

Incentive-Driven Inattention*

Raffaella GIACOMINI[†] Wagner Piazza GAGLIANONE[‡]
João Victor ISSLER[§] Vasiliki SKRETA[¶]

November 10, 2015

Abstract

KEYWORDS: Heterogeneity, Attention degree, Professional forecasters.

JEL CLASSIFICATION: E27, E37, D80, D83.

*The views expressed in the paper are those of the authors and do not necessarily reflect those of the Banco Central do Brasil (BCB) or of Getulio Vargas Foundation (FGV). We are grateful to Valdemar Rodrigues de Pinho Neto for excellent research assistance and to the Investor Relations and Special Studies Department (Gerin) of the BCB for kindly providing the data used in this paper. Any remaining errors are ours.

[†]UCL. r.giacomini@ucl.ac.uk

[‡]Research Department, Banco Central do Brasil. E-mail: wagner.gaglianone@bcb.gov.br

[§]GVF.

[¶]UCL. vskreta@gmail.com

1 Introduction

Mankiw and Reis (2010) argue that the Phillips curve (or the aggregate supply curve) is probably the most important macroeconomic relationship. They state that “it is the key connection between real and nominal variables. It explains why monetary policy, and aggregate demand more broadly, has real effects.” As is well known, the existence of a short-run trade-off between inflation and output gap relies heavily on some sort of inertia in price expectations. Otherwise, as in Lucas (1972), although agents may not correctly predict full-employment equilibrium at any point in time, their forecast errors will still be unpredictable and the economy will fluctuate randomly about full-employment – not a bad outcome.

The work of Mankiw and Reis (2002), Sims (2003) and Woodford (2003) on inattention made the profession aware of the fact that the consensus forecast among agents responded slowly to new information available to them; see, *inter alia*, Coibion and Gorodnichenko (2012) and Andrade and Le Bihan (2013). In some of these models inattention is assumed to be exogenous to the agent, and in others it is assumed to be a function of the limited ability of agents in processing information, while this limitation is itself exogenous once we control for idiosyncratic random shocks. Of course, inattention can generate a short-run trade-off between inflation and output gap and have an important role in the construction of the Phillips curve.

In this paper, we propose a novel approach in which we endogenize inattention while retaining some important features of the previous literature. In our framework, agents respond to incentives to update their predictions based on new information, therefore the title *Incentive-Driven Inattention*. New information flows in continuously and acquiring information is costly. To decide how much new information should be used in predicting future outcomes, agents do cost-benefit analysis. Incentives play an important role

in defining the benefits of updating information. We consider two different regimes regarding incentives. In the first regime, incentives imply that agents have high benefits for updating their forecasts, but in the second they have no incentives and therefore are inattentive as in Mankiw and Reis, using an exogenous updating rule. When they decide to update their information set, they do not observe the true state, but receive noisy-information, as in Sims and Coibion and Gorodnichenko (2014). Once we put together all these ingredients, we have a structural macroeconomic model of incentive-driven inattention, which can be tested at the individual level. Because we can always aggregate individual data emulating a representative-agent framework, we can also test previous results in the literature that were based on a consensus-forecast approach.

To be able to evaluate and test the incentive-driven inattention model we need to have access to high-frequency micro-data at the individual level, where we can observe enough changes in costs and benefits of updating information. While most papers in the literature use consensus forecasts for a representative agent, we are able to exploit the high-frequency micro-data character of the *Focus Survey* gathered by the Brazilian Central Bank (BCB). It has a few unique features that sets it apart from other well-known expectation surveys. For example, its frequency of observations is daily (working days), as opposed to monthly (*Wall Street Journal Forecasting Survey*), quarterly (*Survey of Professional Forecasters*), or semi-annually (*Livingston Survey*). In our sample, dating from January 2nd, 2004 until January, 8th, 2015, at each working day, a survey respondent can *choose* to participate inserting in the data base forecasts of a myriad of important macroeconomic and financial forecasts for different horizons. An important feature of the *Focus-Survey* design is that of incentives to participate at any given day, which is related to reputation and accuracy of forecasts. The *Focus Survey* has periodical *critical dates* in which forecasts for different agents and

horizons are used to compute individual forecast errors and their respective mean-absolute-forecast-errors (MAFEs). Only forecasts at critical dates are used to compute MAFE statistics at the individual level, and all participants know this rule and when the critical dates are. Periodically, the Central Bank compares MAFEs across individual institutions and releases the names of the *top five* forecasting institutions for selected variables at different horizons. Theoretically, since MAFE should not decrease with the forecast horizon, institutions have an incentive to update forecasts at the critical date to increase their chance of being listed as a *top-five* forecaster¹. In this paper, for mathematical convenience, and following the standard approach of the forecasting literature, we adopt as measure of forecast accuracy the mean-squared-forecast-error (MSFE), which shares many properties with the MAFE risk function, such as symmetry and strict monotonicity for non-zero forecast errors.

The characterization, estimation, and testing of the incentive-driven inattention model is done as follows. First, we document the difference in behavior for critical and non-critical dates in forecasting inflation: a significant spike of updating forecasts in the data base – from about 12% on non-critical dates to about 48% in critical dates. This spike coincides with a downward structural break for the behavior of the consensus forecast MSFE across dates. Second, we postulate a parsimonious cost-benefit analysis model where inattention differs endogenously for critical and non-critical dates. Then, after Coibion and Gorodnichenko (2012), we compute the new updating rule for the consensus forecast. To characterize the daily consensus forecast error,

¹An informal incentive for respondents regards their invitation (inclusion) in a quarterly meeting with members of the board of the BCB and other survey respondents. Respondents that frequently participate on the market expectations system and that have logged in the system within the previous 30 days might be invited to the meeting. During this meeting, central bank representatives listen to the opinions of market participants about the economic outlook and the dynamic behavior of the key variables being surveyed, being important for survey respondents to identify what is common knowledge about forecasting.

we have to overcome the problem that it is a function of the daily rational-expectation forecast error, while inflation is sampled at the monthly frequency. Following Amemiya and Wu (1972), we use the fact that the best representation for monthly inflation – an $ARMA(1, 1)$ model – corresponds to an $AR(1)$ model for daily inflation. Coupled with the interpolation method used in Mönch and Uhlig (2005), we back out daily inflation and the daily rational-expectation error in forecasting inflation. Third, as in Sims (2003), and Coibion and Gorodnichenko (2014), we assume that agents never observe the true state, but instead a noisy signal. This yields three structural equations to be estimated by Hansen’s (1982) generalized method of moments (GMM) using daily data: the incentive-driven participation equation, the $AR(1)$ characterization for daily inflation, and the MSFE updating rule for the consensus forecast. The model has six structural parameters: the two participation parameters – critical and non-critical dates, the three $AR(1)$ parameters – constant, $AR(1)$ coefficient and the variance of the error term – and the variance of the noise in the signal-extraction problem. Misspecification testing is done via the usual over-identifying-restriction J -test.

When we take the model to the data, we obtain sensible structural parameters that are able to explain the dynamic behavior of the consensus MSFE across dates, including the critical date. In misspecification tests, the structural model is not rejected using a variety of instruments in GMM estimation. Simulating the model also generates an MSFE profile that matches that of the data when we consider a point-wise 95% confidence interval around it. We also investigate the determinants of updating information using a panel of individual institutions. Our results show that updating is explained mostly by the dummy variable of the critical date, followed by economic uncertainty, measured by the Emerging Market Bond Index (EMBI). Calendar-effect dummies are also relevant, as well as the uncertainty of the individual forecast error forecasting inflation in the previous month.

Equipped with our proposed model, we then discuss two important issues in forecasting. First, using individual and consensus data, we perform a set of rationality tests to investigate whether individuals or the consensus forecast pass the rationality tests pioneered by Mincer and Zarnowitz (1969) and its extensions in Coibion and Gorodnichenko (2014). Second, we discuss survey-design issues, performing counter-factual exercises with our model in order to assess the effects of changing the incentive structure of the survey – a relevant issue – since we believe that the expectation surveys of the future will have the high-frequency structure of the *Focus Survey*, coupled with incentives to update information, i.e., they will be similar to today’s *Focus Survey* on those grounds.

When we apply a set of rationality tests to individuals or the consensus forecast, results show that, at the individual level, at least 90% of individuals are rational when we consider critical dates alone when they have proper incentives to update forecasts. In contrast, for all dates, including non-critical dates, at most 26% of individuals pass rationality tests. Similar tests applied to panel-data and the consensus forecast show that rejection of rationality is the rule, not the exception.

In our survey-design investigation, we show how one can choose the *Critical Date* in order to minimize the MSFE of the consensus forecast. We perform two different counter-factual exercises. In the first, we ask what is the best critical date within a standard month from the point of view of a user of the consensus forecast. In the *Focus Survey*, the critical date is currently in the mid-point of the working month – between two consecutive releases of inflation data.² However, to minimize the MSFE profile, the best critical date is the day after the release of inflation data, since the updating behavior in the model will lower the MSFE profile for the whole month. In the second exercise, we compute what would be the gain of making every

²On average, the critical date is on the 21st day of the calendar month.

day a critical day, starting just after the release of inflation, all the way to the k -th working day within each month. Results show that the minimum average MSFE is the one which every day is a critical date, which reinforces the importance of proper incentives embedded in the critical date.

This paper is organized as follows. The next section discusses the Focus data base and the stylized facts about updating behavior and the resulting MSFE across dates. In Section 3 we present the details of the incentive-driven inattention model proposed here. In Section 4 we present the empirical evidence regarding the model – characterization, estimation, and testing. In Section 5 we present the results of rationality tests and of the survey-design counter-factual exercises for critical dates. In Section 6 we conclude.

2 The Focus Survey and Stylized Facts for Updating Forecasts

2.1 The Focus Survey of the Brazilian Central Bank (BCB)

The *Focus Survey* of expectations of the Central Bank of Brazil (BCB) is unique in which it collects high-frequency data (working days) on 254 individual institutions – about 100 of those are active at any point in time – since 1999, although the number of participants increased substantially in the mid 2000’s. Data collection started with the implementation of the Brazilian *Inflation-Targeting Regime*. Institutions include commercial banks, asset management firms, consulting firms, non-financial institutions and other firms with an assigned economist responsible for the information to be passed on to the BCB. Forecasts are supplied over different forecast horizons and for a large array of economic variables, e.g., inflation using different price

indices, interest and exchange rates, GDP, industrial production, balance of payments accounts, fiscal variables, etc.

Our focus is on Brazilian inflation, measured by the Brazilian Consumer Price Index – IPCA, the official target inflation monitored by the Brazilian *Inflation-Targeting Regime*. Our sample covers daily inflation forecasts collected from January 2nd, 2004 until January, 8th, 2015 (2,751 workdays)³. In each day t , $t = 1, \dots, T$, survey respondent i , $i = 1, \dots, N$, may inform her/his expectations regarding inflation rates (or other variables) all the way up to the next 18 months, as well as for the next 5 years on a year-end basis. For example, market respondent i may inform on February, 2nd, 2004 her/his forecast for the inflation rate of January, 2004 (not yet released), as well as for February, 2004, and the following 16 months. Next working day, the same agent may (or may not) update the forecasts for the same inflation rates of January, 2004, February, 2004,..., up to June, 2005. This way, our sample covers forecasts for the IPCA inflation from January, 2004 to December, 2014 (i.e. 132 months or events). The dataset forms an unbalanced panel ($N \times T$) containing an amount of 234,605 observations. Decomposing our total number of observations into N and T , gives the following breakdown: $T = 2,751$ daily observations (only nowcasts), and an average of $N = 85.3$ forecasters in our sample.

Besides its large size, the survey has other desirable features: (i) the system access can be done at any time by survey participants, with no pre-set schedule for updates; (ii) the confidentiality of information is guaranteed and the anonymity of forecasters is preserved (i.e., there are no reputation concerns); (iii) the *Focus Survey* has specific incentives for participants to update their forecasts (participate). The survey has periodical *critical dates* in which forecasts for different horizons are used to compute individual fore-

³Since the survey had a small cross-sectional coverage (small N) in the beginning of the sample, we start the empirical exercise in 2004.

cast errors and their respective mean-absolute-forecast-error (MAFE). For a given horizon, only forecasts at *critical dates* are used to compute individual MAFE statistics and all participants know when critical dates are. The Central Bank compares MAFEs across institutions and releases the names of the *top five* forecasting institutions (individuals) for selected variables at different horizons. This serves as an incentive scheme for survey respondents to update their forecasts for the critical date; see Carvalho and Minella (2012) and Marques (2013) for further details on the Focus Survey⁴.

