Rational Inattention during an RCT*

Bartosz Maćkowiak
European Central Bank and CEPR

Mirko Wiederholt
LMU Munich and CEPR

September 2024, first version: February 2024

Abstract

We study what rational inattention predicts about the treatment effect of information provided during a randomized control trial. This is a non-trivial question because the model implies optimal attention choice before, during, and after an information treatment. The model is consistent with the empirical negative relationship between the level of inflation and the strength of people's response to information about inflation. We describe the model's implications for cross-sectional variation in treatment effects across agents, persistence of treatment effects, experiments in macro- and microeconomics, and communication policy.

Keywords: rational inattention, information provision experiment, randomized control trial, inflation expectations (JEL: D8, E3, E7).

*Maćkowiak: ECB, 60640 Frankfurt am Main, Germany (e-mail: bartosz.a.mackowiak@gmail.com); Wiederholt: LMU Munich, Ludwigstrasse 28, 80539 Munich, Germany (e-mail: mirko.wiederholt@gmail.com). We thank Ryan Chahrour, Cosmin Ilut, Jennifer La'O, and participants in the 2024 Annual Meetings of the Society for Economic Dynamics and the Central Bank Research Association. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the ECB.

1 Introduction

Surveys with information provision experiments, a form of randomized control trials (RCTs), have become popular in economics.[1] Naturally, economists are interested in comparing the information treatment effects, on beliefs and on actions, measured in this literature with theories of behavior under incomplete information or bounded rationality, including rational inattention (RI) introduced by Sims (2003).[2] This comparison is complicated because RI implies optimal attention choice *before, during, and after* a survey. The goal of this paper is to make this comparison easier.

As a motivation, consider an information provision experiment about inflation. Suppose one would like to know whether to expect a larger or a smaller treatment effect when inflation is high. It is intuitive that as inflation rises, people find it more important to attend to it, and consequently they enter a survey with a sharper prior about inflation, which reduces the treatment effect. But it is likewise intuitive that when people find it more important to attend to inflation, they have more incentive to pay attention during the survey, which tends to increase the treatment effect.

This paper addresses three questions. First, what does the theory of RI predict about the treatment effect in an information provision experiment? Second, are these predictions consistent with findings in the empirical literature on information provision experiments? Third, given the theory, what can one infer from behavior during the survey about behavior outside the survey?

To address these questions, we introduce a survey into a plain-vanilla dynamic RI model. In the model, agents make optimal attention choices before, during, and after a survey. For concreteness, we assume that agents track inflation. Inflation follows a first-order autoregressive process. Agents decide optimally how much attention to devote to inflation. At some point, a subset of the population is unexpectedly called into the survey and is provided with information about inflation. Respondents know that after the survey it will still be useful to be informed about inflation. We derive the predictions for the treatment effect as a function of the parameters of the model.

We obtain four main results. "The importance of attending to inflation" is formalized in the model as "the utility or profit cost of a given mistake in tracking inflation." The first result is that the treatment effect is decreasing in the cost of a mistake and in the variance of innovations in

[1] See Fuster and Zafar (2023) and Haaland, Roth, and Wohlfart (2023) for reviews of the literature on information provision experiments.

[2] See Maćkowiak, Matějka, and Wiederholt (2023) for a review the literature on rational inattention.

inflation. It is plausible that in the real world a higher level of inflation is associated with a greater cost of a mistake in tracking inflation, a larger variance of the innovations in inflation, or both. Hence, the treatment effect is predicted to be smaller when inflation is high.

Next, consider cross-sectional variation in the cost of a given mistake. The second result is that the treatment effect is predicted to be smaller for agents who face a higher cost of a mistake. These two results arise because the first of the two counteracting forces from the motivating example (attention outside the survey) turns out to dominate the second (attention during the survey).

Third, the nature of the information provided during an experiment matters. If agents could have attended to the same information before the experiment (which may be a typical setting in macroeconomics), a strong response to the information during the experiment is a sign of a low cost of a mistake, as already pointed out. By contrast, if agents could not have attended to this or similar information before the experiment (which may be a typical setting in microeconomics), a strong response to the information during the experiment is a sign of a *high* cost of a mistake.

Fourth, the model predicts that the treatment effect may be transient, or, more specifically, the treatment effect will be less persistent than what one would expect based on a model with exogenous imperfect information. Rationally inattentive agents reduce their attention for some time after an information treatment. Consequently, after a typical experiment the treatment group will be absorbing less information for some time than the control group, which will pull their beliefs together.

Weber et al. (2024) conduct surveys with information provision experiments about inflation in different countries over time. They document a negative relationship between the level of inflation and the magnitude of information treatment effects: when inflation is high, households and firms respond less to provided information about inflation. Cavallo, Cruces, and Perez-Truglia (2017) report the same finding based on comparing information provision experiments for individuals from Argentina and the United States. The RI model is consistent with this finding.

Link et al. (2023) document that firms have more accurate expectations of central bank policy rates than households, consistent with policy rates being more important to firms than to households. Link et al. then find that firms update their policy rate expectations less than households when provided with an expert forecast, as predicted by the RI model.[3] Knotek et al. (2024) find

[3]Similarly, Armantier et al. (2016) find that those individuals who are less uncertain about inflation (which

that monetary policy communications have a much stronger effect on people who had not previously heard news about monetary policy and who take sufficient time to read the treatment. The interpretation of this finding based on the RI model is that these are individuals with a low utility cost of a mistake in tracking inflation. A lesson from the model, consistent with this evidence from macroeconomic settings, is that statistically and economically significant treatment effects can be expected most of all from agents for whom the variable of interest is relatively unimportant.

Regarding the persistence of the treatment effects, Haaland, Roth, and Wohlfart (2023), p. 27, write as a summary of the evidence that "people persistently learn from the information provided, but (...) effects in most cases become weaker over time." The RI model gives a reason for why the information treatment effect may dissipate fairly quickly.

Section 2 presents the model. The four main results described above are derived in Section 3 (the comparative statics of the treatment effect, the persistence of the treatment effect, and the role of the nature of information). Additional results are in Section 4. Section 5 provides a summary and implications for communication policy.

