Multivariate Rational Inattention*

Jianjun Miao† Jieran Wu‡ Eric R. Young§

First version: December 2018
This version: August 2021

Abstract

We study optimal control problems in the multivariate linear-quadratic-Gaussian framework under rational inattention. We propose a three-step procedure to solve this problem using semidefinite programming and derive the optimal signal structure without strong prior restrictions. We analyze both the transition dynamics of the optimal posterior covariance matrix and its steady state. We characterize the optimal information structure for some special cases and develop numerical algorithms for general cases. Applying our methods to solve three multivariate economic models, we obtain some results qualitatively different from the literature.

Keywords: Rational Inattention, Endogenous Information Choice, Tracking Problem, Optimal Control, Entropy, Semidefinite Program

JEL Classifications: C61, D83, E21, E22, E31.

*We are extremely grateful to four referees and the editor (Gianluca Violante) whose insightful comments led us to improve our paper significantly. In particular, Chris Sims has guided us throughout the revision of our paper. Many of our results reflect his own ideas or our responses to his insightful comments. We thank Lars Hansen, Yulei Luo, Lin Peng, Giacomo Rondina, Pengfei Wang, Michael Woodford, Wei Xiong, He Yang, and Tao Zha for helpful comments. We also thank the participants at a seminar in Boston University and the AFR Summer Institute of Economics and Finance for comments. Please address all correspondences to Jianjun Miao.

†Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215. Email: miaoj@bu.edu

‡Academy of Financial Research and College of Economics, Zhejiang University, jw5ya@zju.edu.cn

§Department of Economics, University of Virginia, ey2d@virginia.edu, and Research Department, Federal Reserve Bank of Cleveland, eric.young@clev.frb.org

1 Introduction

Humans have limited capacity to process information when making decisions. People often ignore some pieces of information and pay attention to some others. In seminal contributions, Sims (1998, 2003) formalizes limited attention as a constraint on information flow and models decision-making with limited attention as optimization subject to this constraint. Such a framework for rational inattention (RI) has wide applications in economics as surveyed by Sims (2011) and Maćkowiak, Matějka, and Wiederholt (2020). Despite the rapid growth of this literature, most theories and applications have been limited to univariate models.

Multivariate RI models are difficult to analyze both theoretically and numerically, especially in dynamic settings. Because many economic decision problems involve multivariate states and multivariate actions, it is of paramount importance to make progress in this direction as Sims (2011) points out. Our paper contributes to the literature by developing a framework for analyzing multivariate RI problems in a linear-quadratic-Gaussian (LQG) control setup along the lines of Sims (2011).[1] The LQG control setup has a long tradition in economics and can deliver analytical results to understand economic intuition. It is also useful to derive numerical solutions for approximating nonlinear dynamic models (Kydland and Prescott (1982)). We formulate the LQG control problem under RI in both finite- and infinite-horizon setups as a problem of choosing both the control and information structure. The decision maker observes a noisy signal about the unobserved controlled states. The signal vector is a linear transformation of the states plus a noise. The signal dimension, the linear transformation, and the noise covariance matrix are all endogenously chosen subject to period-by-period capacity constraints. Alternatively, the information choice incurs discounted (Shannon entropy) information costs measured in utility units.

Our second contribution is to develop an efficient three-step solution procedure. The first step is to derive the full information solution and the second step is to apply the certainty equivalence principle and the separation principle to derive the optimal control under an exogenous information structure. These two steps follow from the standard control literature. The third step is to solve for the optimal information structure under RI. We focus on the formulation with discounted information costs and the analysis for the formulation with period-by-period capacity constraints is more subtle.

Like Sims (2011), we show that solving for the optimal information structure is equivalent to solving for the sequence of optimal posterior covariance matrices for the state vector. It seems natural to solve this sequence using dynamic programming. The difficulty is that this problem may not be convex and the choice variable must be a positive semidefinite matrix. Moreover, the RI problem involves no-forgetting constraints which are matrix inequality constraints. To

[1] See Sims (2006), Matějka and McKay (2015), and Caplin, Dean, and Leahy (2019) for static non-Gaussian RI models.

tackle these issues, we adopt the semidefinite programming (SDP) approach in the mathematics and engineering literature, which is the mathematical tool to study optimization over positive semidefinite matrices (Vandenberghe and Boyd (1996) and Tanaka et al (2017)). We first transform the original dynamic programming problem into an auxiliary convex dynamic program and then derive an SDP representation. To facilitate an efficient and robust numerical implementation, we construct the representation as a disciplined convex program (DCP) (Grant (2004) and Grant, Boyd and Ye (2006)). A DCP must conform to the DCP ruleset so that it can be easily verified as convex and solvable in a computer. DCPs can be numerically solved using the powerful software CVX (Grant and Boyd (2008) and CVX Research, Inc. (2012)), which is freely available from the internet (http://cvxr.com/cvx/).[2]

The mathematics and engineering literature typically focuses on static SDP. We contribute to the literature by studying dynamic discounted SDP and establishing the convexity of the value function. For the infinite-horizon case, such a dynamic program does not give a contraction mapping. Nevertheless, we use the method of value function iteration (VFI) to show that the sequence of value functions for the truncated finite-horizon problems converges to the infinite-horizon value function. We can then derive the optimal sequence of posterior covariance matrices and the limiting steady state. As is well known, the method of VFI can be numerically slow especially for high dimensional problems. We then modify the basic VFI method in two ways. First, we apply the envelope condition. Second, we solve an equivalent sequence of static RI problems. Both ways speed up computation significantly. We also characterize the first-order conditions and develop efficient algorithms to solve for the steady state and transition dynamics based on these conditions. The first-order conditions based methods are much faster, but the convergence is not guaranteed as a convergence proof is unavailable. The value function based methods are more flexible to incorporate many occasionally binding constraints and nonsmooth objective functions. We develop an efficient Matlab toolbox to implement both types of methods (Miao and Wu (2021)).

We provide some characterization results for the steady state when the discount factor is equal to one and the state transition matrix is diagonal with equal lag coefficients. This includes two special cases: (i) the state vector is serially independently and identically distributed (IID), conditional on a control, and (ii) all states are equally persistent AR(1) processes with correlated innovations. The first special case also gives the solution for the static RI problem, which generalizes the reverse water-filling solution in Theorem 10.3.3 of Cover and Thomas (2006, p. 314) by allowing for general positive semidefinite benefit matrix of information and general positive definite prior covariance matrix. We characterize the optimal signal dimension and show that it does not exceed the minimum of the state dimension and the control dimension and weakly decreases as the

[2]CVX supports two free SDP solvers, SeDuMi (Sturm (1999)) and SDPT3 (Toh, Todd, and Tutuncu (1999)), and a commercial SDP solver, Mosek, which is also free for academic users. These solvers use the interior point methods. We find that Mosek is the fastest and SDPT3 is the most reliable for our examples.

information cost rises. Allowing for nonstationary state processes, we show that RI can make their posterior covariance matrix stationary after acquiring information endogenously.

For pure tracking problems in which all states follow exogenous dynamics and the objective is mean squared error, we prove that the optimal signal is one dimensional if the rank of the benefit matrix of information is equal to one. This case happens when there is only one control variable. The optimal signal is equal to the target under full information plus a noise when the target follows an AR(1) process.

Our third contribution is to apply our results to three economic problems. For all applications, we focus on the steady-state solution for the optimal information structure. Our first application is the price setting problem adapted from Maćkowiak and Wiederholt (2009), in which there are two exogenous state variables representing two sources of uncertainty. We first ignore the general equilibrium price feedback effect and just focus on the decision problem as in Sims (2011). The profit-maximizing price is equal to a linear combination of the two shocks. We then study the general equilibrium model of Maćkowiak and Wiederholt (2009) in which the endogenous aggregate price level affects individual firms' profit-maximizing prices.

We approximate the equilibrium price by an ARMA process as in Maćkowiak, Matějka, and Wiederholt (2018) and derive a state space representation for a firm's tracking problem under RI. We find that the optimal signal is one dimensional, implying that a firm is confused about the sources of shocks and hence there is a volatility spillover effect: An increase in the volatility of one source of the shock causes the firm to raise price responses to both sources of shocks.

Our second application is the consumption/saving problem analyzed by Sims (2003), in which there is an endogenous state variable (wealth) and two exogenous persistent state variables (income shocks). We find that the optimal signal is one dimensional as in Luo (2008). Unlike Sims's (2003) finding, the consumption responses to shocks with different persistence follow similar dynamics.

Our last application is the firm investment problem in which the firm makes both tangible and intangible capital investment. We find that the signal dimension increases from one to two as the information cost parameter declines to a sufficiently small value. Moreover, given a small information cost parameter during the transition phase, the firm does not acquire information initially, then acquires a one-dimensional signal at some time in the future, and finally acquires a two-dimensional signal after some additional time as state innovations arrive in each period. Sims (1998, 2003) argues that RI can substitute for adjustment costs in a dynamic optimization problem. Our numerical results show that RI can generate inertia and delayed responses of investment to shocks, just like capital adjustment costs. Moreover, we find that RI combined with capital adjustment costs can generate hump-shaped investment responses as in Zorn (2018).

We now discuss the related literature. Sims (2003) is the first paper that introduces multivariate

LQG RI models with information-flow constraints.[3] He simplifies the solution for the steady-state posterior covariance matrix by minimizing the steady-state welfare loss. Much of the literature has followed Sims's approach. However, for a general control problem, one must take care of the initial state, which is drawn from an endogenous steady-state distribution. The literature has neglected this initial value problem, which does not arise in pure tracking problems.[4] We tackle this issue and show that Sims's solution coincides with the steady-state solution when the discount factor approaches one. Thus Sims's solution may be viewed as an approximation to the steady-state solution when the discount factor is close to one.

Sims (2011) proposes the formulation with discounted information costs and analyzes the LQG RI problem without explicit reference to the signal structure. His solution procedure consists of two steps. His first step is essentially the same as our first two steps. His second step is to transform the control problem under RI into a problem of choosing a sequence of optimal posterior covariance matrices for the state vector. Sims (2011) proposes to solve for the steady state as the limit point of the optimal sequence. He outlines a method based on first-order conditions when the no-forgetting constraints do not bind and recommends to use the Cholesky decomposition when they bind without providing a detailed analysis.

Maćkowiak, Matějka, and Wiederholt (2018) study a pure tracking problem under information-flow constraints with one control and one exogenous state, which follows a general ARMA process. They assume that the decision maker has chosen the information structure in period minus one and then receives a long sequence of signals such that the prior covariance matrix of the state in period zero equals the steady-state prior covariance matrix of the state. They derive some elegant analytical results and characterize optimal signal form. They also discuss the extension to the case with multiple exogenous states, but still with one control. Consistent with our result, the optimal signal is one dimensional. Relative to their study, our contribution is to study the steady state and transition dynamics for general LQG control problems under RI with multiple controls and endogenous states.

In independent and contemporaneous work, Afrouzi and Yang (2019) study a pure tracking problem with discounted information costs as in our Section 4 and provide characterizations of first-order conditions with no analysis of whether these conditions are sufficient for optimality. They precede us by developing a Julia toolbox to solve for both transition dynamics and steady state. Our paper differs from theirs in four respects.

First, we study the general LQG control problem under RI with endogenous states in both finite- and infinite-horizon setups. We show that this problem can be cast into a form that produces a deterministic dynamic programming problem of the same mathematical structure as the tracking

[3]Luo (2008), Luo and Young (2010), and Luo, Nie and Young (2015) follow Sims's approach closely, but mainly focus on univariate models.

[4]We are extremely grateful to Chris Sims for pointing out this issue to us.

problem that the algorithm of Afrouzi and Yang (2019) solves. Second, we provide a dynamic programming characterization and establish both necessity and sufficiency of first-order conditions. We also develop several value function based methods to numerically solve dynamic RI problems. We show that a dynamic RI problem can be viewed as a sequence of static RI problems as in Steiner, Stewart, and Matějka (2017). We can then apply the static reverse water-filling solution to characterize the first-order conditions for the dynamic RI model. Third, we provide some analytical results for the steady state when the discount factor is equal to one. Finally, we study different applications including both tracking and control problems such as our consumption and investment examples.

Because of the difficulty of solving multivariate RI models, researchers often make simplifying assumptions. For example, Peng (2005), Peng and Xiong (2006), and Van Niewerburgh and Veldkamp (2010) impose the signal independence assumption in static models. An undesirable implication is that initially independent states remain ex post independent. Mondria (2010) and Kacperczyk, Van Niewerburgh, and Veldkamp (2016) remove this assumption in static finance models. The former paper considers only two independent assets (states), while the latter studies the case of many assets given some invertibility restriction on the signal form. Maćkowiak and Wiederholt (2009, 2015) and Zorn (2018) study dynamic tracking models and use Sims's (2003) method of approximating an MA(∞) representation by a finite MA process. They then compute the MA coefficients by brute-force optimization.

As in Maćkowiak, Matějka, and Wiederholt (2018), both the linear transformation and the noise covariance matrix in the signal form must be endogenously chosen under our formulation. In addition to the attention allocation effect emphasized in the literature, the learning effect induced by the linear transformation of states is also important for decision making because the linear transformation determines how the decision maker collects different sources of information by combining different states. Linear combination of states can cause the decision maker to be confused about different sources of uncertainty, thereby generating a spillover effect.

In independent work Fulton (2018) and Kőszegi and Matějka (2019) analyze similar multivariate RI problems in the static case and derives results similar to our generalized reverse water-filling solution. Kőszegi and Matějka (2019) assume that all states are IID ex ante. Fulton (2017) discusses dynamic tracking problems with exogenous states and proposes a different solution method.[5]

Our paper is also related to other studies that are not in the discrete-time LQG framework. This literature is growing. Recent papers include Steiner, Stewart, and Matějka (2017), Dewan (2018), Hébert and Woodford (2018), and Zhong (2019). Miao (2019) studies continuous-time LQG RI problems, which require different mathematical tools. He does not study transitional dynamics and

[5]We would like to thank Gianluca Violante for pointing out Fulton's papers to us, when we presented a preliminary version of our paper in a conference in June 2018.

many economic examples in this paper.

The remainder of the paper proceeds as follows. Section 2 presents the RI problem in the general LQG framework and discusses the three-step solution procedure. Section 3 focuses on the last step of solving for the sequence of optimal posterior covariance matrices using dynamic SDP. We characterize the first-order conditions and derive some analytical results for some special cases. Section 4 studies pure tracking problems and numerically solves an example taken from Sims (2011) to illustrate the effects of some parameters on the steady state and transition dynamics. Section 5 uses our methods to study three applications. Section 6 discusses an alternative solution concept and its relation to Sims's (2003) approach and the steady-state solution. Section 7 concludes. Proofs and technical details are relegated to appendices.

2 LQG Control Problems with Rational Inattention

We start with a finite-horizon linear-quadratic control problem under rational inattention. Let the n_x dimensional state vector x_t follow the linear dynamics

$$x_{t+1} = A_t x_t + B_t u_t + \epsilon_{t+1},\ t = 0, 1, ..., T, \tag{1}$$

where u_t is an n_u dimensional control variable and ϵ_{t+1} is a Gaussian white noise with covariance matrix W_t. The matrix W_t is positive semidefinite, denoted by $W_t \succeq 0$.[6] The state transition matrix A_t and the control coefficient matrix B_t are deterministic and conformable. The state vector x_t may contain both exogenous states such as AR(1) shocks and endogenous states such as capital.

Suppose that the decision maker does not observe the state x_t perfectly, but observes a multi-dimensional noisy signal s_t about x_t given by

$$s_t = C_t x_t + v_t,\ t = 0, 1, ..., T, \tag{2}$$

where C_t is a conformable deterministic matrix and v_t is a Gaussian white noise with covariance matrix $V_t \succ 0$. Notice that we do not impose any other restriction on C_t or V_t.[7] In particular, C_t may not be an identity matrix or invertible. Assume that x_0 is a Gaussian random variable with mean $\bar{x}_0$ and covariance matrix $\Sigma_{-1} \succ 0$. The random variables ϵ_t, v_t, and x_0 are all mutually independent for all t. The decision maker's information set at date t is generated by $s^t = \{s_0, s_1, ..., s_t\}$. The control u_t is measurable with respect to s^t.

Suppose that the decision maker is boundedly rational and has limited information-processing

[6] We use the conventional matrix inequality notations: $X \succ (\succeq) Y$ means that $X - Y$ is positive definite (semidefinite) and $X \prec (\preceq) Y$ means $X - Y$ is negative definite (semidefinite).

[7] As will be clear later, the signal form in (2) is not restrictive and can be recovered from the optimal posterior covariance matrix for the state vector (see Proposition 1).

capacity. They face the following period-by-period capacity constraint[8]

$$I\left(x_t; s_t | s^{t-1}\right) \leq \kappa, \ t = 0, 1, ..., T, \tag{3}$$

where $\kappa > 0$ denotes the information-flow rate or capacity and $I\left(x_t; s_t | s^{t-1}\right)$ denotes the conditional (Shannon) mutual information between x_t and s_t given s^{t-1},

$$I\left(x_t; s_t | s^{t-1}\right) \equiv H\left(x_t | s^{t-1}\right) - H\left(x_t | s^t\right).$$

Here $H(\cdot|\cdot)$ denotes the conditional entropy operator.[9] Let $s^{-1} = \varnothing$. Intuitively, entropy measures uncertainty. At each time t, given past information s^{t-1}, observing s_t reduces uncertainty about x_t. The decision maker can process information by choosing the information structure represented by $\{C_t, V_t\}_{t=0}^T$ for the signal s_t, but the rate of uncertainty reduction in each period is limited by an upper bound κ.