Given its high-frequency character – daily observations – and the incentive mechanism to update, the *Focus Survey* is well suited to investigate the issue of incentive-driven inattention, since we have days in which participants have an incentive to update their forecast and others where the opposite occurs. At lower frequencies, it is hard to think about a similar incentive scheme being operational, since, for example, it makes little sense to have a critical month, quarter, or semester, respectively for the *Wall Street Journal Forecasting Survey*, the *Survey of Professional Forecasters*, or the *Livingston Survey*. Perhaps, for these low-frequency surveys, every observation is a *critical date*, since the cost of updating once a month, a quarter, and a semester is low relative to that of updating daily.

2.2 Stylized Facts

Critical Dates, Updating, and Accuracy

For the nowcast of monthly inflation, the chronology of relevant events is the following: (i) IPCA inflation is usually released on the 8th of the month subsequent to the reference month; (ii) on the 22nd of the reference month there

⁴The *Focus survey* is widely used in the Brazilian economy and its excellence has been internationally recognized. In 2010, the survey received the Certificate of Innovation Statistics, due to a second place in the II Regional Award for Innovation in Statistics in Latin America and Caribbean, offered by the World Bank.

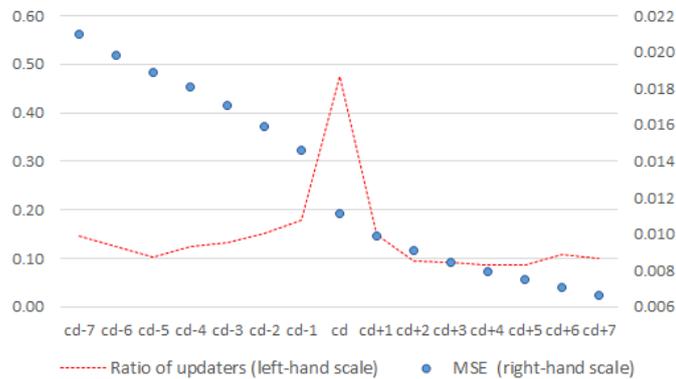
is the release of IPCA-15 inflation, which contains important information about IPCA inflation to be released in about 15 days; (iii) for every reference month, the critical date usually happens one day prior to the release of IPCA-15 inflation. Thus, 17 days prior to the release of IPCA inflation. Therefore, the critical date happens near the mid point between two consecutive monthly IPCA inflation releases.

Inflation is only observed monthly, so choosing a critical day to compute forecast errors and MSFE still gives one forecast error for every (monthly) inflation observation. It also leaves a group of other days in which the forecaster has no incentive to update her/his forecast. So, this scheme introduces variation across days in the micro data, which we can exploit to study inattentiveness. Because the individual MSFE decreases with the forecast horizon as a result of the use of more current information, and individuals want to be among the group of *top-five* forecasters, theory implies that they should update on the critical date, for which we should observe a spike in updating. This is exactly what we observe in Figure 1, where the percentage of individuals updating is displayed on the right scale. Although there is some asymmetry for updating behavior on dates prior and after the critical date, up to a reasonable approximation, the updating behavior is almost constant around 12%, on average, but we observe a spike on the critical date of about 48% of individuals updating their information.

In Figure 1, when we also plot the MSFE of the consensus forecast for current-month inflation (nowcast). We observe the usual (almost linear) decline as the forecast horizon shrinks to zero. This is the result of the acquisition of more and more information as time goes by and we approach the inflation release date. However, for the MSFEs at the critical date and later, we also observe a visible structural break: they are laying below the position they should have been – the projection of the previous almost linear MSFE curve for the days before the critical date (cd).

The stylized fact we want to explain is essentially contained in Figure 1. We want to have a model that links the increase in updating behavior due to incentives with the structural break observed in the MSFE statistic for the consensus forecast. We will endogenize updating behavior and link it to the presence of incentives to update. At the same time, this massive movement in updating behavior that we observe will have a downward effect on the MSFE statistic for the consensus forecast (an increase in accuracy), reducing it in excess of what we would have observed in the absence of these incentives. The model will be parsimonious and we will only distinguish between the critical date and the non-critical dates, although there is some asymmetry prior to and after the critical date that could still be modeled.

Figure 1 - Daily *MSFE* and ratio of updaters around the critical date (cd)



Critical Dates and Disagreement

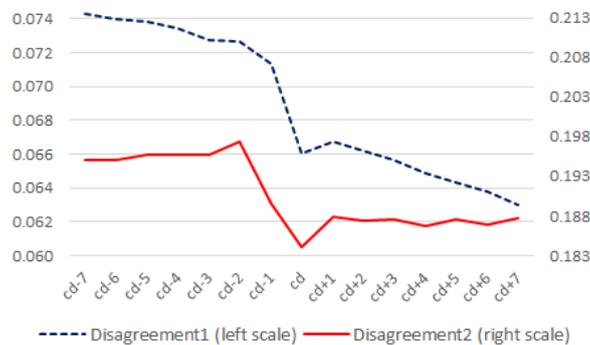
Figure 2 shows the evolution of the disagreement around the critical date based on two empirical measures: (i) Disagreement 1, which is the cross-sectional standard deviation of individual forecasts – a measure of absolute dispersion; and (ii) Disagreement 2, which is the standard deviation divided by mean – a measure of relative dispersion.⁵ Both display a change in behav-

⁵In order to avoid computing a relative dispersion based on a mean forecast very close to zero, we excluded from the disagreement2 measure the observations from June 21st,

ior in the critical date, where we observe a big drop in disagreement, some of which is partially reverted in the next working day.

Overall, we observe a change in behavior either in MSFEs and in Disagreement around the critical date. Given the obvious incentive mechanism to update forecasts in it, we now turn into a model that accounts for incentive-driven inattention.

Figure 2 - Disagreement around the critical date (cd)



Notes: Disagreement 1 is the cross-section standard deviation of individual forecasts, whereas Disagreement 2 is the cross-section standard deviation divided by the cross-section mean forecast.

3 The *Incentive-Driven Inattention* Model

We consider a model of the forecaster’s decision problem that is able to replicate the stylized facts discussed above. The model builds on Mankiw and Reis (2002), Sims (2003), Coibion and Gorodnichenko (2012, 2014), and Andrade and Le Bihan (2013). It provides an endogenous mechanism that can explain the observed variation in attention in our data base. The model is the first to link inattention to the decision problem of the forecaster and to investigate the forecaster’s incentives to provide a forecast in a context of inattention.

2006 to July 7th, 2006.

The decision problem of an individual forecaster can be described as follows. At each day t , during a given month m , a forecaster i has the option of logging into the BCB system to provide a forecast for the current month's inflation level y_m , a nowcast problem. Let t^{CD} represent the critical date and consider a window of days before and after the critical date. At each day t the forecaster decides to update the forecast if the benefits of doing so outweigh the costs. The costs of updating could be due to a variety of reasons, for example, obtaining information, processing information, running a forecasting model, logging into the BCB system, etc. The benefits are due to the higher chance of being in the group of *top five* forecasters – a clear distinction among peers. This has a higher probability to occur if the forecast in the BCB system is currently updated at $t = t^{CD}$, since the more updated a forecast is, the lower should be its *MSFE*.

We model the benefit-cost ratio for each person at each point in time $\frac{B_{it}}{C_{it}}$, where B_{it} denotes benefits and C_{it} denotes costs, as a draw from a *Normal* distribution with mean μ_t and unit variance, so that the proportion of forecasters who update at date t is given by:

$$\lambda_t = P\left(\frac{B_{it}}{C_{it}} > 1\right) = P\left(\frac{B_{it}}{C_{it}} - \mu_t > 1 - \mu_t\right) = 1 - \Phi(1 - \mu_t), \quad (1)$$

where $\Phi(\cdot)$ is the *Standard Normal* CDF.

Consider a model in which agents forecast monthly inflation, defined as $y_m = \log P_m - \log P_{m-1}$ where P_m is the price level at the end of month m . If $y_t = \log P_t - \log P_{t-1}$ is daily inflation, we must have $y_m = \sum_{t=1}^{21} y_t$, where we have assumed for simplicity that there are 21 working days in a standard month.⁶ At each date t a fraction λ_t of agents update their forecast of monthly inflation y_m . Here, we will need to account for both y_m and y_t , since inflation is sampled every month m , but the frequency of the data base

⁶In the empirical exercise we do take into account the fact that each month in our sample has a different number of working days.

is daily $-t$.

The incentive structure of the *Focus Survey* and the stylized facts shown in Figure 1 above are consistent with the assumption that the distribution of the benefit-cost ratio across agents is constant over time, but has a mean shift at the critical date due to the fact that the benefits of updating increase dramatically for every forecaster at the critical date. At $t = t^{CD}$, their accuracy is recorded in order to form the ranking of the *top five* forecasters. This setup implies that the mean of the benefit-cost distribution evolves as:

$$\mu_t = \mu_1 + \mu_2 1 [t = t^{CD}],$$

where $1 [t = t^{CD}]$ is an indicator function of the critical date. This, in turn, using (1), corresponds to a fraction of updaters that evolves as:

$$\lambda_t = \lambda_1 + \lambda_2 1 [t = t^{CD}]. \quad (2)$$

To derive the properties of the consensus forecast implied by the model, we build on Coibion and Gorodnichenko (2012) and adapt their result to our setting of monthly inflation nowcasting with a time-varying proportion of daily updaters. Coibion and Gorodnichenko show that, in models with inattentive agents, the consensus forecast at time t , F_t , can be written as a convex combination of the consensus forecast at the previous period, F_{t-1} , and the current rational expectation of inflation, $\mathbb{E}_t(y_m)$:

$$F_t = \lambda_t \mathbb{E}_t(y_m) + (1 - \lambda_t) F_{t-1}, \quad (3)$$

which implies that the consensus forecast error at time t is

$$y_m - F_t = \frac{1 - \lambda_t}{\lambda_t} (F_t - F_{t-1}) + v_t, \quad (4)$$

where $v_t = y_m - \mathbb{E}_t[y_m]$ is a rational expectation error.

We now derive the implications of this result for the accuracy of the consensus forecast in our setting. Here, the MSFE of the consensus forecast at time t evolves as:

$$MSFE_t = \mathbb{E}[(y_m - F_t)^2] = \left(\frac{1 - \lambda_t}{\lambda_t}\right)^2 \mathbb{E}[(\Delta F_t)^2] + \mathbb{E}(v_t^2), \quad (5)$$

where we have used the fact that, by construction, v_t is uncorrelated with information available at time t .

