Asset Pricing with Misspecified Models*

Pooya Molavi† Alireza Tahbaz-Salehi‡ Andrea Vedolin§

October 2020

Abstract

This paper analyzes how limits to the complexity of statistical models used by market participants can shape asset prices. We consider an economy in which agents can only entertain models with at most k factors, where k may be distinct from the true number of factors that drive the economy's fundamentals. We first characterize the implications of the resulting departure from rational expectations for return predictability at various horizons. We then apply our framework to two applications in asset pricing: (i) violation of uncovered interest rate parity at different horizons and (ii) momentum and reversal in equity returns.

Keywords: model complexity, asset pricing, constrained-rational expectations.
JEL Classification: G4, G12, F31

*We are grateful to Fernando Alvarez, Andrea Buffa, Ignacio Esponda, Xavier Gabaix, Ali Jadbabaie, Lars Hansen, Stefan Nagel, Farzad Saidi, Harald Uhlig, and seminar participants at the University of Chicago, Brandeis University, the Virtual Finance Theory Workshop, and the World Congress of the Econometric Society for helpful comments and suggestions.
†Northwestern University, pmolavi@kellogg.northwestern.edu.
‡Northwestern University and CEPR, alirezat@kellogg.northwestern.edu.
§Boston University, NBER, and CEPR, avedolin@bu.edu.

1 Introduction

The rational expectations framework maintains that agents have a complete understanding of their economic environment: they know the structural equations that govern the relationship between endogenous and exogenous variables, have full knowledge of the stochastic processes that determine the evolution of shocks (if not the shocks' actual realizations), and are capable of forming and updating beliefs about as many variables as necessary. These assumptions are imposed irrespective of how complex the actual environment is. In reality, however, limits to agents' cognitive and computational abilities mean that market participants are bound to rely on simplified models that may not fully account for the complexity of their environment. Thus, to the extent that agents employ such simplified (and potentially misspecified) models, their decisions—and more generally, any endogenous variable that depends on those decisions—would depart from predictions obtained under rational expectations.

In this paper, we study how limits to the complexity of statistical models used by market participants shape asset prices. We consider a framework in which the stochastic process that governs the evolution of economic variables may not have a simple representation, and yet, agents are only capable of entertaining statistical models with a certain level of complexity. As a result, they may end up with a low-dimensional approximation to the true data-generating process. We show that this form of model misspecification generates systematic deviations from rational expectations with sharp predictions for the extent and nature of return predictability at different horizons.[1]

We present our results in the context of a simple asset pricing environment, in which a sequence of exogenous fundamentals (say, an asset's dividends) are generated by an n-factor model. While agents can observe the sequence of realized fundamentals, they neither observe nor know the underlying factors that drive them. As a result, they rely on their past observations to estimate a hidden-factor model that would allow them to make predictions about the future. As our main behavioral assumption, we assume that agents can only entertain models with at most k factors, where k may be distinct from the true number of factors, n. This assumption captures the idea that there is a limit to the complexity of statistical models that agents are able to consider, with a larger k corresponding to a more sophisticated agent who can entertain a richer class of models.

Subject to the constraint on the number of factors k, agents optimally choose the statistical model that minimizes the disparity between forecasts obtained under their subjective expectations and the true data-generating process. More specifically, we assume that they choose the k-factor model that minimizes the Kullback–Leibler (KL) divergence between the two. This modeling choice has two immediate, but important, implications. First, it implies that, when the agents' set of models is flexible enough to contain the true data-generating process (that is, when $k \geq n$), they can forecast future realizations of the fundamentals as if they knew the true process. In that case, our framework's

[1] While we focus on limits on the complexity of agents' statistical models, one can also consider other dimensions along which decision makers are constrained, such as limited memory (Nagel and Xu, 2019; da Silveira, Sung, and Woodford, 2020), limited capacity for processing information (as in models of rational inattention), an incomplete understanding of general equilibrium effects (e.g., due to level-k reasoning as in Farhi and Werning (2019) or limited "depth of knowledge" as in Angeletos and Sastry (2020)), or restriction on the number of variables to pay attention to (Gabaix, 2014, 2017).

predictions coincide with those of the rational expectations. However, when the constraint on model complexity binds (i.e., when $k < n$), agents *endogenously* select the best k-factor approximation to the true n-factor model.

With our behavioral framework in hand, we then establish that, irrespective of the constraint on the number of factors in agents' model, their subjective expectations are always internally consistent, in the sense that they satisfy the law of iterated expectations. We also provide a series of micro-foundations for KL divergence as the proper notion of model fit. In particular, we show that, in our context, whether agents choose their k-factor model to minimize the mean-squared error of their forecasts, use a maximum likelihood estimator, or engage in Bayesian learning, they end up with a model that minimizes the KL divergence with the true data-generating process.

We then turn to the asset pricing implications of our framework by providing a series of results relating the asset's return predictability to the number of factors k in agents' models. We measure the extent of return predictability by relying on two families of linear regressions, which we refer to as the *Fama* and *momentum regressions*. These regressions measure, respectively, the extent to which current fundamentals and excess returns predict future returns at different horizons. Our main theoretical results provide a characterization of the slope coefficients of the aforementioned regressions in terms of the complexity of agents' models and the autocorrelation function of the process that drives the fundamentals. Since all parameters in the agents' models, other than the number of factors k, are endogenously determined, our results provide sharp predictions for the extent of return predictability at different horizons, with no remaining degrees of freedom. We also extend our framework to a heterogeneous-agent economy in which only a fraction of agents face a constraint in the number of factors in their model, while the remaining fraction can entertain models of any order. We show how, in the presence of such heterogeneity, higher-order expectations play a central role in shaping the extent of return predictability at various horizons.

We then apply our framework to two asset pricing applications. As a first application, we study the implications of our behavioral assumption for violations of the uncovered interest rate parity (UIP) condition in foreign exchange. We show that the constraint on the number of factors in agents' model can generate return predictability patterns that are simultaneously consistent with two well-known, but seemingly contradictory violations of UIP, namely, the forward discount and the predictability reversal puzzles. The forward discount puzzle, which dates back to Fama (1984), is the robust empirical finding that, in short time horizons ranging from a week to a quarter, high interest rate currencies tend to have positive excess returns. The predictability reversal puzzle, more recently documented by Bacchetta and van Wincoop (2010) and Engel (2016), refers to the fact that high interest rate currencies tend to have negative excess returns over longer horizons, that is, the violation of UIP reverses sign after some point. The seemingly contradictory implications of these puzzles for the relationship between currency excess returns and interest rate differentials has led some to argue for the inadequacy of existing models for explaining UIP violations (Engel, 2016).

To test our model's implications for the violation of UIP at different horizons, we use a large cross-section of currency returns and, following Engel (2016), run predictive regressions from a trade-weighted average currency return on the corresponding interest rate differential. We then compare the resulting slope coefficients to those implied by our theoretical framework when agents are constrained in the number of factors in their models. We find that when investors are constrained to using single-factor models, the model-implied slope coefficients line up with the ones from the data: at short horizon, deviations from rational expectations generate UIP violations that imply positive excess returns for high interest rate currencies, whereas at longer horizons the pattern reverses, with high interest rate currencies earning a negative excess return. In other words, our framework generates return predictability patterns that are simultaneously consistent with the forward discount and predictability reversal puzzles. Crucially, the pattern of our model-implied Fama slope coefficients matches its empirical counterpart without using the data on exchange rates or excess returns: the model-implied slope coefficients are constructed solely from the autocorrelation of interest rate differentials.[2] We also show that our findings remain mostly unchanged even if only a relatively small fraction of agents are subject to our behavioral constraint.

As a further illustration, we study violations of UIP in the cross-section of developed and emerging market currencies by running the Fama regressions on a currency-by-currency basis. We find positive and statistically significant relationships between the empirically-estimated slope coefficients and their model-implied counterparts at various horizons.

As a second application of our framework, we study short-run momentum and long-run reversal in equity returns. Time-series momentum implies that past returns are predictors of future returns. This pattern usually persists for a year and reverses over the longer term. To estimate model-implied slope coefficients, we again rely on the autocorrelation function of fundamentals, in this case, dividend-price ratios. Using the closed-form expressions for the one-factor model case, we find that model-implied slope coefficients track the slope coefficients in the data, simultaneously generating short-term momentum and long-run reversal.

Related Literature Our paper belongs to the literature that studies the asset pricing implications of deviations from rational expectations. The most related strand of this literature considers agents who have a misspecified view of the true data-generating process as a result of behavioral biases. For instance, Barberis, Shleifer, and Vishny (1998) assume agents mistakenly believe the innovations in earnings are drawn from a regime with excess reversals or excess streaks. Similarly, Rabin and Vayanos (2010) show that gambler's fallacy—the belief that random sequences should exhibit systematic reversals, even in small samples—can generate momentum and reversal in equity returns. More recently, Guo and Wachter (2019) focus on a model in which investors believe that returns are predictable when in fact they are not. As in Rabin and Vayanos (2010) and Guo and Wachter (2019), we consider a framework in which agents form expectations using misspecified factor models. In contrast to these papers, where misspecification manifests itself as constraints on the parameters of agents' model, agents in our framework are constrained only in the number of factors they can entertain.