2 Model: RI before, during, and after the survey

We introduce an information provision experiment into a dynamic RI model. Agents maximize a quadratic payoff. The payoff-relevant variable, which we refer to as "inflation," follows a serially correlated process, for simplicity a Gaussian first-order autoregression:

$$\pi_t = \rho \pi_{t-1} + \varepsilon_t \tag{1}$$

where $\rho \in [0, 1)$ and $\varepsilon_t \sim i.i.d.N\left(0, \sigma_\varepsilon^2\right)$. There are three stages:

- stage 1: before the survey (periods $t = 0, -1, -2, ...$);
- stage 2: during the survey (end of period $t = 0$);
- stage 3: after the survey (periods $t = 1, 2, ...$).

One can think of a period as a quarter, a month, a week, or a day. The survey with an information provision experiment occurs at the end of period 0. At this point, some agents ("the treatment

indicates that inflation may be more important to them) respond less to the information treatment.

group") are unexpectedly called into the survey. Think of agents as being asked during the survey to report their conditional expectation of the current inflation rate, before and after an information treatment. The treatment consists of showing agents some information relevant for updating their belief. Those agents who are not called into the survey ("the control group") move directly from stage 1 to stage 3.[4] Our model of attention is rational inattention, as in Sims (2003), in all three stages. "Rational inattention" means that information is in principle available, but absorbing information requires paying attention which is costly; and the cost of attention is linear in entropy reduction.[5]

Stage 3: Agents solve the following RI problem at the beginning of period $t = 1$:

$$V_1\left(\Sigma_{1|0}\right) = \underset{\{\Sigma_{t|t}\}_{t=1}^{\infty}}{min} \left\{ \sum_{t=1}^{\infty} \beta^{t-1} \left(\gamma \Sigma_{t|t} + \mu ln \left(\frac{\Sigma_{t|t-1}}{\Sigma_{t|t}} \right) \right) \right\} \tag{2}$$

subject to

$$\Sigma_{t|t} \leq \Sigma_{t|t-1} \tag{3}$$

$$\Sigma_{t+1|t} = \rho^2 \Sigma_{t|t} + \sigma_{\varepsilon}^2 \tag{4}$$

$$\Sigma_{1|0} \ \textit{given.} \tag{5}$$

$\Sigma_{t|t-1}$ is the prior variance of inflation in period t, and $\Sigma_{t|t}$ is the posterior variance of inflation in period t (i.e., the conditional variance of inflation after paying attention to inflation). Agents enter period $t = 1$ with some prior variance of inflation $\Sigma_{1|0}$ and choose the posterior variance $\Sigma_{t|t}$ for each period $t \geq 1$. The prior variance of inflation in a future period is given by equation (4) and the posterior variance cannot exceed the prior variance (inequality (3)). The first term in the per-period loss function in (2) is the loss from being uninformed about inflation. The second term is the cost of attention. $V_1\left(\Sigma_{1|0}\right)$ is the value function in period $t = 1$. One could add a subscript i for an individual agent i in the parameters γ, μ, β as well as in the prior and posterior variances.

[4] In a real-world survey, the control group are typically agents who are called into the survey but do not receive the information treatment during the survey. In the model, such agents move directly from stage 1 to stage 3.

[5] Regarding inattention during the survey, Haaland, Roth, and Wohlfart (2023), p. 25, write: "One concern in online surveys is that respondents are inattentive and speed through the surveys (...)." Knotek et al. (2024), p. 3, write: "(...) while our RCTs explicitly communicate information to the treatment groups, it is hard to know how much of this information is really read and/or absorbed by respondents." Knotek et al. measure attention during the survey, which they refer to as compliance, through the time that respondents choose to spend reading the information treatment in the online survey.

We suppress the subscript for ease of exposition. Likewise for ease of exposition, we express the per-period loss directly in terms of the conditional variance of inflation. One can think of the agent as ultimately choosing an action where the period t optimal action depends on the period t inflation rate (e.g., the optimal action is proportional to the inflation rate). The parameter γ reflects both the cost of a given mistake in the action and the sensitivity of the optimal action to inflation. With this understanding, we refer to γ simply as the cost of a mistake in tracking inflation.

Stage 1: Agents solve the same problem as in stage 3 because they are in the same environment as in stage 3; hence: (i) the same γ, μ, and β as in stage 3, (ii) the same inflation process (the same ρ and σ_ε^2 as in stage 3), and (iii) the horizon is infinite because agents in stage 1 did not expect to be called into the survey. Next, assuming that agents have been in stage 1 for a sufficiently long time, the effect of the initial prior variance of inflation has disappeared and in stage 1 one can focus on the steady-state information structure, denoted $\bar{\Sigma}_{t+1|t}$ and $\bar{\Sigma}_{t|t}$. Computationally, this steady-state information structure is a by-product of solving the problem in stage 3. Intuitively, in stage 1 agents have converged to the same information structure that they will converge to in stage 3, because of the symmetry of stages 1 and 3.

Stage 2: Agents who are called into the survey and receive the information treatment solve the following problem at the end of period $t = 0$ (i.e., during the survey):

$$\underset{\Sigma_{0|0,s}}{min} \left\{ \left(\gamma \Sigma_{0|0,s} + \lambda ln \left(\frac{\Sigma_{0|0}}{\Sigma_{0|0,s}} \right) \right) + \beta V_1 \left(\Sigma_{1|0} \right) \right\} \tag{6}$$

subject to

$$\Sigma_{1|0} = \rho^2 \Sigma_{0|0,s} + \sigma_\varepsilon^2 \tag{7}$$

$$\Sigma_{0|0} = \bar{\Sigma}_{t|t}\,. \tag{8}$$

$\Sigma_{0|0}$ is the prior variance of inflation at the beginning of the survey. $\Sigma_{0|0,s}$ is the posterior variance of inflation after having paid attention during the survey. Equation (7) maps the posterior variance after the survey into the prior variance in period 1, the first period after the survey. Equation (8) states that the prior variance at the beginning of the survey equals the steady-state posterior variance (from stage 1). The term $\beta V_1 \left(\Sigma_{1|0} \right)$ captures the continuation value of information processed during the survey. The marginal cost of attention during the survey, λ, may differ from the marginal cost of attention outside the survey, μ. We assume that $\lambda \leq \mu$. It is plausible that the marginal cost of attention is smaller during the survey than outside of it (i.e., the opportunity cost

of paying attention to the variable of interest declines conditional on having accepted to participate in the survey about this variable).