To characterize $v_t = y_m - \mathbb{E}_t[y_m]$, and compute the last term in (5) – $\mathbb{E}(v_t^2)$ – we need a model for $\mathbb{E}_t[y_m]$ in a daily basis. The main problem here is that we do not observe daily inflation, since it is sampled monthly. However, this is simply an *interpolation* problem, which has been dealt with previously by several authors, e.g., Mönch and Uhlig (2005), *inter alia*, in a multivariate setting. In the simpler context of *ARMA* models, Amemiya and Wu (1972) have general results for these processes, which we can apply to our case. Their results show that, if the daily process is an *AR*(1), then, the monthly process should be an *ARMA*(1, 1), in a one-to-one mapping. Jumping to our empirical section, this is exactly the process that best describes monthly inflation, therefore we assume that daily inflation follows an *AR*(1) process:

$$y_t = \phi y_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim i.i.d. \mathcal{N}(0, \sigma_\varepsilon^2), \quad (6)$$

which implies that the rational expectation of monthly inflation at day t is:

$$\mathbb{E}_t[y_m] = \mathbb{E}_t \left[\sum_{j=1}^{21} y_j \right] = \sum_{j=t+1}^{21} \phi^{j-t} y_t, \quad (7)$$

which corresponds to a rational expectation error:

$$\begin{aligned} v_t &= y_m - \mathbb{E}_t[y_m] = \sum_{j=1}^{21} y_j - \sum_{j=t+1}^{21} \phi^{j-t} y_t \\ &= \varepsilon_{21} + (1 + \phi)\varepsilon_{20} + \dots + (1 + \phi + \dots + \phi^{21-t-1})\varepsilon_{t+1}, \end{aligned}$$

and its variance is:

$$\mathbb{E}(v_t^2) = [1 + (1 + \phi)^2 + \dots + (1 + \phi + \dots + \phi^{21-t-1})^2] \sigma_\varepsilon^2. \quad (8)$$

The *ARMA*(1, 1) process for monthly inflation y_m as follows:

$$(1 - \phi^k L) y_m = (1 + \theta L) \varepsilon_t^*,$$

where, in our case $k = 21$, since we assume that there are 21 working days in a standard month. In the empirical implementation we will impose the proper monthly restriction to each month, but for exposition purposes we keep $k = 21$ here. Because $y_m = \sum_{t=1}^{21} y_t$, we are able to write a relationship between ε_t^* and ε_t that will depend on θ and ϕ :

$$\begin{aligned} (1 + \theta L) \varepsilon_t^* &= (1 + \phi L + \dots + \phi^{k-1} L^{k-1}) (1 + L + \dots + L^{k-1}) \varepsilon_t \\ &= \varepsilon_t \left[\sum_{i=0}^{k-1} L^i \left(\sum_{j=0}^i \phi^j \right) + \sum_{i=k}^{2(k-1)} L^i \left(\sum_{j=0}^{2(k-1)-i} \phi^{k-1-j} \right) \right]. \end{aligned}$$

Given consistent estimates for θ and ϕ , we can solve for σ_ε^2 using either of

the following equations:⁷

$$\begin{aligned}
(1 + \theta^2)\sigma_{\varepsilon^*}^2 &= \sigma_\varepsilon^2 \left[\sum_{i=0}^{k-1} \left(\sum_{j=0}^i \phi^j \right)^2 + \sum_{i=k}^{2(k-1)} \left(\sum_{j=0}^{2(k-1)-i} \phi^{k-1-j} \right)^2 \right], \text{ or,} \quad (9) \\
\theta\sigma_{\varepsilon^*}^2 &= \sigma_\varepsilon^2 \left[\sum_{i=1}^{k-1} \left(\left(\sum_{j=0}^i \phi^j \right) \left(\sum_{j=0}^{i-1} \phi^j \right) \right) \right] + \sigma_\varepsilon^2 \left[\left(\sum_{j=0}^{k-1} \phi^j \right) \left(\sum_{j=1}^{k-1} \phi^j \right) \right] \\
&\quad + \sigma_\varepsilon^2 \left[\sum_{i=k+1}^{2(k-1)} \left(\left(\sum_{j=0}^{2(k-1)-i} \phi^{k-1-j} \right) \left(\sum_{j=0}^{2(k-1)-i-1} \phi^{k-1-j} \right) \right) \right]. \quad (10)
\end{aligned}$$

We now introduce into the model a “lower bound” for MSFE by employing the “noisy-information” approach, where agents never fully observe the true state, but instead receive a noisy signal as in Sims (2003) and Coibion and Gorodnichenko (2014). First, recall that the monthly (observed) inflation y_m is modeled as the sum of the daily (latent) inflation y_t , such that $y_m = \sum_{j=1}^{21} y_j$. Each survey participant i does not observe the (latent) daily inflation y_t , but, instead, an idiosyncratic signal $y_{i,t}$, such that $y_{i,t} = y_t + \eta_{i,t}$, where $\eta_{i,t} \sim i.i.d.(0, \sigma_{\eta,i}^2)$. If, as before, daily inflation follows an $AR(1)$ model, note that $\mathbb{E}_{i,t}(y_{i,t+h}) = \mathbb{E}_{i,t}(y_{t+h})$, since $\mathbb{E}_{i,t}(\eta_{i,t+h}) = 0$ for all $h \geq 1$, where $\mathbb{E}_{i,t}(\cdot)$ denotes the conditional expectation operator of participant i at time t . Since the observed monthly inflation is the sum of the latent daily inflation, it also follows that the sum of the daily signals $y_{i,t}$ must be equal to y_m . Thus, $y_m = \sum_{j=1}^{21} y_j = \sum_{j=1}^{21} y_{i,j}$ and, therefore, the following restriction holds: $\sum_{j=1}^{21} \eta_{i,j} = 0$.

Let $\eta_t = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^t \eta_{i,j}$ be the cross-section average of accumulated

⁷I one wishes to rely solely on the estimate of ϕ , then the following set of equations can be viewed as a system of two equations with two unknowns that can be solved for θ and σ_ε^2 .

$\eta_{i,t}$. Thus, the individual conditional expectation of y_m can be written as:

$$\begin{aligned}\mathbb{E}_{i,t}(y_m) &= \mathbb{E}_{i,t}\left(\sum_{j=1}^{21} y_{i,j}\right) = \mathbb{E}_{i,t}\left(\sum_{j=1}^{21} y_j + \eta_{i,j}\right) \\ &= \sum_{j=1}^t (y_j + \eta_{i,j}) + \sum_{j=t+1}^{21} \phi^{j-t} y_t = \sum_{j=1}^t y_j + \sum_{j=t+1}^{21} \phi^{j-t} y_t + \sum_{j=1}^t \eta_{i,j}.\end{aligned}$$

The cross-section average rational expectation of y_m is, then, given by:

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{i,t}(y_m) = \mathbb{E}_t(y_m) + \eta_t.$$

Therefore, equation (3) becomes:

$$F_t = \lambda_t (\mathbb{E}_t(y_m) + \eta_t) + (1 - \lambda_t) F_{t-1}, \quad (11)$$

where $\eta_t \sim i.i.d.(0, \sigma_\eta^2)$. Note that if $\sigma_\eta^2 = 0$, then, we are back to the Coibion and Gorodnichenko (2012) updating formula as a special case.

Equation (11) implies now that a percentage λ_t of attentive survey participants update their forecasts in a (not fully) rational way, but instead including some noise, due to the uncertainty regarding latent daily inflation. By assuming that η_t is uncorrelated with v_t and $((1 - \lambda_t) / \lambda_t) \Delta F_t$, the aggregated forecast error is given by:

$$y_m - F_t = ((1 - \lambda_t) / \lambda_t) \Delta F_t + (v_t - \eta_t),$$

and the respective MSFE is given by:

$$MSFE_t = \mathbb{E} [(y_m - F_t)^2] = \left(\frac{1 - \lambda_t}{\lambda_t}\right)^2 \mathbb{E} [(\Delta F_t)^2] + \mathbb{E} (v_t^2) + \mathbb{E} (\eta_t^2). \quad (12)$$

As long as the forecast horizon decreases, from (8) it follows that $\mathbb{E} (v_t^2)$

approaches σ_ε^2 . Thus, for very short forecast horizons, even if $\mathbb{E}[(\Delta F_t)^2]$ and σ_ε^2 are close to zero, there is now a “residual” MSFE (or lower bound for the MSFE) due to $\mathbb{E}(\eta_t^2) = \sigma_\eta^2$, which is not time-dependent, applies to the whole term structure of $MSFE_t$, and might help reconciling the model outcome with the empirical evidence from data, as discussed below.

4 Empirical Analysis

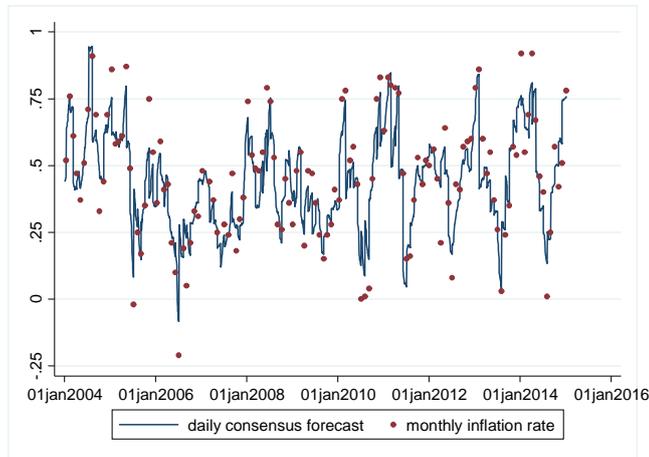
We first assess the plausibility of the $AR(1)$ model for inflation by estimating an $ARMA(1, 1)$ model for monthly inflation and backing out the corresponding parameters ϕ and σ_ε^2 for the daily model as discussed in the previous section. Then, we employ a state-space approach to construct the series y_t for daily inflation. Next, we compute the empirical counterparts of λ_1 , and λ_2 , which, together with ϕ , σ_ε^2 and y_t , allow us to uncover the right hand side of equation (12), which can then be confronted with the actual MSFEs reported in Figure 1. Our structural model has two other structural equations: equation (2), describing the evolution of the fraction of updaters, and equation (6), describing the evolution of daily inflation. Once we consider the use of proper instruments, the structural model can be estimated and tested using the generalized method of moments (GMM).

To the best of our knowledge, this paper is the first to address such issues in terms of incentive-driven inattention. The cost-benefit model proposed here should be viewed as complementary to the rational inattention literature, since it brings to it new and relevant topics, including survey design and the dynamics of participants’ behavior. The model outcomes might also be useful for policymakers interested in better understanding the evolution of market expectations, which could help in designing an optimal-incentive rule. Some of these issues are discussed below.

Modeling Monthly Inflation y_m

Figure 3 shows the observed monthly inflation rate y_m and the respective daily consensus nowcast F_t . Table 1 presents the Akaike's information criterion (AIC) and the Bayesian information criterion (BIC) for models of the *ARMA* class estimated for monthly Inflation measured by IPCA. In both cases, the *ARMA*(1, 1) belongs to the set of the three best models, the other competitors being the *AR*(1), the *AR*(2), and the *ARMA*(2, 1) model. Table 2 presents the estimated *ARMA*(1, 1) parameters. We discard the *AR*(1) alternative given that the *MA*(1) coefficient is significant at the usual levels, which could also be a problem for the *AR*(2) model, and discard the *ARMA*(2, 1) because the *ARMA*(1, 1) shows no signs of misspecification in its error structure, whereas the *ARMA*(1, 1) is a more parsimonious model.

Figure 3 - Monthly inflation rate (y_m) and daily consensus forecast (F_t)



Notes: Monthly inflation rate sample: January 2004-December 2014 (132 observations).

Daily consensus forecast sample: January 2nd, 2004 to January 8th, 2015 (2,751 workdays).

Consensus forecasts shown above are "nowcasts" with forecast horizons ranging from $h=1, \dots, 21$ workdays.

Table 1 - Information Criteria

	AIC	BIC
AR(1)	-65.84	-57.19
AR(2)	-66.19	-54.66
AR(3)	-64.24	-49.83
AR(4)	-65.25	-47.96
MA(1)	-59.07	-50.42
MA(2)	-63.44	-51.91
MA(3)	-64.91	-50.50
MA(4)	-63.37	-46.07
ARMA(1,1)	-66.16	-54.63
ARMA(2,1)	-65.87	-51.45
ARMA(1,2)	-64.17	-49.76
ARMA(2,2)	-63.87	-46.57
ARIMA(1,1,1)	-59.36	-47.86
ARIMA(2,1,2)	-58.18	-40.93

Note: The best three models according to each information criterion are marked in blue.

Table 2 - Model for monthly inflation rate: IPCA (y_m)

$$\text{ARMA}(1,1): (1 - \phi^k L) y_m = (1 + \theta L) \varepsilon_t^*$$

$\widehat{\phi}^k$	$\widehat{\theta}$	$\widehat{\sigma}_{\varepsilon^*}^2$
0.4544 (0.1261)	0.2249 (0.1220)	0.0333 (0.0046)

Note: Sample: January 2004-December 2014 (132 observations).

Robust standard errors in parentheses.