Within this literature, our paper is also related to Hong, Stein, and Yu (2007), who focus on

[2]We contrast these findings to factor models of higher order and find that as agents become more sophisticated—and in particular, as we increase the number of factors in agents' model to $k = 3$—UIP violations mostly disappear at all horizons.

an economy where agents rely on univariate models of the world, even when the true underlying model is multivariate. The set of univariate models entertained by agents in Hong, Stein, and Yu (2007) are exogenously given, with the transition probabilities between those models treated as free parameters. In contrast, agents in our framework *endogenously* select the k-factor model that best describes their observations. In other words, the only free parameter in our framework is the maximum number of factors in agents' models.

Our behavioral framework is also related to the natural expectations model of Fuster et al. (2010, 2012), in which decision makers are assumed to forecast a time series by modeling it as a low-order autoregressive process with k lags. Our paper builds on and generalizes these works by allowing agents to entertain any k-factor model without restrictions. Therefore, rather than making an arbitrary assumption about the kind of statistics the agents can keep track of, we allow them to optimally choose the k-factor model that best fits their past observations.[3]

Another strand of literature studies the asset pricing implications of deviations from Bayesian updating. The extrapolative expectations models of Hirshleifer, Li, and Yu (2015), Barberis, Greenwood, Jin, and Shleifer (2015), and Choi and Mertens (2019) and diagnostic expectations models of Bordalo, Gennaioli, and Shleifer (2018), Bordalo, Gennaioli, La Porta, and Shleifer (2019), and Bordalo, Gennaioli, Ma, and Shleifer (2020) are examples of such non-Bayesian models of updating. Relatedly, Adam, Marcet, and Nicolini (2016) and Nagel and Xu (2019) consider a setting in which agents update their subjective beliefs using a constant gain, which induces fading memory. In a departure from this literature, we maintain the assumption of Bayesian updating and instead consider agents whose priors assign zero probability to complex models with a large number of factors.

Our paper is also related to the recent work of Martin and Nagel (2020), who study an environment in which the number of factors potentially relevant for prediction is of the same order of magnitude as the number of assets with available cashflow data. They show that, as a consequence of this high-dimensional prediction problem, Bayesian investors use regularization to trade off the costs of downweighting certain pieces of information against the benefit of reduced parameter estimation error. In contrast, we focus on an environment in which agents have access to abundant data but instead are restricted in the number of factors they can use in their models.

More generally, our paper builds on the broader literature that studies the theoretical implications of misspecification as an expression of bounded rationality, as such as the prior works of Esponda and Pouzo (2016, 2019) as well as Molavi (2019), who studies the implications of misspecified factor models for business-cycle fluctuations. We instead focus on how restricting agents to factor models with a small number of endogenously-chosen factors shapes the extent and nature of return predictability in asset pricing applications.

Methodologically, our paper is also related to the literature on model order reduction in control theory, which is concerned with characterizing lower-order approximations to large-scale dynamical systems. See Antoulas (2005) and Sandberg (2019) for textbook treatments of the subject. In contrast

[3]Note that the family of k-factor models includes all $\mathrm{ARMA}(p,q)$ processes such that $\max\{p,q+1\} \leq k$ and thus nests the family of $\mathrm{AR}(k)$ processes as a special case.

to this literature, which mostly focuses on optimal Hankel-norm approximations, agents in our model use the KL divergence as their notion of model fit. This choice is motivated by the fact that, in our context, choosing the k-factor model with the smallest KL divergence to the true data-generating process is equivalent to the outcome of Bayesian learning when agents assign a zero prior belief to models with more than k factors. Therefore, our framework allows for imposing constraints on the complexity of models that agents can entertain, while at the same time preserving other properties of rational expectations.

Finally, we also contribute to the literature that studies UIP violations and in particular, the reversal in currency return predictability. Engel (2016) develops a model to reconcile the forward discount and predictability reversal puzzles by introducing a non-pecuniary liquidity return on assets. Valchev (2020) proposes a different mechanism based on endogenous fluctuations in bond convenience yields, while Bacchetta and van Wincoop (2019) argue that delayed portfolio adjustments can explain various anomalies in exchange rate dynamics. Different from these papers, we show that behavioral constraints on the complexity of investors' statistical models—and the resulting departure from rational expectations—generate patterns that are simultaneously consistent with the forward discount and predictability reversal puzzles.

Outline The rest of the paper is organized as follows. Section 2 presents the environment and specifies our behavioral assumption. Section 3 contains our main theoretical results and explores the implications of constraints on the complexity of agents' models for asset prices. Section 4 presents our two empirical applications. All proofs and some additional mathematical details are provided in the Appendix.

2 Framework

Consider a discrete-time economy consisting of a unit mass of identical agents. Agents form subjective expectations about an exogenous sequence $\{x_t\}_{t=-\infty}^{\infty}$ of variables, which we refer to as the economy's *fundamentals*. Depending on the context, the fundamentals may correspond to an asset's dividend stream over time, interest rates differentials between two countries, or any sequence of variables that can be treated as exogenous in that specific application.

The sequence of fundamentals $\{x_t\}_{t=-\infty}^{\infty}$ is generated by a stationary n-factor model given by

$$\begin{aligned} z_t &= A^* z_{t-1} + \epsilon_t \\ x_t &= b^{*\prime} z_t, \end{aligned} \tag{1}$$

where $z_t \in \mathbb{R}^n$ denotes the vector of factors that drive the dynamic of fundamentals, $A^* \in \mathbb{R}^{n\times n}$ is a square matrix that governs the evolution of factors over time, and $b^* \in \mathbb{R}^n$ is a vector of constants that captures the fundamental's loading on each of the n factors.[4] The noise terms $\epsilon_t \in \mathbb{R}^n$ are independent over time and are distributed according to $\epsilon_t \sim \mathcal{N}(0, \Sigma_\epsilon^*)$. The process that generates the fundamentals can thus be summarized by the n underlying factors $(z_1, z_2, \ldots, z_n)$

[4]To ensure stationarity, we assume that A^* is a stable matrix, i.e., all its eigenvalues are within the unit circle.

as well as the collection of parameters $\theta^* = (A^*, b^*, \Sigma_\epsilon^*)$. Throughout, we assume that the underlying factors remain unobservable to agents, who, at any given time t, can only observe the realization of fundamentals $\{x_\tau\}_{\tau \leq t}$ up to that time. We also assume that there are no redundant factors in (1), in the sense that the dynamics of the fundamentals cannot be represented by a factor model with fewer factors than n.

Whereas the fundamental sequence $\{x_t\}_{t=-\infty}^{\infty}$ is assumed to be exogenous, we are interested in how the fundamentals and agents' subjective expectations about them jointly determine an endogenous sequence of variables $\{y_t\}_{t=-\infty}^{\infty}$, which we refer to as *prices*. More specifically, we assume that the price at time t satisfies the recursive equation

$$y_t = x_t + \delta \mathbb{E}_t[y_{t+1}], \tag{2}$$

where $\mathbb{E}_t[\cdot]$ denotes agents' time t subjective expectation and $\delta \in [0, 1]$ is a constant. Equation (2), which serves as the basis of our subsequent analysis, has a natural interpretation. For example, if the sequence $\{x_t\}_{t=-\infty}^{\infty}$ represents an asset's dividend stream, then y_t corresponds to the price of the asset at time t, with equation (2) simply capturing the fact that the asset's price is the sum of its dividend at that time and its expected future price, discounted at some rate δ.

Given the sequence of fundamentals and prices $\{(x_t, y_t)\}_{t=-\infty}^{\infty}$, we define *excess returns* at time $t + 1$ as the sum of the fundamental and the change in price between time t and $t + 1$, properly discounted:

$$\mathrm{rx}_{t+1} = \delta y_{t+1} - y_t + x_t. \tag{3}$$

Using equation (2), we can also express excess returns as the difference between the realized and the expected price: $\mathrm{rx}_{t+1} = \delta(y_{t+1} - \mathbb{E}_t[y_{t+1}])$.

Together with the specification of how agents' expectations are formed, equations (1)–(3) fully describe our environment. As we will discuss in Section 4, the framework represented by these equations is general enough to capture various asset pricing applications. We also note that, thus far, we have not made any assumptions about how agents' expectations are formed. In particular, it may be the case that agents' expectations $\mathbb{E}_t[\cdot]$ are distinct from the expectations arising from the true data-generating process (1), which we denote by $\mathbb{E}_t^*[\cdot]$. Nonetheless, as we will discuss in further detail in subsequent sections, our framework ensures that agents' expectations are internally consistent, in the sense that they satisfy the law of iterated expectations.