Notice that the choice of $\{C_t, V_t\}_{t=0}^T$ implies that the dimension of the signal vector s_t and the correlation structure of the noise v_t are endogenous and may vary over time. The decision maker makes decisions sequentially. They first choose the information structure $\{C_t, V_t\}_{t=0}^T$ and then select a control $\{u_t\}_{t=0}^T$ adapted to $\left\{s^t\right\}$ to maximize an objective function. Suppose that the objective function is quadratic. We are ready to formulate the decision maker's problem as follows:

Problem 1 *(Finite-horizon LQG problem under RI with period-by-period capacity constraints)*

$$\max_{\{u_t\},\{C_t\},\{V_t\}} -\mathbb{E}\left[\sum_{t=0}^{T} \beta^t \left(x_t' Q_t x_t + u_t' R_t u_t + 2x_t' S_t u_t\right) + \beta^{T+1} x_{T+1}' P_{T+1} x_{T+1}\right]$$

subject to (1), (2), and (3), where $\beta \in (0,1]$ and the expectation is taken with respect to the joint distribution induced by the initial distribution for x_0 and the state dynamics (1).

The parameter $\beta \in (0,1]$ represents the discount factor. The deterministic matrices Q_t, R_t, and S_t for all t and P_{T+1} are conformable and exogenously given. In applications it may be more convenient to consider the following relaxed problems with discounted information costs.

Problem 2 *(Finite-horizon LQG problem under RI with discounted information costs)*

$$\begin{aligned}\max_{\{u_t\},\{C_t\},\{V_t\}} \quad & -\mathbb{E}\left[\sum_{t=0}^{T} \beta^t \left(x_t' Q_t x_t + u_t' R_t u_t + 2x_t' S_t u_t\right) + \beta^{T+1} x_{T+1}' P_{T+1} x_{T+1}\right] \\ & -\lambda \sum_{t=0}^{T} \beta^t I\left(x_t; s_t | s^{t-1}\right)\end{aligned}$$

subject to (1) and (2), where $\beta \in (0,1]$ and the expectation is taken with respect to the joint distribution induced by the initial distribution for x_0 and the state dynamics (1).

[8] We do not adopt the capacity constraint on the total information flows across periods because this formulation causes the dynamic inconsistency issue.

[9] See Cover and Thomas (2006) or Sims (2011) for the definitions of entropy, conditional entropy, mutual information, and conditional mutual information.

In this problem $\lambda > 0$ can be interpreted as the shadow price (cost) of the information flow. For the infinite-horizon stationary case, we set $T \to \infty$ and remove the time index for all exogenously given matrices A_t, B_t, Q_t, R_t, S_t, and W_t. Under some stability conditions, the posterior distribution for x_t will converge to a long-run stationary distribution.

For simplicity we focus our analysis on Problem 2 and its infinite-horizon limit as $T \to \infty$. We discuss how we solve Problem 1 in Online Appendix C.

2.1 Full Information Case

Before analyzing Problem 2, we first present the solution in the full information case, in which the decision maker observes x_t perfectly. The solution can be found in the textbooks by Ljungqvist and Sargent (2004) and Miao (2014). If all states are endogenous, suppose that $P_{T+1} \succeq 0$, $R_t \succ 0$, and

$$\begin{bmatrix} Q_t & S_t \\ S_t' & R_t \end{bmatrix} \succeq 0$$

for all $t = 0, 1, ..., T$. Then the value function given a state x_t takes the form

$$v_t^{FI}(x_t) = -x_t' P_t x_t - \sum_{\tau=t}^{T} \beta^{\tau-t+1} \text{tr}(W_\tau P_{\tau+1}), \tag{4}$$

where $P_t \succeq 0$ and satisfies the Riccati equation

$$\begin{aligned} P_t &= Q_t + \beta A_t' P_{t+1} A_t \\ &\quad - \left(\beta A_t' P_{t+1} B_t + S_t\right) \left(R_t + \beta B_t' P_{t+1} B_t\right)^{-1} \left(\beta B_t' P_{t+1} A_t + S_t'\right), \end{aligned} \tag{5}$$

for $t = 0, 1, ..., T$. Here $\text{tr}(\cdot)$ denotes the trace operator.

The optimal control is

$$u_t = -F_t x_t, \tag{6}$$

where

$$F_t = \left(R_t + \beta B_t' P_{t+1} B_t\right)^{-1} \left(S_t' + \beta B_t' P_{t+1} A_t\right). \tag{7}$$

If some states are exogenous, we can augment the state vector and derive similar results.

For the infinite horizon case, all exogenous matrices are time invariant. As $T \to \infty$, we obtain the infinite-horizon solution under some standard stability conditions. The value function becomes

$$v^{FI}(x_t) = -x_t' P x_t - \frac{\beta}{1-\beta} \text{tr}(WP),$$

where P satisfies the Riccati equation

$$P = Q + \beta A' P A - \left(\beta A' P B + S\right) \left(R + \beta B' P B\right)^{-1} \left(\beta B' P A + S'\right). \tag{8}$$

The optimal control is given by

$$u_t = -F x_t, \tag{9}$$

where

$$F = \left(R + \beta B' P B\right)^{-1} \left(S' + \beta B' P A\right).$$

2.2 Control under Exogenous Information Structure

We solve Problem 2 in three steps. In the first step we derive the full-information solution as in Section 2.1. In the second step we observe that Problem 2 is a standard LQG problem under partial information when the information structure $\{C_t, V_t\}_{t=0}^T$ is exogenously fixed. Thus the usual separation principle and certainty equivalence principle hold. This implies that the optimal control is given by

$$u_t = -F_t \widehat{x}_t, \tag{10}$$

where $\widehat{x}_t \equiv \mathbb{E}\left[x_t | s^t\right]$ denotes the estimate of x_t given information s^t. Notice that the matrix F_t is determined by (7) in the full information case, which is independent of the information structure.

The state under the optimal control satisfies the dynamics

$$x_{t+1} = A_t x_t - B_t F_t \widehat{x}_t + \epsilon_{t+1}. \tag{11}$$

By the Kalman filter formula, $\widehat{x}_t$ follows the dynamics

$$\widehat{x}_t = \widehat{x}_{t|t-1} + \Sigma_{t|t-1} C_t' \left(C_t \Sigma_{t|t-1} C_t' + V_t\right)^{-1} \left(s_t - C_t \widehat{x}_{t|t-1}\right), \tag{12}$$

$$\widehat{x}_{t+1|t} = \left(A_t - B_t F_t\right) \widehat{x}_t, \ t \geq 0, \tag{13}$$

where $\widehat{x}_{t|t-1} \equiv \mathbb{E}\left[x_t | s^{t-1}\right]$ with $\widehat{x}_{0|-1} = \bar{x}_0$ and $\Sigma_{t|t-1} \equiv \mathbb{E}\left[\left(x_t - \widehat{x}_{t|t-1}\right)\left(x_t - \widehat{x}_{t|t-1}\right)' | s^{t-1}\right]$ with $\Sigma_{0|-1} = \Sigma_{-1}$ exogenously given. Moreover,

$$\Sigma_{t+1|t} = A_t \Sigma_t A_t' + W_t, \tag{14}$$

$$\Sigma_t = \left(\Sigma_{t|t-1}^{-1} + \Phi_t\right)^{-1}, \tag{15}$$

for $t = 0, 1, ..., T$, where $\Sigma_t \equiv \mathbb{E}\left[\left(x_t - \widehat{x}_t\right)\left(x_t - \widehat{x}_t\right)' | s^t\right]$ denotes the posterior covariance matrix given s^t and Φ_t denotes the signal-to-noise ratio (SNR) defined by $\Phi_t = C_t' V_t^{-1} C_t \succeq 0$, $t = 0, 1, ..., T$.

We need the following lemma to derive the optimal information structure. Its proof is given in Appendix A.

Lemma 1 *Under the optimal control policy in (10) for fixed information structure $\{C_t, V_t\}_{t=0}^T$, we have*

$$\begin{aligned} & \mathbb{E}\left[\sum_{t=0}^T \beta^t \left(x_t' Q_t x_t + u_t' R_t u_t + 2 x_t' S_t u_t\right) + \beta^{T+1} x_{T+1}' P_{T+1} x_{T+1}\right] \\ = \ & \mathbb{E}\left[x_0' P_0 x_0\right] + \sum_{t=0}^T \beta^{t+1} \mathrm{tr}\left(W_t P_{t+1}\right) + \sum_{t=0}^T \beta^t \mathrm{tr}\left(\Omega_t \Sigma_t\right), \end{aligned}$$

where

$$\Omega_t = F_t' (R_t + \beta B_t' P_{t+1} B_t) F_t \succeq 0. \tag{16}$$

Notice that the matrix Ω_t is positive semidefinite because $R_t \succ 0$ and $B_t' P_{t+1} B_t \succeq 0$. The matrix Ω_t translates estimation error Σ_t into welfare loss, and measures the marginal benefit of information (the reduction of uncertainty). Since F_t is an n_u by n_x dimensional matrix, the rank of Ω_t, denoted by rank(Ω_t), does not exceed the minimum of the state dimension n_x and the control dimension n_u. Thus it is possible that Ω_t is singular. If $n_x \geq n_u$ and F_t has full column rank, then rank$(\Omega_t) = n_u$. If $n_x < n_u$ and F_t has full row rank, then rank$(\Omega_t) = n_x$.

2.3 Optimal Information Structure

In the final step of our solution procedure, we solve for the optimal information structure $\{C_t, V_t\}$. In doing so, we compute the mutual information[10]

$$\begin{aligned} I\left(x_t; s_t | s^{t-1}\right) &= H\left(x_t | s^{t-1}\right) - H\left(x_t | s^t\right) \\ &= \frac{1}{2}\log\det\left(A_{t-1}\Sigma_{t-1}A_{t-1}' + W_{t-1}\right) - \frac{1}{2}\log\det\left(\Sigma_t\right) \end{aligned}$$

for $t = 1, 2, .., T$, and

$$I\left(x_0; s_0 | s^{-1}\right) = H\left(x_0\right) - H\left(x_0 | s_0\right) = \frac{1}{2}\log\det\left(\Sigma_{-1}\right) - \frac{1}{2}\log\det\left(\Sigma_0\right)$$

for $t = 0$, where the functions $H(\cdot)$ and $H(\cdot|\cdot)$ denote the entropy and conditional entropy operators, and $\det(\cdot)$ denotes the determinant operator.

Since $\{P_t\}$ is independent of the information structure and $\mathbb{E}\left[x_0' P_0 x_0\right]$ is determined by the exogenous initial prior distribution, it follows from Lemma 1 that solving for the optimal information structure in Problem 2 is equivalent to solving for the optimal sequence of posterior covariance matrices for the state vector:

Problem 3 *(Optimal information structure for Problem 2)*

$$\min_{\{\Sigma_t\}_{t=0}^T} \sum_{t=0}^{T} \beta^t \left[\operatorname{tr}\left(\Omega_t \Sigma_t\right) + \lambda I\left(x_t; s_t | s^{t-1}\right)\right]$$

subject to

$$I\left(x_t; s_t | s^{t-1}\right) = \frac{1}{2}\log\det\left(A_{t-1}\Sigma_{t-1}A_{t-1}' + W_{t-1}\right) - \frac{1}{2}\log\det\left(\Sigma_t\right),$$

$$I\left(x_0; s_0 | s^{-1}\right) = \frac{1}{2}\log\det\left(\Sigma_{-1}\right) - \frac{1}{2}\log\det\left(\Sigma_0\right),$$

$$\Sigma_t \preceq A_{t-1}\Sigma_{t-1}A_{t-1}' + W_{t-1}, \tag{17}$$

$$\Sigma_0 \preceq \Sigma_{-1}, \tag{18}$$

for $t = 1, 2, ..., T$.

[10] The usual base for logarithm in the entropy formula is 2, in which case the unit of information is a “bit.” In this paper we adopt natural logarithm, in which case the unit is called a “nat.”

It follows from Lemma 1 and (4) that the expression $\sum_{t=0}^{T} \beta^t \text{tr}(\Omega_t \Sigma_t)$ represents the expected welfare loss due to the limited information (i.e., the difference between the expected discounted utilities under full information and under limited information). The optimal information structure under RI minimizes the welfare loss plus the discounted information cost. Sims (2011) formulates an essentially identical problem for the infinite-horizon case as $T \to \infty$, except that there is a difference in constraints at date zero. The matrix inequalities (17) and (18) are called the no-forgetting constraints (Sims (2003, 2011)). They can be derived from (14) and (15) as the SNR Φ_t is positive semidefinite. After obtaining $\{\Sigma_t\}$, we can recover $\{\Phi_t\}$ and hence $\{C_t\}$ and $\{V_t\}$ from the following result. Its proof and the proofs of all other results in the main text are collected in Online Appendix B.

Proposition 1 *Given an optimal sequence $\{\Sigma_t\}_{t=0}^T$ determined from Problem 3, the optimal SNR is given by*

$$\Phi_0 = \Sigma_0^{-1} - \Sigma_{-1}^{-1},\ \Phi_t = \Sigma_t^{-1} - \left(A_{t-1}\Sigma_{t-1}A_{t-1}' + W_{t-1}\right)^{-1},\ t \geq 1.$$

An optimal information structure $\{C_t, V_t\}_{t=0}^T$ satisfies $\Phi_t = C_t' V_t^{-1} C_t$. A particular solution is that $V_t = diag\left(\varphi_{it}^{-1}\right)_{i=1}^{m_t}$ and the m_t columns of $n_x \times m_t$ matrix C_t' are orthonormal eigenvectors for all positive eigenvalues of Φ_t, denoted by $\{\varphi_{it}\}_{i=1}^{m_t}$. The optimal dimension of the signal vector s_t is equal to rank $(\Phi_t) = m_t \leq n_x$.

This proposition shows that the optimal information structure $\{C_t, V_t\}_{t=0}^T$ is not unique and can be computed by the singular-value decomposition. The optimal signal can always be constructed such that the components in the noise vector v_t of the signal s_t are independent. Throughout the paper we focus on the signal structure such that V_t is diagonal for each t. In this case C_t is unique up to a scalar constant and up to an interchange of rows. When C_t is scaled by a constant b, V_t is scaled by b^2. By the Kalman filter, the impulse responses to structural shocks to all state variables do not change, but the responses to noise shocks are scaled by $1/b$. Notice that optimal signals are in general not independent in the sense that the matrix C_t may not be diagonal or invertible.

3 Dynamic Semidefinite Programming

In this section we focus on the analysis of Problem 3, which is not a trivial dynamic problem because the choice variables are positive semidefinite matrices and the constraints are matrix inequalities. We extend the SDP approach recently proposed by Tanaka et al (2017) for static programs to the dynamic case. We also provide some characterization results for some special cases.

3.1 Finite-Horizon Case

We use dynamic programming to study Problem 3 (Stokey and Lucas with Prescott (1989) and Miao (2014)). Let $\mathcal{V}_0(\Sigma_{-1})$ be the value function for Problem 3. Let $\mathcal{V}_t(\Sigma_{t-1})$ be the value function for the continuation problem in period $t \geq 1$ defined as

$$\mathcal{V}_t\left(\Sigma_{t-1}\right) = \min_{\{\Sigma_\tau\}_{\tau=t}^T} \sum_{\tau=t}^T \beta^{\tau-t}\left[\operatorname{tr}\left(\Omega_\tau \Sigma_\tau\right) + \lambda I\left(x_\tau ; s_\tau | s^{\tau-1}\right)\right]$$

subject to

$$I\left(x_\tau ; s_\tau | s^{\tau-1}\right) = \frac{1}{2}\log\det\left(A_{\tau-1}\Sigma_{\tau-1}A'_{\tau-1} + W_{\tau-1}\right) - \frac{1}{2}\log\det\left(\Sigma_\tau\right),$$

$$\Sigma_\tau \preceq A_{\tau-1}\Sigma_{\tau-1}A'_{\tau-1} + W_{\tau-1},$$

for $\tau = t, t+1, ..., T$.

The sequence of value functions $\mathcal{V}_t\left(\Sigma_{t-1}\right)$ for $t \geq 0$ satisfies Bellman equations. But $\mathcal{V}_t\left(\Sigma_{t-1}\right)$ may not be convex, as will become clear later. We thus solve an auxiliary convex problem. Specifically, in the last period T, consider

$$J_T\left(\Sigma_{T-1}\right) \equiv \min_{\Sigma_T \succ 0} \operatorname{tr}\left(\Omega_T \Sigma_T\right) - \frac{\lambda}{2}\log\det\left(\Sigma_T\right) \tag{19}$$

subject to (17) for $t = T$. Since the log-determinant function is strictly concave and (17) is a linear matrix inequality, the problem in (19) is a convex program and hence $J_T\left(\Sigma_{T-1}\right)$ is also strictly convex in Σ_{T-1}.