Based on the estimation of the $ARMA(1,1)$ model for the monthly inflation, we compute the corresponding estimates $\widehat{\phi} = 0.9631$ and $\widehat{\sigma}_{\varepsilon}^2 = 1.1268E - 05$ of the $AR(1)$ model for daily inflation (considering that each month has a fixed number of $k = 21$ workdays).⁸

Constructing a daily inflation series

In order to generate a daily time series of inflation, we interpolate the monthly inflation based on the method employed by Mönch and Uhlig (2005),

⁸For forecasting purposes, if one considers a month with $k = 31$ calendar days, then, it follows that the corresponding estimates become $\widehat{\phi} = 0.9749$ and $\widehat{\sigma}_{\varepsilon}^2 = 3.5475E - 06$.

which implements the state-space approach of Bernanke, Gertler and Watson (1997). They consider an *interpolation* (or rather, *distribution*) problem with mixed frequencies. Their complete model applied to unobserved daily inflation y_t would be:

$$(1 - \phi L) y_t = x_t \beta + u_t, \text{ with } u_t = \rho u_{t-1} + \varepsilon_t, \quad (13)$$

where y_t is the daily inflation, x_t is a vector of covariates (observables) and u_t is an $AR(1)$ error term, and observed monthly inflation y_m is the variable being interpolated. They impose the restriction that the daily interpolated series y_t exactly add up to the monthly observed inflation series y_m in the following way:⁹

$$y_m = \begin{cases} \sum_{t=1}^k y_t, & t = k, 2k, 3k, \dots, T \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

Monthly inflation can only be observed on days $t = k, 2k, \dots, T$, and will be the sum of the corresponding daily inflation rates in any given month. Otherwise, it is just set to a fictional value of zero. Notice that setting $y_m = 0$ for the days we do not observe it is a way of making monthly inflation

⁹Since the interpolation methodology requires the use of balanced data, we considered a broader case here, with $k = 31$ as the fixed number of days within every month. This procedure is adopted just to generate the filtered daily inflation series. In the remainder of the paper, a month with $k = 21$ workdays is employed, as previously mentioned.

observable at the daily frequency. The corresponding state-space model is:¹⁰

$$y_m = \mathbf{H}'_t \xi_t, \text{ where,} \quad (15)$$

$$\xi_t = \begin{pmatrix} y_t \\ y_{t-1} \\ y_{t-2} \\ \vdots \\ y_{t-k} \\ u_t \end{pmatrix} = \begin{pmatrix} \phi & 0 & 0 & \dots & 0 & \rho \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ \vdots & & \ddots & & & \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho \end{pmatrix} \begin{pmatrix} y_{t-1} \\ y_{t-2} \\ y_{t-3} \\ \vdots \\ y_{t-k-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} x_t \beta \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} \varepsilon_t \\ 0 \\ 0 \\ \vdots \\ 0 \\ \varepsilon_t \end{pmatrix}, \quad (16)$$

and (15) and (16) represent, respectively, the observation and state equations.

The matrix \mathbf{H}'_t is time-varying with the following format:

$$\mathbf{H}'_t = \begin{cases} [1 & 1 & 1 & \dots & 1 & 0], & t = k, 2k, 3k, \dots, T \\ [0 & 0 & 0 & \dots & 0 & 0], & \text{otherwise.} \end{cases} \quad (17)$$

To apply Mönch and Uhlig's setup to our problem of *interpolating* monthly inflation, we set ϕ to 0.9631 (which is the estimate from the *ARMA*(1, 1) – *AR*(1) approach of Amemiya and Wu, 1972) and set $\beta = \rho = 0$ in the system

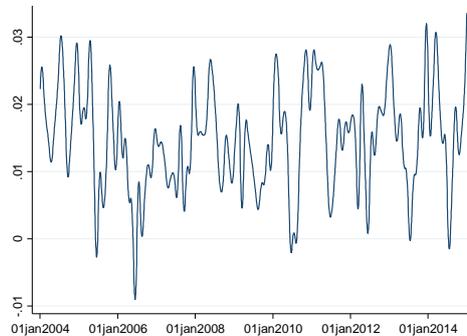
¹⁰In the Kalman-filter literature for mixed frequency models (e.g., Giannone, Reichlin and Small, 2008) a fictional value is usually assumed for missing observations (zero is the most frequent choice). The crucial step is to impose that the fictional observed data has a very large variance, so that the zero value is discounted and overwritten by the Kalman-filter technique. This is exactly how Mönch and Uhlig proceed.

(15), (16) and (17). This yields our corresponding daily $AR(1)$ model:

$$y_t = \phi y_{t-1} + \varepsilon_t,$$

with an embedded interpolating restriction, which is consistent with the $ARMA(1,1)$ model estimated with monthly data. Results are presented in Figure 4.

Figure 4 - Daily inflation rate (y_t) series based on the state-space model with $\phi = 0.9631$



GMM Estimation of the *Incentive-Driven Inattention* model

We propose the joint estimation the structural parameters of our model using the generalized method of moments (GMM) based on three moment conditions derived from the cost-benefit model of incentive-driven inattention. The first is the equation fitting the observed time-varying degree of attention λ_t using equation (2), with two structural parameters – λ_1 and λ_2 . Since we observe the proportion of agents updating in each point in time (days), we have an empirical measure of λ_t from the data, which we confront to our model in equation (2). The second equation is the square of equation (6) for daily inflation y_t (including an intercept c), which was estimated using the techniques discussed in the previous section. The third equation describes

the evolution of the $MSFE_t$ using (12), where we can compute the daily MSFE in nowcasting inflation from the data, which can then be confronted with our model. The three structural equations are:

$$\begin{aligned} \lambda_t - \lambda_1 - \lambda_2 \mathbf{1}_{t=t^{CD}} &= 0, \\ (y_t - c - \phi y_{t-1})^2 - \varepsilon_t^2 &= 0, \\ MSFE_t - \left(\frac{1 - \lambda_t}{\lambda_t} \right)^2 \mathbb{E} [(\Delta F_t)^2] - \mathbb{E}(v_t^2) - \mathbb{E}(\eta_t^2) &= 0. \end{aligned} \quad (18)$$

From (8), it follows that $\mathbb{E}(v_t^2) = w_t \sigma_\varepsilon^2$, where $w_t \equiv [1 + (1 + \phi)^2 + \dots + (1 + \phi + \dots + \phi^{21-t-1})^2]$. Taking the conditional expectation $\mathbb{E}(\cdot | \mathcal{F}_{t-1})$ on both sides of the three equations in (18), where \mathcal{F}_{t-1} is the information set available at time $t - 1$, and using respectively for each equation the set of valid instruments $z_{1,t-1}$, $z_{2,t-1}$, and $z_{3,t-1}$, the system (18) can be written as a system of restrictions on the set of structural parameters $(\lambda_1, \lambda_2, c, \phi, \sigma_\varepsilon^2, \sigma_\eta^2)$, as follows:

$$0 = \mathbb{E} [(\lambda_t - \lambda_1 - \lambda_2 \mathbf{1}_{t=t^{CD}}) \otimes z_{1,t-1}], \quad (19)$$

$$0 = \mathbb{E} [((y_t - c - \phi y_{t-1})^2 - \sigma_\varepsilon^2) \otimes z_{2,t-1}], \quad (20)$$

$$0 = \mathbb{E} \left[\left(\begin{array}{c} MSFE_t - \left(\frac{1 - \lambda_1 - \lambda_2 \mathbf{1}_{t=t^{CD}}}{\lambda_1 + \lambda_2 \mathbf{1}_{t=t^{CD}}} \right)^2 \mathbb{E} [(\Delta F_t)^2] \\ - w_t \sigma_\varepsilon^2 - \sigma_\eta^2 \end{array} \right) \otimes z_{3,t-1} \right], \quad (21)$$

where the vector of instruments include an intercept and other observables dated on $t-1$ and older, and $\dim(z_{j,t-1}) = n_j \geq 1$, $j = \{1, 2, 3\}$. There are six structural parameters $(\lambda_1, \lambda_2, c, \phi, \sigma_\varepsilon^2, \sigma_\eta^2)$ to be estimated and $(n_1 + n_2 + n_3)$ moment conditions. Over-identification requires that $(n_1 + n_2 + n_3) > 6$.

In GMM estimation, we are treating the following as observables: λ_t , $MSFE_t$ and $\mathbf{1}_{t=t^{CD}}$. In addition, we also treated as observables the following variables: y_t , w_t and F_t , where y_t comes from the (Kalman-filtered) daily inflation series, w_t is computed from an initial estimate of ϕ (i.e., $\hat{\phi} = 0.9631$)

and $w_t \equiv [1 + (1 + \phi)^2 + \dots + (1 + \phi + \dots + \phi^{2^{1-t}-1})^2]$, and F_t is constructed using the Coibion and Gorodnichenko (2012) updating formula, which, in turn, relies on the observed series λ_t and on the constructed series $E_t(y_m)$, generated from y_t and $\hat{\phi}$.¹¹ We also did some robustness analysis regarding F_t , since we can get a daily observed consensus forecast from the *Focus Survey*.

Table 5 presents the GMM estimates¹² from two different sets of instruments.¹³ The over-identifying-restrictions (OIR) are tested using the J -test due to Hansen (1982). The joint significance of the variance terms σ_ε^2 and σ_η^2 is also verified.

¹¹In addition, as an initial value, for every month (or event), the first consensus forecast from the Coibion and Gorodnichenko formula is set equal to the respective consensus forecast from data.

¹²The "iterative" procedure of Hansen et al. (1996) is employed in the GMM estimation and the initial weight matrix is the identity.

¹³Set of instruments I: $z_{1,t-1} = [1, \Delta F_{t-2}, \dots, \Delta F_{t-4}]'$, $z_{2,t-1} = [1, \lambda_{t-2}, \dots, \lambda_{t-5}]'$ and $z_{3,t-1} = [1, y_{t-2}, \dots, y_{t-5}]'$. Set of instruments II: $z_{1,t-1} = [1, \Delta^2 F_{t-1}, \dots, \Delta^2 F_{t-3}]'$, $z_{2,t-1} = [1, \Delta \lambda_{t-1}, \dots, \Delta \lambda_{t-3}]'$ and $z_{3,t-1} = [1, \Delta y_{t-1}, \dots, \Delta y_{t-3}]'$. This way, for the first set of instruments, we have 14 moment conditions and 6 structural parameters, resulting in a GMM setup with 8 degrees of freedom. In the second set of instruments, we have 12 moment conditions and, thus, a GMM setup with 6 degrees of freedom.

Table 5 - GMM estimates of the structural parameters
of the *Incentive-Driven Inattention* model

Panel A: Set of instruments I

$\widehat{\lambda}_1$	$\widehat{\lambda}_2$	\widehat{c}	$\widehat{\phi}$	$\widehat{\sigma}_\varepsilon^2$	$\widehat{\sigma}_\eta^2$
0.1149*** (0.0206)	0.3746* (0.2073)	0.0007*** (0.0002)	0.9545*** (0.0123)	$2.77e - 07$ *** ($6.80e-08$)	0.0079*** (0.0006)

OIR test (8 degrees of freedom): p-value = 0.8022

$H_0 : [\sigma_\varepsilon^2 = 0; \sigma_\eta^2 = 0]$ p-value = 0.0000

Panel B: Set of instruments II

$\widehat{\lambda}_1$	$\widehat{\lambda}_2$	\widehat{c}	$\widehat{\phi}$	$\widehat{\sigma}_\varepsilon^2$	$\widehat{\sigma}_\eta^2$
0.1158*** (0.0089)	0.3915*** (0.1047)	0.0036** (0.0015)	0.7638*** (0.0970)	$3.53e - 06$ ($2.76e-06$)	0.0070*** (0.0015)

OIR test (6 degrees of freedom): p-value = 0.8348

$H_0 : [\sigma_\varepsilon^2 = 0; \sigma_\eta^2 = 0]$ p-value = 0.0000

Notes: (i) Robust Newey-West SE in parentheses. (ii) ***, **, * indicate significance at 1%, 5% and 10% levels, respectively. (iii) OIR denotes the Over-Identifying Restriction J-test due to Hansen (1982). (iv) F_t used here is constructed using the Coibion and Gorodnichenko (2012) updating formula.

In both estimation setups, the over-identifying-restriction test is not rejected at the usual levels of significance. The joint test for the zero variances involving σ_ε^2 and σ_η^2 is rejected at any admissible level in both instances, although the σ_ε^2 is not significant in Panel B above. By comparing the point estimates from the two sets of instruments, one should note that the estimated parameters $\widehat{\lambda}_1$, $\widehat{\lambda}_2$ and $\widehat{\sigma}_\eta^2$ remained quite stable, although the point estimates \widehat{c} , $\widehat{\phi}$ and $\widehat{\sigma}_\varepsilon^2$, have changed.¹⁴

Next, we perform the estimation of the system in (19), (20), and (21),

¹⁴In this case, note that for a given time series $y_t = c + \phi y_{t-1} + \varepsilon_t$, and its respective unconditional mean $E(y_t) \equiv \mu = c/(1 - \phi)$, it follows that a decrease in the autoregressive parameter ϕ (e.g., due to a change in the set of instruments in a GMM estimation) might lead to an increase in the variance of residuals (σ_ε^2) and, at the same time, imply an increase of the intercept c (provided that μ remains unchanged). This is exactly the observed change in point estimates \widehat{c} , $\widehat{\phi}$ and $\widehat{\sigma}_\varepsilon^2$ when one goes from the GMM estimation based on the first set of instruments to the second set of instruments.

using observed F_t – the daily observed consensus forecast from the *Focus Survey* – instead of that obtained using Coibion and Gorodnichenko’s (2012) updating formula. Results are presented in Table 6.