As a final remark, we note that one can write the price and excess returns in terms of the fundamentals and agents' subjective expectations. In particular, as long as subjective expectations satisfy law of iterated expectations, iterating on equation (2) implies that

$$y_t = \sum_{\tau=0}^{\infty} \delta^\tau \mathbb{E}_t[x_{t+\tau}] \tag{4}$$

$$\mathrm{rx}_{t+1} = \sum_{\tau=1}^{\infty} \delta^\tau \left(\mathbb{E}_{t+1}[x_{t+\tau}] - \mathbb{E}_t[x_{t+\tau}] \right). \tag{5}$$

While the recursive representations in (2) and (3) are more compact, the above expressions illustrate how prices and excess returns depend on agents' expectations about the entire future path of fundamentals. For instance, going back to the example discussed earlier, the representation in (4) simply illustrates that the asset's price at time t is equal to the discounted sum of its expected future dividends. Similarly, equation (5) captures the intuitive idea that the asset's excess return at time $t+1$ is equal to the change in agents' forecasts (from t to $t+1$) about its dividend stream (Campbell, 1991).

2.1 Return-Predictability Regressions

To assess the implications of our (yet-to-be-specified) behavioral framework for asset prices, we start with a measure to capture deviations from the rational expectations benchmark. One natural choice for such a measure is the extent of (excess) return predictability. This is because the benchmark rational expectations framework implies that future returns should be unpredictable.

We measure the extent of return predictability by relying on two families of linear regressions. The first family of regressions, which we refer to as *Fama regressions*, measures the extent to which current fundamentals predict future excess returns at different horizons. More specifically,

$$\mathrm{rx}_{t+h} = \alpha_h^{\mathrm{Fama}} + \beta_h^{\mathrm{Fama}} x_t + \epsilon_{t,h}, \tag{6}$$

where $h \geq 1$. If returns are unpredictable—as would be the case under rational expectations—then $\beta_h^{\mathrm{Fama}} = 0$ for all horizons h. The family of slope coefficients $(\beta_1^{\mathrm{Fama}}, \beta_2^{\mathrm{Fama}}, \dots)$ thus provides not only a measure for departures from the rational expectations benchmark, but also the extent to which such departures vary with the prediction horizon.[5]

Our second family of regressions, which we refer to as *momentum regressions*, measures the extent to which current returns predict future returns:

$$\mathrm{rx}_{t+h} = \alpha_h^{\mathrm{mom}} + \beta_h^{\mathrm{mom}} \mathrm{rx}_t + \epsilon_{t,h}. \tag{7}$$

Once again, the unpredictability of future returns under rational expectations requires that $\beta_h^{\mathrm{mom}} = 0$ for all $h \geq 1$. Hence, the term structure of coefficients $(\beta_1^{\mathrm{mom}}, \beta_2^{\mathrm{mom}}, \dots)$ provides us with a natural measure for assessing the implications of our behavioral framework at different time horizons compared to the rational expectations benchmark, with a positive (negative) β_h^{mom} corresponding to the extent of time-series momentum (reversal) in returns at horizon h.

We conclude by providing a simple result that relates the slope coefficients in regressions (6) and (7) to agents' subjective expectations about the fundamental as well as the properties of the true data-generating process.

Proposition 1. *If agents' expectations satisfy the law of iterated expectations, the slope coefficients in*

[5] We refer to equation (6) as the Fama regression because of its similarity to Fama's (1984) regression specification for testing deviations from the uncovered interest rate parity condition. In that context, the fundamental x_t corresponds to log interest rate differential between two countries, y_t is the log of the foreign exchange rate, and rx_t is the currency risk premium at time t. See Section 4 for a more detailed discussion of the application of our framework to foreign exchange.

regressions (6) *and* (7) *are given by*

$$\beta_h^{\text{Fama}} = \frac{1}{\mathbb{E}^*[x_t^2]} \sum_{\tau=0}^{\infty} \delta^{\tau+1} \mathbb{E}^*[x_t(\mathbb{E}_{t+h}[x_{t+h+\tau}] - \mathbb{E}_{t+h-1}[x_{t+h+\tau}])] \tag{8}$$

$$\beta_h^{\text{mom}} = \frac{\sum_{\tau,s=1}^{\infty} \delta^{\tau+s} \mathbb{E}^*[(\mathbb{E}_{t+1}[x_{t+\tau}] - \mathbb{E}_t[x_{t+\tau}])(\mathbb{E}_{t+h+1}[x_{t+h+s}] - \mathbb{E}_{t+h}[x_{t+h+s}])]}{\mathbb{E}^*[\sum_{\tau=1}^{\infty} \delta^{\tau} (\mathbb{E}_{t+1}[x_{t+\tau}] - \mathbb{E}_t[x_{t+\tau}])]^2}, \tag{9}$$

respectively, where $\mathbb{E}^*[\cdot]$ *is the expectation with respect to the true data-generating process* (1).

The above result has a few immediate, but important implications. First, it is straightforward to verify that, under rational expectations, $\beta_h^{\text{Fama}} = \beta_h^{\text{mom}} = 0$ for all $h \geq 1$. Thus, as expected, in the rational expectations benchmark, returns are unpredictable at any horizon, irrespective of whether agents condition on the current value of the fundamental, x_t, or the current excess returns, rx_t. Second, when agents' subjective expectations do not coincide with the expectation under the true data-generating process, the coefficients in the return-predictability regressions may in general be different from zero, even when agents' expectations are internally consistent. Finally, equations (8) and (9) illustrate that the wedge between agents' subjective expectations and rational expectations may have differential impacts on the slope coefficients at different horizons h. How departures from rational expectations shape the term structures of β_h^{Fama} and β_h^{mom} will be the focus of our subsequent theoretical and empirical analyses.

2.2 Behavioral Assumption

So far, we have remained agnostic with respect to how agents' subjective expectations are formed. We now specify the behavioral model that governs agents' expectation formation process.

Recall that the fundamentals $\{x_t\}_{t=-\infty}^{\infty}$ are generated by the n-factor model in (1). While at any given time t, agents observe the sequence of realized fundamentals up to that time, we assume that (i) they neither observe nor know the underlying factors $(z_1, \ldots, z_n)$ that drive the fundamentals and (ii) they do not know the collection of parameters $\theta^* = (A^*, b^*, \Sigma_\epsilon^*)$ that govern the true data-generating process.

In addition, we assume that agents can only entertain models with at most k factors, where k may be distinct from the true number of factors, n. This assumption, which serves as our main behavioral assumption, captures the intuitive idea that there is a limit to the complexity of statistical models that agents are able to consider. The number of factors, k, indexes the degree of agents' sophistication, with a larger k corresponding to agents who can entertain a richer class of models.

Formally, we assume that agents can only entertain models of the form

$$\begin{aligned} \omega_t &= A\omega_{t-1} + \epsilon_t \\ x_t &= b'\omega_t, \end{aligned} \tag{10}$$

where $\omega_t \in \mathbb{R}^k$ denotes the vector of k underlying factors, $A \in \mathbb{R}^{k\times k}$ and $b \in \mathbb{R}^k$ capture the process that govern factors' evolution and the fundamental's loading on each of the factors, respectively, and the noise terms $\epsilon_t \sim N(0, \Sigma_\epsilon)$ are independent and identically distributed over time. We summarize such a k-factor model with the collection of parameters $\theta = (A, B, \Sigma_\epsilon)$ and use Θ_k to denote the set of all k-factors models in the form of (10). Note that, with some abuse of notation, we can write $\Theta_k \subseteq \Theta_{k+1}$, thus capturing the fact that agents with a higher k can contemplate a larger class of models.

A few remarks are in order. First, note that we impose no restrictions, other than the number of factors, on the agents' model: the k factors $(\omega_1, \ldots, \omega_k)$ in the agents' model may overlap with a subset of the n factors $(z_1, \ldots, z_n)$ that drive the fundamental, may be linear combinations of the underlying n factors, or can be constructed in an entirely different way altogether. Second, as we discuss in further detail below, the agents' choice of which k-factor model to use is endogenous: they pick the model from set Θ_k that best fits their past observations. Third, when $k < n$, the set of models entertained by the agents does not contain the true n-factor data-generating process in (1), i.e., $\theta^* \notin \Theta_k$. In such a case, our behavioral assumption implies that irrespective of their choice, agents will end up with a misspecified model of the world. This observation also clarifies the bite of our behavioral assumption: whereas more sophisticated agents with $k \geq n$ can recover the model that generates the fundamentals (at least in principle), those with $k < n$ can at best construct lower dimensional approximations to the true data-generating process.

As already mentioned, we assume that agents choose their model by picking the k-factor model that best fits their observations. Following Esponda and Pouzo (2016, 2019) and Molavi (2019), we use the Kullback–Leibler divergence as our notion of fit:

Definition 1. The *Kullback–Leibler (KL) divergence* of model $\theta \in \Theta_k$ from the true process is given by

$$\mathrm{KL}(\theta^* \| \theta) = \mathbb{E}^*[-\log f^\theta(x_{t+1}|x_t, \ldots)] - \mathbb{E}^*[-\log f^*(x_{t+1}|x_t, \ldots)], \tag{11}$$

where f^* is the density of the fundamental under the true data-generating process, f^θ is agents' subjective density under model θ, and $\mathbb{E}^*[\cdot]$ is the expectation with respect to the true process.[6]

The KL divergence measures the disparity between agents' subjective expectations, as captured by density f^θ, and the true process f^*. It is always non-negative and obtains its minimum value of zero if and only if the two densities coincide almost everywhere. To see the interpretation of $\mathrm{KL}(\theta^* \| \theta)$ as a goodness-of-fit measure, note that $\mathbb{E}^*[-\log f^\theta(x_{t+1}|x_t, \ldots)]$ and $\mathbb{E}^*[-\log f^*(x_{t+1}|x_t, \ldots)]$ are entropy-like terms that capture the extent of uncertainty regarding one-step-ahead predictions conditional on the history of past observations under, respectively, model θ and the true model, θ^*. Therefore, the right-hand side of (11) measures the additional uncertainty—and hence, the resulting degradation in prediction quality—when agents use model θ as opposed to the true model.