In any period $t = 0, 1, ..., T-1$, consider the Bellman equation:

$$J_t\left(\Sigma_{t-1}\right) = \min_{\Sigma_t \succ 0} \operatorname{tr}\left(\Omega_t \Sigma_t\right) + \frac{\lambda}{2}\left[\beta\log\det\left(A_t\Sigma_tA'_t + W_t\right) - \log\det\left(\Sigma_t\right)\right] + \beta J_{t+1}\left(\Sigma_t\right) \tag{20}$$

subject to (17) for $t \geq 1$ and (18) for $t = 0$.

It is straightforward to verify that

$$\mathcal{V}_t\left(\Sigma_{t-1}\right) = J_t\left(\Sigma_{t-1}\right) + \frac{\lambda}{2}\log\det\left(A_{t-1}\Sigma_{t-1}A'_{t-1} + W_{t-1}\right) \tag{21}$$

for $t \geq 1$ and

$$\mathcal{V}_0\left(\Sigma_{-1}\right) = J_0\left(\Sigma_{-1}\right) + \frac{\lambda}{2}\log\det\left(\Sigma_{-1}\right). \tag{22}$$

Moreover, the optimal solution $\{\Sigma_t\}_{t=0}^T$ for (19) and (20) also gives the optimal solution to Problem 3 by the dynamic programming principle.

Our dynamic programming formulation allows us to interpret the dynamic RI Problem 3 as an investment problem.[11] Specifically, the state variable in period t is the posterior covariance matrix Σ_{t-1} determined in the previous period, which can be interpreted as a stock of knowledge. The decision maker invests in the stock by acquiring information to reduce uncertainty given the prior covariance matrix $\Sigma_{t|t-1} = A_{t-1}\Sigma_{t-1}A'_{t-1} + W_{t-1}$. That is, the decision maker chooses $\Phi_t = C'_tV_t^{-1}C_t \succeq 0$ to determine Σ_t such that (15) is satisfied. Moving to period $t+1$, the state variable becomes Σ_t and the prior covariance matrix $\Sigma_{t+1|t}$ evolves according to (14).

The following lemma is nontrivial and important for the convexity of the value function J_t.

[11] We thank an anonymous referee for suggesting this interpretation.

Lemma 2 *Suppose that* $W \succeq 0$ *and* $AA' + W \succ 0$. *Then the function*

$$\mathcal{F}(\Sigma) = \beta \log\det\left(A\Sigma A' + W\right) - \log\det\Sigma \tag{23}$$

is convex in $\Sigma \succ 0$ *for* $\beta \in (0, 1]$ *and is strictly convex for* $\beta \in (0, 1)$.

It is not obvious whether $\mathcal{F}$ is convex as it is the difference of two concave functions. Sims (2003) establishes the convexity of $\mathcal{F}$ for $\beta = 1$ assuming A is invertible, while we assume W is invertible in a previous version of this paper. Afrouzi and Yang (2019) introduce the weaker assumption in this lemma to ensure the invertibility of $A\Sigma A' + W$. But they do not establish the convexity of $\mathcal{F}$. Its proof is quite involved as shown in Online Appendix B.

Since the software CVX cannot recognize whether the difference of two concave functions is convex by its ruleset, we need to transform the dynamic programming problem (20) into a DCP form. To achieve this goal, the following proposition derives a dynamic SDP representation.

Proposition 2 *Suppose that* $W_t \succeq 0$ *and* $A_t A_t' + W_t \succ 0$ *for* $t = 0, 1, ..., T-1$. *(a) The value function* $J_t(\Sigma_{t-1})$ *is strictly convex in* $\Sigma_{t-1} \succ 0$ *for* $\beta \in (0, 1)$. *(b) If* $W_t \succ 0$, $J_t(\Sigma_{t-1})$ *satisfies the dynamic SDP for* $t = 0, 1, ..., T-1$:

$$\begin{aligned} J_t(\Sigma_{t-1}) = \min_{\Pi_t \succ 0, \Sigma_t \succ 0} & \operatorname{tr}(\Omega_t \Sigma_t) - \frac{\lambda}{2}(1-\beta)\log\det(\Sigma_t) \\ & + \frac{\lambda\beta}{2}(\log\det W_t - \log\det \Pi_t) + \beta J_{t+1}(\Sigma_t) \end{aligned} \tag{24}$$

subject to (17) and

$$\begin{bmatrix} \Sigma_t - \Pi_t & \Sigma_t A_t' \\ A_t \Sigma_t & W_t + A_t \Sigma_t A_t' \end{bmatrix} \succeq 0,$$

where $J_T(\Sigma_{T-1})$ *satisfies (19) and is also strictly convex. For* $t = 0$, *(17) is replaced by (18).*

Since $J_t(\Sigma_{t-1})$ is strictly convex for $t = 0, 1, ..., T$ by this proposition and since the log-determinant function is strictly concave, the objective function in (24) as the sum of four convex functions is convex in Σ_t and Π_t. The dynamic programming problem (24) is a DCP as the constraints are linear matrix inequalities. We can then apply the software CVX to derive numerical solutions. Notice that $\mathcal{V}_t(\Sigma_{t-1})$ also satisfies a dynamic programming equation. But we do not solve it directly because $\mathcal{V}_t(\Sigma_{t-1})$ may not be convex as it is equal to the sum of a convex function $J_t(\Sigma_{t-1})$ and a concave function by (21).

The additional assumption of $W_t \succ 0$ in Proposition 2 can be relaxed. In Online Appendix D we assume that A_t is invertible when W_t is singular. Both assumptions ensure an SDP representation of the dynamic RI problem and the latter is satisfied for VAR(p) and ARMA(p, q) processes ($p > q$). But both are violated for MA processes that satisfy only the weaker invertibility assumption of $A_t A_t' + W_t$. In Online Appendix D we discuss how our approach can work under this weaker assumption and apply our approach to general ARMA and MA processes. In Online Appendix G we develop related algorithms. In particular, Algorithm 5 is our preferred value function based method, which allows for the weakest assumption.

We next derive the first-order conditions and characterize the solution. Afrouzi and Yang (2019) use a different method to derive these results for a subset of our general control problems (pure tracking problems in Section 4). Our approach applies to the case with endogenous state variables. We also provide a different proof and a new interpretation based on a sequence of static RI problems.

Because the dynamic semidefinite programs in (19) and (20) are convex, the following Kuhn-Tucker conditions are necessary and sufficient for optimality:

$$\frac{\lambda}{2}\Sigma_t^{-1} = \Theta_t + \Lambda_t, \tag{25}$$

$$\Lambda_t \bullet \left(\Sigma_{t|t-1} - \Sigma_t\right) = 0, \quad \Sigma_{t|t-1} \succeq \Sigma_t, \tag{26}$$

for $t = 0, 1, ..., T$, where $\bullet$ denotes the trace inner product for the space of positive semidefinite matrices,[12] $\Lambda_t \succeq 0$ is the Lagrange multiplier associated with the no-forgetting constraint in period t, $\Sigma_{t|t-1}$ is the prior covariance matrix satisfying (14), and

$$\Theta_t \equiv \Omega_t + \frac{\beta\lambda}{2} A_t' \Sigma_{t+1|t}^{-1} A_t + \beta \frac{\partial J_{t+1}\left(\Sigma_t\right)}{\partial \Sigma_t}, \ 0 \le t \le T-1, \ \Theta_T = \Omega_T. \tag{27}$$

The envelope condition gives

$$\frac{\partial J_t\left(\Sigma_{t-1}\right)}{\partial \Sigma_{t-1}} = -A_{t-1}' \Lambda_t A_{t-1}. \tag{28}$$

Plugging (28) into (27) yields

$$\Theta_t = \Omega_t + \frac{\beta\lambda}{2} A_t' \Sigma_{t+1|t}^{-1} A_t - \beta A_t' \Lambda_{t+1} A_t, \ 0 \le t \le T-1, \ \Theta_T = \Omega_T. \tag{29}$$

The left side of the first-order condition (25) represents the marginal cost of information acquisition and the right side represents the associated marginal benefit. The Lagrange multiplier Λ_t is related to the shadow value by (28) and satisfies the complementary slackness condition (26). The term Θ_t incorporates both the current benefit Ω_t and the future benefit from the reduction of uncertainty as the future prior belief will be revised given currently acquired information.

To characterize the above system, we notice that the dynamic programming problem (20) in each period t can be viewed as a static RI problem like (19), where the marginal benefit of information is given by Θ_t and the prior covariance matrix is given by $\Sigma_{t|t-1}$. In a previous version of our paper circulated in 2018, we developed a generalized reverse water-filling solution for the static case, which is also a special case of Proposition 4. We restate this result in Lemma 3 of Online Appendix B. Applying this lemma, we immediately obtain the following result:[13]

[12] The trace inner product is defined as $A \bullet B = \mathrm{tr}(AB)$ for any positive semidefinite matrices A and B.

[13] Let $X^{1/2}$ denote the square root of any positive semidefinite matrix X. Following the Matlab operation, $\max(A, B)$ ($\min(A, B)$) for any equal sized matrices A and B denotes the matrix with the largest (smallest) elements taken from A and B.

Proposition 3 *Suppose that* $W_t \succeq 0$ *and* $A_t A_t' + W_t \succ 0$ *for all* $t = 0, 1, ..., T-1$. *Perform the eigendecomposition*

$$\Sigma_{t|t-1}^{\frac{1}{2}} \Theta_t \Sigma_{t|t-1}^{\frac{1}{2}} = U_t D_t U_t', \tag{30}$$

where U_t *is an orthogonal matrix and* D_t *is a diagonal matrix of eigenvalues. Then the sequence of optimal posterior covariance matrices for the finite-horizon RI problem satisfies*

$$\Sigma_t = \Sigma_{t|t-1}^{\frac{1}{2}} U_t \left[\max \left(\frac{2}{\lambda} D_t, I \right) \right]^{-1} U_t' \Sigma_{t|t-1}^{\frac{1}{2}}, \quad \Sigma_{0|-1} = \Sigma_{-1} \succ 0 \text{ given}, \tag{31}$$

for $t = 0, 1, ..., T$, *where* $\Sigma_{t+1|t}$ *satisfies (14) and* Θ_t *satisfies*

$$\Theta_t = \Omega_t + \beta A_t' \Sigma_{t+1|t}^{-\frac{1}{2}} U_{t+1} \min \left(D_{t+1}, \frac{\lambda}{2} I \right) U_{t+1}' \Sigma_{t+1|t}^{-\frac{1}{2}} A_t, \quad \Theta_T = \Omega_T, \tag{32}$$

for $t = 0, 1, ..., T-1$. *Moreover, these conditions are sufficient for optimality.*

The sufficiency part follows from Lemma 2 and Proposition 2. Using (25) and (31), we can immediately derive the Lagrange multiplier

$$\Lambda_t = \Sigma_{t|t-1}^{-\frac{1}{2}} U_t \max \left(\frac{\lambda}{2} I - D_t, 0 \right) U_t' \Sigma_{t|t-1}^{-\frac{1}{2}}.$$

Thus the no-forgetting constraint binds (i.e., $\Lambda_t \succ 0$) whenever all eigenvalues in D_t are less than 0.5λ. In this case the decision maker will not acquire any information so that $\Sigma_t = \Sigma_{t|t-1}$. When all eigenvalues in D_t are greater than 0.5λ, we have $\Lambda_t = 0$ and $\Sigma_t = 0.5\lambda\Theta_t^{-1}$ by (30) and (31). In this case the no-forgetting constraint does not bind.

To better understand intuition, suppose that the prior covariance matrix $\Sigma_{t|t-1}$ is diagonal and Θ_t is an identity matrix I. Then $D_t = \Sigma_{t|t-1}$ and $U_t = I$ by (30). We have $\Sigma_t = \min\left(\Sigma_{t|t-1}, 0.5\lambda I\right)$ by (31). The decision maker acquires information to reduce any prior variance greater than 0.5λ to a posterior variance of 0.5λ. They do not pay attention to the components of prior variances lower than 0.5λ so that the corresponding posterior variances remain the same. To see how information signals are determined, we apply Proposition 1 to derive the following result:

Corollary 1 *The optimal information structure for the finite-horizon RI problem satisfies*

$$C_t' V_t^{-1} C_t = \Sigma_{t|t-1}^{-\frac{1}{2}} U_t \max \left(0, \frac{2}{\lambda} D_t - I \right) U_t' \Sigma_{t|t-1}^{-\frac{1}{2}}.$$

The signal dimension in period t *is given by the number of eigenvalues in* D_t *greater than* $\lambda/2$.

This corollary shows that the signal dimension may change over time as eigenvalues in D_t are time varying. Starting from any initial prior $\Sigma_{0|-1} \succ 0$, the decision maker revises their beliefs over time by acquiring signals endogenously. We will illustrate this point using numerical examples in Sections 4 and 5.

The system in Proposition 3 facilitates numerical methods for the first-order conditions because the Lagrange multipliers and the complementary slackness conditions are eliminated. In Online Appendix G we propose a backward-forward shooting algorithm. Intuitively, equation (31) is a backward-looking equation for Σ_t that can be solved forward given any initial prior $\Sigma_{0|-1} = \Sigma_{-1}$. Equation (32) is a forward-looking equation for Θ_t that can be solved backward given the terminal value $\Theta_T = \Omega_T$. The system of these two equations can be solved jointly by starting with an initial guess for $\{\Sigma_t\}_{t=0}^{T-1}$ and iterating until convergence.

3.2 Infinite-horizon Case

In the infinite-horizon case, all exogenous matrices A_t, B_t, Q_t, R_t, S_t, and W_t are time invariant. We can derive the solution for the infinite-horizon case by taking the limit of the finite-horizon solution as $T \to \infty$. For numerical implementation, we can apply the VFI method. We present a formal analysis in Online Appendix E, where Proposition 9 establishes a convergence result. Here we sketch the key idea.

Under some stability conditions in the standard control theory, P_t and F_t converge to P and F given in Section 2.1 as $T \to \infty$. By (16), Ω_t converges to

$$\Omega = F'(R + \beta B'PB)F \succeq 0. \tag{33}$$

Moreover, the value functions $J_t(\Sigma_{t-1})$ and $\mathcal{V}_t(\Sigma_{t-1})$ also converge to some time-invariant functions $J(\Sigma_{t-1})$ and $\mathcal{V}(\Sigma_{t-1})$ for any fixed $t \geq 1$ as $T \to \infty$. Let the optimal policy function for problem (20) be $\Sigma_t = h_t(\Sigma_{t-1})$ for a finite T. As $T \to \infty$, h_t converges to a time-invariant function h for any fixed $t \geq 1$. Since the initial no-forgetting constraint (18) is different from (17) for $t \geq 1$, the initial policy function h_0 is different from h.

The infinite-horizon solution can also be characterized by the first-order conditions in Proposition 3 except that the model parameters are replaced by their time-invariant counterpart and the terminal condition $\Theta_T = \Omega_T$ is replaced by a transversality condition. As $t \to \infty$, the solution may converge to a steady state in which $\lim_{t\to\infty} \Sigma_t = \Sigma$, $\lim_{t\to\infty} \Sigma_{t|t-1} = \Sigma_p$, and $\lim_{t\to\infty} \Theta_t = \Theta$. These limits satisfy the following time-invariant system:

$$\Sigma = \Sigma_p^{\frac{1}{2}} U \left[\max\left(\frac{2}{\lambda}D, I\right)\right]^{-1} U'\Sigma_p^{\frac{1}{2}}, \tag{34}$$

$$\Theta = \Omega + \beta A'\Sigma_p^{-\frac{1}{2}} U \min\left(D, \frac{\lambda}{2}I\right) U'\Sigma_p^{-\frac{1}{2}} A, \tag{35}$$

$$\Sigma_p^{\frac{1}{2}}\Theta\Sigma_p^{\frac{1}{2}} = UDU',\ \Sigma_p = A\Sigma A' + W. \tag{36}$$

We can then recover the steady-state information structure (C, V) using

$$C'V^{-1}C = \Sigma_p^{-\frac{1}{2}} U \max\left(0, \frac{2}{\lambda}D - I\right) U'\Sigma_p^{-\frac{1}{2}} \succeq 0. \tag{37}$$

The optimal signal s_t takes the form $s_t = Cx_t + v_t$, where v_t is a Gaussian white noise with covariance matrix V.

By the steady-state version of the Kalman filter (11), (12), and (13), we have

$$\begin{aligned} \widehat{x}_t &= (I - KC)(A - BF)\widehat{x}_{t-1} + K(Cx_t + v_t), && (38) \\ x_{t+1} &= Ax_t - BF\widehat{x}_t + \epsilon_{t+1}, \ t \geq 0, && (39) \end{aligned}$$

where the matrix K is the Kalman gain

$$K \equiv \left(A\Sigma A' + W\right) C' \left[C\left(A\Sigma A' + W\right) C' + V\right]^{-1}. \tag{40}$$

The posterior covariance matrix Σ_t of x_t will stay at Σ for all $t \geq 0$ by (14) and (15), whenever x_0 is drawn from the prior Gaussian distribution with covariance matrix $\Sigma_p = A\Sigma A' + W$. Notice that equations (38) and (39) can be used to generate impulse response functions to an innovation shock at time 1 starting from $x_0 = v_0 = \widehat{x}_{-1} = 0$.