For an agent who receives the information treatment during the survey, the posterior (after-the-survey) mean of period-0 inflation, $\pi_{0|0,s}$, is a convex combination of the prior (before-the-survey) mean, $\pi_{0|0}$, and the signal absorbed from the treatment, S:

$$\pi_{0|0,s} = \pi_{0|0} + G\left(S - \pi_{0|0}\right). \tag{9}$$

During the survey, an agent updates their belief as if they observed a realization of $S = \pi_0 + \eta$, where $\eta \sim N\left(0, \sigma_\eta^2\right)$ with σ_η^2 chosen optimally (stage-2 problem). This is a simple way to model attention choice during the survey; furthermore, a signal of this form is optimal under RI.[6] Think of the treatment as showing the agent the current inflation rate. The agent is "speeding through the survey" (footnote 5) with the speed that reflects their environment and preferences, and therefore the agent absorbs a noisy signal where the variance of the noise is endogenous.[7]

From Bayesian updating, $G = \Sigma_{0|0} / \left(\Sigma_{0|0} + \sigma_\eta^2\right) = 1 - \Sigma_{0|0,s} / \Sigma_{0|0}$. We refer to the updating coefficient $G \in [0, 1]$ as the treatment effect. Note that the treatment effect G is increasing in *the ratio* of the steady-state posterior variance, $\Sigma_{0|0} = \bar{\Sigma}_{t|t}$, to the after-the-survey posterior variance, $\Sigma_{0|0,s}$, both of which are endogenous in the RI model. Think of the incentives to pay attention outside the survey and during the survey, respectively.

Link to the empirical literature. The treatment effect, also referred to as the learning rate, is a key object of interest in the literature on information provision experiments. As a recent example, Weber et al. (2024) use the following regression specification to estimate the treatment effect (the subscript i is for individuals):

$$Posterior_i = \beta_0 + \beta_1 Prior_i + \beta_2 I_i + \beta_3 I_i Prior_i + \epsilon_i \tag{10}$$

[6] In the Gaussian AR(1) case with a quadratic objective, the optimal signal under RI has the form "current inflation plus Gaussian noise" (e.g., Maćkowiak, Matějka, and Wiederholt, 2018, Miao, Wu, and Young, 2022, and Jurado, 2023). It follows that outside the survey an agent updates their belief as if in each period t they observed a realization of $\pi_t + \psi_t$, where ψ_t is distributed normally and independently over time with mean zero and variance $\sigma_{\psi_t}^2$ which is chosen optimally (stage-1 and stage-3 problems).

[7] One can also think of the treatment as showing the agent some non-quantitative information (e.g., a news report) from which the agent extracts the signal S. In Section 4, we consider the case where the treatment consists of showing a forecast of a future inflation rate and the case where the treatment consists of showing a noisy measure of the current inflation rate (with noise that is beyond the agent's control).

where $Posterior_i$ and $Prior_i$ are the respondent's posterior and prior about the variable of interest, and I_i is the indicator variable for being in the treatment group as opposed to being in the control group. Comparing equations (9) and (10), the model predicts $\beta_3 = -G$ (also, $\beta_1 = 1$ and $\beta_2 = GS$). Weber et al. estimate the treatment effect as $\hat{\beta}_3/\hat{\beta}_1$, allowing for $\beta_1 \neq 1$ since priors and posteriors are measured using different questions in their data (they also report $\hat{\beta}_3$ separately). Hence, by studying G we focus on the key coefficient of interest in the empirical literature.[8]

3 Main results

If the constraint (3) is non-binding in all periods, the model can be solved analytically. The economic meaning of "the constraint (3) is non-binding in all periods" is that agents pay some strictly positive amount of attention, however small, to inflation in each period, which seems reasonable in a quarterly model. In the RI literature, the constraint (3) is called the "no-forgetting" constraint because it states that posterior uncertainty cannot exceed prior uncertainty.

Analytical solution of the model. Let's start from stage 3. Substituting constraints (4)-(5) into objective (2), the problem reads

$$V_1\left(\Sigma_{1|0}\right) = \underset{\{\Sigma_{t|t}\}_{t=1}^{\infty}}{min} \left\{ \gamma\Sigma_{1|1} + \mu ln\left(\frac{\Sigma_{1|0}}{\Sigma_{1|1}}\right) + \sum_{s=2}^{\infty} \beta^{s-1} \left[\gamma\Sigma_{s|s} + \mu ln\left(\frac{\rho^2\Sigma_{s-1|s-1} + \sigma_\varepsilon^2}{\Sigma_{s|s}}\right)\right]\right\} \tag{11}$$

subject to the no-forgetting constraint

$$\forall t = 1, 2, \ldots : \Sigma_{t|t} \leq \Sigma_{t|t-1}\,. \tag{12}$$

If the no-forgetting constraint (12) is non-binding for all $t \geq 1$, the first-order condition for any $\Sigma_{t|t}$ with $t \geq 1$ reads

$$\gamma - \mu\frac{1}{\Sigma_{t|t}} + \beta\mu\rho^2\frac{1}{\rho^2\Sigma_{t|t} + \sigma_\varepsilon^2} = 0. \tag{13}$$

Note that $\Sigma_{t|t}$ is the same for all $t \geq 1$: If the no-forgetting constraint is non-binding for all $t \geq 1$, convergence to the steady-state information structure is *immediate* in period $t = 1$.[9] From equation

[8] Haaland, Roth, and Wohlfart (2023) include in their survey article a table with the estimated G's from the literature (their Table 1). Note that in some experiments priors are unobserved and one cannot estimate the treatment effect G (one can see how ex post beliefs differ between the control group and the treatment group).

[9] Afrouzi and Yang (2021), Proof of Proposition 3.1 in the Appendix, and Miao, Wu, and Young (2022), Proof of

(13), choosing a marginally lower posterior variance has three effects: It reduces the current loss from inattention (the first term), it increases the current attention cost (the second term), and it reduces the future attention cost (the third term). The third term captures that, in a dynamic environment, current attention reduces future prior uncertainty thereby saving future attention. The third term drops out if inflation follows a white noise process, $\rho = 0$ (or the discount factor β equals zero).