Table 6 - GMM estimates of the structural parameters
of the *Incentive-Driven Inattention* model

Panel A: Set of instruments I

$\hat{\lambda}_1$	$\hat{\lambda}_2$	\hat{c}	$\hat{\phi}$	$\hat{\sigma}_\varepsilon^2$	$\hat{\sigma}_\eta^2$
0.1125*** (0.0211)	0.3997* (0.2128)	0.0007*** (0.0002)	0.9530*** (0.0117)	$2.86e - 07$ *** ($6.65e-08$)	0.0097*** (0.0004)

OIR test (8 degrees of freedom):

p-value = 0.8887

$H_0 : [\sigma_\varepsilon^2 = 0; \sigma_\eta^2 = 0]$

p-value = 0.0000

Panel B: Set of instruments II

$\hat{\lambda}_1$	$\hat{\lambda}_2$	\hat{c}	$\hat{\phi}$	$\hat{\sigma}_\varepsilon^2$	$\hat{\sigma}_\eta^2$
0.1090*** (0.0105)	0.4739*** (0.1248)	0.0038** (0.0015)	0.7543*** (0.0961)	$3.79e - 06$ ($2.85e-06$)	0.0092*** (0.0014)

OIR test (6 degrees of freedom):

p-value = 0.7779

$H_0 : [\sigma_\varepsilon^2 = 0; \sigma_\eta^2 = 0]$

p-value = 0.0000

Notes: (i) Robust Newey-West SE in parentheses. (ii) ***, **, * indicate significance at 1%, 5% and 10% levels, respectively. (iii) OIR denotes the Over-Identifying Restriction J-test due to Hansen (1982).

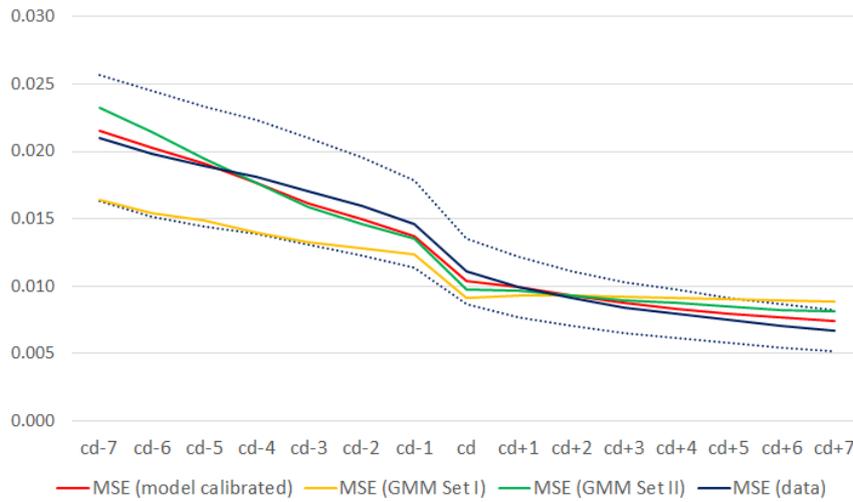
(iv) F_t used here is the daily observed Consensus Forecast from the Focus Survey.

The robustness exercise presented in Table 6 shows very similar results to those in Table 5. All in all, we obtained sensible and mostly significant structural parameters that are able to explain the dynamic behavior of the consensus MSFE across dates, including the critical date. In misspecification tests, the structural model is not rejected using a variety of instruments in GMM estimation. We now turn to the motivating picture of this paper – Figure 1 – asking whether or not the estimated structural model can reproduce the MSFE profile depicted there.

Figure 6 shows a comparison of the MSFE observed in our dataset with

the one implied by the *Incentive-Driven Inattention* model, when GMM point estimates are used to construct the MSFE profile. We also include OLS-estimated parameters for comparison; see the Appendix.

Figure 5 - MSFE (Data *versus* GMM-estimated and OLS-estimated) and 95% confidence interval



Note: Confidence interval based on (pointwise) asymptotic values: $\text{std. error} = \text{sample std.dev.} / \sqrt{T}$, where $T = 132$ events (or months) for each investigated date $CD + j$; $j = -7, \dots, 7$.

First, note that all three *MSFE* curves implied by the model fall inside the asymptotic 95% confidence interval for the observed *MSFE* – a small exception for GMM using set of instruments I, at $cd+6$ and $cd+7$. This confirms the usefulness of the *Incentive-Driven Inattention* model in explaining what we observe for the MSFE profile.

Second, note that the MSFE curve from the second set of instruments (in the GMM setup) provides better results compared to the first set of instruments. A possible explanation is the point estimate $\hat{\sigma}_\varepsilon^2$, which is more than ten times higher in the second set of instruments (compared to the first one), seems to impact the overall slope of the MSFE curve.

Finally, note that all three curves implied by the model capture quite well the sharp decrease in MSFE at the critical date (cd), which is due to the fact that the point estimates for the attention parameters λ_1 and λ_2 are similar across these three approaches (i.e., OLS-two-step, GMM set I, GMM set II) and reflect reasonably well the observed series λ_t through equation $\lambda_t = \lambda_1 + \lambda_2 1_{t=t^{CD}}$.

Modeling the Probability of Updating (Logit/Probit)

This section presents the estimates of an econometric model of the probability of updating in a panel-data context. It is a reduced-form model for the dependent variable $\lambda_{i,t}^*$, the unobserved willingness to update information. We do not observe $\lambda_{i,t}^*$. However, we do observe $\lambda_{i,t}$, which shows whether or not agent i has updated information in period t :

$$\lambda_{i,t} = \begin{cases} 1 & \text{if } \lambda_{i,t}^* > 0 \\ 0 & \text{if } \lambda_{i,t}^* \leq 0. \end{cases}$$

The reduced form for $\lambda_{i,t}^*$ is the following:

$$\lambda_{i,t}^* = \alpha_i + x_t \beta + z_{i,t-1} \gamma + e_{i,t}^2 \delta + \varepsilon_{i,t}, \quad (22)$$

where α_i is the unobserved heterogeneity, captured by either fixed or random effects; the $x_t' = [d_t^{mon}; d_t^{tue}; d_t^{thu}; d_t^{fri}; d_t^{CD}; d_t^{ipca}; d_t^{ipca15}; d_{t-1}^{MPC}]'$ vector contains the following variables: $d_t^{mon}; d_t^{tue}; d_t^{thu}$ and d_t^{fri} are dummies for the days of the week (workdays only), d_t^{CD} is a dummy for the critical date, d_t^{ipca} and d_t^{ipca15} are dummies for the days of release of the IPCA and IPCA15, respectively, and d_{t-1}^{MPC} is a dummy for the (lagged) day of release of the minutes of the meeting of the Monetary Policy Committee (MPC) of the Central Bank of Brazil; $z_{i,t-1}$ is a one-day lagged measure of uncertainty

captured by the Emerging Market Bond Index (EMBI) for Brazil – the daily risk premium on Brazilian sovereign debt; $e_{i,t}^2$ is the squared forecast error made by respondent i in the previous month. The parameters in β , γ , and δ are estimated using a panel Probit or Logit model,¹⁵ based on fixed effect (only Logit¹⁶) or random effects (both Logit and Probit). This amounts to choosing the parametric form of $G(\cdot)$ in:

$$\Pr(\lambda_{i,t} = 1 | x_t, z_{i,t-1}, e_{i,t}^2) = G(\alpha_i + x_t\beta + z_{i,t-1}\gamma + e_{i,t}^2\delta),$$

where $G(\cdot)$ can take the form of a conditional Normal CDF for the Probit model or the conditional Logistic CDF for the Logit model.

In order to get comparable estimated coefficients, we standardized all regressors (i.e., zero mean and unit variance, by demeaning and dividing each regressor by its respective sample standard deviation). The estimation results are presented below in Table 8:

¹⁵In order to estimate the referred panel models we assume strict exogeneity, in which the explanatory variables in each time period are uncorrelated with the idiosyncratic error in each time period; which is a much stronger assumption than simply assuming no contemporaneous correlation.

¹⁶Performing fixed effects Probit estimation is more complicated than its Logit analog, since there does not exist a sufficient statistic allowing the fixed effects to be conditioned out of the likelihood. Nonetheless, unconditional fixed-effects Probit models could be estimated with indicator variables for the panel dimension, although resulting in (possible) biased estimates. See Wooldridge (2002) for further details.

Table 8 - Panel Logit/Probit model of the probability of update
(dependent variable: $\lambda_{i,t}$)

Regressors	Logit-FE	Logit-RE	Probit-RE
d_t^{mon}	0.112 *** (0.008)	0.112 *** (0.008)	0.061 *** (0.005)
d_t^{tue}	-0.013 (0.009)	-0.013 (0.009)	-0.006 (0.005)
d_t^{thu}	-0.007 (0.009)	-0.008 (0.009)	-0.004 (0.005)
d_t^{fri}	0.180 *** (0.009)	0.180 *** (0.009)	0.097 *** (0.005)
d_t^{CD}	0.493 *** (0.005)	0.493 *** (0.005)	0.286 *** (0.003)
d_t^{IPCA}	0.117 *** (0.006)	0.117 *** (0.006)	0.065 *** (0.003)
d_t^{IPCA15}	0.057 *** (0.006)	0.057 *** (0.006)	0.031 *** (0.003)
d_{t-1}^{MPC}	0.038 *** (0.006)	0.038 *** (0.006)	0.021 *** (0.003)
Embi_{t-1}	0.037 *** (0.007)	0.038 *** (0.007)	0.022 *** (0.004)
$e_{i,t}$	0.010 (0.007)	0.008 (0.007)	0.005 (0.004)
Log likelihood	-77749.5	-78772.2	-78741.0

Notes: FE means fixed effects and RE random effects. Robust standard errors in parentheses.

Number of observations = 223,685. *** indicates 1% significance level.

The main result in Table 8 is that the dummy for the *Critical Date* – d_t^{CD} – is the most important variable in explaining the probability of agent i updating her/his forecast in day t . Given that we normalized all regressors to have zero mean and unit variance, we can compare coefficients in the Logit and Probit regressions. Note that, in absolute value, the estimated coefficient of d_t^{CD} is at least 2.5 times that of the second largest coefficient – d_t^{fri} .

It is also worth noting that the impact of the IPCA release is higher in comparison to that of the release of IPCA15 and of the Central-Bank minutes. Moreover, the coefficient for the past squared forecast error is not significant, but the coefficient for uncertainty (EMBI) is positive and significant, although with a lower magnitude in comparison to the significant

coefficients for the calendar-effect dummies.

With respect to the dummies for the weekdays, the highest estimated coefficient was for the Friday dummy. This can be related to a market readout published on every Monday morning by the Central Bank of Brazil (Focus-Market Readout). It releases key aggregate statistics from the *Focus Survey*, such as the consensus forecast, the median, the standard deviation, the coefficient of variation, the maximum and minimum, etc., based on data collected up to 5 PM of the previous Friday. See Marques (2013, p. 305) for further details.

Finally, note that the Logit coefficients obtained from fixed effects are very close to the ones from the random effects model. In comparing the coefficients estimated by the Probit and Logit models, Amemiya (1985) suggests multiplying a Logit estimate by 0.625 to get an estimate of the corresponding Probit estimate, since their respective distributions have zero mean, but their variances are different (1 for the Normal and $\pi^2/3$ for the logistic). In this respect, our empirical results are matched in most cases.

5 Rationality Tests and Survey Design

In this section, we provide evidence that results of rationality tests are severely affected if one disregards the different behavior of agents in critical and non-critical dates. Moreover, we also provide evidence that test results based on a time-series of individual forecasts are very different from those using the time-series of the aggregate measure of consensus forecasts or panel-data on individuals under homogeneity restrictions.