A steady state for $\{\Sigma_t\}$ is a fixed point of the policy function h and can be numerically solved by starting from an initial guess of Σ_p and Θ and iterating equations (34), (35) and (36) until convergence. We can use the backward-forward shooting algorithm to solve for the transition dynamics numerically starting from any initial prior to the steady state by assuming a sufficiently large transition period T and setting the terminal value Θ_T to its steady-state value.

3.3 Some Analytical Results

Analytical results for general multivariate dynamic RI problems are rarely available even in the steady state (see Maćkowiak, Matějka, and Wiederholt (2018) for an exception). While equations (34), (35), and (36) facilitate numerical solution, they are not useful to derive analytical results. It is even unclear whether there exists a solution to these equations. To better understand the nature of the RI solution, in this subsection we provide some analytical steady-state results in the limit as β approaches 1 for the multivariate case and then derive an explicit solution for both the steady state and transition dynamics in the univariate case with $\beta \in (0, 1)$.

Consider the following static problem

$$\min_{\Sigma \succ 0} \ \operatorname{tr}(\Omega\Sigma) + \frac{\lambda}{2}\left[\log\det\left(A\Sigma A' + W\right) - \log\det(\Sigma)\right] \tag{41}$$

subject to

$$\Sigma \preceq A\Sigma A' + W. \tag{42}$$

By Lemma 2, this is a convex program if $AA' + W$ is invertible. We can easily check that the first-order conditions for this problem are the same as the steady-state version of equations (25), (26), and (29) when $\beta = 1$. Thus the solution to this static problem is the same as the steady-state solution to the infinite-horizon RI problem in the limit as β tends to 1. We will provide more discussions on this static problem in Section 6.

Instead of using equations (34), (35), and (36) to characterize the steady-state solution, we apply tools from semidefinite programming to problem (41) to derive analytical results.

Proposition 4 *Suppose that $\Omega \succeq 0$, $W \succ 0$, and $A = \rho I$ in the infinite-horizon RI problem. Perform the eigendecomposition $W^{\frac{1}{2}}\Omega W^{\frac{1}{2}} = U\Omega_d U'$, where U is an orthogonal matrix and $\Omega_d \equiv$ diag$(d_1, ..., d_{n_x})$ is a diagonal matrix of eigenvalues. Then the steady-state posterior covariance matrix for x_t in the $\beta \to 1$ limit is given by*

$$\Sigma = W^{\frac{1}{2}} U \widehat{\Sigma} U' W^{\frac{1}{2}}, \tag{43}$$

where $\widehat{\Sigma} \equiv diag\left(\widehat{\Sigma}_i\right)_{i=1}^{n_x}$ with

$$\widehat{\Sigma}_i = \min\left(\frac{1}{1-\rho^2}, \widehat{\Sigma}_i^*\right), \quad \widehat{\Sigma}_i^* = \frac{1}{2\rho^2}\left(\sqrt{1+\frac{2\rho^2\lambda}{d_i}} - 1\right), \tag{44}$$

if $|\rho| < 1$; and $\widehat{\Sigma}_i = \widehat{\Sigma}_i^*$, if $|\rho| \geq 1$ and $\Omega \succ 0$.

This proposition gives a closed-form solution in the sense that all expressions on the right-hand side equation (43) are in terms of the primitive parameters Ω, A, and W. In the IID case with $\rho = 0$ or $A = 0$, the solution is reduced to that for the static problem with $T = 0$, $\Omega = \Omega_0$, and $W = \Sigma_{-1}$. The optimal posterior covariance matrix is given by (43) with $\widehat{\Sigma}_i = \min\left(1, \lambda/\left(2d_i\right)\right)$. This static solution generalizes the standard reverse water-filling solution analyzed by Cover and Thomas (2006) and Kőszegi and Matějka (2019) by allowing for general $\Omega \succeq 0$ and $W \succ 0$. The diagonal matrix $\widehat{\Sigma}$ can be interpreted as a scaling factor for the eigenvalues $\{d_i\}$ of the weighted innovation covariance matrix $W^{\frac{1}{2}}\Omega W^{\frac{1}{2}}$. The attention is allocated according to a decreasing order of $\{d_i\}$, instead of innovation variances. High eigenvalues d_i are scaled down by the factor $\widehat{\Sigma}_i$ for sufficiently small information costs.

Notice that the additional assumption $\Omega \succ 0$ for $|\rho| \geq 1$ is important. The intuition is best understood for the univariate case. If $\Omega = 0$, then there is no benefit from reducing uncertainty. The decision maker will not acquire information so that the posterior variance will explode without violating the no-forgetting constraint for $|\rho| \geq 1$.

What kind of signal structure can generate the optimal covariance matrix Σ in Proposition 4? Let the signal be $s_t = Cx_t + v_t$, where v_t is a Gaussian white noise with covariance matrix V. Using equation (37), we can recover Φ, C, and V. The following result characterizes the signal structure.

Proposition 5 *Suppose that $\Omega \succeq 0$, $W \succ 0$, and $A = \rho I$ in the infinite-horizon RI problem. Then the steady-state information structure (C, V) in the $\beta \to 1$ limit satisfies*

$$C'V^{-1}C = W^{-\frac{1}{2}} U \ diag\left\{\max\left(0, \frac{2d_i}{\lambda}\left[1 - \left(1-\rho^2\right)\widehat{\Sigma}_i^*\right]\right)\right\}_{i=1}^{n_x} U' W^{-\frac{1}{2}},$$

where U, d_i, and $\widehat{\Sigma}_i^*$ are given in Proposition 4. If $|\rho| < 1$, the signal dimension is equal to the number of d_i such that $0 < \lambda < 2d_i / \left(1-\rho^2\right)^2$ and weakly decreases as λ increases. If $|\rho| \geq 1$ and $\Omega \succ 0$, then the signal dimension is equal to n_x.

This proposition shows that the maximal dimension does not exceed the rank of the matrix Ω (i.e., the number of positive of eigenvalues d_i), which does not exceed the minimum of the state dimension n_x and the control dimension n_u because $\Omega = F'(R + \beta B'PB)F$ in the infinite-horizon LQG control problem, where F is an n_u by n_x matrix. If all states are nonstationary, then the signal dimension must be equal to the state dimension.

In the special case in which Ω has rank 1, we have the following closed-form solution. This case can arise in tracking problems analyzed in the next section.

Proposition 6 *Let G be an n_x-dimensional row vector. Suppose that $\Omega = G'G$, $W \succ 0$, and $A = \rho I$ ($|\rho| < 1$) in the infinite-horizon RI problem. If $\lambda \geq 2\left\|W^{1/2}G'\right\|^2 / \left(1-\rho^2\right)^2$, then no information is acquired in the steady state limit as $\beta \to 1$ and the optimal posterior covariance matrix is given by $\Sigma = W/\left(1-\rho^2\right)$.[14] If $0 < \lambda < 2\left\|W^{1/2}G'\right\|^2 / \left(1-\rho^2\right)^2$, then the steady-state signal in the $\beta \to 1$ limit is one dimensional and can be normalized as*

$$s_t = Gx_t + \left\|W^{1/2}G'\right\| v_t. \tag{45}$$

The variance V of v_t satisfies

$$V^{-1} = \frac{2\left\|W^{1/2}G'\right\|^2}{\lambda}\left[1 - \left(1-\rho^2\right)\widehat{\Sigma}_1^*\right] > 0,$$

and the steady-state posterior covariance matrix Σ for x_t in the $\beta \to 1$ limit is given by

$$\Sigma = \frac{W}{1-\rho^2} - \frac{W\Omega W}{\left\|W^{1/2}G'\right\|^2}\left[\left(1-\rho^2\right)^{-1} - \widehat{\Sigma}_1^*\right],$$

where

$$\widehat{\Sigma}_1^* = \frac{1}{2\rho^2}\left(\sqrt{1 + \frac{2\rho^2\lambda}{\left\|W^{1/2}G'\right\|^2}} - 1\right).$$

Maćkowiak, Matějka, and Wiederholt (2018) use a different approach to derive a similar result for a pure tracking model with information-flow constraints. When $\rho = 0$, Proposition 6 is reduced to the IID case, which is also the static case studied by Fulton (2018).

For the univariate case, we can derive an explicit solution for both the steady state and transition dynamics. This case is studied by Sims (2011) for $|\rho| < 1$. Afrouzi and Yang (2019) solve a similar example with $\rho = 1$. Here we consider a general ρ.

[14] We use $\|\cdot\|$ to denote the Euclidean norm.

Proposition 7 *For the univariate case with $\beta \in (0,1)$, $A = \rho$, $\Omega = 1$, and $W > 0$. Let Σ^* be the unique positive solution to the equation*

$$2\rho^2\Sigma^2 + \left(2W - (1-\beta)\lambda\rho^2\right)\Sigma - \lambda W = 0.$$

Then the optimal posterior variance Σ_t follows the dynamics

$$\Sigma_t = \min\left(\Sigma_{t|t-1}, \widehat{\Sigma}\right), \quad t \geq 0,$$

where $\Sigma_{0|-1} > 0$ is given, $\Sigma_{t|t-1} = \rho^2\Sigma_{t-1} + W$ for $t \geq 1$, and the steady state is given by

$$\widehat{\Sigma} = \begin{cases} \min\left(W/\left(1-\rho^2\right), \Sigma^*\right) & \text{if } |\rho| < 1 \\ \Sigma^* & \text{if } |\rho| \geq 1 \end{cases}.$$

The decision maker acquires a signal $s_t = x_t + v_t$ for $t \geq t_0$ where t_0 is the first time t such that $\widehat{\Sigma} < \Sigma_{t|t-1}$ and v_t is a Gaussian white noise with variance V_t satisfying

$$V_t^{-1} = \widehat{\Sigma}^{-1} - \Sigma_{t|t-1}^{-1}.$$

In the case of $|\rho| < 1$, we can check that if the information cost parameter λ is sufficiently small, then the steady-state optimal posterior variance is reduced from the stationary prior variance $W/\left(1-\rho^2\right)$ to a smaller variance $\widehat{\Sigma}$. But if λ is sufficiently large, then no information is collected and $\widehat{\Sigma} = W/\left(1-\rho^2\right)$. If $|\rho| \geq 1$, the state process x_t is nonstationary ex ante. But its estimate has a stationary variance after the decision maker acquires costly information to reduce uncertainty. Starting from a small prior variance $\Sigma_{0|-1}$, the no-forgetting constraint binds and the posterior Σ_t grows with the prior $\Sigma_{t|t-1} = \rho^2\Sigma_{t-1} + W$ over time as $|\rho| \geq 1$ and the state innovation variance W is added to the prior. In this case the decision maker does not acquire information. After some transition periods, the decision maker acquires information to reduce uncertainty. Eventually Σ_t stays at the steady state Σ^* forever and the no-forgetting constraint never binds as $\rho^2\Sigma^* + W > \Sigma^*$.

As discussed earlier, the information acquisition problem can be interpreted as an investment problem. The discount factor β is important to determine the intertemporal cost and benefit tradeoff. Intuitively, a more patient decision maker (higher β) has a higher incentive to invest in the stock of knowledge and hence they acquire information earlier. Formally, it is straightforward to show that Σ^* decreases with β. The steady-state limit as $\beta \to 1$ coincides with Proposition 4.

4 Tracking Problems

We now turn to the special case of pure tracking problems similar to that in Sims (2011). In addition to being interesting in their own right, these problems provide a simplified case that illustrates our general framework.

Suppose that the state vector x_t and the target y_t have a state space representation:

$$x_{t+1} = Ax_t + \eta_{t+1}, \quad y_t = Gx_t,$$

where G is a conformable matrix, x_0 is Gaussian with mean $\bar{x}_0$ and covariance matrix $\Sigma_{0|-1} = \Sigma_{-1} \succ 0$, and η_{t+1} is a Gaussian white noise with covariance matrix W. Unlike the general control problem, the state vector x_t follows an exogenous process. The decision maker does not observe x_t and wants to keep an action z_t close to y_t with a quadratic loss, given their observation of histories of signals s^t. The signal s_t satisfies (2) with $T = \infty$. The decision maker selects an optimal information structure before choosing z_t by paying an information cost of λ per nat.

Let Σ_t denote the posterior covariance matrix of x_t given information s^t. We formulate the tracking problem with discounted information costs as follows:

Problem 4 *(Tracking problem with discounted information costs)*

$$\min_{\{z_t\},\{\Sigma_t\}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[(y_t - z_t)' (y_t - z_t) + \lambda I \left(x_t; s_t | s^{t-1} \right) \right]$$

subject to (18),

$$I \left(x_0; s_0 | s^{-1} \right) = \frac{1}{2} \log \det \left(\Sigma_{-1} \right) - \frac{1}{2} \log \det \left(\Sigma_0 \right),$$

$$I \left(x_t; s_t | s^{t-1} \right) = \frac{1}{2} \log \det \left(A \Sigma_{t-1} A' + W \right) - \frac{1}{2} \log \det \left(\Sigma_t \right),$$

$$\Sigma_t \preceq A \Sigma_{t-1} A' + W, \tag{46}$$

for $t \geq 1$.

As is well known, it is optimal to set $z_t = G\mathbb{E}\left[x_t | s^t\right]$. Thus $\mathbb{E}\left[(y_t - z_t)'(y_t - z_t)\right] = \operatorname{tr}(G'G\Sigma_t)$ and this problem becomes an infinite-horizon version of Problem 3 with $\Omega = G'G$. The analysis in Section 3 applies. Afrouzi and Yang (2019) study Problem 4 and derive first-order conditions and steady-state solution similar to our Proposition 3. They also develop related algorithms that precede ours. Our characterizations apply to a much wider class of control problems. Unlike ours, their derivation relies on the simultaneous diagonalization of Λ_t and $\Sigma_{t|t-1} - \Sigma_t$. Like us, they also apply the eigendecomposition of a special weighted prior covariance matrix in (30). They do not show the convexity of the minimization problem and the sufficiency of the first-order conditions for optimality. Maćkowiak, Matějka, and Wiederholt (2018) propose a different approach to solve tracking problems with one control under information-flow constraints for general ARMA processes.

In the special case in which G is an n_x-dimensional row vector, the rank of $\Omega = G'G$ is one. Proposition 6 provides an explicit solution for the steady state when $\beta = 1$ and when all states have the same persistence parameter ρ, but innovations are arbitrarily correlated. If these assumptions are relaxed, we are unable to derive analytical results.

Numerical example. We now study an example taken from Sims (2011) using numerical methods. This example can be interpreted as a single firm's price setting problem adapted from Maćkowiak and Wiederholt (2009).[15] We use this example to illustrate our computation tools and show how these tools allow exploring an agent's behavior in tracking problems, as the information cost, shock persistence, or the discount factor varies.

Let x_t represent a vector of exogenous aggregate and idiosyncratic shocks, $y_t = [1,1]\, x_t$ the full information profit-maximizing price, and z_t the optimal price under RI. We use the same baseline parameter values as in Sims (2011):[16]

$$A = \begin{bmatrix} 0.95 & 0 \\ 0 & 0.4 \end{bmatrix}, \; W = \begin{bmatrix} 0.0975 & 0 \\ 0 & 0.84 \end{bmatrix}, \; G = [1,1]\,, \; \beta = 0.9, \; \lambda = 2. \tag{47}$$

Given these parameter values, the two shocks have identical stationary unconditional variance of 1. Proposition 6 shows that, if the two shocks have equal persistence, then the steady-state optimal signal for $\beta = 1$ is one dimensional and takes the form of the profit-maximizing price y_t plus a noise, independent of the innovation covariance matrix. Using numerical methods, we find that this result still holds for the steady-state solution with $\beta \in (0,1)$. We will show that if the two shocks have different persistence, then the steady-state optimal signal is still one dimensional, but takes a different form.

Computing the solution. Using the algorithm described in the first paragraph of Online Appendix G.3 based on the first-order conditions (34), (35), and (36), it takes about 0.038 seconds for a PC with Intel Core i5-9500 CPU and 16GB memory to compute the steady-state posterior covariance matrix within error 10^{-8} :

$$\Sigma = \begin{bmatrix} 0.3571 & -0.1725 \\ -0.1725 & 0.7828 \end{bmatrix}. \tag{48}$$

The steady-state signal takes the normalized form $s_t = [1.3778, 1]\, x_t + v_t$, where v_t is a Gaussian white noise with variance 2.6149. The optimal signal puts more weight on the slow-moving component (aggregate shock). Because the signal weights on the two shocks are positive, conditional on a given signal value, a positive shock to one state must be associated with a negative shock to the other state. Thus the two states are negatively correlated conditional on the optimal signal.

We use the VFI method (Algorithm 2 of Online Appendix G.2) to solve this example and find the same numerical solution. It takes about 160 seconds for the same PC to get convergence of Σ_t within error 10^{-6}, starting from an initial prior covariance matrix of $0.5W$. In Online Appendix G.2 we propose two modified VFI methods (Algorithms 3 and 5) which take about 6 seconds to get the steady-state solution within error 10^{-6}. As is well known in the literature, value function based

[15] See Woodford (2003, 2009) for related pricing models.