Definition 2. A *k-factor constrained-rational expectations equilibrium (CREE)* consists of a process for prices and a model $\theta_k^{\text{CREE}} \in \Theta_k$ for agents such that

(i) prices satisfy equation (2),

[6] Note that, by equation (2), y_t is measurable with respect to the information set generated by $\{x_\tau\}_{\tau \leq t}$. Therefore, irrespective of whether past prices are observed by agents or not, it is sufficient to consider expectations conditioned on the past realizations of the fundamental.

(ii) agents' subjective expectations are generated by the k-factor model θ_k^{CREE},

(iii) $\theta_k^{\text{CREE}} \in \operatorname{argmin}_{\theta\in\Theta_k} \text{KL}(\theta^*\|\theta)$.

Under constrained-rational expectations, agents' subjective expectations are consistent with the k-factor model (10) that minimizes the KL divergence to the true process (1). Importantly, even though agents are restricted to models with at most k factors, they face no constraints in choosing what these factors are, how they evolve, or how the fundamental loads on each of the factors.

Note that, whenever $k \geq n$, the k-factor constrained-rational expectations reduce to rational expectations. This is, of course, not surprising: when agents can contemplate models that are more complex than the true model, their subjective expectations will coincide with the underlying objective expectations once they choose the model that best fits the data. More specifically, since $\theta^* \in \Theta_k$, they can achieve the lower bound of $\text{KL} = 0$. In contrast, when $k < n$, the constraint imposed on the number of factors binds, in the sense that $\text{KL}(\theta^*||\theta) > 0$ for all $\theta \in \Theta_k$. As a result, agents end up with a misspecified model of the world and subjective expectations that may diverge from rational expectations.

With our notion of goodness-of-fit in hand, our next result establishes that a constrained-rational expectations equilibrium always exists and generates subjective expectations that are internally consistent.

Theorem 1. *A constrained-rational expectations equilibrium exists for all k. Furthermore, agents' expectations in any CREE satisfy the law of iterated expectations.*

Besides its intuitive appeal as a measure of additional uncertainty arising from potential model misspecification, using the KL divergence in (11) as the notion of fit can also be rationalized in a number of other ways. Our next result provides a series of micro-foundations for our equilibrium notion by illustrating that whether agents choose their model to minimize the mean-squared error of their one-step-ahead predictions (statement (a)), use a maximum likelihood estimator (statement(b)), or engage in Bayesian learning (statement (c)), they end up with a model in the set of constrained-rational expectations equilibria.

Theorem 2. *Let $\Theta_k^{\text{CREE}} \subseteq \Theta_k$ denote the set of all models that are part of a k-factor constrained-rational expectations equilibrium. Then,*

(a) $\Theta_k^{\text{CREE}} = \operatorname{argmin}_{\theta\in\Theta_k} \mathbb{E}^*[(x_{t+1} - \mathbb{E}_t^\theta[x_{t+1}])^2]$,

(b) $\arg\max_{\theta\in\Theta_k} f^\theta(x_t|x_{t-1},\dots x_0) \to \Theta_k^{\text{CREE}}$ *as $t \to \infty$ with $\mathbb{P}^*$-probability one,*

(c) if $\mu_0 \in \Delta\Theta_k$ is a prior with full support and $\mu_t \in \Delta\Theta_k$ is the corresponding Bayesian posterior at time t, then $\lim_{t\to\infty} \mu_t(\mathcal{U}) = 1$ with $\mathbb{P}^*$-probability one for any open set $\mathcal{U} \supset \Theta_k^{\text{CREE}}$.

We conclude this discussion by noting that even when the constraint on the complexity of models entertained by the agents binds (i.e., when $k < n$), the fact that they choose the k-factor model that fits the data best means that they may still recover important features of the data-generating process, even if not the dynamics of the underlying factors or the factors themselves. The following example illustrates this point.

Example 1. Suppose the underlying data-generating process is governed by n independent-evolving factors $(z_1, \ldots, z_n)$. More specifically, suppose $z_{it+1} = a_i^* z_{it} + \epsilon_{it+1}$ and $x_t = \sum_{i=1}^n b_i^* z_{it}$, where $a_i^* \in (-1, 1)$ is the persistence of the i-th factor, b_i^* is the fundamental's loading on that factor, and the shocks $(\epsilon_{1t}, \ldots, \epsilon_{nt})$ are independent from one another and have unit variance.[7] In the language of equation (1), the underlying model is represented by the set of parameters $\theta^* = (A^*, b^*, \Sigma^*)$, where Σ^* is the identity matrix and the matrix that governs the dynamics of the data-generating process is given by

$$A^* = \begin{bmatrix} a_1^* & 0 & \cdots & 0 \\ 0 & a_2^* & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & a_n^* \end{bmatrix}.$$

Now, suppose agents can only entertain single-factor models, that is, $k = 1$. Thus, as long as $n \geq 2$, the agents can never recover the dynamics that govern the fundamental. Instead, they form their subjective expectations about future fundamentals from the single-factor model that best fits their observations in the sense of Definition 1. As we show in the appendix, the model θ_1^{CREE} corresponding to the single-factor constrained-rational expectations equilibrium is given by $\omega_{t+1} = a^{\text{CREE}} \omega_t + \epsilon_t$ and $x_t = \omega_t$, where $\omega_t \in \mathbb{R}$ is the factor that drives the process, $\epsilon_t \sim N(0, 1)$, and

$$a^{\text{CREE}} = \left(\sum_{i=1}^n \frac{b_i^{*2} a_i^*}{1 - a_i^{*2}} \right) \Big/ \left(\sum_{i=1}^n \frac{b_i^{*2}}{1 - a_i^{*2}} \right). \tag{12}$$

Note that the single-factor model constructed by the agents depends on the dynamics of all n underlying factors that drive the fundamental. However, the corresponding persistence parameter a^{CREE} puts a lower weight on factors with smaller loadings: a lower $|b_i^*|$ translates into a smaller weight on the persistence a_i^* of the i-th factor. So, for example, if all but one of the b_i^* are very small in magnitude—so that the model behaves similar to a single-factor model—the agents' model closely approximates that specific factor's dynamics.

Another observation that emerges from (12) is that if one of the factors follows a near unit-root process, in the sense that $|a_i^*| \approx 1$, then agents' estimated single-factor model also behaves like a near unit-root process, with $|a^{\text{CREE}}| \approx 1$. This is the case even if the fundamental's loading b_i^* on that factor is small.

Taken together, these observations illustrates that, even though agents are restricted to models that may be significantly less complex than the true data-generating process and have no way of identifying the true underlying factors, they nonetheless recover a lower dimensional representation of the data-generating process that may approximate salient features of the true model.

3 Asset Pricing Implications

In this section, we study the implications of constrained-rational expectations for return predictability. In particular, we characterize the slope coefficients of return-predictability

[7] To ensure that there are no redundant factors in the model, we assume that $a_i^* \neq a_j^*$ for all $i \neq j$ and $b_i^* \neq 0$ for all i.

regressions (6) and (7) when agents can only entertain models consisting of at most k factors, where, in general, $k \neq n$.

Recall from Definition 2 that all parameters corresponding to agents' models, other than the number of factors k, are endogenously determined by minimizing the KL divergence from the true data-generating process. Consequently, once we specify the number of factors in the agents' model, there are no more degrees of freedom on how they form their expectations. This, in turn, implies that the coefficients of return-predictability regressions can be expressed only in terms of the number of factors k in the agents' model and the statistical properties of the true data-generating process. We have the following result:

Theorem 3. *Suppose agents are constrained to k-factor models. Then, the slope coefficients of the Fama and momentum regressions* (6) *and* (7) *are, respectively, given by*

$$\beta_h^{\text{Fama}} = \delta u'(I - \delta M)^{-1} u \left(\xi_h^* - \sum_{\tau=1}^{\infty} \phi_\tau \xi_{h-\tau}^* \right) \tag{13}$$

$$\beta_h^{\text{mom}} = \frac{\xi_h^* - \sum_{s=1}^{\infty} \phi_s (\xi_{h-s}^* + \xi_{h+s}^*) + \sum_{s,\tau=1}^{\infty} \phi_s \phi_\tau \xi_{h+s-\tau}^*}{1 - 2\sum_{s=1}^{\infty} \phi_s \xi_s^* + \sum_{s,\tau=1}^{\infty} \phi_s \phi_\tau \xi_{s-\tau}^*}, \tag{14}$$

where $\xi_\tau^* = \mathbb{E}^*[x_t x_{t+\tau}]/\mathbb{E}^*[x_t^2]$ is the autocorrelation of the fundamental, $\phi_s = u'[M(I - uu')]^{s-1} M u$, and M and u are, respectively, a $k \times k$ weakly stable matrix and a k-dimensional unit vector that minimize

$$H(M, u) = 1 - 2\sum_{s=1}^{\infty} \phi_s \xi_s^* + \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \phi_s \phi_\tau \xi_{\tau-s}^*. \tag{15}$$

This theorem, which is the main characterization result of the paper, relates the extent of return predictability under constrained-rational expectations to the number of factors in the agents' model and the statistical properties of the underlying data-generating process. Importantly, it illustrates that the coefficients of the Fama and momentum regressions happen to only depend on the underlying process through the autocorrelation function of the fundamental, $\xi_\tau^* = \mathbb{E}^*[x_t x_{t+\tau}]/\mathbb{E}^*[x_t^2]$. This is a consequence of the fact that, under constrained-rational expectations, agents optimally scale the volatility of the factors in their model with the volatility of the fundamental.