When testing individuals separately, at the critical dates alone, we find compelling evidence of rationality, whereas we find the opposite in different contexts: non-critical dates, all dates, time-series of the consensus forecasts, and panel-data on individuals under homogeneity restrictions. This shows

the importance of understanding and modelling *incentive-driven inattention*. The latter is not merely a curiosity found in an exquisite data-base, but something that must be understood as we gather more frequent time-series observations in panels of forecasts. It also has a bearing on survey design, where surveys are taken in high frequency (weekly, daily, or tic by tic) and we consider changes in the incentive structure of the survey in order to reduce the MSFE for the survey user. Since we believe that the expectation surveys of the future will be of that kind, this paper provides a step forward in designing their structure.

Rationality tests

Let π_{t+h} be the inflation rate as measured by the IPCA, $f_{i,t+h|t}$ be the respective inflation forecast of survey participant i formed at period t , and $F_{t+h|t}$ be the consensus forecast, i.e., $F_{t+h|t} = \frac{1}{N} \sum_{i=1}^N f_{i,t+h|t}$. We considered six sampling schemes, due to the possibility of using aggregated or disaggregated data in both dimensions. The time dimension can be computed in two ways: daily frequency ($T = 2,751$ workdays), where we use all days in the sample, or in monthly frequency ($T = 132$ months), where we use only critical dates. The cross-section dimension can be considered in three ways: (i) disaggregated data, with an individual OLS regression for each $i = 1, \dots, N$ ¹⁷; (ii) disaggregated data, with panel data linear regression (random-effects or fixed-effects models); and (iii) aggregated data (consensus forecast), with OLS regression. On the other hand, we considered three types of rationality tests: (1) Mincer-Zarnowitz (MZ) regressions; (2) FIRE (Full-Information Rational Expectations) tests, from Coibion and Gorodnichenko (2014, equation 10); and (3) FIRE tests, weaker version, from Coibion and Gorodnichenko (2014, equation 12).

¹⁷In each test with disaggregated data (individual OLS regressions), we only considered survey participants with more than 10 available nowcasts.

Mincer-Zarnowitz Tests

First we consider disaggregate data, with an individual OLS regression for each $i = 1, 2, \dots, N$. The individual-forecasts OLS regressions are given as follows:

$$\pi_{t+h} = \alpha_i + \beta_i f_{i,t+h|t} + \varepsilon_{i,t+h}, \quad (23)$$

where, π_{t+h} denotes observed inflation on period $t + h$, and $f_{i,t+h|t}$ denotes individual i forecast of π_{t+h} using only information up to period t . The Mincer and Zarnowitz (1969) test is based on a Wald test of the joint null hypothesis of rationality, $H_0 : [\alpha_i = 0; \beta_i = 1]$ for each i . We then compute the proportion of agents for which we do not reject the null of rationality.

Next, we consider disaggregate data in a panel-data linear-regression (random-effects or fixed-effects models) context:

$$\pi_{t+h} = \alpha + \beta f_{i,t+h|t} + v_i + \varepsilon_{i,t+h}, \quad (24)$$

where v_i is the unobserved heterogeneity and $\varepsilon_{i,t+h}$ is the error term, and, due to the curse of dimensionality, we impose the homogeneity restriction that $\beta_i = \beta$ for all $i = 1, 2, \dots, N$. Here, the Mincer-Zarnowitz test of rationality is based on a Wald test of the joint null hypothesis $H_0 : [\alpha = 0; \beta = 1]$.

At last we consider aggregate data (consensus forecast), with OLS regression in time-series data. The consensus-forecast OLS regression is given by:

$$\pi_{t+h} = \alpha + \beta F_{t+h|t} + \varepsilon_{t+h}, \quad (25)$$

where the rationality test is based on a Wald test of the joint null hypothesis $H_0 : [\alpha = 0; \beta = 1]$. In testing all the null hypotheses listed above we have employed robust standard errors. Results are summarized in Table 9.

Table 9 - Mincer-Zarnowitz (MZ) regressions

(i) Individual Forecasts, OLS	N	NR	% Rational= NR/N
All Days – daily frequency	203	28	14%
Critical Days only – monthly frequency	176	161	91%
(ii) Individual Forecasts, Panel	$\hat{\alpha}$	$\hat{\beta}$	p-value
All Days - daily (FE), $N \times T = 234, 605$	0.02097 (0.00228)	0.98264 (0.00515)	0.0009
All Days - daily (RE), $N \times T = 234, 605$	0.02135 (0.00362)	0.98261 (0.00514)	0.0000
Critical Days only - monthly (FE), $N \times T = 11, 290$	0.01711 (0.00249)	0.98414 (0.00557)	0.0048
Critical Days only - monthly (RE), $N \times T = 11, 290$	0.01839 (0.00317)	0.98269 (0.00557)	0.0000
(iii) Consensus Forecast, OLS	$\hat{\alpha}$	$\hat{\beta}$	p-value
All Days – daily, $T = 2, 751$	-0.04243 (0.00548)	1.12749 (0.0114)	0.0000
Critical Days only – monthly, $T = 132$	-0.04523 (0.02508)	1.12517 (0.05308)	0.0363

Note: Robust standard error in parenthesis. FE means fixed-effects and RE means random-effects.

In the second column, N indicates the number of survey participants with at least 10 observations. In the third column, NR indicates the number of rational participants (p-value>0.05 in the rationality test).

Results in Table 9 show compelling evidence that incentives are important to understand the results of rationality tests. First, when we consider all the days in our sample, tests using Mincer-Zarnowitz regressions with disaggregate data and individual OLS regressions for each i show that only 14% of the individuals pass rationality tests at the 5% significance level, i.e., for only 14% of the individuals we cannot reject the null hypothesis. However, if we perform the same test using only the *Critical Dates*, results are com-

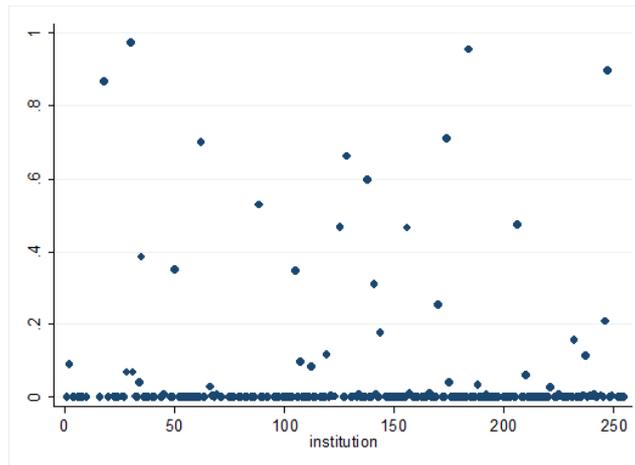
pletely different in which we cannot reject the null of rationality for 91% of the individuals at the 5% level. A natural conclusion is that incentives matter for rationality-test results. As argued above, in critical dates, individuals become more attentive and this additional attentiveness reduces inertia in inflation expectations, changing the results of rationality tests. For critical dates, rational behavior is the rule, not the exception.

In Table 9, we also consider disaggregate data in a panel-data linear-regression context, where we impose homogeneity restrictions ($\beta_i = \beta, \forall i$) in testing the joint null hypothesis $H_0 : [\alpha = 0; \beta = 1]$. Here, in testing the null, we have the auxiliary hypothesis of homogeneity. Hence, rejection can also be due to homogeneity itself. Regardless of whether we consider critical or non-critical dates, the null is always rejected at the 5% level, or even at the 1% level.

Finally, we analyze test results using the consensus forecasts with OLS regressions. The null is always rejected at the 5% level, regardless of whether we consider critical or non-critical dates, which is consistent with results previously found in the literature, where rationality was rejected overwhelmingly in Mincer-Zarnowitz tests.

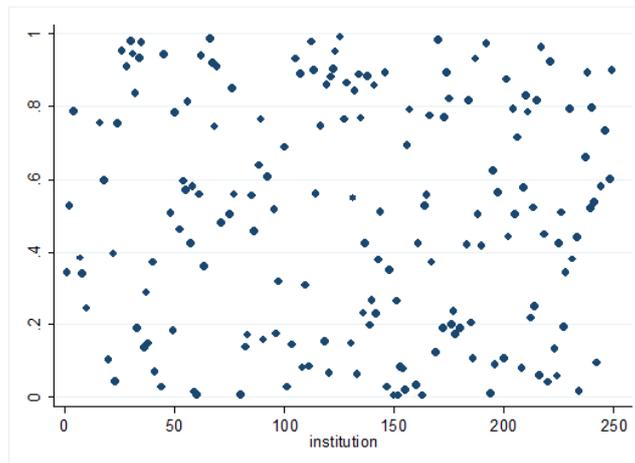
Next, we provide a summarizing picture of the results of rationality tests using individual time-series data for non-critical and critical dates, respectively.

Figure 10 - P-values of the joint test $H_0 : [\alpha_i = 0; \beta_i = 1]$ for each i
 Case (i) Individual Forecasts, OLS (daily frequency)



Notes: The p-value (vertical axis) of the rationality test based on the MZ regression is plotted by each survey participant (horizontal axis).

Figure 11 - P-values of the joint test $H_0 : [\alpha_i = 0; \beta_i = 1]$ for each i
 Case (i) Individual Forecasts, OLS (only critical dates, monthly freq.)



Notes: The p-value (vertical axis) of the rationality test based on the MZ regression is plotted by each survey participant (horizontal axis).

FIRE Tests

In this section we consider full-information rational-expectation (FIRE) tests as used in Coibion and Gorodnichenko (2014, Equation 10). They are a variant of the Mincer-Zarnowitz tests. For individual forecasts, OLS regressions, we run:

$$(\pi_{t+h} - f_{i,t+h|t}) = \alpha_i + \delta_i (f_{i,t+h|t} - f_{i,t+h|t-1}) + \mu_{i,t+h}, \quad (26)$$

where the FIRE test is based on the null of rationality, $H_0 : [\alpha_i = \delta_i = 0]$ for each i . We then compute the proportion of agents for which we do not reject the null of rationality. For individual forecasts, panel-data regression, we run:

$$(\pi_{t+h} - f_{i,t+h|t}) = \alpha + \delta (f_{i,t+h|t} - f_{i,t+h|t-1}) + v_i + \mu_{i,t+h}, \quad (27)$$

where v_i is the unobserved heterogeneity, and the FIRE test is based on the null $H_0 : [\alpha = \delta = 0]$. For the consensus forecast, OLS regression, we run:

$$(\pi_{t+h} - F_{t+h|t}) = \alpha + \delta (F_{t+h|t} - F_{t+h|t-1}) + \mu_{t+h}, \quad (28)$$

where the FIRE test is based on the null $H_0 : [\alpha = \delta = 0]$. Results of FIRE tests are presented in Table 10.

Table 10 - FIRE Test Results
(Full-Information Rational-Expectation Tests)

(i) Individual Forecasts, OLS	N	NR	% Rational= NR/N
All Days – daily frequency	203	52	26%
Critical Days only – monthly frequency	171	157	92%
(ii) Individual Forecasts, Panel	$\hat{\alpha}$	$\hat{\delta}$	p-value
All Days - daily (FE), $N \times T = 233, 120$	0.01325 ($2.00E-7$)	-0.10445 (0.00635)	0.0000
All Days - daily (RE), $N \times T = 233, 120$	0.01351 (0.00271)	-0.10447 (0.00635)	0.0000
Critical Days only - monthly (FE), $N \times T = 10, 710$	0.00956 ($8.51E-07$)	-0.02925 (0.00639)	0.0000
Critical Days only - monthly (RE), $N \times T = 10, 710$	0.01058 (0.00155)	-0.02908 (0.00640)	0.0000
(iii) Consensus Forecast, OLS	$\hat{\alpha}$	$\hat{\delta}$	p-value
All Days – daily, $T = 2, 750$	0.01382 (0.00209)	0.28923 (0.11122)	0.0000
Critical Days only – monthly, $T = 131$	0.00987 (0.00905)	0.14912 (0.07264)	0.0411

Notes: Standard error in parenthesis. FE means fixed-effects and RE means random-effects.

In the second column, N indicates the number of participants with at least 10 observations.

In the third column, NR indicates the number of rational participants (p-value>0.05 in the rationality test). The lagged forecast is related to the previous day (in daily frequency)

and to the previous month (i.e., previous critical date, in the monthly frequency).