Multiplying equation (13) by $\Sigma_{t|t}\left(\rho^2\Sigma_{t|t} + \sigma_\varepsilon^2\right)$ yields a quadratic equation with the unique positive solution given by

$$\Sigma_{t|t} = \frac{-\left(\gamma\sigma_\varepsilon^2 - \mu\rho^2 + \beta\mu\rho^2\right) + \sqrt{\left(\gamma\sigma_\varepsilon^2 - \mu\rho^2 + \beta\mu\rho^2\right)^2 + 4\gamma\rho^2\mu\sigma_\varepsilon^2}}{2\gamma\rho^2}. \tag{14}$$

Hence, stage 3 (and thereby stage 1) is solved analytically.

Next, let us turn to stage 2. The no-forgetting constraint is never binding during the survey (because $\Sigma_{0|0}$ is given by the steady-state information structure and $\lambda \leq \mu$) and the first-order condition for $\Sigma_{0|0,s}$ reads

$$\gamma - \lambda\frac{1}{\Sigma_{0|0,s}} + \beta\mu\rho^2\frac{1}{\rho^2\Sigma_{0|0,s} + \sigma_\varepsilon^2} = 0. \tag{15}$$

Compare the second term in (15) with the second term in (13). If the marginal cost of attention is strictly lower during the survey than outside of it ($\lambda < \mu$), then it is optimal to pay some attention during the survey and reduce the posterior variance of inflation further, $\Sigma_{0|0,s} < \Sigma_{0|0}$ and $G > 0$. In addition, if inflation follows an autocorrelated process, $\rho > 0$ (and $\beta > 0$), paying some attention during the survey further reduces the prior variance of inflation in the period after the survey (the third term in equation (15)).

Multiplying equation (15) by $\Sigma_{0|0,s}\left(\rho^2\Sigma_{0|0,s} + \sigma_\varepsilon^2\right)$ yields a quadratic equation with the unique positive solution given by

$$\Sigma_{0|0,s} = \frac{-\left(\gamma\sigma_\varepsilon^2 - \lambda\rho^2 + \beta\mu\rho^2\right) + \sqrt{\left(\gamma\sigma_\varepsilon^2 - \lambda\rho^2 + \beta\mu\rho^2\right)^2 + 4\gamma\rho^2\lambda\sigma_\varepsilon^2}}{2\gamma\rho^2}. \tag{16}$$

It follows that the treatment effect $G = 1 - \Sigma_{0|0,s}/\Sigma_{0|0}$ is determined analytically, with $\Sigma_{0|0}$ given by equation (14) and $\Sigma_{0|0,s}$ given by equation (16).

Proposition 7 in the Supplement, also note that convergence to the steady-state information structure is immediate in the AR(1) case with the non-binding no-forgetting constraint.

Comparative statics. We now focus on the comparative statics of G with respect to γ and with respect to σ_ε^2. It is plausible that in the real world a higher level of inflation is associated with a greater cost of a given mistake in tracking inflation, a larger variance of the innovations in inflation, or both. In a time-series context, one can think of a change in γ or σ_ε^2 or both over time.[10] The steady-state assumption in stage 1 implies that the parameter change occurred sufficiently in advance of the survey for agents to have converged to the steady-state information structure in advance of the survey. As we already know, "sufficiently in advance" means simply "at least *one period* in advance." In a cross-sectional context, one can think that some agents have a higher γ than others (there is a distribution of γ across agents) or one can think that there is a distribution of γ and σ_ε^2 across countries.

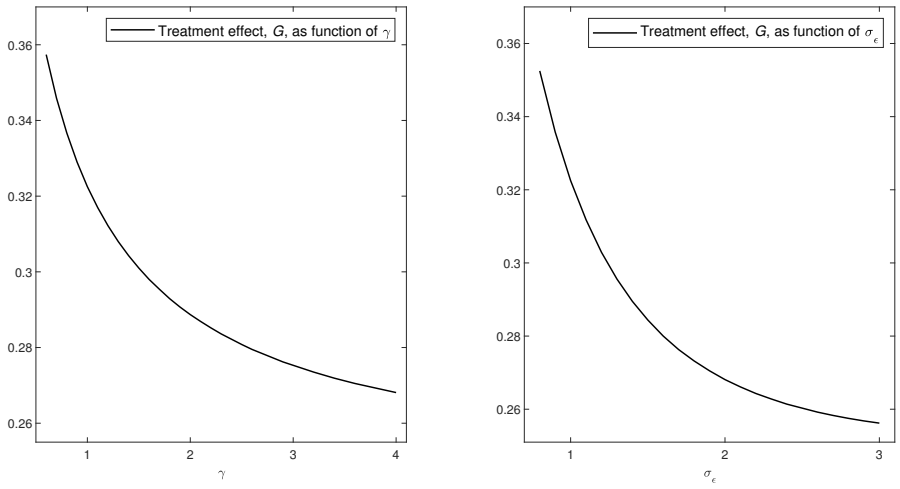
It is helpful to start with the case when inflation follows a white noise process, $\rho = 0$ (or $\beta = 0$). Then current attention has value only in the current period (equations (13) and (15)). The posterior variance is inversely proportional to γ and independent of σ_ε^2, $\Sigma_{0|0} = \mu/\gamma$ and $\Sigma_{0|0,s} = \lambda/\gamma$, which implies that the treatment effect depends only on the ratio of the marginal attention costs, $G = 1 - \lambda/\mu$. The treatment effect is independent of γ and σ_ε^2 in the static (or independent-over-time) version of the model. The two counteracting forces described in the introduction balance out exactly.

The treatment effect does vary with γ and σ_ε^2, however, in the empirically relevant case when the inflation process has some persistence, $\rho > 0$ (and $\beta > 0$). Then current attention has a continuation value. The third term in both first-order conditions (13) and (15), which reflects the continuation value, implies that: (i) each posterior variance is no longer inversely proportional to γ; and (ii) each posterior variance is increasing in σ_ε^2. The third term also implies that G *declines monotonically* with γ and with σ_ε^2. Figure 1 illustrates this result, starting from a baseline parameterization with $\rho = 0.9$, $\sigma_\varepsilon^2 = 1$, $\gamma = 1$, $\beta = 0.99$, $\mu = 2/ln(2)$, and $\lambda = 0.75\mu$, which yields a treatment effect of $G = 0.32$. When γ rises, (i) the steady-state posterior variance $\Sigma_{0|0}$ and the after-survey posterior variance $\Sigma_{0|0,s}$ fall; (ii) agents pay less attention during the survey so that *the ratio* $\Sigma_{0|0}/\Sigma_{0|0,s}$ and the updating coefficient G decline (left panel of Figure 1). When σ_ε^2

[10] Weber et al. (2024) report a flat or weakly rising agents' uncertainty about inflation in the COVID high-inflation period, which as the authors point out is not what one expects after an increase in γ. However, a flat or weakly rising uncertainty is consistent with a combination of an increase in γ and an increase in σ_ε^2.

increases, (i) the steady-state posterior variance $\Sigma_{0|0}$ and the after-survey posterior variance $\Sigma_{0|0,s}$ rise; (ii) agents pay less attention during the survey so that *the ratio* $\Sigma_{0|0}/\Sigma_{0|0,s}$ and the updating coefficient G decline (right panel of Figure 1).