The characterization result in Theorem 3 serves two distinct purposes in our analysis. First, it allows us to directly apply our framework to various asset pricing applications. In particular, in any context in which endogenous and exogenous variables are related to one another via equation (2), we can use the autocorrelation function of the exogenous variables and the expressions in (13) and (14) to compute the implied coefficients of the Fama and momentum regressions under constrained-rational expectations. This is the approach we take in the next section. Second, Theorem 3 also enables us to perform comparative statics analyses with respect to the primitives and to compare return predictability at different horizons, as we do in the remainder of this section.

3.1 Comparative Statics

We start with a simple result that considers the case where agents' models are sufficiently rich to fully capture the statistical properties of the data-generating process.

Theorem 4. *Suppose $\delta > 0$ and let k and n denote the number of factors in, respectively, the agents' model and the true data-generating process. Then,*

(a) If $k \geq n$, then $\beta_h^{\text{Fama}} = \beta_h^{\text{mom}} = 0$ for all horizons h.

(b) If $k < n$, then there exists h and $\tilde{h}$ such that $\beta_h^{\text{Fama}} \neq 0$ and $\beta_{\tilde{h}}^{\text{mom}} \neq 0$.

Statement (a) of Theorem 4 establishes that if the number of factors agents can entertain is at least equal to the number of true factors driving the fundamental, then there is no return predictability at any horizon. This result is a consequence of two key assumptions embedded in the definition of a CREE. First, agents choose their k factors optimally, by minimizing the divergence between their forecasts and their observations. Second, in a CREE, agents behave as if they have used an infinitely long sequence of observations to discipline the parameters of their models. As a result, agents do not suffer from finite-sample problems such as overfitting: they recover the underlying data-generating process even if they use models that have too many parameters relative to the true model (i.e., when $k > n$). Statement (b) of the theorem then shows that the converse implication is also true: if there is no return predictability, the set of models entertained by the agents has to be rich enough to contain the true underlying model. Consequently, under constrained-rational expectations, return predictability and model misspecification are one and the same.

Our next result concerns the extent of return predictability in long horizons.

Theorem 5. *Under constrained-rational expectations, excess returns are not predictable in the long run:*

$$\lim_{h\to\infty} \beta_h^{\text{Fama}} = \lim_{h\to\infty} \beta_h^{\text{mom}} = 0.$$

This result is a direct consequence of the assumptions that (i) the true data-generating process is stationary and (ii) agents only entertain stationary factor models. These stationarity assumptions imply that the effect of time-t variables (including the fundamental x_t and excess returns rx_t) on returns at time $t + h$ die out eventually as h increases.

3.2 Single-Factor Models

While there does not exist a closed-form representation of agents' expectations for a general k, our next result provides such a characterization for the case in which agents are restricted to using single-factor models. Despite being only a special case, this closed-form characterization is a transparent and easy-to-use result that is informative about how the extent of return predictability varies with the horizon h as a function of the autocorrelation function of the true data-generating process.

Proposition 2. *If agents are constrained to single-factor models, then the slope coefficients of the Fama and momentum regressions are given by*

$$\beta_h^{\text{Fama}} = \frac{\delta}{1-\delta\xi_1^*}(\xi_h^* - \xi_{h-1}^*\xi_1^*) \tag{16}$$

$$\beta_h^{\text{mom}} = \frac{(1+\xi_1^{*2})\xi_h^* - \xi_1^*(\xi_{h-1}^* + \xi_{h+1}^*)}{1-\xi_1^{*2}}, \tag{17}$$

respectively, where ξ_h^* is the autocorrelation of the fundamental with lag h.

Using the above closed-form representation, we can obtain the following result:

Proposition 3. *Suppose agents are constrained to single-factor models. Furthermore, suppose $\xi_1^* > 0$. Then, there exist h and h' such that $\beta_h^{\text{mom}}\beta_{h'}^{\text{mom}} < 0$.*

Therefore, as long as the underlying process that drives the fundamental exhibits some persistence in the very short run—and irrespective of any of its other characteristics—subjective expectations induced by agents' misspecified model generate momentum and reversal simultaneously.

3.3 Heterogenous-Agent Economy

Our results thus far relied on the assumption that the economy consists of a unit mass of identical agents, all of whom are restricted to using models with the same maximum number of factors, k. In this subsection, we extend our previous results by assuming that only a fraction $1-\lambda$ of the agents are subject to our behavioral constraint, while the remaining λ fraction can entertain models with any number of factors. For simplicity, we refer to the two groups of agents as behavioral and rational agents, respectively.[8] We assume that parameter λ is common knowledge in the economy.

The heterogeneity in agents' ability to entertain statistical models of different complexities results in heterogenous subjective expectations, with λ and $1-\lambda$ fractions of agents having rational and constrained-rational expectations, respectively. As a result, the relationship between endogenous prices and exogenous fundamentals is given by the following generalization of equation (2):

$$y_t = x_t + \delta\bar{\mathbb{E}}_t[y_{t+1}],$$

where $\bar{\mathbb{E}}[\cdot] = \lambda\mathbb{E}^*[\cdot] + (1-\lambda)\mathbb{E}[\cdot]$ denotes the cross-sectional average of agents' expectations. Iterating on the above, we can also obtain the following counterpart to equation (5) for excess returns in terms of agents' expectations:

$$\text{rx}_{t+1} = \sum_{\tau=1}^{\infty}\delta^\tau\left(\bar{\mathbb{E}}_{t+1}\bar{\mathbb{E}}_{t+2}\cdots\bar{\mathbb{E}}_{t+\tau}[x_{t+\tau}] - \bar{\mathbb{E}}_t\bar{\mathbb{E}}_{t+1}\cdots\bar{\mathbb{E}}_{t+\tau}[x_{t+\tau}]\right). \tag{18}$$

[8] Despite the terminology, note that, in view of Theorem 2 all agents in this economy are Bayesian, with the only difference between the two groups being that behavioral agents assign zero prior probability to models consisting of more than k factors. Furthermore, note that when $k \geq n$, agents in both groups end up with the same exact subjective expectations. As a result, the subjective expectations of the two groups differ only when $k < n$.

The key observation is that, even though subjective expectations of each group of agents satisfy the law of iterated expectations, the cross-sectional average expectation $\bar{\mathbb{E}}[\cdot]$ may not. Therefore, unlike the representative-agent framework of Section 2, excess returns in the heterogenous-agent economy also depend on higher-order expectations whenever $k < n$.

Before presenting our result, we note that the failure of law of iterated expectations with respect to $\bar{\mathbb{E}}[\cdot]$ in our framework resembles a similar phenomenon in differential-information economies, such as Allen, Morris, and Shin (2006), Barillas and Nimark (2017), Angeletos and Lian (2018) and Angeletos and Huo (2019), where agents have access to private signals about fundamentals. In contrast to these papers, however, all information in our framework is public and it is the heterogeneity in the maximum number of factors in agents' models that results in heterogenous expectations and the potential violation of law of iterated expectations. Another important consequence of this assumption is how agents' higher-order expectations are formed: while rational agents can recover the model used by behavioral agents, the latter behave as if they live in a representative-agent economy only consisting of agents with k-factor models. This is because behavioral agents are convinced—mistakenly so when $k < n$—that a k-factor model is sufficient to capture the process that drives the fundamental.

Our next result characterizes the slope coefficient of the Fama regression (6) in the heterogenous-agent economy in terms of the corresponding family of coefficients in the representative-agent economy consisting of only behavioral agents (i.e., $\lambda = 0$). Let $\beta_h^{\text{Fama}}(\lambda)$ denote the slope coefficient of the Fama regression in an economy with λ and $1-\lambda$ fraction of rational and behavioral agents, respectively. We have the following result:

Proposition 4. *The slope coefficient of the Fama regression in the heterogenous-agent economy is given by*

$$\beta_h^{\text{Fama}}(\lambda) = (1-\lambda)\sum_{s=0}^{\infty}(\delta\lambda)^s\beta_{h+s}^{\text{Fama}}(0), \tag{19}$$

where $\beta_h^{\text{Fama}}(0)$ is the slope coefficient of the representative-agent economy and is given by (13).