Results in Table 10 are roughly equal to those in Table 9: at critical dates only, we have compelling evidence that individuals are rational when tested separately – we cannot reject the null of rationality for 92% of the individuals at the 5% level. When they are tested imposing homogeneity restrictions ($\delta_i = \delta, \forall i$), or using a consensus forecast, the null of rationality is always rejected at the 5% level (critical dates, consensus forecasts), and sometimes

even at the 1% level (all remaining cases). For the consensus forecast, note that, under the sticky-information model, it follows that $\lambda = \delta/(1 + \delta)$, whereas under the noisy-information model, it follows that $G = 1/(1 + \delta)$, where G is the Kalman-filter gain. Hence, for daily data (all dates), it follows that $\hat{\lambda} = 0.22$ and $\hat{G} = 0.78$ whereas, for monthly data (critical dates only), it follows that $\hat{\lambda} = 0.13$ and $\hat{G} = 0.87$.

Weak Version of FIRE Tests

In this section we consider the weaker version of the FIRE tests, as used in Coibion and Gorodnichenko (2014, Equation 12). They are a variant of the FIRE tests presented above. For individual forecasts, OLS regressions, we run:

$$(\pi_{t+h} - f_{i,t+h|t}) = \alpha_i + \beta_{1,i}f_{i,t+h|t} + \beta_{2,i}f_{i,t+h|t-1} + \omega_{i,t+h}, \quad (29)$$

where the test is based on the null of rationality, $H_0 : [\alpha_i = 0; \beta_{1,i} + \beta_{2,i} = 0]$ for each i . We then compute the proportion of agents for which we do not reject the null of rationality. For individual forecasts, panel-data regression, we run:

$$(\pi_{t+h} - f_{i,t+h|t}) = \alpha + \beta_1 f_{i,t+h|t} + \beta_2 f_{i,t+h|t-1} + v_i + \omega_{i,t+h}, \quad (30)$$

where v_i is the unobserved heterogeneity, and the test is based on the null $H_0 : [\alpha = 0; \beta_1 + \beta_2 = 0]$. For the consensus forecast, OLS regression, we run:

$$(\pi_{t+h} - F_{t+h|t}) = \alpha + \beta_1 F_{t+h|t} + \beta_2 F_{t+h|t-1} + \omega_{t+h}, \quad (31)$$

where the test is based on the null $H_0 : [\alpha = 0; \beta_1 + \beta_2 = 0]$. As argued by Coibion and Gorodnichenko, in models with informational rigidities, one should expect that $\beta_1 > 0$ and $\beta_2 < 0$.

Table 11 - FIRE Test Results, Weak Version
(Full-Information Rational-Expectation Tests)

(i) Individ.Forec., OLS	N	NR	% Rational NR/N	
All Days – daily frequency	203	29	14%	
Critical Days Only – monthly freq.	171	154	90%	
(ii) Individ.Forec., Panel	$\hat{\alpha}$	$\hat{\beta}_1$	$\hat{\beta}_2$	p-value
All Days - daily (FE), $N \times T = 233, 120$	0.01969 (0.00228)	-0.11182 (0.00801)	0.09729 (0.00552)	0.0000
All Days - daily (RE), $N \times T = 233, 120$	0.02011 (0.00364)	-0.11186 (0.00801)	0.09730 (0.00551)	0.0000
Critical Days only - monthly (FE), $N \times T = 10, 710$	0.01285 (0.00262)	-0.03293 (0.00736)	0.02554 (0.00670)	0.0000
Critical Days only - monthly (RE), $N \times T = 10, 710$	0.01555 (0.00279)	-0.03470 (0.00742)	0.02323 (0.00659)	0.0000
(iii) Consensus, OLS	$\hat{\alpha}$	$\hat{\beta}_1$	$\hat{\beta}_2$	p-value
All Days - daily frequency, $T = 2, 750$	-0.04080 (0.00549)	0.34895 (0.10344)	-0.22525 (0.10442)	0.0000
Critical Days only - monthly, $T = 131$	-0.02942 (0.02677)	0.19294 (0.07509)	-0.10421 (0.07746)	0.2113

Notes: Standard error in parenthesis. FE means fixed-effects and RE means random-effects.

In the second column, N indicates the number of participants with at least 10 observations.

In the third column, NR indicates the number of rational participants (p-value>0.05 in the rationality test). The lagged forecast is related to the previous day (in daily frequency) and to the previous month (i.e., previous critical date, in the monthly frequency).

When we perform separate individual tests of rationality, results in Table 11 are essentially the same as those in Tables 9 and 10 – we have strong evidence that individuals are rational when tested separately – we cannot

reject the null of rationality for 90% of the individuals at the 5% level. When we employ homogeneity restrictions ($\beta_i = \beta, \forall i$) in panel-data tests, the null of rationality is always rejected at the 5% and at the 1% level. However, when we test rationality using consensus forecasts, results do change vis-a-vis those in Tables 9 and 10. When we consider all days, rationality tests reject the null at the 5% and at the 1% levels. But, when we consider critical dates alone, tests using consensus forecasts support rationality at the usual significance levels. This is major difference vis-a-vis the results in the previous sections.

Counter-Factual Exercises and Survey Design

We now propose a discussion of survey design using the *incentive-driven inattention* model. It is based on counter-factual exercises for the mean-squared-forecast error (MSFE) of the consensus forecast, under the assumption that there is a final user of the consensus forecast operating under an MSFE *risk function*. The counter-factual analyses are aimed at assessing the effects of changing the incentive structure of the survey. We consider different scenarios (or survey designs) for the benefit-cost ratio and produce counter-factual consensus forecasts and their corresponding MSFEs. Each scenario is attached to a different time path for λ_t , and we consider how this change in the time path for λ_t will affect the updating behavior of agents in constructing the counter-factual scenario. The counter-factual consensus forecasts are generated according to the Coibion and Gorodnichenko (2012) updating formula for F_t , with a time-varying λ_t , where the initial forecast F_1 is the first consensus nowcast available (for each event) in the dataset¹⁸, and $\mathbb{E}_t(y_m)$ is the rational expectation forecast in (7).

¹⁸The first nowcast is available (on average, in our sample) at the 8th calendar day of the month, which corresponds to the day of release of the IPCA inflation rate of the previous month.

Survey Design (A): Change in the Single Critical Date for the Standard Month

Here we consider a simple change in the design of the survey. Currently, in the *Focus Survey*, the critical date for each month is in the mid-point of the working month, i.e., between two consecutive releases of monthly inflation data. From the point of view of a final user of the consensus forecast, a valid question is what is the *optimal* date to have the critical date in a standard month of 21 working days? In our model, incentives change the individual updating behavior, which, in turn, changes the individual daily MSFE and thus the daily profile MSFE of the consensus forecast.

The dynamics of inattention is given by $\lambda_t = \lambda_1 + \lambda_2 1 [t = t^{CD}]$. We consider the same benefit-cost ratio distribution as in the estimated model. Hence, λ_1 and λ_2 are the same as in the estimated ones, but the critical date can be different, i.e., it can be any date: $t^{CD} = 1, 2, \dots, 21$. This is compatible with our assumption of constant accuracy improvement during the month, which implies that the additional attention from shifting the benefit/cost distribution is also constant.

To perform the counter-factual analysis we construct a benchmark scenario. Since we want to ask what is the optimal single-date to be the critical date, we need a benchmark with no critical dates in the standard month. Hence, for the benchmark, we set $\lambda_2 = 0$ in the dynamics of λ_t , making $\lambda_t = \lambda_1$ for all t . Based on that, we construct $MSFE_t$ for the consensus forecast under the Coibion and Gorodnichenko (2012) updating formula for F_t .

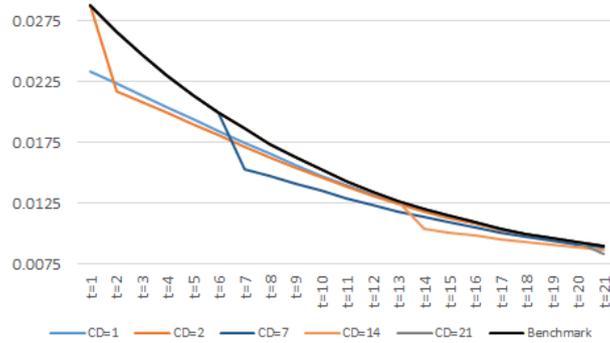
We change the benchmark case by constructing alternative scenarios in which there are a single critical date: $t^{CD} = 1, 2, \dots, 21$. We then compute the ratio of the MSFE for the benchmark case and for each of the 21 counter-factual critical dates, averaging these ratios across the days in a standard

month, i.e., we compute $\sum_{t=1}^{21} \frac{MSFE_t^{Benchmark}}{MSFE_t^{t^{CD}}}$.

Figure 7 and Table 6 present the results of the counter-factual analysis. Figure 7, Panel B, presents the $MSFE$ time-evolution for all possible designs based on a single-critical date. Table 6 presents the relative average $MSFE_t$ across the month for selected survey designs. The results show that the highest gain for the $MSFE_t$ profile is obtained when we choose the critical date for $t^{CD} = 1$, i.e., the day after the release of the inflation rate of the previous month. Having the critical date in the $t^{CD} = 1$ lowers the $MSFE_t$ profile for the whole month as a consequence of anticipating the benefits of being attentive since the first day of the new inflation cycle.

Figure 7 - $MSFE_t$ along a $t = 21$ -day-month of different survey designs

Panel A: Selected single-CD designs and the benchmark (no CD)



Panel B: Comparison of all single-CD designs ($t^{CD} = 1, \dots, 21$)

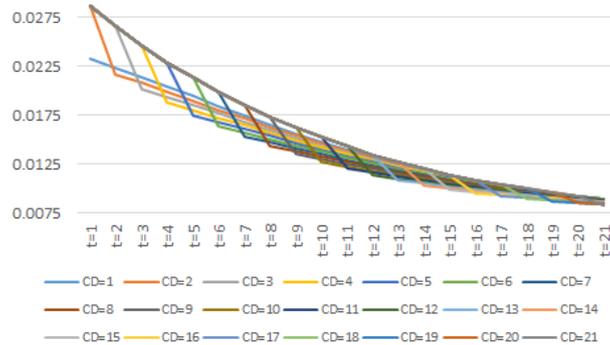


Table 6 - Relative average $MSFE_t$ for selected survey designsSingle CD: t^{CD}

Critical Date	$\sum_{t=1}^{21} \frac{MSFE_t^{Benchmark}}{MSFE_t^{t^{CD}}}$
$t^{CD} = 1$	1.074
$t^{CD} = 2$	1.070
$t^{CD} = 3$	1.066
$t^{CD} = 7$	1.050
$t^{CD} = 14$	1.022
$t^{CD} = 21$	1.002

Survey Design (B): Multiple Critical Dates (k) in the Month

The second survey-design exercise considers the same benchmark as before, i.e., we set $\lambda_2 = 0$ in the dynamics of λ_t , making $\lambda_t = \lambda_1$ for all t . We want to introduce multiple critical dates for the standard month – from the day after the release of the inflation rate of the previous month all the way to the k -th day, $k = 1, \dots, 21$, i.e., we consider multiple critical dates continuously for the first k days of the standard month.

In order to construct the multiple critical-date scenario, we first back out the mean shift in the benefit-cost distribution that corresponds to the increase in attention on the critical date, λ_2 , which equals $\mu_2 = 1 - \Phi^{-1}(1 - \lambda_2)$. We then consider spreading this additional benefit over the k days in the month. Because we are now updating more often, the costs also increase. We assume that the cost of updating k times is kC_i . Therefore, for the counter-factual of updating every day (i.e., $k = 21$) we use $\mu_2/21$.

For a fixed number k of critical dates and for each day t in the month

($t = 1, \dots, 21$) it follows that:

$$\lambda_t = \begin{cases} \lambda_1 + 1 - \Phi(1 - \mu_2/k); & \text{if } t = 1, \dots, k \\ \lambda_1 & ; \text{ if } t = k + 1, \dots, 21 \end{cases} . \quad (32)$$

We employ the estimates of λ_1 and λ_2 , we get $\hat{\mu}_2 = 1 - \Phi^{-1}(1 - \hat{\lambda}_2) = 0.632$. Note that the particular case in which $k = 21$ corresponds to the counter-factual where forecasters are continuously evaluated over the entire month. In this particular case, $\hat{\lambda}_t = \hat{\lambda}_1 + 1 - \Phi(1 - \hat{\mu}_2/21) = 0.285$ for all $t = 1, \dots, 21$. With these results we can then construct the the $MSFE_t$ profile for the whole month and compare all profiles. Figure 9 presents average relative $MSFE$ results for all designs. Table 7 presents relative (average) MSFEs for selected designs. Its last column shows that having multiple critical dates is beneficial in terms of its reduction of the $MSFE_t$ profile. Moreover, if every day is a critical day, i.e., $k = 21$, the MSFE gains are roughly 22% compared to the benchmark case, being the best choice within this comparison. This also shows that, if the BCB let every day to be a critical day it will lower the most the $MSFE_t$ profile.