Figure 1: **Treatment effect in the rational inattention model**



The monotonicity with respect to γ and σ_ε^2 arises because of the continuation value of attention (the third term in first-order conditions (13) and (15)). This term is decreasing and convex in the respective posterior variance. In a high-γ environment, the extra updating during the survey happens in a range where the posterior variance is already small; in that range, a small decrease in the posterior variance yields a large increase in the continuation value. This gives intuition for why the ratio $\Sigma_{0|0}/\Sigma_{0|0,s}$ falls (why $\Sigma_{0|0,s}$ declines by less than $\Sigma_{0|0}$). In a high-σ_ε^2 environment, the extra updating during the survey happens in a range where the posterior variance is large; in that range, a large increase in the posterior variance yields a small decrease in the continuation value. This gives intuition for why the ratio $\Sigma_{0|0}/\Sigma_{0|0,s}$ again falls (why $\Sigma_{0|0,s}$ rises by more than $\Sigma_{0|0}$).[11]

We also find that the treatment effect G is increasing in ρ (holding constant the unconditional variance of inflation) and β, because agents pay more attention during the survey when inflation is more persistent or when agents are more patient. The treatment effect is decreasing in the ratio of the marginal attention costs λ/μ (G falls from 1 toward 0 when λ rises from near 0 toward μ),

[11] We find that as γ or σ_ε^2 keeps rising, G tends to $1-\lambda/\mu$ which is the value of G in the static version of the model, 0.25 in this parameterization.

and the treatment effect is increasing in μ for given λ/μ (so long as $\rho, \beta > 0$).[12] The comparative statics may be helpful in time-series and cross-sectional contexts. The key prediction is that the treatment effect G will be lower in a high-γ or high-σ_ε^2 situation.

Alternative models. The result "the treatment effect declines monotonically with γ and σ_ε^2" does not naturally arise in other models of costly information acquisition or in models with exogenous imperfect information. As an example, the same result does not even arise in the model "rational inattention only outside the survey" where the survey itself is modeled without attention choice. Suppose agents solve the same RI problem in stage 1, and during the survey in stage 2 each agent updates their prior based on a realization of the signal "current inflation rate plus exogenous noise," with noise η drawn from a normal distribution with mean zero and variance σ_η^2. In this model, the treatment effect is equal to $G = \Sigma_{0|0} / \left(\Sigma_{0|0} + \sigma_\eta^2\right)$ with an exogenous σ_η^2. Since $\Sigma_{0|0}$ is decreasing in γ and increasing in σ_ε^2, in this model G is decreasing in γ and G is increasing in σ_ε^2. The difference between the two models is that in the RI model σ_η^2 is endogenous and agents reduce their attention during the survey when γ or σ_ε^2 increases.

As another example, consider an exogenous-imperfect-information model in which: (i) outside the survey, in every period each agent updates their prior based on a realization of the signal "current inflation rate plus exogenous noise," with noise ψ drawn from a normal distribution with mean zero and variance σ_ψ^2, (ii) the survey is modeled as in the previous paragraph. Here the noise outside the survey and the noise during the survey are exogenous. In this model, the treatment effect is invariant to γ and increasing in σ_ε^2.

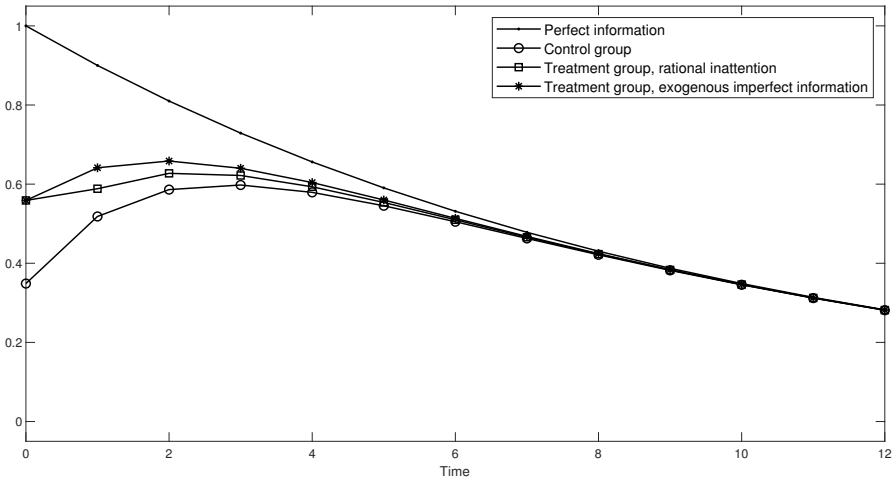
Persistence of the treatment effect. In the RI model, the treatment effect of an information provision experiment is less persistent than in an exogenous-imperfect-information model like the one from the previous paragraph. The reason is that rationally inattentive agents reduce their attention for some time after an information treatment.