[16] The parameter λ in our paper corresponds to 2λ in Sims (2011). Chris Sims informed us that the variance of the slow-moving component should be 0.84 and the value of 0.86 in Sims (2011) is a typo.

methods are much slower than first-order conditions based methods. But the former methods are more reliable as we have a convergence proof in Proposition 9 of Online Appendix E. They are also more flexible and can incorporate many occasionally binding constraints and nonsmooth objective functions. See Online Appendix G.4 for an example that can be solved by our VFI methods, but cannot be solved by the first-order conditions (34), (35), and (36).

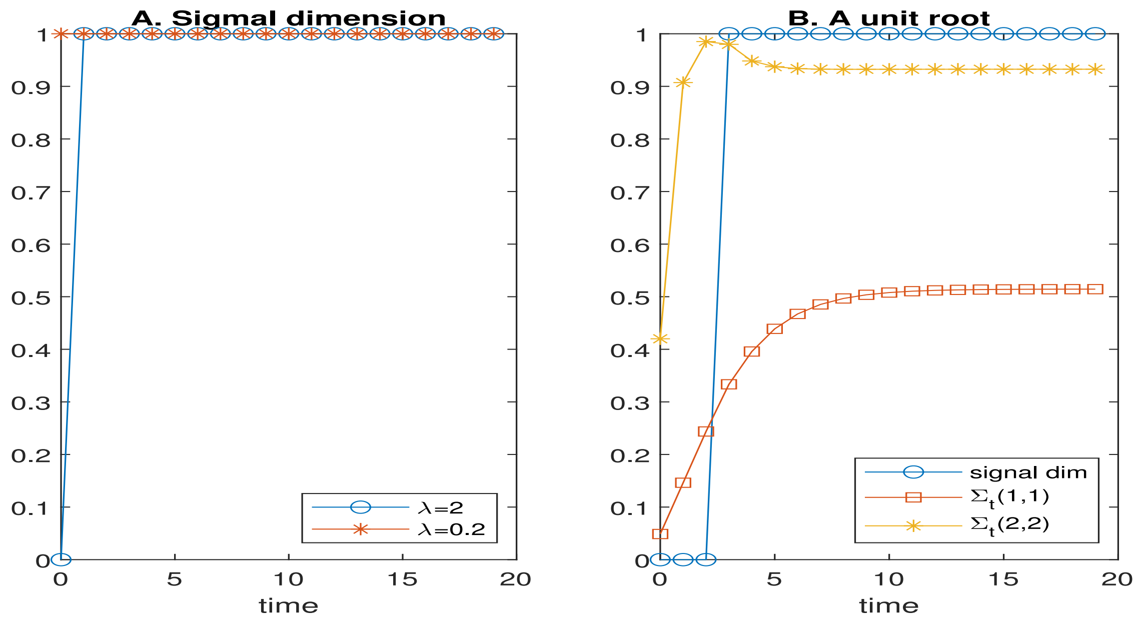

Figure 1: Panel A displays the dynamics of signal dimension for two values of λ. Panel B displays the dynamics of signal dimension and posterior variances of the two state components when the first component has a unit root.

It takes about 0.056 seconds for our backward-forward shooting method (Algorithm 7 of Online Appendix G.3) based on the first-order conditions to find the transition dynamics within error 10^{-6}, starting from an initial prior covariance matrix of $0.5W$. Figure 1 panel A displays the optimal signal dimension at time $t = 0, 1, ..., 19$, which is determined by the number of the eigenvalues in D_t greater than 0.5λ by Proposition 3, or the rank of the SNR $\Phi_t = \Sigma_t^{-1} - \Sigma_{t|t-1}^{-1}$ by Proposition 1. This figure shows that the firm does not acquire information initially and acquires a one-dimensional signal from period 1 on.

Information cost. For $\lambda = 0.2$ and other parameter values held fixed as in (47), Figure 1 panel A shows that the firm acquires a one-dimensional signal immediately at $t = 0$. Intuitively, a smaller information cost λ induces the firm to acquire information earlier. The optimal posterior covariance matrix Σ_t arrives at the steady state in 26 periods starting from the prior $0.5W$:

$$\Sigma = \begin{bmatrix} 0.3161 & -0.3001 \\ -0.3001 & 0.3819 \end{bmatrix}.$$

The normalized signal weight vector becomes $[1.0314, 1]$ and the signal is less noisy with its innovation variance of 0.1091. Compared with (48), the posterior variance of the slow-moving component does not change much, but the posterior variance of the fast-moving component is reduced significantly. As in Sims (2011), news about the fast-moving component is perceived fairly promptly when λ declines, while there is little immediate reaction to news about the slow-moving component.

Shock persistence. Next we consider the impact of the shock persistence. We raise the persistence of the slow-moving component from 0.95 to 0.98 and decrease the innovation variance from 0.0975 to 0.0396 to keep the unconditional variance at the same value of 1. We also fix other parameter values as in (47). We find the steady-state posterior covariance matrix

$$\Sigma = \begin{bmatrix} 0.2488 & -0.1197 \\ -0.1197 & 0.7882 \end{bmatrix},$$

the normalized signal weight vector $[1.4842, 1]$, and the signal noise variance 3.0764. As the signal weight on the slow-moving component is higher, the firm allocates more attention to the component when its persistence becomes higher, even though the prior unconditional variance remains the same. Compared with (48), the steady-state posterior variance of the slow-moving component decreases sharply from 0.3571 to 0.2488, while the steady-state posterior variance of the fast-moving component does not change much. This implies that the price responds faster to the more persistent shock when the unconditional variances of the two shocks are the same. The price responses are the same in the symmetric case when they have the same persistence.

The analytical solutions in Propositions 4, 5, and 7 for equally persistent states show that RI can make the posterior covariance matrix of a nonstationary process stationary conditional on endogenously acquired information if $\Omega \succ 0$. For the pricing model here, $\Omega = G'G$ is singular. We can numerically check that there is no steady-state solution if both states have a unit root. We then assume that only the aggregate shock (slow-moving component) contains a unit root. Setting $\lambda = 4$ and keeping other parameter values fixed as in (47), we numerically compute the steady state and transition dynamics given the initial prior covariance matrix $\Sigma_{0|-1} = 0.5W$. Figure 1 panel B presents the signal dimensions and posterior variances of the two states against time. We find that the no-forgetting constraints bind in that $\Sigma_t = \Sigma_{t|t-1}$ for $t = 0, 1, 2$. During these periods, the firm does not acquire any information. As $\Sigma_{t+1|t} = A\Sigma_t A' + W$, both Σ_t and $\Sigma_{t+1|t}$ grow over time. In period $t = 3$, the cost of uncertainty becomes so high that the firm has an incentive to acquire a one-dimensional signal to reduce uncertainty. Then the growth of Σ_t and $\Sigma_{t+1|t}$ is stabilized and Σ_t reaches the steady state at $t = 20$.

Discount factor. Finally, we consider the impact of the discount factor β. As shown in Proposition 7 for the univariate case, a higher β leads to earlier information acquisition. For the two

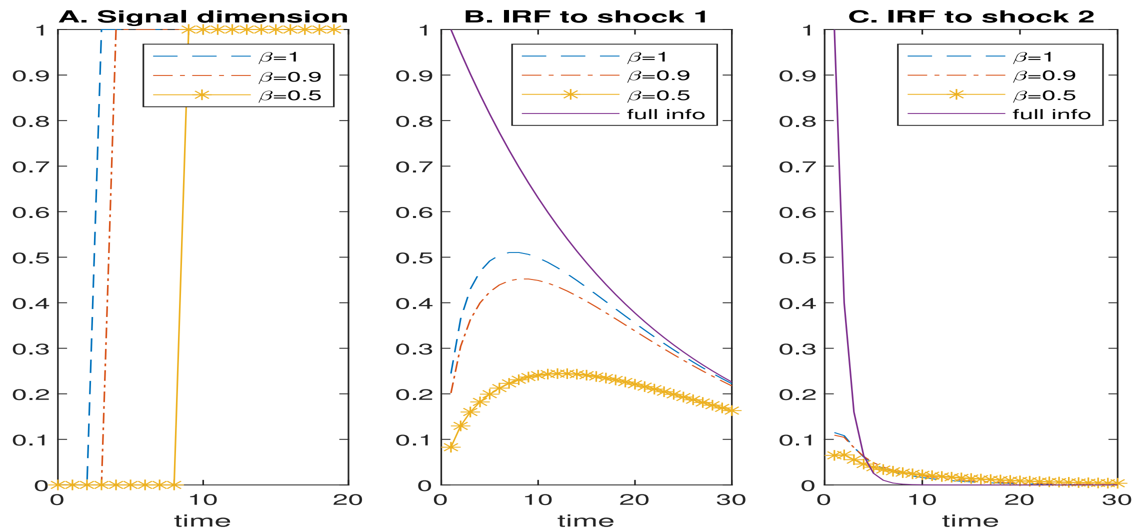

Figure 2: Impact of the discount factor. Panel A plots signal dimension against time. Panels B and C plot impulse response functions of the price to a unit size innovation shock to the slow-moving (shock 1) and fast-moving (shock 2) components, respectively.

dimensional example here, we solve for the transition dynamics for $\beta = 1$, 0.9, and 0.5 given an initial prior covariance matrix $\Sigma_{0|-1} = 0.5W$. We set $\lambda = 4$ and fix other parameter values as in (47). Figure 2 panel A shows that the firm acquires a one-dimensional signal at $t = 3$ for $\beta = 1$, $t = 4$ for $\beta = 0.9$, and $t = 9$ for $\beta = 0.5$. By the steady-state Kalman filter, the impulse responses of price y_t to an innovation shock of one unit size to each of the two state components jump higher on impact and peak at a higher value for a larger value of β (see Figure 2 panels B and C).[17] Compared to the solution under full information, RI generates dampened and delayed responses. Moreover, the initial price responds faster to a slow-moving shock, even though its innovation variance is smaller. Intuitively, the firm pays more attention to the slow-moving component as discussed above, because learning about it is more useful to predict the future value of the shock.

5 Applications

In this section we study three applications to illustrate our results using our RI Matlab toolbox. We analyze a pure tracking problem in an equilibrium setting in the first application and dynamic control problems in the other two. In the first application there are two exogenous states and one control. In the second application there are one endogenous and two exogenous states and one control. In the last application there are two endogenous and two exogenous states and two controls. For all applications we focus on the steady-state solution for the optimal information structure

[17] All impulse response functions in this paper are computed using the steady-state Kalman filter.

discussed in Section 3.2.[18] We also study the transition dynamics for the last two applications.

5.1 Equilibrium Sticky Prices

We extend the pricing problem in Section 4 to an equilibrium setting as in Maćkowiak and Wiederholt (2009). Here we present the key equilibrium conditions directly and refer the reader to their paper for detailed derivations and interpretations. We drop their signal independence assumption. They argue that this assumption is more realistic for firms and can help their model match data reasonable well. They also discuss ways to relax it. Our purpose is to illustrate our numerical methods without this assumption and highlight additional insights.[19]

Consider an economy with a continuum of firms indexed by $j \in [0,1]$. Firm j sells good j and sets its prices to maximize the present discounted value of profits. The full-information profit-maximizing price is given by

$$p_{jt}^{*} = (1-\alpha_2)\, p_t + \alpha_2 q_t + \alpha_3 z_{jt}, \tag{49}$$

where p_t is the aggregate price level, q_t is nominal aggregate demand, and z_{jt} represents an idiosyncratic shock. The parameter $\alpha_2 \in (0,1]$ describes the degree of strategic complementarity. Suppose that z_{jt} and q_t follow exogenous AR(1) processes

$$\begin{aligned} z_{jt} &= \rho_i z_{j,t-1} + \epsilon_{jt}, \quad 0 < \rho_i < 1, \\ q_t &= \rho_a q_{t-1} + \epsilon_{at}, \quad 0 < \rho_a < 1, \end{aligned}$$

where ϵ_{jt} and ϵ_{at} are independent Gaussian white noise processes with variances σ_i^2 and σ_a^2. Assume that z_{jt} is also independent across firms $j \in [0,1]$ such that $\int \epsilon_{jt} dj = 0$.

Each firm j does not observe q_t and z_{jt}. It acquires an optimal signal vector s_{jt} about a vector x_{jt} of unobserved states subject to discounted entropy information costs. To fit in the framework of Section 4, assume that the vector of states x_{jt} and the target p_{jt}^{*} have a state space representation. We will specify the state vector x_{jt} later.

Firm j sets price p_{jt} to track p_{jt}^{*} subject to entropy information costs. It solves the following tracking problem under RI:

$$\min_{\{p_{jt}, C_{jt}, V_{jt}\}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \alpha_1 \left[\left(p_{jt} - p_{jt}^{*}\right)^2 \right] + \lambda \sum_{t=0}^{\infty} \beta^t I\left(x_{jt}; s_{jt} | s_j^{t-1}\right), \tag{50}$$

subject to no-forgetting constraints, where $\alpha_1 > 0$, $s_{jt} = C_{jt} x_{jt} + v_{jt}$, and v_{jt} is a Gaussian white noise with covariance matrix V_{jt}. Then the optimal price under RI is given by $p_{jt} = \mathbb{E}\left[p_{jt}^{*} | s_j^t\right]$.

[18] We have used several different methods to solve all applications and get the same results up to small numerical errors. The first-order conditions based method is the fastest. It takes about 2, 2, and 0.2 seconds on average to respectively solve for the equilibrium pricing problem, the consumption problem, and the investment problem for a wide range of parameter values. For the first problem A is invertible, while for the last two problems W is invertible.

[19] Maćkowiak, Matějka, and Wiederholt (2018) solve a similar model of Woodford (2003) without idiosyncratic shocks.

Assume that v_{jt} is independent of all other shocks, and is independent across firms $j \in [0,1]$ such that $\int v_{jt} dj = 0$. The model is closed by the equilibrium condition:

$$p_t = \int_0^1 p_{jt} dj. \tag{51}$$

In the analysis below, we normalize $\alpha_1 = 1$. Let Σ_{jt} denote the posterior covariance matrix of the state x_{jt}. We focus on the steady-state symmetric equilibrium in which $\Sigma_{jt} = \Sigma$, $C_{jt} = C$, and $V_{jt} = V$ for all j and t.

5.1.1 No Strategic Complementarity

When there is no strategic complementarity ($\alpha_2 = 1$), we have $p_{jt}^* = q_t + \alpha_3 z_{jt}$. Then there is no equilibrium price feedback to individual pricing decisions. After defining the state vector as $x_{jt} = (z_{jt}, q_t)'$, we obtain the state space representation: $p_{jt}^* = G x_{it}$, $G = (\alpha_3, 1)$,

$$x_{jt} = A x_{j,t-1} + \begin{bmatrix} \epsilon_{jt} \\ \epsilon_{at} \end{bmatrix}, \; A = \begin{bmatrix} \rho_i & 0 \\ 0 & \rho_a \end{bmatrix}, \; W = \begin{bmatrix} \sigma_i^2 & 0 \\ 0 & \sigma_a^2 \end{bmatrix}.$$

The problem (50) becomes a single firm's pricing problem under RI studied in Section 4.

Firm j's optimal price under RI is given by

$$p_{jt} = \mathbb{E}\left[p_{jt}^* | s_j^t\right] = G\mathbb{E}\left[x_{jt} | s_i^t\right] = G\widehat{x}_{jt}, \tag{52}$$

where $\widehat{x}_{jt}$ satisfies the Kalman filter:

$$\widehat{x}_{jt} = (I - KC)\, A\widehat{x}_{j,t-1} + K\,(Cx_{jt} + v_{jt})\,, \tag{53}$$

for $t \geq 0$, with $\widehat{x}_{j,-1} = 0$, where K satisfies (40). Unlike (38) and (39) in the optimal control case, there is no control feedback in (53).

Equations (52) and (53) show that individual price responses p_{jt} to shocks through s_{jt} are determined by two effects for a given G: (i) the learning effect reflected by the term KC, and (ii) the attention allocation effect reflected by the optimal choice of information structure Σ or (C, V).

In Online Appendix F we show that the equilibrium aggregate price satisfies

$$p_t = \int_0^1 p_{jt} dt = G\int_0^1 \widehat{x}_{jt} dj = G\left[I - (I - KC)A\mathbf{L}\right]^{-1} KC(I - A\mathbf{L})^{-1}\,[0,1]'\,\epsilon_{at},$$

where $\mathbf{L}$ represents the lag operator and $KC = I - \Sigma\,(A\Sigma A' + W)^{-1}$.

When $\rho_i = \rho_a$ and $\beta = 1$, Proposition 6 applies and the steady-state optimal signal can be normalized as the profit-maximizing price plus a noise (i.e., $s_{jt} = p_{jt}^* + v_{jt}$). This signal form implies that the impulse responses of individual prices to the idiosyncratic shock z_{jt} are larger than to the aggregate shock q_t if and only if it carries a larger weight α_3 as shown in equations (52) and

(53). The individual price responses are the same when $\alpha_3 = 1$. This result is independent of the dimension of states and the innovation covariance matrix W.

