Wiederholt (2016): "The Rational Inattention Filter," Tech. rep., CEPR Discussion Papers.

Mankiw, N. G. and R. Reis (2002): "Sticky Information versus Sticky Prices: A proposal to replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Mankiw, N. G., R. Reis, and J. Wolfers (2004): "Disagreement about inflation expectations," in *NBER Macroeconomics Annual 2003, Volume 18*, The MIT Press, 209–270.

Manso, G. (2011): "Motivating innovation," *The Journal of Finance*, 66, 1823–1860.

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," .

Murphy, K. J. (1999): "Executive compensation," *Handbook of labor economics*, 3, 2485–2563.

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.

Wiederholt, M. (2010): "Rational inattention," *The New Palgrave Dictionary of Economics (Online Edition ed.)*.

Woodford, M. (2001): "Imperfect common knowledge and the effects of monetary policy," Tech. rep., National Bureau of Economic Research.

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