To illustrate this result, Figure 2 compares the impulse response of the conditional expectation of inflation, $\pi_{t|t}$, $t \geq 0$, to an innovation in inflation in period $t = 0$ for three groups of agents: (i) agents from the RI model who did not participate in the survey ("Control group"), (ii) agents from the RI model who participated in the survey ("Treatment group, rational inattention"), and (iii) agents

[12] That an increase in γ and a decrease in μ have equivalent effects on optimal attention outside the survey is apparent from first-order condition (13).

from the exogenous-imperfect-information model who participated in the survey ("Treatment group, exogenous imperfect information"). The figure also reports the impulse response for hypothetical agents with perfect information. The parameter values are the same as before with $\sigma_\eta^2 = 2.8954$ and $\sigma_\psi^2 = 3.9529$ in the exogenous-imperfect-information model. With this parameterization, both treatment groups absorb the same amount of information during the survey. Furthermore, the treatment group from the exogenous-imperfect-information model absorbs the same amount of information in each period after the survey as does the control group.[13]

Figure 2: **Impulse response of $\pi_{t|t}$ to an innovation in π_0**



The key insight is that in period $t = 1$, the period after the survey, the treatment group from the RI model absorbs less information than the treatment group from the exogenous-imperfect-information model. The extra information from the survey pushes the conditional variance below its steady-state, making it optimal to reduce attention in order to let the conditional variance rise back to the steady-state. As a result, the Kalman gain is lower in period 1 than in the steady-state. In Figure 2, both treatment groups start out together in period 0, closer to the perfect-information benchmark than the control group (the initial treatment effect is the same for both treatment groups). In period 1, the treated agents from the exogenous-imperfect-information model are

[13] Agents in the RI model choose $\sigma_\eta^2 = 2.8954$ during the survey and $\sigma_\psi^2 = 3.9529$ in the steady state; in the steady state, agents update their belief as if in each period t they observed a realization of $\pi_t + \psi$, where $\psi \sim i.i.d.N\left(0, \sigma_\psi^2\right)$ and $\sigma_\psi^2 = 3.9529$.

about 75 percent further from the control group than the treated agents from the RI model, which is a non-trivial difference in the persistence of the treatment effect.

Nature of the provided information. From an RI perspective, the interpretation of an information provision experiment depends on the nature of the provided information. If agents could have attended to the same information before the experiment (which may be a typical setting in macroeconomics), a strong response to the information treatment is a sign of a low γ. If agents could not have attended to this or similar information before the experiment (which may be a typical setting in microeconomics), a strong response to the information treatment is a sign of a high γ.

Let us expand on the latter case. Either the relevant information was not publicly available, or the opportunity cost of paying attention to the relevant information outside the survey was prohibitively high (a large μ, so that agents pay no attention in the steady state of the model). Either way, the prior variance at the beginning of the survey equals the unconditional variance, $\Sigma_{0|0} = \sigma_\varepsilon^2 / \left(1 - \rho^2\right)$, and the model predicts that attention during the survey and the treatment effect G increase with γ (and with σ_ε^2).

4 Additional results

The case of a binding no-forgetting constraint. So far we have studied the case where agents pay strictly positive attention in equilibrium in the period after the survey. However, an agent may exit the survey so well informed about inflation that the individual chooses to pay no attention to inflation in the period after the survey, or even in multiple periods. This case seems unlikely in the current circumstances, but it is a theoretical possibility.

In the period after the survey, an agent pays either strictly positive attention or no attention. Recall that $\bar{\Sigma}_{t|t}$ is the steady-state posterior uncertainty (equation (14)). When $\Sigma_{1|0} > \bar{\Sigma}_{t|t}$, the prior uncertainty in period 1 exceeds the steady-state posterior uncertainty and the agent pays strictly positive attention in period 1. This is the case studied in Section 3. When $\Sigma_{1|0} \leq \bar{\Sigma}_{t|t}$, the agent pays no attention in period 1. The agent then pays no attention until the prior uncertainty in some future period exceeds $\bar{\Sigma}_{t|t}$. To solve for equilibrium in the case of a binding no-forgetting constraint, we express the value function in stage 3 analytically. This expression is given in the

online appendix.

Figure 1 showed the treatment effect G as a function of γ and σ_ε^2 when the no-forgetting constraint was non-binding in all periods in equilibrium. The lowest value of γ was 0.6 (Figure 1, left panel) and the lowest value of σ_ε was 0.8 (Figure 1, right panel). Figure 3 shows the equilibrium for the same parameterization except that the lowest value of γ is 0.3 (Figure 3, left panel) and the lowest value of σ_ε is 0.6 (Figure 3, right panel). The steady-state attention is increasing in γ and σ_ε^2. As one lowers γ or σ_ε^2, the steady-state attention falls and at some point the agent chooses to pay no attention in period 1: the no-forgetting constraint becomes binding in period 1 in equilibrium. It turns out that, in the range of γ where the no-forgetting constraint is binding, G is an increasing function of γ: when γ rises in that range, the agent pays more attention during the survey so that the ratio $\Sigma_{0|0}/\Sigma_{0|0,s}$ and G increase (left panel of Figure 3). Similarly, in the range of σ_ε^2 where the no-forgetting constraint is binding, G is an increasing function of σ_ε^2 (right panel of Figure 3). The peak in G in each panel in Figure 3 corresponds to the parameter value such that, to the left, the agent is at the corner solution for attention in period 1, and, to the right, the no-forgetting constraint is non-binding in all periods.

We emphasized before that the economic meaning of "the no-forgetting constraint is non-binding in all periods" is that agents pay some strictly positive amount of attention, however small, to inflation in every period, which seems sensible in a quarterly model, and very likely for most individuals in the COVID high-inflation period. The model then predicts that the treatment effect will decline monotonically with γ and σ_ε^2.[14]

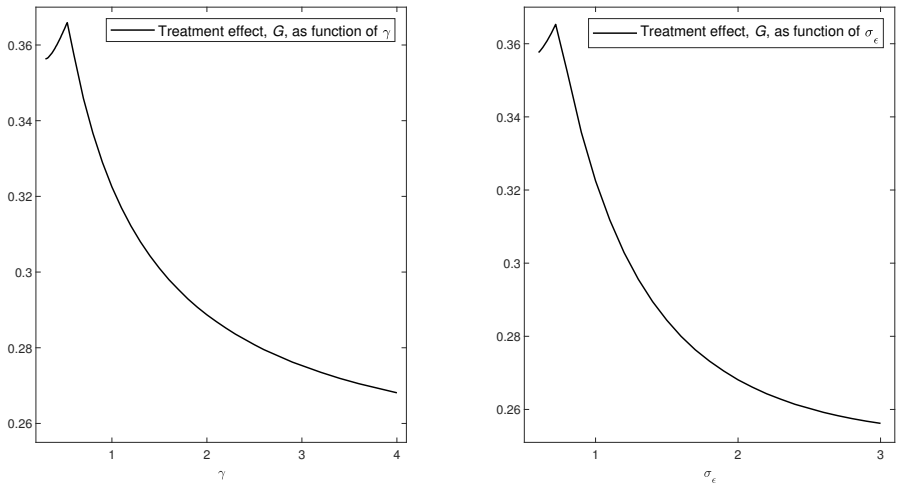
Forecast treatment. During the survey agents may be shown a forecast of inflation, and they may be asked to report their forecast of inflation. One may wonder if in the model it makes a difference if in stage 2 agents are constrained to learn from a forecast of inflation. In this case, the signal absorbed during the survey is a realization of $S = \rho\pi_0 + \eta$, with $\rho > 0$ and σ_η^2 chosen in the same stage-2 problem. Note that $\rho\pi_0$ is the best forecast of next period's inflation at the time of the survey. The updating equation reads:

$$\pi_{1|0,s} = \pi_{1|0} + \frac{\Sigma_{0|0}}{\Sigma_{0|0} + \sigma_\eta^2/\rho^2}\left(S - \pi_{1|0}\right)$$

where $\pi_{1|0,s}$ is the posterior (after-the-survey) mean of period-1 inflation, and $\pi_{1|0}$ is the prior