When $\rho_i \neq \rho_a$ and $\beta \in (0,1)$, based on numerical solutions for a wide range of parameter values, we find that the steady-state optimal signal is still one dimensional, but it does not take the normalized form of the profit-maximizing price plus a noise. If $\alpha_3 = 1$, then the relative size of the initial individual price responses to the two shocks is determined by the attention allocation effect, i.e, the signal weight vector C, as in the pricing example of Section 4. The comparative statics analysis in Section 4 applies to individual firm prices. Instead of repeating it here, we turn to the more interesting case with strategic complementarity.

5.1.2 Strategic Complementarity

When there is strategic complementarity, i.e., $\alpha_2 \in (0,1)$, there is equilibrium price feedback in (49). The equilibrium solution becomes more involved due to higher-order beliefs. The state vector x_{jt} must contain endogenous variables that incorporate the equilibrium aggregate price information. We present the technical details in Online Appendix F.

We focus on the equilibrium in which the aggregate price p_t follows a causal stationary process, which has an MA(∞) representation. As in Maćkowiak, Matějka, and Wiederholt (2018), we approximate such an equilibrium by a stationary ARMA(r,m) process $p_t = \Psi(\mathbf{L})\epsilon_{at}$ for a large enough $r \geq m+1$,[20] where

$$\Psi(\mathbf{L}) \equiv \frac{b_0 + b_1\mathbf{L} + b_2\mathbf{L}^2 + \cdots + b_m\mathbf{L}^m}{1 - a_1\mathbf{L} - a_2\mathbf{L}^2 - \cdots - a_r\mathbf{L}^r}. \tag{54}$$

All coefficients in the rational function Ψ and the order (r and m) are endogenous with $a_r \neq 0$ and $b_m \neq 0$. Notice that the equilibrium aggregate price p_t contains only aggregate innovations ϵ_{at}, because idiosyncratic innovations ϵ_{jt} wash out in the aggregate.

We adopt the following state space representation (Hamilton (1994)):

$$x_{jt} = \begin{bmatrix} \rho_i & 0 & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \rho_a & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & a_1 & a_2 & \cdots & \cdots & a_{r-1} & a_r \\ 0 & 0 & 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix} x_{j,t-1} + \begin{bmatrix} \epsilon_{jt} \\ \epsilon_{at} \\ \epsilon_{at} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \tag{55}$$

$$p^*_{jt} = Gx_{jt}, \quad G = [\alpha_3, \ \ \alpha_2, \ \ (1-\alpha_2)D], \ D = \begin{bmatrix} b_0 & b_1 & \cdots & b_{r-2} & b_{r-1} \end{bmatrix}, \tag{56}$$

where the state vector $x'_{jt} = \left[z_{jt}, q_t, \xi'_t\right]$ consists of the exogenous states z_{jt}, q_t, and an endogenous r-dimensional state (column) vector ξ_t such that we can write $p_t = D\xi_t$. Moreover, we set $b_{m+1} =$

[20] This assumption ensures the state transition matrix A constructed in equation (55) is invertible.

$b_{m+2} = ... = b_{r-1} = 0$. Let the $(r+2) \times 1$ noise vector be $\eta_{jt} \equiv [\epsilon_{jt}, \epsilon_{at}, \epsilon_{at}, 0, ..., 0]'$. Then η_{jt} is a Gaussian white noise and its covariance matrix W is singular. Let A denote the $(r+2) \times (r+2)$ transition matrix in equation (55). We can check that A is invertible.

We solve individual pricing problem under RI with $\Omega = G'G$ and derive the steady-state information structure. After aggregating individual optimal prices using (51) and (52), we obtain a fixed point problem for the coefficients $(a_1, a_2, ..., a_r, b_0, b_1, ..., b_m)$. In Online Appendix F we describe an algorithm to solve this fixed point problem and determine the endogenous r and m. Then we can determine the equilibrium aggregate price function and individual pricing rules.

We set baseline parameter values as follows: $\beta = 0.95$, $\lambda = 0.002$, $\rho_i = \rho_a = 0.95$, $\sigma_i = 10\%$, $\sigma_a = 1\%$, $\alpha_1 = \alpha_3 = 1$, and $\alpha_2 = 0.15$. For these parameter values we find that an ARMA(2,1) process is a good approximation of the equilibrium aggregate price p_t. Then the state vector x_{jt} is $r + 2 = 4$ dimensional. We find that the steady-state optimal signal vector s_{jt} is one dimensional and takes the form

$$s_{jt} = 0.8552 z_{jt} + 0.1283 q_t + 0.5021 \xi_{1t} - 0.0110 \xi_{2t} + v_{jt},$$

where v_{jt} is a Gaussian white noise with variance 0.0741. The signal assigns weights to the endogenous state $\xi_t = (\xi_{1t}, \xi_{2t})'$ contained in the equilibrium aggregate price. We also find that the optimal signal takes a similar one-dimensional form for all parameter values considered below. This signal form implies that the exogenous aggregate and idiosyncratic shocks (q_t and z_{jt}) are confounded. We will show below that this feature has interesting economic implications.

Now we consider the impact of the information cost λ on the impulse responses of the aggregate equilibrium price to a unit innovation shock to the nominal aggregate demand, shown in Figure 3 Panel A. Under full information, the aggregate price moves one-to-one with the nominal aggregate demand shock so that real output does not change. The responses under RI are dampened and delayed. The higher the information cost λ, the less responsive the aggregate price is.

Panel B of Figure 3 shows the impact of the degree of strategic complementarity α_2. The case with $\alpha_2 = 1$ corresponds to the solution without strategic complementarity studied earlier. As in Maćkowiak and Wiederholt (2009), when the profit-maximizing price is less sensitive to real aggregate demand (i.e., when α_2 is lower), the response of the price level to a nominal demand shock is more dampened. The reason is that the price feedback effects are stronger.

Next we study the impact of innovation volatilities presented in Figure 4. Under the signal independence assumption, Maćkowiak and Wiederholt (2009) find when the innovation variance of a shock increases, firms shift attention toward that shock, and away from the other shock. By contrast, Figure 4 shows that when the innovation variance of a shock increases, the individual price responses to both aggregate and idiosyncratic shocks rise. Thus there is a spillover effect similar to that in Mondria (2010). The intuition is that the optimal signal structure implies that

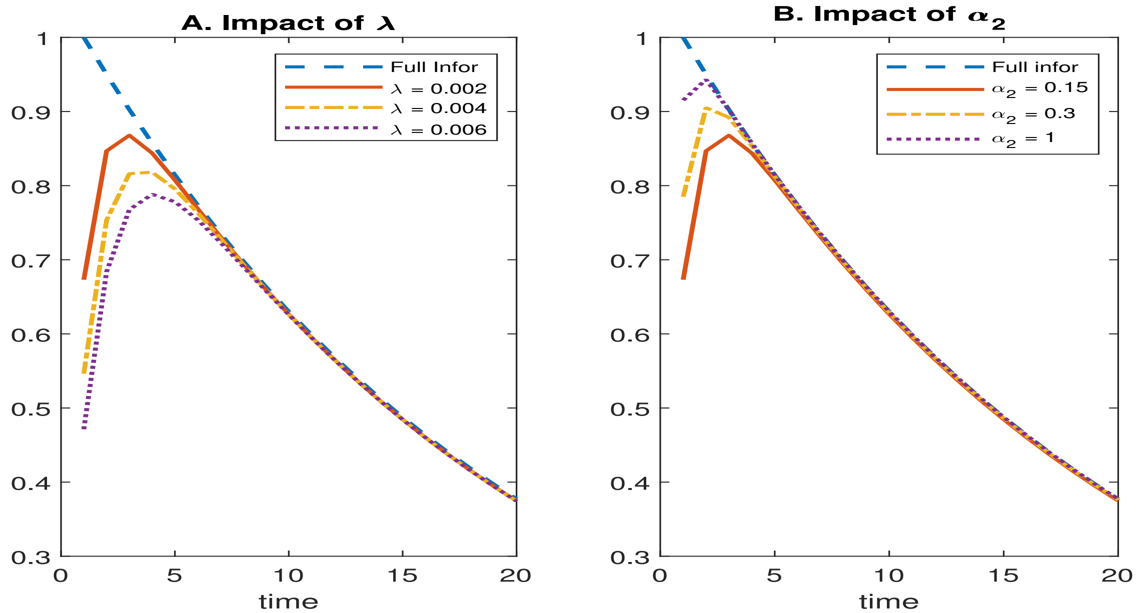

Figure 3: Impulse responses of the aggregate price to a unit size innovation in nominal aggregate demand for the case with strategic complementarity. Panel A shows the impact of information cost. Panel B shows the impact of strategic complementarity.

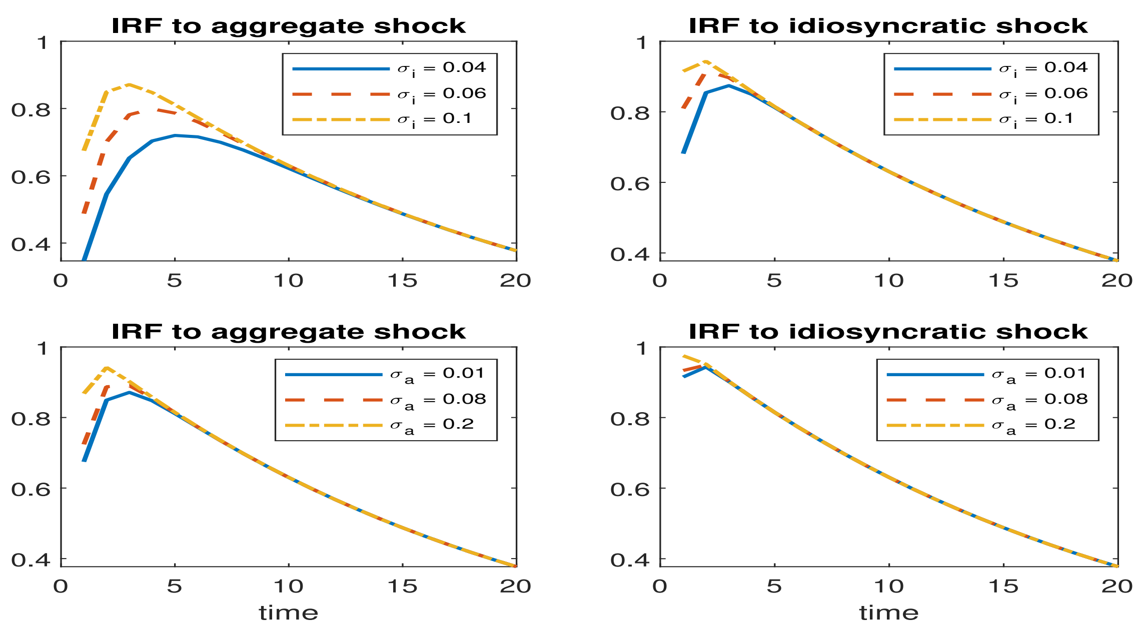

Figure 4: Impulse responses of the individual price to a unit size innovation in nominal aggregate demand and idiosyncratic productivity for different innovation variances in the case with strategic complementarity.

aggregate and idiosyncratic shocks are confounded. The impact of an increase in the innovation variance of one shock is transmitted to the other shock due to the learning effect via the term KC.

Given the signal independence assumption, Maćkowiak and Wiederholt (2009) can match the empirical finding that prices respond much faster to idiosyncratic shocks than to aggregate shocks, while our model without this assumption has difficulty matching this fact quantitatively as shown in Figure 4.

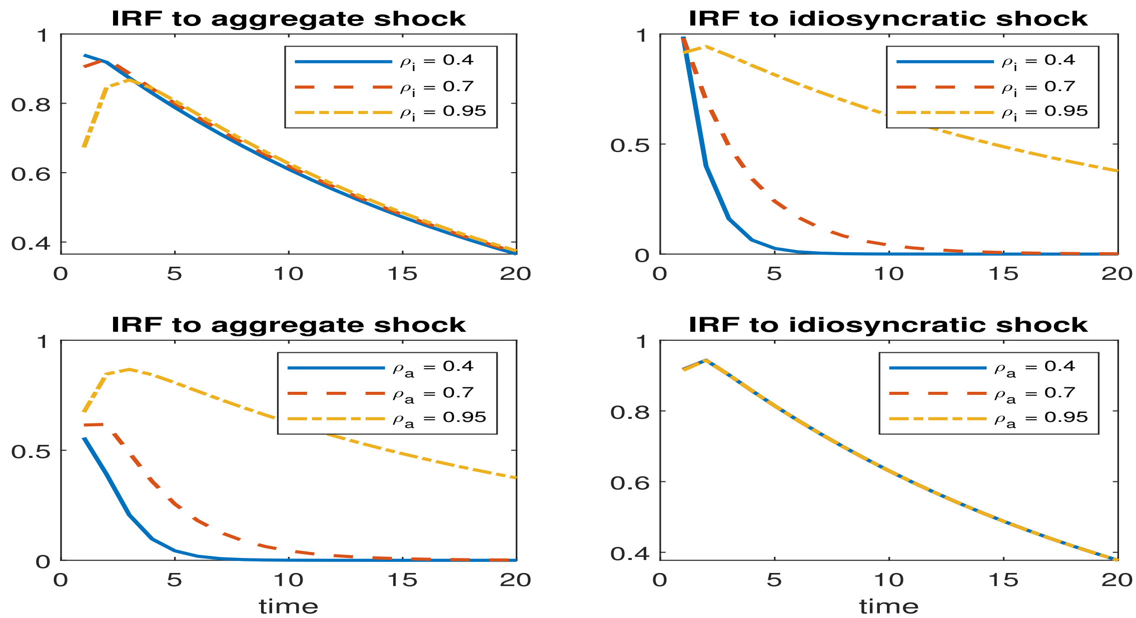

Figure 5: Impulse responses of the individual price to a unit size innovation in nominal aggregate demand and idiosyncratic productivity for different persistence of shocks in the case with strategic complementarity.

We finally study the impact of the shock persistence presented in Figure 5. When we change one persistence parameter ρ_i or ρ_a, we adjust the innovation variance to hold the unconditional variance fixed as in Maćkowiak and Wiederholt (2009). We also keep other parameters fixed at the baseline values. We find that the impact of persistence on individual price responses is ambiguous, a result similar to Maćkowiak and Wiederholt (2009). One reason is that the unconditional variances of the two shocks are different, unlike in the pricing example of Section 4. Another reason is that there is strategic complementarity in the model here. Figure 5 shows that individual prices respond faster to an idiosyncratic shock because its innovation variance is much larger, even though it is less persistent than an aggregate shock for many parameter values.

5.2 Consumption/Saving

In this subsection we study a consumption/saving problem similar to those in Hall (1978), Sims (2003), and Luo (2008). A household has quadratic utility over a consumption process $\{c_t\}$,

$$-\frac{1}{2}\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(c_t-\bar{c}\right)^2\right]$$

and faces the budget constraints

$$w_{t+1}=(1+r)\left(w_t-c_t\right)+y_{t+1},\ t\geq 0,$$

where $\bar{c}$ is a bliss level of consumption, w_t is wealth, and y_t is labor income. For simplicity let $\beta(1+r)=1$. We also impose a standard no-Ponzi game condition.

Suppose that income y_t consists of two persistent components and a transitory component:

$$\begin{aligned} y_t &= \bar{y}+z_{1,t}+z_{2,t}+\epsilon_{y,t}, \\ z_{i,t} &= \rho_i z_{i,t-1}+\eta_{i,t},\ i=1,2, \end{aligned}$$

where $\bar{y}$ is average income and innovations $\epsilon_{y,t}$, $\eta_{1,t}$, and $\eta_{2,t}$ are mutually independent Gaussian white noises with variances σ_y^2, σ_1^2, and σ_2^2. The two persistent components $z_{1,t}$ and $z_{2,t}$, and the transitory component $\epsilon_{y,t}$ may capture aggregate, local, and individual income uncertainties. The state vector is $x_t=(w_t,z_{1,t},z_{2,t})'$ plus a constant state 1. Suppose that the household does not observe the state vector x_t and solves the optimal consumption/saving problem under RI with discounted information costs.

By the certainty equivalence principle, it is straightforward to show that optimal consumption under RI is given by

$$c_t=\frac{\bar{y}}{1+r}+\frac{r}{1+r}\left(\widehat{w}_t+\frac{\rho_1}{1+r-\rho_1}\widehat{z}_{1,t}+\frac{\rho_2}{1+r-\rho_2}\widehat{z}_{2,t}\right),$$

where $\widehat{x}_t=\mathbb{E}\left[x_t|s^t\right]$. As is well known in the literature (e.g., Luo (2008)), optimal consumption under full information is linear in permanent income

$$m_t\equiv w_t+\frac{\rho_1}{1+r-\rho_1}z_{1,t}+\frac{\rho_2}{1+r-\rho_2}z_{2,t},$$

which is equal to the sum of financial wealth and expected present value of labor income. Then the RI problem can be equivalently solved using the permanent income as the only state variable. Thus the optimal signal is one dimensional and can be written as the permanent income m_t plus a noise (Luo (2018)).