It is worth mentioning that the results of the counter-factual exercise, especially those based on multiple critical dates, should be interpreted here with a bit of caution due to the Lucas critique (and the fact that agents' optimal behavior is not design-invariant and would necessarily change whenever the survey incentives change).

Nonetheless, these results point out to the importance of a proper discussion of survey design and evaluation mechanisms, from a microeconomic perspective, *vis-à-vis* the responses of survey participants (captured here by a time-varying degree of attention and respective individual forecast revisions).

Figure 9 - $MSFE$ for different survey designs: Multiple CDs $t^{CD} = 1, \dots, k$ and the Benchmark (no CD)

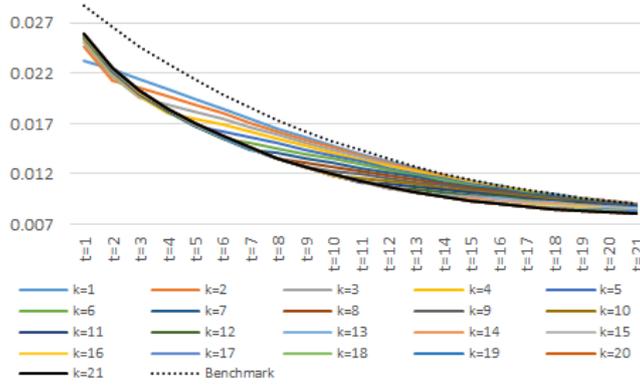


Table 7 - Relative (average) $MSFE_t$ for selected survey designs

Multiple CDs: $t^{CD} = 1, \dots, k$

k	$\sum_{t=1}^{21} \frac{MSFE_t^{t^{CD}=1, \dots, k}}{MSFE_t^{t^{CD}=1, \dots, 21}}$	$\sum_{t=1}^{21} \frac{MSFE_t^{Benchmark}}{MSFE_t^{t^{CD}=1, \dots, k}}$
1	1.135	1.074
2	1.121	1.088
3	1.105	1.103
7	1.054	1.156
14	1.009	1.208
21	1.000	1.219

6 Conclusions and Further Research

In this paper, we propose a novel approach to inattention, in which we endogenize it while retaining some important features of the previous literature. Agents respond to incentives to update their predictions based on new information. To decide how much new information should be used in predicting future outcomes, agents do cost-benefit analysis, with incentives playing an important role in defining the benefits of updating information. When they

decide to update their information set, they receive noisy-information. The final result of using these ingredients is what we have labelled the *incentive-driven inattention* model.

We test the model using the *Focus Survey* gathered by the Brazilian Central Bank (BCB). Its frequency of observations is daily (working days), as opposed to monthly (*Wall Street Journal Forecasting Survey*), quarterly (*Survey of Professional Forecasters*), or semi-annually (*Livingston Survey*). At each working day, a survey respondent can *choose* to participate inserting in the data base forecasts of a myriad of important macroeconomic and financial forecasts for different horizons. There is an incentive to participate in the so called *critical dates*, since only forecasts at critical dates are used to compute MAFE statistics, which are compared across agents, and are partially released by the BCB in the form as a list of the *top five* forecasters. Thus, the *Focus Survey* is well suited to study the updating behavior of individual agents: it has high-frequency micro-data at the individual level, where we can observe enough changes in costs and benefits for the individual updating behavior.

When we take the model to the data, we obtain sensible structural parameters that are able to explain the dynamic behavior of the consensus MSFE across dates, including the critical date. In misspecification tests, the structural model is not rejected using a variety of instruments in GMM estimation. Simulating the model also generates an MSFE profile that matches that of the data when we consider a point-wise 95% confidence interval around it. We also investigate the determinants of updating information using a panel of individual institutions. Our results show that updating is explained mostly by the dummy variable of the critical date, followed by economic uncertainty, measured by the Emerging Market Bond Index (EMBI). Calendar-effect dummies are also relevant, as well as the uncertainty of the individual forecast error forecasting inflation in the previous month.

Finally, equipped with the *incentive-driven inattention* model, we discuss two important issues in forecasting. First, using individual and consensus data, we perform a set of rationality tests pioneered by Mincer and Zarnowitz (1969) and its extensions in Coibion and Gorodnichenko (2014). The results show that, at the individual level, at least 90% of individuals are rational when we consider critical dates alone when they have proper incentives to update forecasts. In contrast, for all dates, including non-critical dates, at most 26% of individuals pass rationality tests. Similar tests applied to panel-data and the consensus forecast show that rejection of rationality is the rule, not the exception. Second, we discuss survey-design issues, performing counter-factual exercises with our model in order to assess the effects of changing the incentive structure of the survey, making every day a critical day has a big potential in lowering the MSFE profile across the standard month – a benefit to users of the consensus forecasts as a forecasting tool.

A Appendix: Reduced-Form Estimation of the *Incentive-Driven Inattention* Model:

λ_1 , λ_2 , σ_ε^2 and σ_η^2

In this section, we compare the $MSFE_t$ from actual data (Figure 1) with that of the *Incentive-Driven Inattention* model, where its parameters are estimated by a two-step approach based on two sequential OLS regressions. The first OLS estimation (called OLS1) is a regression of the observed series λ_t onto a constant (to capture λ_1) and a dummy variable for the critical dates ($1_{t=t^{CD}}$) to get an estimate of λ_2 , where ϵ_t is the residual:

$$OLS1 : \lambda_t = \lambda_1 + \lambda_2 1_{t=t^{CD}} + \epsilon_t. \quad (33)$$

Table A1 - OLS1 estimation

$\widehat{\lambda}_1$	$\widehat{\lambda}_2$
0.1192 (0.0017)	0.3566 (0.0094)

Notes: Sample: January 2nd, 2004 to January 8th, 2015 (2,751 workdays).

Newey-West standard errors in parentheses.

It is worth noting that a coefficient restriction test $H_0 : \lambda_1 = \lambda_2$, strongly rejects the null hypothesis, at usual significance levels (p-value = 0.000). We also reject the null $H_0 : \lambda_2 = 0$ with high confidence. Thus, the data clearly indicates that the fraction of updaters λ_t differs for critical and non-critical dates, as already expected from the dynamics of λ_t around the critical date depicted in Figure 1.

The second regression (called OLS2) stems from equation (12). The dependent variable is $MSFE_t - \left(\frac{1-\widehat{\lambda}_t}{\widehat{\lambda}_t}\right)^2 \mathbb{E}[(\Delta F_t)^2]$, where $\widehat{\lambda}_t = \widehat{\lambda}_1 + \widehat{\lambda}_2 1_{t=t^{CD}}$, $[\widehat{\lambda}_1; \widehat{\lambda}_2]$ are estimates from OLS1, and F_t is given by the Coibion and Gorodnichenko (2012) updating formula:¹⁹

$$OLS2 : MSFE_t - \left(\frac{1-\widehat{\lambda}_t}{\widehat{\lambda}_t}\right)^2 \mathbb{E}[(\Delta F_t)^2] = w_t \sigma_\varepsilon^2 + \sigma_\eta^2 + v_t, \quad (34)$$

where v_t is the residual. The regressors are an intercept (to get an estimate of σ_η^2) and the series w_t (to get the slope as an estimate of σ_ε^2), where from equation (8) it follows that $\mathbb{E}(v_t^2) = w_t \sigma_\varepsilon^2$, with

$w_t \equiv [1 + (1 + \phi)^2 + \dots + (1 + \phi + \dots + \phi^{21-t-1})^2]$, where w_t is computed

¹⁹Note that, from the Coibion and Gorodnichenko (2012) formula applied to our setup, F_t relies on the observed series λ_t and on the constructed series $E_t(y_m)$, generated from equation (7). The latter, in turn, depends on the Kalman-filtered daily inflation series y_t and on an initial estimate of ϕ (that is, $\widehat{\phi} = 0.9631$) obtained from the $ARMA(1,1)-AR(1)$ approach of Amemiya and Wu (1972). In addition, as an initial value for F_t , for every event (i.e. forecasted monthly inflation rate), the first Coibion and Gorodnichenko consensus forecast is set equal to the respective consensus forecast from data.

from an initial estimate of ϕ (i.e., $\widehat{\phi} = 0.9631$).²⁰

Table A2 - OLS2 estimation

$\widehat{\sigma}_\varepsilon^2$	$\widehat{\sigma}_\eta^2$
1.92e - 06 (2.55e-07)	0.0059 (0.0003)

Note: Sample: January 2nd, 2004 to January 8th, 2015.

Newey-West robust standard errors in parentheses.

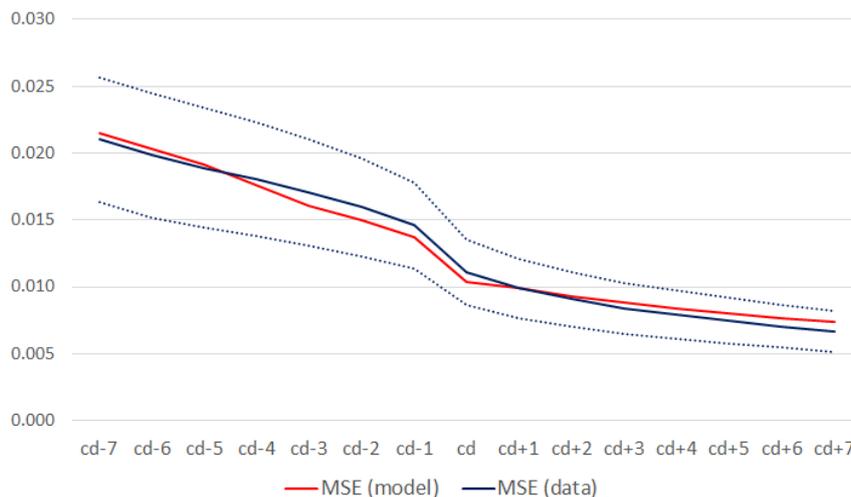
Based on the four estimated parameters from the two-step approach, λ_1 , λ_2 , σ_ε^2 and σ_η^2 , we can construct a reduced-form based $MSFE_t$ estimate in the following way:

$$\widehat{MSFE}_t = \left(\frac{1 - \widehat{\lambda}_1 - \widehat{\lambda}_2 \mathbf{1}_{t=t^{CD}}}{\widehat{\lambda}_1 + \widehat{\lambda}_2 \mathbf{1}_{t=t^{CD}}} \right)^2 E [(\Delta F_t)^2] + w_t \widehat{\sigma}_\varepsilon^2 + \widehat{\sigma}_\eta^2. \quad (35)$$

Figure A1 shows the comparison between the data and calibrated $MSFE_s$. Note that the model-based $MSFE$ is quite close to the data one and inside the 95% confidence interval all over the considered range around the critical date (cd).

²⁰In this calibration exercise, we could not estimate ϕ as a parameter in regression OLS2, because the w_t series depends also on the day t of the month under analysis and, since each month m in our dataset has (in fact) a different number of workdays k_m , also the number of terms in the sum $[1 + (1 + \phi)^2 + \dots + (1 + \phi + \dots + \phi^{k_m - t - 1})^2]$ varies from one month to another. This way, we first calibrate ϕ with an initial estimate $\widehat{\phi}_0$, compute the w_t series based on the actual number of workdays k_m for each month and, then, estimate the remaining parameters of the model.

Figure A1 - *MSFE* (data *versus* calibrated) and 95% confidence interval



Note: Confidence interval based on (pointwise) asymptotic values: $\text{std. error} = \text{sample std.dev.} / \sqrt{T}$, where $T = 132$ events (or months) for each investigated date $CD + j$; $j = -7, \dots, 7$.

It is worth noting that it looks like the fit of the reduced-form estimates for the MSFE is better than that obtained using structural-form estimates based on GMM; see Figure 5. This is not surprising, since, on the one hand, the estimates of in (35) are constructed to match the MSFE profile when equation OLS2 is estimated. The estimates of the structural form, on the other hand, are constructed to match the system with equations (19), (20), and (21) as a whole, and not a single equation for the MSFE. In addition, an advantage of using a structural-form approach is the opportunity to tests the misspecification of the model using a over-identifying-restriction test.