4 Applications

In this section, we apply our framework to two asset pricing applications: the violation of uncovered interest rate parity in foreign exchange and time-series momentum and reversal in equity returns. More specifically, we use the characterization result in Theorem 3 to test our model's predictions for the slope coefficients of return predictability regressions (6) and (7) in these contexts.

4.1 Reversal of Uncovered Interest Rate Parity

One of the central tenets of international finance is the uncovered interest rate parity (UIP) condition, which maintains that high interest rate currencies should depreciate vis-à-vis those with low interest rates. Yet—in what has become known as the "forward discount puzzle"—a vast

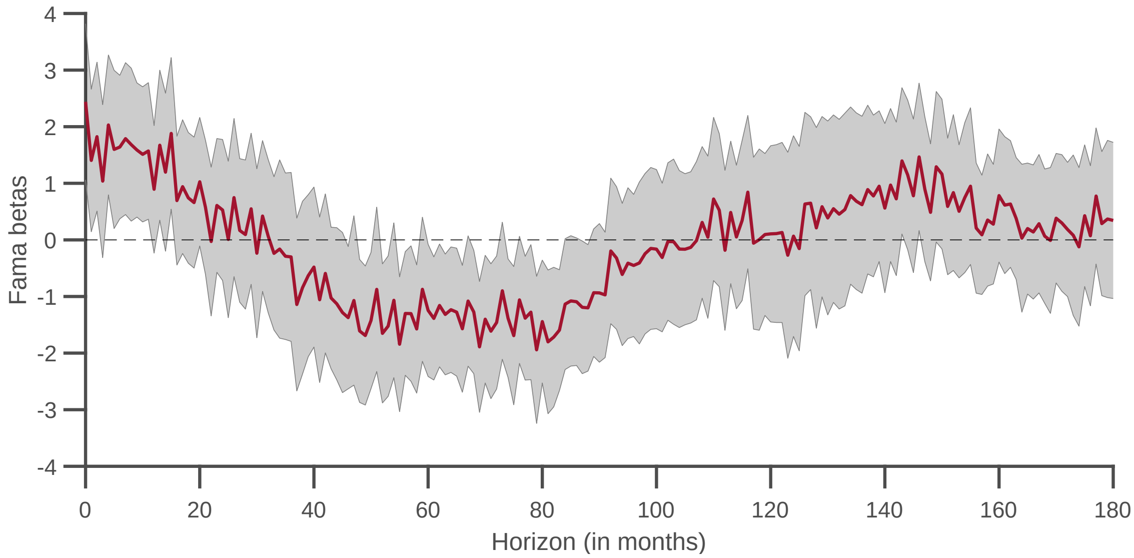

Figure 1. Uncovered Interest Rate Parity Regression: Fama Coefficient

Notes: This figure plots estimated slope coefficients together with 90 percent confidence intervals of the Fama regression: $\mathrm{rx}_{t+h} = \alpha_h^{\text{Fama}} + \beta_h^{\text{Fama}}(i_t^* - i_t) + \epsilon_{t,h}$, where $h = 0, \ldots, 180$ denotes the horizon in months, rx_{t+h} is the currency risk premium on a trade weighted average of seven currencies vis-à-vis the U.S. dollar, and $i_t^* - i_t$ is the log interest rate differential between the trade-weighted average interest rate and the U.S. interest rate. Data is monthly and runs from January 1985 to October 2018. Confidence intervals are calculated using Newey and West (1987) standard errors.

empirical literature documents that, over short time horizons (ranging from a week to a quarter), high interest rate currencies tend to appreciate. In other words, short-term deposits of high-interest rate currencies tend to earn a predictively positive excess return.

More recently, however, Bacchetta and van Wincoop (2010) and Engel (2016) document a distinct but related puzzle, known as the "predictability reversal puzzle." They find that UIP violations reverse sign over longer horizons, with high interest rate currencies earning negative excess returns at horizons from four to seven years. The seemingly contradictory implications of the forward discount and predictability reversal puzzles for the relationship between currency excess returns and interest rate differentials has led some to argue for the inadequacy of existing models for explaining UIP violations. More specifically, Engel (2016) argues that, risk-based explanations of the forward discount puzzle—which attribute the violations of UIP to the relative riskiness of holding short-term deposits in the high-interest rate country—cannot account for the predictability reversal puzzle.

In this subsection, we apply our theoretical framework to study the implications of model misspecification for the pattern of UIP violations at different horizons and investigate the extent to which constrained-rational expectations can jointly explain the forward discount and predictability reversal puzzles.

We map this context to our framework in Section 2 by letting the fundamental denote the log interest rate differential between the U.S. and a foreign country, i.e., $x_t = i_t^* - i_t$, where i_t and i_t^* are nominal interest rates on deposits held in U.S. dollars and the foreign currency, respectively. We also let y_t denote the log of the foreign exchange rate, expressed as the U.S. dollar price of the foreign currency. With the discount rate set to $\delta = 1$, recursive equation (2) then captures the interest rate parity condition, according to which an increase in the U.S. to foreign short-term interest rate differential is associated with an exchange rate appreciation, whereas a higher expected future exchange rate implies a depreciation. Note that, as in Section 2, the expectation in (2) denotes agents' subjective expectations, which may differ from those arising from the true data-generating process. Finally, equation (3) simply expresses the currency risk premium from period t to period $t+1$:

$$\mathrm{rx}_{t+1} = y_{t+1} - y_t + (i_t^* - i_t).$$

We start by reproducing the empirical findings on UIP violations at different horizons. Following Engel (2016), we build a trade-weighted average exchange rate and interest rate differential relative to the U.S. for the following countries: Australia, Canada, Euro (Germany before its introduction), New Zealand, Japan, and the United Kingdom. The weights are constructed as the value of each country's exports and imports as a fraction of the average value of trade over the six countries. Monthly exchange rate data is from Datastream and interest rate differentials are calculated using covered interest rate parity from forward rates, $i_t^* - i_t = f_t - y_t$, also available from Datastream. We then run the following regression:

$$\mathrm{rx}_{t+h} = \alpha_h^{\text{Fama}} + \beta_h^{\text{Fama}}(i_t^* - i_t) + \epsilon_{t+h}, \qquad h = 0, \ldots, 180,$$

where h is the horizon measured in months. This regression is, of course, identical to the Fama regression (6) in Section 2.

Figure 1 plots the estimated slope coefficients at various horizons. For $h = 1$, we find the slope coefficient to be positive, thus recovering the classic forward discount puzzle: at short time horizons, higher interest rate differentials (relative to the U.S.) lead to higher risk premia. This pattern remains the same up to a horizon of three years, but then reverses its sign, illustrating the predictability reversal puzzle: for horizons between four to seven years, higher interest rates predict a lower currency risk premium. Finally, as the figure indicates, the estimated coefficient of the Fama regression becomes indistinguishable from zero at even longer horizons.

Turning to our framework's predictions, we first calculate the autocorrelation function $\{\xi_\tau^*\}_{\tau \geq 1}$ of the interest rate differential for the trade-weighted basket of currencies against the U.S. dollar. Taking this autocorrelation as our primitive, we then calculate the term structure of the model-implied coefficients of the Fama regression using expression (13) in Theorem 3 for different number of factors k in agents' model.

As our first exercise, we consider the case in which agents can only entertain single-factor models, i.e., $k = 1$. Recall that in this special case, we can use the closed-form expression (16) in Proposition 2 to calculate the model-implied slope coefficients of the Fama regression. Figure 2 plots the term structure of the model-implied coefficient for $k = 1$ together with the coefficients obtained from the data from Figure 1. As the figure indicates, the pattern of model-implied coefficients tracks the pattern observed from the data fairly closely. Most importantly, we see a reversal in the slope coefficient: model-implied coefficients are positive for horizons up to 30 months and reverse to