We can verify this result using our numerical methods to solve for the optimal information structure (C,V) for the signal vector $s_t=Cx_t+v_t$. Set the same parameter values as in Sims

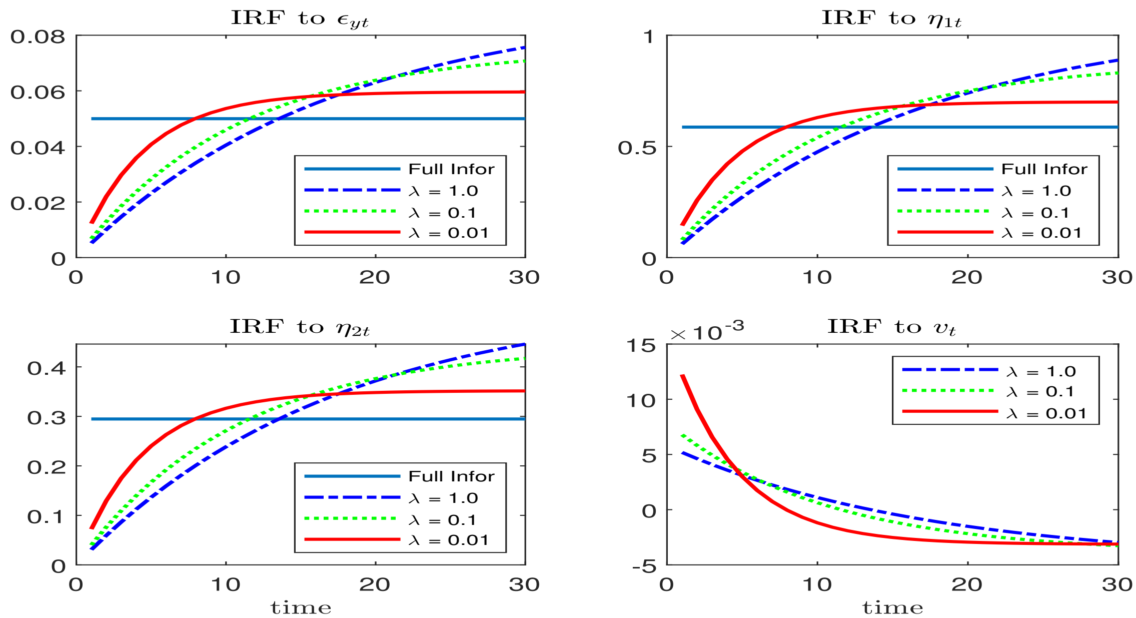

Figure 6: Impulse responses of consumption to a unit size innovation in various shocks for different information cost parameter values.

(2003): $\beta = 0.95$, $\bar{y} = 0$, $\rho_1 = 0.97$, $\rho_2 = 0.90$, $\sigma_y^2 = 0.01$, $\sigma_1^2 = 0.0001$, and $\sigma_2^2 = 0.003$. Unlike Sims (2003) and Luo (2018), we focus on the steady-state solution with discounted information costs, instead of capacity constraints.[21]

For the information cost parameter $\lambda = 0.01$, we find that the steady-state optimal signal vector s_t is one dimensional and can be normalized as $C = [1, \;\; 11.7433, \;\; 5.8978]$ and $V = 3.1876$. As can be verified that

$$\frac{\rho_1}{1 + r - \rho_1} = 11.7433, \; \frac{\rho_2}{1 + r - \rho_2} = \; 5.8978,$$

we have $Cx_t = e_t$ as in Luo (2008). As λ increases, the normalized signal weight vector C remains unchanged, but the signal noise variance increases significantly. Intuitively, the signal becomes more noisy when the information cost is larger.

Figure 6 plots the impulse response functions for consumption to a unit size innovation shock to each of the three true income components and the signal noise, starting from zero consumption. The flat lines correspond to the responses for the full information case. Under RI, the consumption responses to all three true component income shocks are damped initially, and then gradually rise permanently to high levels. Intuitively, the rationally inattentive household responds to shocks sluggishly. Lower consumption early leads to higher wealth. The extra savings earn a return $1 + r$ and allow the household to accumulate higher wealth to fund higher consumption later. We also

[21] In a previous version of the paper we solved the case with capacity constraints. The impulse response functions are qualitatively similar.

find that the initial response is larger for a more persistent income shock given the same λ. And the initial responses to all true income shocks are larger when λ is smaller. Unlike the income shocks, the noise shock causes consumption to rise immediately and then gradually decreases over time.

Our numerical results are different from those reported by Sims (2003). His Figures 7 and 8 show that consumption responses to shocks with different persistence display very different dynamics. By contrast, we find that they follow similar dynamics. The impact of persistence is reflected mainly by the magnitude of the initial response. The intuition is that within the LQG framework the multivariate permanent income model with general income processes can be reduced to a univariate model with IID innovations to permanent income (Luo (2008)).

Unlike Sims (2003) and Luo (2008), we also solve for the transition dynamics starting from the innovation covariance matrix as the initial prior for the state. We find that the household waits for the uncertainty to grow and then acquires a one-dimensional signal to reduce uncertainty at $t = 3$, 11, and 26 for $\lambda = 0.01$, 0.1, and 1, respectively. Intuitively, the household acquires information later when the information cost parameter λ is larger.

5.3 Firm Investment

We finally solve a firm's investment problem subject to convex adjustment costs under RI. Under full information, the firm chooses two types of capital investment to maximize its discounted present value of dividends:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t d_t\right]$$

subject to

$$\begin{aligned} d_t &= \exp\left(z_t + e_t\right) k_{1,t}^{\alpha} k_{2,t}^{\theta} - I_{1,t} - I_{2,t} - \frac{\phi_1}{2}\left(\frac{I_{1t}}{k_{1,t}} - \delta_1\right)^2 k_{1,t} - \frac{\phi_2}{2}\left(\frac{I_{2t}}{k_{2,t}} - \delta_2\right)^2 k_{2,t} \\ & \quad -\tau\left(\exp\left(z_t + e_t\right) k_{1,t}^{\alpha} k_{2,t}^{\theta} - \chi I_{2,t}\right), \end{aligned}$$

where d_t, $k_{1,t}, k_{2,t}, I_{1,t}$, and $I_{2,t}$ denote dividends, tangible capital, intangible capital, tangible capital investment, and intangible capital investment, respectively. The parameters satisfy $\delta_1, \delta_2, \alpha, \theta, \tau \in (0,1)$, $\alpha + \theta < 1$, and $\phi_1, \phi_2 > 0$.

The variables z_t and e_t represent persistent and temporary Gaussian TFP shocks, $z_t = \rho z_{t-1} + \epsilon_{z,t}$. We include taxation of corporate profits because a key distinction between the two types of capital is that a fraction χ of intangible investment is expensed and therefore exempt from taxation. The capital evolution equations are

$$k_{i,t+1} = \left(1 - \delta_i\right) k_{i,t} + I_{i,t} + \epsilon_{i,t+1}, \ i = 1, 2,$$

where $\epsilon_{i,t+1}$ represents depreciation or capital quality shocks. Suppose that $\epsilon_{z,t}$, e_t, $\epsilon_{1,t}$, and $\epsilon_{2,t}$ are mutually independent Gaussian white noises with variances σ_z^2, σ_e^2, σ_1^2, and σ_2^2.

To solve the problem under RI numerically, we first approximate the firm's objective function by a quadratic function in the neighborhood of the nonstochastic steady state. We then obtain a linear-quadratic control problem with the state vector $x_t = \left(z_t, e_t, \tilde{k}_{1,t}, \tilde{k}_{2,t}\right)'$ plus a constant state 1, where $\tilde{k}_{i,t}$, $i = 1, 2$, denotes the deviation from the steady state. From this problem we can derive the decision rules and the benefit matrix Ω in the control problem in which the relevant state vector is x_t. For the no adjustment cost case under full information, the linearized optimal decision rules are given by

$$\tilde{k}_{i,t+1} = \frac{k_i \rho}{1 - \alpha - \theta} z_t + \epsilon_{i,t+1},$$

where k_i is the steady-state capital stock. Notice that the optimal capital and investment choice is independent of transitory shocks e_t.

We now solve for the steady-state information structure. We set baseline parameter values as in McGrattan and Prescott (2010): $\alpha = 0.26$, $\theta = 0.076$, $\delta_1 = 0.126$, $\delta_2 = 0.05$, $\tau = 0.35$, and $\chi = 0.5$. Set $\rho = 0.91$, $\sigma_z = \sigma_1 = \sigma_2 = 0.01$, and $\sigma_e = 0.1$. We choose $\beta = 0.9615$ to generate a 4 percent steady-state interest rate. Following Saporta-Eksten and Terry (2018), we set the capital adjustment cost parameter values as $\phi_1 = 0.46$ and $\phi_2 = 1.40$. For these parameter values, the steady-state levels of capital are $k_1 = 0.98$ and $k_2 = 0.64$.

Since this model features two control variables and four state variables, we can study the nontrivial determination of the information structure. As shown in Proposition 5, the steady-state signal dimension for $\beta = 1$ does not exceed the minimum of the state dimension and the control dimension when all states are equally persistent. Using numerical examples, we find that this result holds true more generally. We also find that the steady-state signal dimension for $\beta \in (0, 1)$ can decrease from 2 to 1 when the information cost λ increases. Here we display the steady-state signal structure for two values of λ with adjustment costs: For $\lambda = 0.0001$,

$$s_t = \begin{bmatrix} -0.8985 z_t + 0.4204 \tilde{k}_{1,t} + 0.1265 \tilde{k}_{2,t} \\ -0.0406 z_t - 0.3665 \tilde{k}_{1,t} + 0.9295 \tilde{k}_{2,t} \end{bmatrix} + v_t,$$

where the covariance matrix of v_t is $\text{diag}(0.0008, 0.0042)$, but for $\lambda = 0.0005$,

$$s_t = -0.8977 z_t + 0.3397 \tilde{k}_{1,t} + 0.2806 \tilde{k}_{2,t} + v_t,$$

where the variance of v_t is 0.0456. Without adjustment costs, the steady-state signal is two dimensional for these values of λ.

Note that in neither case does the signal depend on e_t, the purely transitory productivity shock; since e_t does not affect the value-maximizing level of investment under full information, there is no point using information capacity to learn about it. Thus rational inattention does not explain why investment responds to transitory shocks in the data documented by Saporta-Eksten and Terry (2018). If the information structure is exogenously given as in a standard signal extraction problem,

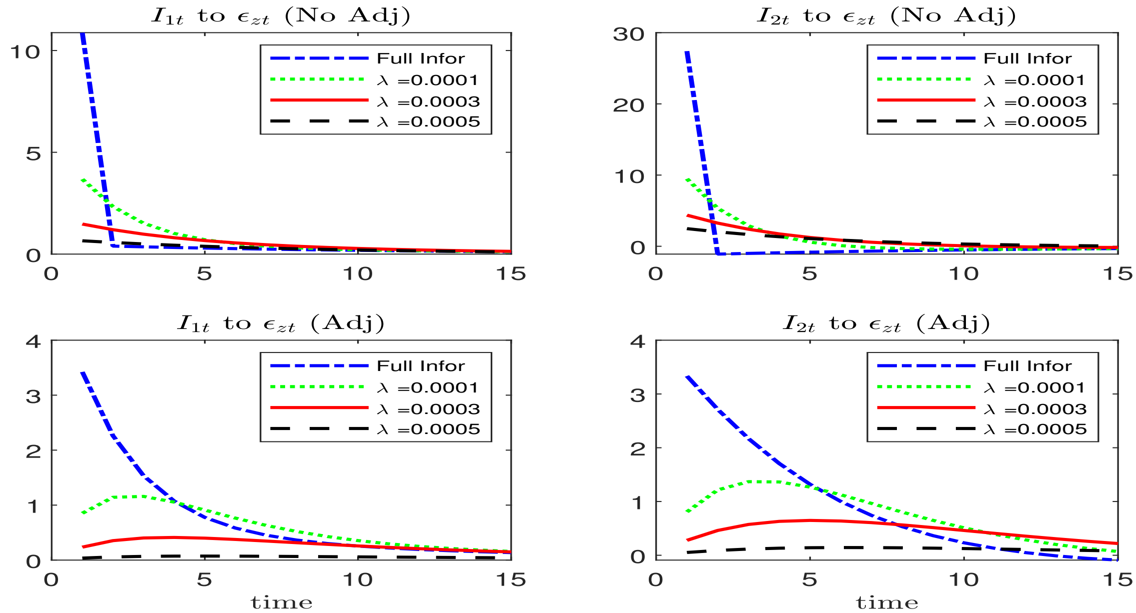

Figure 7: Impulse responses of tangible and intangible investment to a unit size innovation in the persistent TFP shock for different information costs. The vertical axis presents percentage deviations from the steady state.

then firms would be confused about the source of a productivity change; as a result, they would respond to transitory shocks.

We now turn to the impulse responses of two types of capital investment to a positive one unit innovation shock to the persistent TFP component displayed in Figure 7. The top two panels show the case without adjustment costs. Under full information, in response to a positive persistent TFP shock, investment rises too much on impact. As the information cost λ rises, the investment responses under RI become dampened and delayed – investment rises less on impact and remains above the steady state longer.

In the case with adjustment costs displayed in the bottom two panels, investment responses under RI are delayed further, and can become hump-shaped, a pattern not present in the full information case. The reason for the hump-shape is a horse race between two effects. Consider the response of tangible investment to a positive TFP shock z_t (bottom left panel). Value-maximizing investment under full information rises on impact and then gradually falls back to the steady state, but at a slower rate than the case without adjustment costs. Under rational inattention, since the firm does not know z_t with certainty, exactly how much investment has risen is unknown. Since the firm learns slowly and the capital adjustment is costly, it takes several periods before the firm knows the investment level it should have chosen on impact, which leads to a rising investment path. On the other hand, since z_t is mean reverting the value-maximizing level of investment is falling over time. Thus optimal investment under RI will eventually fall back to the steady state.

Without adjustment costs, mean reversion is sufficiently fast such that learning is always behind, leading to monotonic but delayed responses. With adjustment costs, but without information cost, there is no hump-shaped investment response either.

Our results are similar to Zorn's (2018) findings, while his model has only one type of capital and assumes there is no capital quality shock. He documents evidence that investment at the sectoral level displays a hump-shaped response to aggregate shocks and a monotonic response to sectoral shocks. He shows that a model with both rational inattention and capital adjustment costs can deliver the two different types of responses. In contrast, models with just capital adjustment costs, models with just investment adjustment costs, and models with just rational inattention cannot match both types of impulse responses.

Unlike Zorn's (2018) model, our model with four states and two controls allows us to study the nontrivial dynamics of information acquisition during the transition phase. We consider the case with adjustment costs and suppose that the firm starts with the innovation covariance matrix as the initial prior for the state. We find that the firm does not acquire any information initially for $\lambda = 0.0001$. As additional innovations arrive in each period, the firm starts to acquire a one-dimensional signal at $t = 2$ to reduce uncertainty, and then a two-dimensional signal at $t = 5$. The steady state is reached at $t = 30$. For $\lambda = 0.0005$, the information cost is so high that the firm starts to acquire a one-dimensional signal at $t = 14$ and never raises the signal dimension thereafter. It takes a longer time to reach the steady state at $t = 52$.

6 Discussions

In this section we discuss another solution concept related to Sims (2003), which has been followed by much of the literature. We also discuss the relation to the steady-state solution analyzed earlier.

Sims (2003) studies an infinite-horizon version of Problem 1. He considers the long-run situation in which $\Sigma_t = \Sigma$ is constant for $t \geq 0$. He then solves the following problem for optimal Σ :

$$\min_{\Sigma \succ 0} \ \mathrm{tr}\left(\Omega\Sigma\right) \tag{57}$$

subject to (42) and

$$\log\det\left(A\Sigma A' + W\right) - \log\det\left(\Sigma\right) \leq 2\kappa.$$

He interprets the objective as the long-run expected welfare loss under limited information relative to full information. However, minimizing the long-run expected welfare loss may not be equivalent to maximizing long-run expected utility for the control problem under limited information.

To see this point, we use Lemma 1 to write the negative of discounted expected utility under optimal control in the infinite-horizon case when $\Sigma_t = \Sigma$ for all t as

$$\mathbb{E}\left[x_0' P x_0\right] + \frac{1}{1-\beta}\mathrm{tr}\left(WP\right) + \frac{1}{1-\beta}\mathrm{tr}\left(\Omega\Sigma\right).$$

By the Kalman filter, for the posterior distribution Σ_t to stay at Σ for all $t \geq 0$, the initial state x_0 must be drawn from a Gaussian distribution with the long-run prior covariance matrix $\Sigma_{0|-1} = A\Sigma A' + W$. We can then compute

$$\mathbb{E}\left[x_0' P x_0\right] = \bar{x}_0' P \bar{x}_0 + \operatorname{tr}\left(P\Sigma_{0|-1}\right) = \bar{x}_0' P \bar{x}_0 + \operatorname{tr}\left(P\left(A\Sigma A' + W\right)\right), \tag{58}$$

where $\bar{x}_0$ is an exogenous mean of x_0. Because P is independent of Σ by the certainty equivalence principle (see (8)), choosing Σ to maximize expected utility is equivalent to

$$\min_{\Sigma \succ 0} \operatorname{tr}\left(A'PA\Sigma\right) + \frac{1}{1-\beta}\operatorname{tr}\left(\Omega\Sigma\right).$$

We can see that the first term in the above objective is missing in (57).