[14] As one lowers γ or σ_ε^2 further than shown in Figure 3, the no-forgetting constraint becomes binding in period 2 in addition to period 1. In that range for each parameter, G is also an increasing function of the parameter.

Figure 3: **Treatment effect in the rational inattention model**



(before-the-survey) mean. A comparison with equation (9) reveals that in the denominator σ_η^2/ρ^2 has replaced σ_η^2. The treatment effect, $\Sigma_{0|0}/\left(\Sigma_{0|0}+\sigma_\eta^2/\rho^2\right)$, is smaller for a given σ_η^2 than in the baseline model. However, a rationally inattentive agent chooses a lower σ_η^2 here than in the baseline model. In fact, the agent chooses a σ_η^2 such that σ_η^2/ρ^2 is equal to σ_η^2 in the baseline. The reason is that during the survey the agent wants to pay the same amount of attention to the variable they care about, which is the current inflation rate, as in the baseline. As a result, in equilibrium the treatment effect is equal to the treatment effect in the baseline. It does not matter if the provided information is the current realization (π_0) or the forecast ($\rho\pi_0$).

Outside the AR(1) case, we no longer expect the current-realization and the best-forecast treatments to yield exactly the same learning rate. Being provided with the current realization or the best forecast are no longer equivalent options for a rationally inattentive agent, and both options are suboptimal (the optimal signal is on the state vector, see Maćkowiak, Matějka, and Wiederholt, 2018). As an example, we solved numerically the model in the case when inflation follows an AR(2) process, $\pi_t = \rho_1\pi_{t-1} + \rho_2\pi_{t-2} + \varepsilon_t$, assuming $\rho_1 = 1.2$, $\rho_2 = -0.3$ with the rest of the baseline parameterization unchanged.[15] We obtained very similar treatment effects: 0.288 for the current-realization treatment, and 0.293 for the best-forecast treatment. Coibion, Gorodnichenko, and

[15] We used the code of Afrouzi and Yang (2021) because the stage-3 problem is a special case of the decision problem that they study.

Weber (2022) elicit households' inflation expectations and provide different pieces of information regarding inflation during the survey. They find very similar treatment effects when the provided information is the most recent inflation rate or the FOMC's inflation forecast (or the Federal Reserve's inflation target). The inflation-forecast and the inflation-target treatments appear to result in slightly more persistent belief revisions.

"The rate of change in my cost-of-living is different from the official inflation rate." Suppose the agent cares about an inflation rate that is different from the official inflation rate, because the agent's consumption basket differs from the consumption basket underlying the official inflation rate or because there is measurement error in the official inflation rate. Let us now interpret π_t as the inflation rate that the agent cares about. Suppose that the official inflation rate is given by $\pi_t + z_t$, where z_t follows a Gaussian AR(1) process independent of π_t. During the survey agents are constrained to learn from the current official inflation rate and decide how much attention to pay to that variable. The signal absorbed during the survey is a realization of $S = \pi_0 + z_0 + \eta$, with σ_η^2 chosen in the same stage-2 problem. The updating equation reads:

$$\pi_{0|0,s} = \pi_{0|0} + \frac{\Sigma_{0|0}}{\Sigma_{0|0} + \Sigma_z + \sigma_\eta^2} \left(S - \pi_{0|0}\right)$$

where Σ_z is the unconditional variance of z_t.[16] The treatment effect, $\Sigma_{0|0} / \left(\Sigma_{0|0} + \Sigma_z + \sigma_\eta^2\right)$, is smaller than in the baseline model because the signal during the survey is suboptimal for the agent, implying less attention during the survey compared with the baseline. We solve this version of the model numerically making use of the code of Afrouzi and Yang (2021). We find qualitatively the same comparative statics of the treatment effect with respect to γ and σ_ε^2. It does not matter for the comparative statics if the agent cares about an inflation rate that is different from the official inflation rate.

[16] In principle, it is the prior variance of z_t at the beginning of the survey, equal to the steady-state posterior variance from stage 1, but the agent had no incentive to pay attention to z_t in stage 1. Furthermore, the updating equation for the official inflation rate (as opposed to the inflation rate the agent cares about) is very similar; the numerator reads $\Sigma_{0|0} + \Sigma_z$ instead of $\Sigma_{0|0}$.

5 Conclusions

Rational inattention helps interpret information provision experiments and provides an explanation for existing evidence. We would like to argue that RI is a useful starting point for understanding the evidence, not that it will explain all of the evidence. The main insights from this paper can be summarized as follows:

- A standard RI model predicts that the information treatment effect will decline with γ and σ_ε^2, if the variable of interest is autocorrelated and agents expect to pay positive attention after the survey.
- If agents could have attended to the same information before the experiment, a strong response to the treatment is a sign of a low γ. If agents could not have attended to this or similar information before the experiment, a strong response to the treatment is a sign of a high γ.
- The treatment effect will be less persistent than what one expects based on a model with exogenous imperfect information.

Viewed through the lens of the RI model, the RCT evidence indicates that people pay more attention to inflation when inflation is high, which in the model implies that actions will then respond faster to inflationary shocks (Maćkowiak and Wiederholt, 2009). Thus, the Phillips-curve trade-off faced by policymakers changes with the level of inflation.[17] The model also has implications for communication policy. It suggests that communication about inflation may have less effect in a high-inflation environment than usual, so long as communication involves, similar to a typical macro RCT, merely restating information that has already been publicly available elsewhere. The model also implies that communication about inflation may be more effective among agents for whom inflation is relatively unimportant. If communication involves expanding the set of publicly available information, then it is predicted to be more effective when inflation is high, and more effective among agents with a high cost of a mistake.