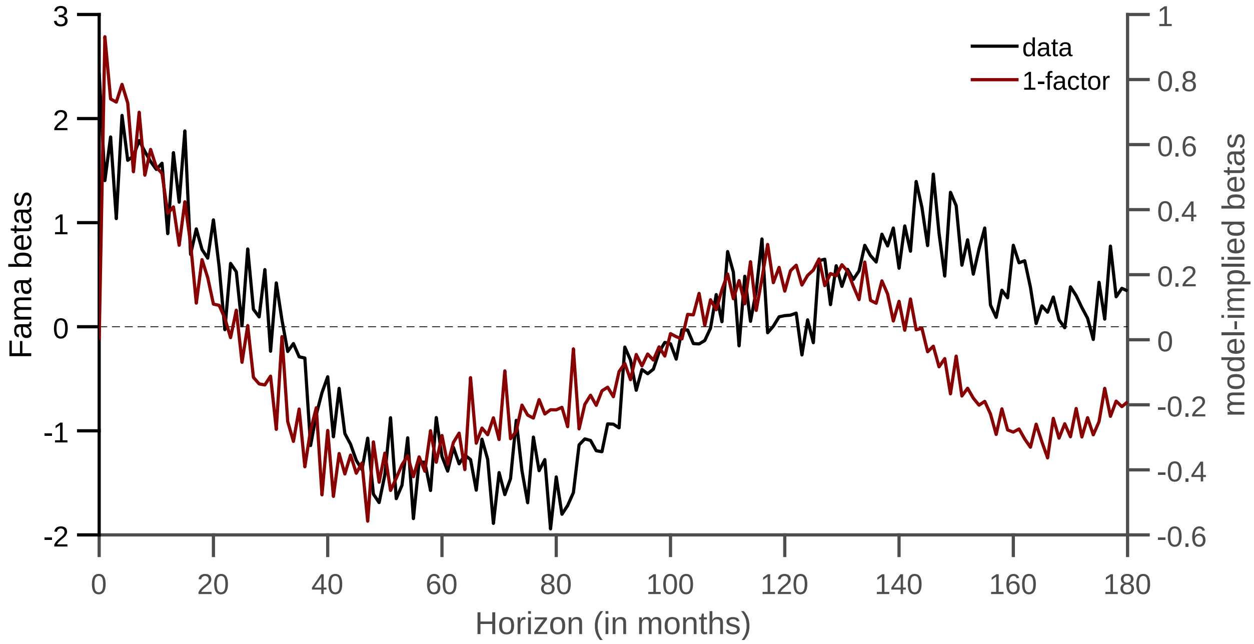

Figure 2. One-Factor Model-Implied Fama Coefficient

Notes: This figure plots estimated Fama slope coefficients from Figure 1 (left axis) together with model-implied (right axis) betas from a one-factor model given in Proposition 2 for horizons 0 to 180 months.

a negative sign thereafter. The pattern in Figure 2 thus suggests that single-factor constrained-rational expectations generate return predictability patterns that are simultaneously consistent with the forward discount and predictability reversal puzzles.

It is important to emphasize that while the estimated coefficients in Figure 1 are obtained from regressing returns on interest rate differentials, the model-implied coefficients in Figure 2 do not use the data on exchange rates or excess returns. Rather, they are simply obtained by plugging the autocorrelation of interest rate differential into equation (16).

Next, we investigate how increasing the number of factors k entertained by the agents impacts the term structure of model-implied slope coefficients. To this end, we once again use the empirical autocorrelation function of the interest rate differential as an input to calculate β_h^{Fama} for $k = 2, 3$. However, when $k > 1$, there is no closed-form expression for the model-implied slope coefficients. As a result, we use equation (13) and the characterization result in Theorem 3 to solve for β_h^{Fama} numerically. Importantly, as a by-product, we also obtain the model-implied autocorrelation function (ACF), i.e., the autocorrelation function as perceived from the (potentially misspecified) k-factor model used by the agents.[9]

The results are reported in Figure 3. The upper panel depicts the model-implied ACF for $k = 1, 2, 3$, together with the empirical ACF of the trade-weighted average interest rate differential for horizons 0 to 180 months. As the figure indicates, the model-implied ACFs differ quite substantially

[9]See Appendix B.1 for the details of how the model-implied autocorrelation function can be calculated in terms of the solution of optimization problem (15).

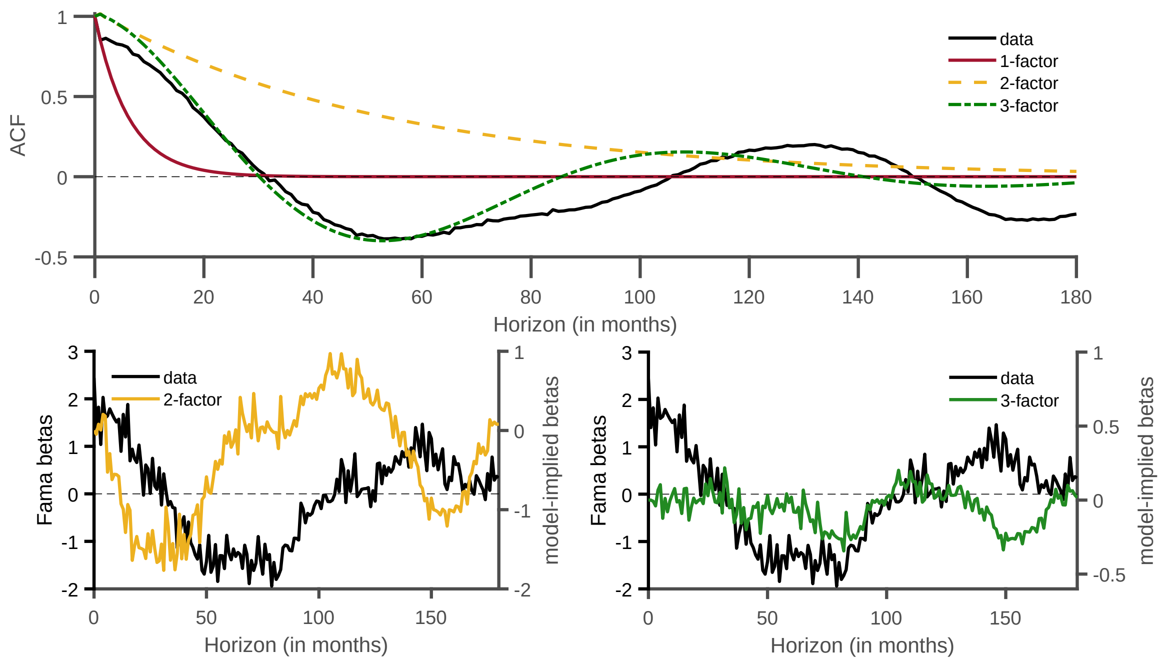

Figure 3. Autocorrelation Function and k-Factor-Model-Implied Fama Coefficients

Notes: The upper panel plots the autocorrelation function of the trade-weighted interest rate differential (data) together with one-, two-, and three-factor-model-implied ACFs. The lower two panels plot the two-factor (lower left panel) and three-factor (lower right panel) model-implied Fama betas for horizons 0 to 180 months.

across various levels of agents' sophistication. For example, while the three-factor model almost perfectly matches the empirical ACF, the ACF implied by the single-factor model looks significantly different. Crucially, this is reflected in the model-implied slope coefficients as illustrated in the lower right panel of Figure 3: the three-factor model, which generates a model-implied ACF that tracks the empirical ACF very closely, also results in model-implied slope coefficients that are significantly smaller at all horizons. This, of course, is to be expected in view of our results in Section 3. As agents are able to entertain richer and more complex statistical models, they end up with models that better fit the empirical ACF, which in turn results in less significant deviations from the rational expectations benchmark and hence less return predictability.

To further validate our model's predictions, we calculate the model-implied slope coefficients of the Fama regression for a larger cross-section of countries and plot them against the corresponding slope coefficients estimated from the data.[10] Figure 4 depicts the results for $k = 1$ at different horizons. Recall from Theorem 3 and Proposition 2 that model-implied slope coefficients depend on the shape of the ACF of the fundamental (in this case, the interest rate differential between the corresponding country and the U.S.). Therefore, as ACFs differ in the cross-section of currencies, so should the model-implied betas. Nonetheless, as Figure 4 illustrates, model-implied and empirically-estimated slope coefficients have a positive and statistically-significant relationship.

[10] We use currency pairs for the following countries (all against the U.S. dollar): Australia, Canada, Hong Kong, Denmark, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, Singapore, Sweden, Switzerland, the United Kingdom, Czech Republic, Kuwait, Mexico, the Philippines, Spain, and Turkey.