For the infinite-horizon version of Problem 2, we have to include the discounted information costs. When $\Sigma_t = \Sigma$ for all $t \geq 0$, we have

$$\sum_{t=0}^{\infty} \beta^t I\left(x_t; s_t | s^{t-1}\right) = \frac{1}{2\left(1-\beta\right)}\left[\log\det\left(A\Sigma A' + W\right) - \log\det\left(\Sigma\right)\right].$$

Thus maximizing the long-run discounted expected utility minus discounted information costs is equivalent to minimizing

$$\operatorname{tr}\left(A'PA\Sigma\right) + \frac{1}{1-\beta}\operatorname{tr}\left(\Omega\Sigma\right) + \frac{\lambda}{2\left(1-\beta\right)}\left[\log\det\left(A\Sigma A' + W\right) - \log\det\left(\Sigma\right)\right].$$

Multiplying by $(1-\beta)$ yields the objective function of the following problem:

Problem 5 *(Golden-rule information structure)*

$$\min_{\Sigma \succ 0} \left(1-\beta\right)\operatorname{tr}\left(A'PA\Sigma\right) + \operatorname{tr}\left(\Omega\Sigma\right) + \frac{\lambda}{2}\left[\log\det\left(A\Sigma A' + W\right) - \log\det\left(\Sigma\right)\right]. \tag{59}$$

subject to (42).

We call the solution to this problem the golden-rule information structure. We offer three comments on this solution concept. First, for the pure tracking Problem 4, we can easily check that the objective function when $\Sigma_t = \Sigma$ for all $t \geq 0$ becomes (41), or (59) without the first term. Thus there is no initial value problem for the pure tracking model. Second, in the $\beta \to 1$ limit, the golden-rule Problem 5 is the same as problem (41), which gives the steady-state solution for $\beta = 1$. Thus the golden-rule solution can be viewed as an approximation to the steady-state solution when β is sufficiently close to 1.

Third, the steady-state solution is a fixed point of the optimal policy function for the posterior covariance matrix discussed in Section 3.2. By contrast, the golden-rule solution is based on the assumption that the decision maker has already received a long sequence of signals before time zero so that the initial prior covariance matrix is the same as the long-run prior $A\Sigma A' + W$.

This assumption follows from Maćkowiak and Wiederholt (2009) and Maćkowiak, Matějka, and Wiederholt (2018). The weakness of this assumption is that it abstracts away from transition dynamics of Σ_t and also ignores some interesting intertemporal tradeoffs of information acquisition as discussed in Sections 4, 5.2, and 5.3. The strength is that it allows researchers to derive some analytical results as demonstrated by Maćkowiak and Wiederholt (2009), Maćkowiak, Matějka, and Wiederholt (2018), and our analysis in Section 3.3.

Moreover, the golden rule can be reliably solved by the powerful software CVX or other SDP software as the optimization problem is static and convex if $AA' + W \succ 0$. It applies to general ARMA processes and is robust to initial guess. By contrast, there is no theory to guarantee the convergence of the brute force iteration algorithm for the steady-state solution discussed in Section 3.2 and in Afrouzi and Yang (2019). This algorithm may be sensitive to the initial guess for high dimensional problems. When solving the equilibrium pricing example in Section 5.1, we indeed encounter the nonconvergence issue. We find that using the golden-rule solution with $\beta = 1$ as the initial guess for the steady-state solution for $\beta \in (0, 1)$ helps convergence.

7 Conclusion

We have developed a framework to analyze multivariate RI problems in a LQG setup. We have proposed a three-step solution procedure to theoretically analyze and numerically solve these problems based on SDP. We have provided generalized reverse water-filling solutions for some special cases and developed both value function based and first-order conditions based numerical methods for the general case. We have also applied our approach to three economic examples. Our analysis of the steady state and transition dynamics of the optimal signal structure generates some new insights such as different roles of the shock persistence and the innovation variance, information spillover, price comovement, and the timing of information acquisition. Our approach provides researchers a useful toolkit to solve multivariate RI problems without simplifying assumptions and will find wide applications in economics and finance.

References

Afrouzi, H., and C. Yang (2019): Dynamic Rational Inattention and the Phillips Curve, working paper, Columbia University.

Caplin, A., M. Dean, and J. Leahy (2019): Rational Inattention, Optimal Consideration Sets, and Stochastic Choice, *Review of Economic Studies,* 86, 1061-1094.

Cover, T. M., and J. A. Thomas (2006): *Elements of Information Theory*, 2th edition, John Wiley & Sons, Inc.

CVX Research, Inc. (2011): CVX: Matlab Software for Disciplined Convex Programming, version 2.0. http://cvxr.com/cvx.

Dewan, A. (2018): Costly Multidimensional Information, working paper, Columbia University.

Fulton, C. (2017): Mechanics of Linear Quadratic Gaussian Rational Inattention Tracking Problems, Finance and Economics Discussion Series 2017-109. Washington: Board of Governors of the Federal Reserve System.

Fulton, C. (2018): Mechanics of Static Quadratic Gaussian Rational Inattention Tracking Problems, working paper, Board of Governors of the Federal Reserve System.

Grant, M. C. (2004): Disciplined Convex Programming. PhD thesis, Information Systems Laboratory, Department of Electrical Engineering. Stanford University.

Grant, M. C., and S. P. Boyd (2008): Graph Implementations for Nonsmooth Convex Programs. In *Recent Advances in Learning and Control* (a tribute to M. Vidyasagar), V. Blondel, S. Boyd, and H. Kimura (eds), pp. 95-110, Lecture Notes in Control and Information Sciences, Springer.

Grant, M. C., S. P. Boyd, and Y. Ye (2006): Disciplined Convex Programming. In *Global Optimization: From Theory to Implementation,* L. Liberti and N. Maculan (eds), pp. 155-210, in the book series *Nonconvex Optimization and its Applications*, Springer.

Horn, R. A., and C. R. Johnson (2013): *Matrix Analysis*, 2nd edition, Cambridge University Press, NY.

Hamilton, J. D. (1994): *Time Series Analysis*, Princeton University Press, Princeton NJ.

Han, Z., F. Tan, and J. Wu (2019): Analytic Policy Function Iteration, Working Paper, St. Louis University

Harville, D. A. (1997): *Matrix Algebra from a Statistician's Perspective*, Springer, 1997, vol. 1.

Hébert, B., and M. Woodford (2018): Rational Inattention in Continuous Time, working paper, Columbia University.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp (2016): A Rational Theory of Mutual Funds' Attention Allocation, *Econometrica,* 84, 571-626.

Kőszegi, B. and F. Matějka (2019): Choice Simplification: A Theory of Mental Budgeting and Naive Diversification, forthcoming in *Quarterly Journal of Economics.*

Kydland, F. E., and E. C. Prescott (1982): Time to Build and Aggregate Fluctuations, *Econometrica,* 50, 1345-70.

Lipp, T., and S. Boyd (2016): Variations and Extension of the Convex–Concave Procedure, *Optimization and Engineering,* 17, 263–287.

Ljungqvist L., and T. J. Sargent (2004): *Recursive Macroeconomic Theory*, 2nd, the MIT Press, Cambridge, MA.

Stokey, N. L. and R. E. Lucas Jr., with E. C. Prescott (1989): *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge MA.

Luo, Y. (2008): Consumption Dynamics under Information Processing Constraints, *Review of Economic Dynamics,* 11, 366-385.

Luo, Y., J. Nie, and E. R. Young (2015): Slow Information Diffusion and the Inertial Behavior of Durables Consumption, *Journal of the European Economic Association,* 13, 805-840.

Luo, Y., and E. R. Young (2010): Risk-sensitive Consumption and Savings under Rational Inattention, *American Economic Journal: Macroeconomics,* 2, 281-325.

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

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

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

Maćkowiak, B., F. Matějka, and M. Wiederholt (2020): Rational Inattention: A Review, working paper, Sciences Po., CEPR DP1540.

Matějka, F., and A. McKay (2015): Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model, *American Economic Review,* 105, 272-298.

McGrattan, E. R., and E. C. Prescott (2010): Unmeasured Investment and the Puzzling US Boom in the 1990s, *American Economic Journal: Macroeconomics,* 2, 88-123.

Miao, J. (2014): *Economic Dynamics in Discrete Time*, the MIT Press, Cambridge, MA.

Miao, J. (2019): Multivariate LQG Control under Rational Inattention in Continuous Time, working paper, Boston University.

Miao, J., and J. Wu (2021): A Matlab Toolbox to Solve Dynamic Multivariate Rational Inattention Problems in the LQG Framework, working paper, Boston University.

Mondria, J. (2010): Portfolio Choice, Attention Allocation, and Price Comovement, *Journal of Economic Theory,* 145, 1837-1864.

Peng, L. (2005): Learning with Information Capacity Constraints, *Journal of Financial and Quantitative Analysis,* 40, 307-329.

Peng, L., and W. Xiong (2006): Investor Attention, Overconfidence and Category Learning, *Journal of Financial Economics,* 80, 563-602.

Saporta-Eksten, I., and S. J. Terry (2018): Short-term Shocks and Long-term Investment, working paper, Boston University.

Sims, C. A. (1998): Stickiness, *Carnegie-Rochester Conference Series On Public Policy,* 49, 317–356.

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

Sims, C. A. (2006): Rational Inattention: Beyond the Linear-Quadratic Case, *American Economic Review,* 96, 158–163.

Sims, C. A. (2011): Rational Inattention and Monetary Economics, in *Handbook of Monetary Economics*, V. 3A, Benjamin M. Friedman and Michael Woodford (ed.), North-Holland.

Steiner, J., C. Stewart, and F. Matějka (2017): Rational Inattention Dynamics: Inertia and Delay in Decision-Making, *Econometrica,* 85, 521-553.

Sturm, J. F. (1999): Using SeDuMi 1.02, A MATLAB Toolbox for Optimization over Symmetric Cones. In *Optimization Methods and Software*, 11-12: 625-633. Special issue on Interior Point Methods (CD supplement with software).

Tanaka, T., K. K. Kim, P. A. Parrio, and S. K. Mitter (2017): Semidefinite Programming Approach to Gaussian Sequential Rate-Distortion Trade-Offs, *IEEE Transactions on Automatic Control,* 62, 1896-1910.

Toh, K., M. J. Todd, and R. H. Tutuncu (1999): SDPT3 — A Matlab Software Package for Semidefinite Programming, *Optimization Methods and Software,* 11, 545–581.

Van Nieuwerburgh, S., and L. Veldkamp (2010): Information Acquisition and Under-diversification, *Review of Economic Studies,* 77, 571-626.

Vandenberghe, L., and S. Boyd (1996): Semidefinite Programming, *SIAM Review,* 38, 49-95.

Wang, B., B. Xi, and F. Zhang (1999): Some Inequalities for Sum and Product of Positive Semidefinite Matrices, *Linear Algebra and its Applications* 293, 39-49.

Woodford, M. (2003): Imperfect Common Knowledge and the Effects of Monetary Policy, In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, 25–58. Princeton, NJ: Princeton University Press.

Woodford, M. (2009): Information-Constrained State-Dependent Pricing, *Journal of Monetary Economics* 56, S100-S124.

Zhong, W. (2019): Optimal Dynamic Information Acquisition, working paper, Columbia University.

Zorn P. (2018): Investment under Rational Inattention: Evidence from US Sectoral Data, working paper, University of Munich.

A Appendix: Proof of Lemma 1

Fix the information structure $\{C_t, V_t\}$. Consider the control problem:

$$\widehat{v}_t \equiv \min_{\{u_\tau\}} \mathbb{E}\left[\sum_{\tau=t}^{T} \beta^{\tau-t}\left(x_\tau' Q_\tau x_\tau + u_\tau' R_\tau u_\tau + 2x_\tau' S_\tau u_\tau\right) + \beta^{T+1} x_{T+1}' P_{T+1} x_{T+1} \middle| s^t\right]$$

subject to (1) and (2) from period t on. Claim that

$$\widehat{v}_t = \mathbb{E}\left[x_t P_t x_t | s^t\right] + \sum_{\tau=t}^{T} \beta^{\tau-t+1} \mathrm{tr}\left(W_\tau P_{\tau+1}\right) + \sum_{\tau=t}^{T} \beta^{\tau-t} \mathrm{tr}\left(\Omega_\tau \Sigma_\tau\right), \quad \text{(A.1)}$$

where P_t and Ω_t satisfy (5) and (16). We prove this claim using backward induction. In the last period T, we compute the objective function as

$$\begin{aligned} & \mathbb{E}\left[\left(x_T' Q_T x_T + u_T' R_T u_T + 2x_T' S_T u_T\right) + \beta x_{T+1}' P_{T+1} x_{T+1} | s^T\right] \\ = \; & \mathbb{E}\left[\left(x_T' Q_T x_T + u_T' R_T u_T + 2x_T' S_T u_T\right) | s^T\right] \\ & + \beta \mathbb{E}\left[\left(A_T x_T + B_T u_T + \epsilon_{T+1}\right)' P_{T+1} \left(A_T x_T + B_T u_T + \epsilon_{T+1}\right) | s^T\right]. \end{aligned} \tag{A.2}$$

Rewrite the above expression as

$$\begin{aligned} & \mathbb{E}\left[\left(x_T' Q_T x_T + u_T' R_T u_T + 2x_T' S_T u_T\right) | s^T\right] \\ & + \beta \mathbb{E}\left[x_T' A_T' P_{T+1} A_T x_T + u_T' B_T' P_{T+1} B_T u_T + \epsilon_{T+1}' P_{T+1} \epsilon_{T+1} | s^T\right] \\ & + 2\beta \mathbb{E}\left[x_T' A_T' P_{T+1} B_T u_T | s^T\right] \\ = \; & \beta \mathrm{tr}\left(W_T P_{T+1}\right) + \mathbb{E}\left[x_T' Q_T x_T | s^T\right] + \beta \mathbb{E}\left[x_T' A_T' P_{T+1} A_T x_T | s^T\right] \\ & + \mathbb{E}\left[u_T'\left(R_T + \beta B_T' P_{T+1} B_T\right) u_T + 2x_T'\left(S_T + \beta A_T' P_{T+1} B_T\right) u_T | s^T\right] \end{aligned}$$

Taking the first-order condition gives the optimal control $u_T = -F_T \widehat{x}_T$, where F_T satisfies (7) for $t = T$. Substituting this equation back into the objective function yields

$$\widehat{v}_T = \mathbb{E}\left[x_T' P_T x_T | s^T\right] + \beta \mathrm{tr}\left(W_T P_{T+1}\right) + \mathrm{tr}\left(\Omega_T \Sigma_T\right),$$

where P_T satisfies (5) for $t = T$ and where we notice that x_T conditional on s^T is Gaussian with mean $\hat{x}_T$ and covariance matrix Σ_T.

Suppose that (A.1) holds for $\widehat{v}_{t+1}$ in period $t+1$. By dynamic programming, we have

$$\widehat{v}_t = \min_{u_t} \mathbb{E}\left[\left(x_t' Q_t x_t + u_t' R_t u_t + 2x_t' S_t u_t\right) + \beta \widehat{v}_{t+1} | s^t\right].$$

Rewriting the objective function by the induction hypothesis yields

$$\begin{aligned} & \mathbb{E}\left[\left(x_t' Q_t x_t + u_t' R_t u_t + 2x_t' S_t u_t\right) + \beta \widehat{v}_{t+1} | s^t\right] \\ = \; & \mathbb{E}\left[\left(x_t' Q_t x_t + u_t' R_t u_t + 2x_t' S_t u_t\right) | s^t\right] + \beta \mathbb{E}\left[x_{t+1} P_{t+1} x_{t+1} | s^t\right] \\ & + \sum_{\tau=t+1}^{T} \beta^{\tau-t+1} \mathrm{tr}\left(W_\tau P_{\tau+1}\right) + \sum_{\tau=t+1}^{T} \beta^{\tau-t} \mathrm{tr}\left(\Omega_\tau \Sigma_\tau\right). \end{aligned}$$

The expression on the second line has the same form as in (A.2). By the previous analysis, we deduce that the optimal policy is given by $u_t = -F_t \widehat{x}_t$, where F_t satisfies (7). Substituting this policy back into the preceding objective function, we find that the resulting objective function equals

$$\begin{aligned} & \mathbb{E}\left[x_t' P_t x_t | s^t\right] + \beta \mathrm{tr}\left(W_t P_{t+1}\right) + \mathrm{tr}\left(\Omega_t \Sigma_t\right) \\ & + \sum_{\tau=t+1}^{T} \beta^{\tau-t+1} \mathrm{tr}\left(W_\tau P_{\tau+1}\right) + \sum_{\tau=t+1}^{T} \beta^{\tau-t} \mathrm{tr}\left(\Omega_\tau \Sigma_\tau\right), \end{aligned}$$

where P_t satisfies (5). Thus $\widehat{v}_t$ takes the form in (A.1), completing the induction proof. Finally, letting $t = 0$ and taking unconditional expectations, we obtain the desired result. Q.E.D.