[17] Non-RCT evidence also indicates that people pay more attention to inflation when inflation is high; e.g., Bracha and Tang (2023), Korenok, Munro, and Chen (2023), and Pfäuti (2023).

References

Afrouzi, H. and C. Yang (2021): "Dynamic Rational Inattention and the Phillips Curve," Working paper.

Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar (2016): "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment," *Review of Economics and Statistics*, 98, 503–523.

Bracha, A. and J. Tang (2023): "Inflation Levels and (In)Attention," *Review of Economic Studies*, forthcoming.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017): "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments," *American Economic Journal: Macroeconomics*, 9, 1–35.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022): "Monetary Policy Communications and Their Effects on Household Inflation Expectations," *Journal of Political Economy*, 130, 1537–1584.

Fuster, A. and B. Zafar (2023): "Survey Experiments on Economic Expectations," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, Academic Press, 107–130.

Haaland, I., C. Roth, and J. Wohlfart (2023): "Designing Information Provision Experiments," *Journal of Economic Literature*, 61, 3–40.

Jurado, K. (2023): "Rational Inattention in the Frequency Domain," *Journal of Economic Theory*, 208, 10.1016/j.jet.2022.105604.

Knotek, E. S., J. Mitchell, M. Pedemonte, and T. Shiroff (2024): "The Effects of Interest Rate Increases on Consumers' Inflation Expectations: The Roles of Informedness and Compliance," Working Paper 24-01, Federal Reserve Bank of Cleveland.

Korenok, O., D. Munro, and J. Chen (2023): "Inflation and Attention Thresholds," *Review of Economics and Statistics*, forthcoming.

Link, S., A. Peichl, C. Roth, and J. Wohlfart (2023): "Information Frictions among Firms and Households," *Journal of Monetary Economics*, 135, 99–115.

Maćkowiak, B., F. Matějka, and M. Wiederholt (2018): "Dynamic Rational Inattention: Analytical Results," *Journal of Economic Theory*, 176, 650–692.

——— (2023): "Rational Inattention: A Review," *Journal of Economic Literature*, 61, 226–273.

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

Miao, J., J. Wu, and E. R. Young (2022): "Multivariate Rational Inattention," *Econometrica*, 90, 907–945.

Pfäuti, O. (2023): "The Inflation Attention Threshold and Inflation Surges," Working paper.

Sims, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

Weber, M., B. Candia, H. Afrouzi, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, G. Kenny, and J. Ponce (2024): "Tell Me Something I Don't Already Know: Learning in Low and High-Inflation Settings," Working paper.

Online Appendix to "Rational Inattention during an RCT"

Bartosz Maćkowiak
European Central Bank and CEPR

Mirko Wiederholt
LMU Munich and CEPR

September 2024

This online appendix provides an analytical expression for the value function in stage 3. We use this expression to solve for equilibrium in the case of a binding no-forgetting constraint in Section 4 of the paper.

As a preliminary step, let $\bar{V}$ denote the value function if the agent enters period 1, the period after the survey, with the steady-state prior uncertainty, $\rho^2\bar{\Sigma}_{t|t}+\sigma_\varepsilon^2$:

$$\bar{V}\equiv\sum_{t=1}^{\infty}\beta^{t-1}\left(\gamma\bar{\Sigma}_{t|t}+\mu ln\left(\frac{\rho^2\bar{\Sigma}_{t|t}+\sigma_\varepsilon^2}{\bar{\Sigma}_{t|t}}\right)\right)=\frac{1}{1-\beta}\left(\gamma\bar{\Sigma}_{t|t}+\mu ln\left(\rho^2+\frac{\sigma_\varepsilon^2}{\bar{\Sigma}_{t|t}}\right)\right).$$

Let $x\equiv\Sigma_{1|0}$ denote the prior uncertainty in period 1. The value function in stage 3, for any prior uncertainty in period 1, can be expressed as:

$$V_1(x)=\begin{cases}\mu ln(x)+\gamma\bar{\Sigma}_{t|t}-\mu ln\left(\bar{\Sigma}_{t|t}\right)+\beta\bar{V} & if \quad x\geq\bar{\Sigma}_{t|t}\\ \gamma x+\beta\left[\mu ln\left(\rho^2x+\sigma_\varepsilon^2\right)+\gamma\bar{\Sigma}_{t|t}-\mu ln\left(\bar{\Sigma}_{t|t}\right)\right]+\beta^2\bar{V} & if \quad x\in\left[x_1,\bar{\Sigma}_{t|t}\right)\\ \gamma x+\beta\gamma\left(\rho^2x+\sigma_\varepsilon^2\right)+\beta^2\left[\mu ln\left(\rho^4x+\left(1+\rho^2\right)\sigma_\varepsilon^2\right)+\gamma\bar{\Sigma}_{t|t}-\mu ln\left(\bar{\Sigma}_{t|t}\right)\right]+\beta^3\bar{V} & if \quad x\in[x_2,x_1)\\ \vdots & \vdots & \vdots\end{cases}$$

If $x\geq\bar{\Sigma}_{t|t}$, the no-forgetting constraint is non-binding in period 1 and the agent chooses the steady-state posterior uncertainty in period 1, entering period 2 with the steady-state prior uncertainty.

If $x\in\left[\left(\bar{\Sigma}_{t|t}-\sigma_\varepsilon^2\right)/\rho^2,\bar{\Sigma}_{t|t}\right)$, the no-forgetting constraint is binding in period 1 and the agent chooses the steady-state posterior uncertainty in period 2, entering period 3 with the steady-state prior uncertainty.

If $x \in \left[\left(\bar{\Sigma}_{t|t} - \sigma_\varepsilon^2\left(1+\rho^2\right)\right)/\rho^4, \left(\bar{\Sigma}_{t|t} - \sigma_\varepsilon^2\right)/\rho^2\right)$, the no-forgetting constraint is binding for two periods and the agent chooses the steady-state posterior uncertainty in period 3, entering period 4 with the steady-state prior uncertainty, and so on.

The stage-2 problem is given by expressions (6)-(8) in the paper with $V_1\left(\Sigma_{1|0}\right)$ given in the previous equation.