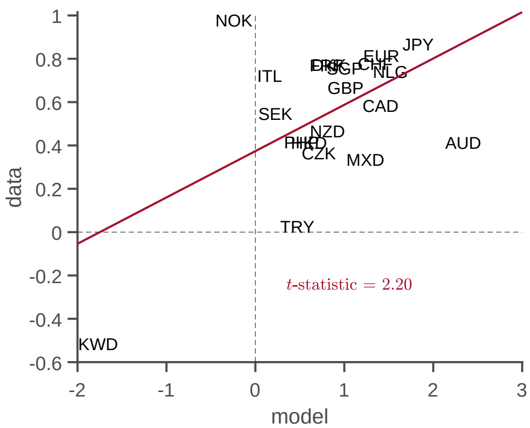

(a) $h = 2$

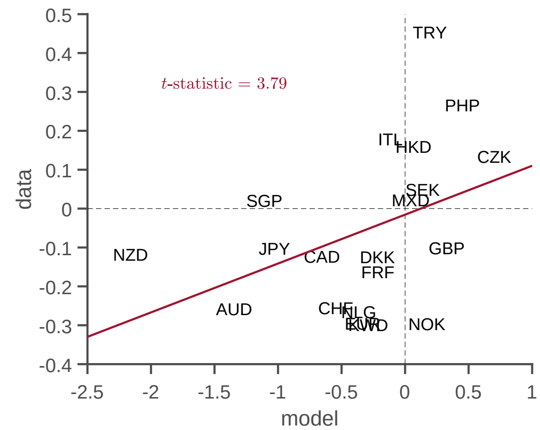

(b) $h = 40$

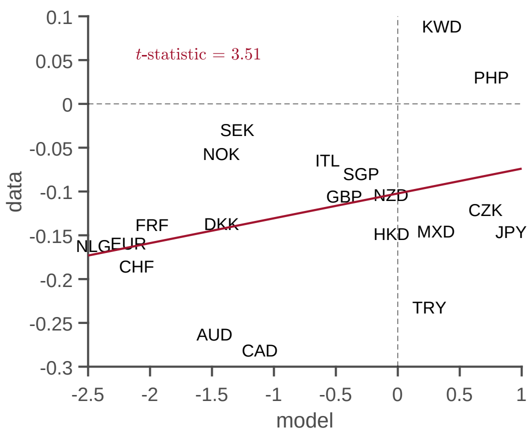

(c) $h = 80$

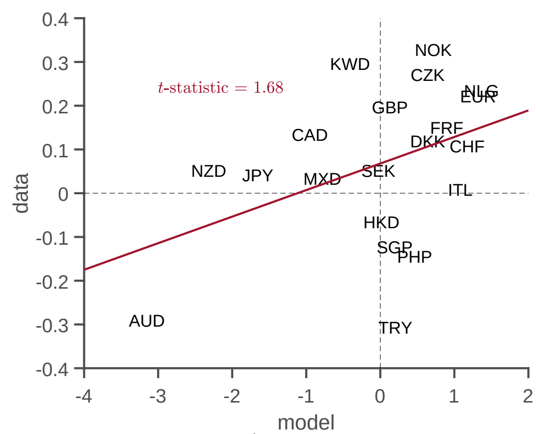

(d) $h = 120$

Figure 4. Fama Coefficient Cross-Section

Notes: This figure plots the empirically-estimated Fama regression coefficients against the corresponding model-implied coefficients for 21 currencies against the U.S. dollar at $h = 2$, 40, 80, and 120 month horizons. The red line in each panel indicates the corresponding least-square fit. The currencies are Australia, Canada, Hong Kong, Denmark, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, Singapore, Sweden, Switzerland, the United Kingdom, Czech Republic, Kuwait, Mexico, the Philippines, Spain, and Turkey.

This positive relationship holds both at short horizons (such as two months) when most Fama coefficients are positive, as well as for the longer horizons (such as 80 months) when coefficients tend to be mostly negative. We therefore conclude that in a world where agents are constrained by their ability to entertain high-dimensional models, constrained-rational expectations can generate patterns of UIP violations that are jointly consistent with the forward discount and predictability reversal puzzles.

As a final exercise, we test whether the above findings are robust to the introduction of heterogeneity in the number of factors in agents' models. To this end, we use the characterization in Proposition 4 to calculate the model-implied slope coefficients of the Fama regression in a heterogeneous-agent economy, in which a fraction $1 - \lambda$ of agents are constrained to using a

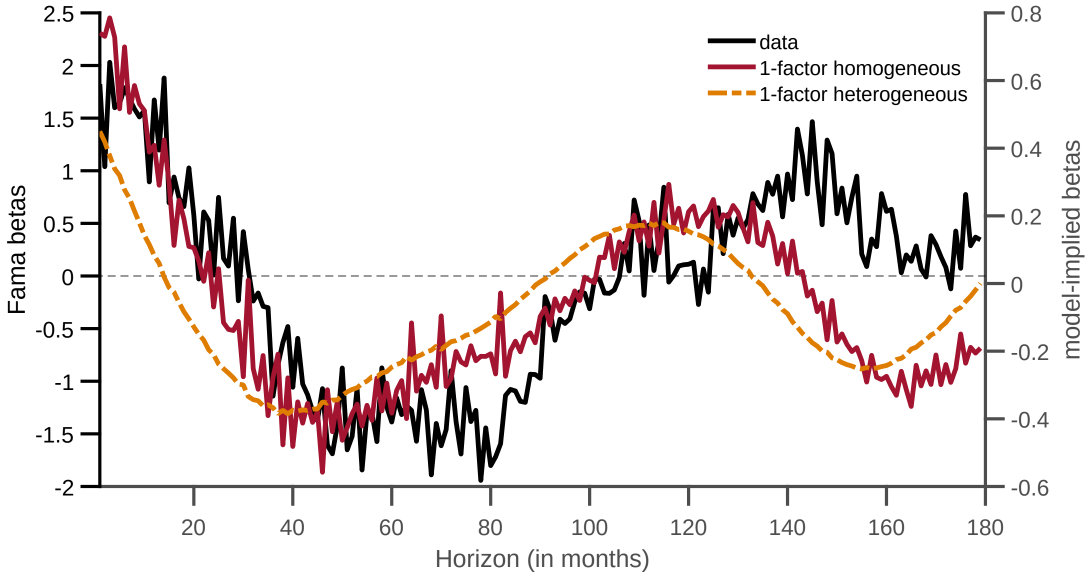

Figure 5. 1-Factor-Model-Implied Fama Coefficients Homogeneous and Heterogeneous Economy

Notes: This figure plots the single-factor model-implied slope coefficients of the Fama regression for horizons 0 to 180 months in a representative- and heterogeneous-agent economies (right axis) together with the corresponding coefficients estimated from the data (left axis). The fraction of behavioral agents in the heterogenous-agent economy is 10%, i.e., $\lambda = 0.9$.

single-factor model, while the remaining λ fraction can entertain models with any number of factors. Figure 5 plots estimated coefficients in an economy populated by 90% rational and 10% behavioral agents, i.e., for $\lambda = 0.9$. As the figure illustrates, the model-implied slope coefficients in the heterogeneous-agent economy look very similar to those in the representative-agent economy consisting of only behavioral agents ($\lambda = 0$). This indicates that even small fractions of behavioral agents can lead to notable deviations from the rational expectations benchmark, generating patterns that are consistent with the slope coefficients in the data.

4.2 Time-Series Momentum and Reversal in Equity Returns

One of the starkest challenges to the "random walk hypothesis" of asset prices is the existence of time-series momentum and reversal, whereby past returns predict future returns. For example, Moskowitz, Ooi, and Pedersen (2012) document that returns of a diverse set of futures and forward contracts exhibit persistence for one to 12 months, an effect that partially reverses over longer horizons.

As a second illustration of our framework, we focus on time-series momentum and reversal in equity returns. As in the previous subsection, we start by reproducing the empirical findings that show the existence of return predictability. We then apply our theoretical results from Section 3 and compare the degree of return predictability in the data to that implied by our framework.

We collect data on MSCI price and total return indices for Australia, Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, the United Kingdom, and the United States available from Datastream. Since volatility varies across the different country indices, we scale

Table 1. Autocorrelations for Equity Excess Returns

	Horizon (in months)			
	1–12	13–24	25–36	37–48
Australia	−0.0098	−0.0238	0.0072	−0.0001
Belgium	0.0325	−0.0014	−0.0029	−0.0153
Canada	0.0073	−0.0346	−0.0040	0.0026
France	0.0104	−0.0217	−0.0016	−0.0095
Germany	0.0036	−0.0154	−0.0056	0.0028
Italy	0.0241	−0.0149	−0.0089	−0.0197
Japan	0.0539	−0.0132	−0.0081	−0.0211
The Netherlands	−0.0019	−0.0104	0.0018	0.0014
Sweden	0.0148	−0.0338	−0.0008	0.0009
Switzerland	0.0114	−0.0100	−0.0022	−0.0080
United Kingdom	0.0047	−0.0063	−0.0153	−0.0154
United States	0.0107	−0.0096	0.0076	−0.0085

Notes: This table reports average autocorrelations of MSCI country excess returns for the 12 months indicated time period.

excess returns by their lagged volatility as in Moskowitz, Ooi, and Pedersen (2012).[11] From the raw return series, we also calculate dividend-price ratios.[12]

As a first exercise, we follow Cutler, Poterba, and Summers (1991) and calculate average autocorrelations of returns across the different equity indices, with the results reported in Table 1. The first column illustrates the well-known pattern of positive serial correlation over horizons shorter than one year (with the exception of the Netherlands and Australia), indicating short-term time-series momentum. The average autocorrelations, however, turn negative over the horizon of 12–24 months, as is evident from the second column of Table 1, pointing towards reversals in excess returns at longer horizons.

We also test for the extent of return predictability by running a pooled panel regression, as in Moskowitz, Ooi, and Pedersen (2012), of the form

$$\mathrm{rx}^s_{t+h} = \alpha^{\mathrm{mom}}_h + \beta^{\mathrm{mom}}_h \mathrm{rx}^s_t + \epsilon^s_{t,h}, \tag{20}$$

where rx^s_t is the excess return of equity index s at time t. Figure 6 plots the slope coefficients of the pooled regression for $h = 1, 2, \ldots, 40$ months. The results echo the findings in Table 1: the estimated slope coefficients are positive up to ten months (thus indicating short-term momentum) and turn negative and significant at longer horizons (indicating long-term reversal).

Turning to our framework's predictions, we start by noting that, in this context, the fundamental x_t corresponds to the dividend-price ratio of an equity, while rx_t is the corresponding equity excess

[11] The conditional return volatilities are calculated using a GARCH(1,1) model. Our results remain unchanged without this normalization.

[12] Because dividends feature a strong seasonal component, we follow Ang and Bekaert (2006) and construct dividend time-series by summing up monthly dividends over the past year.

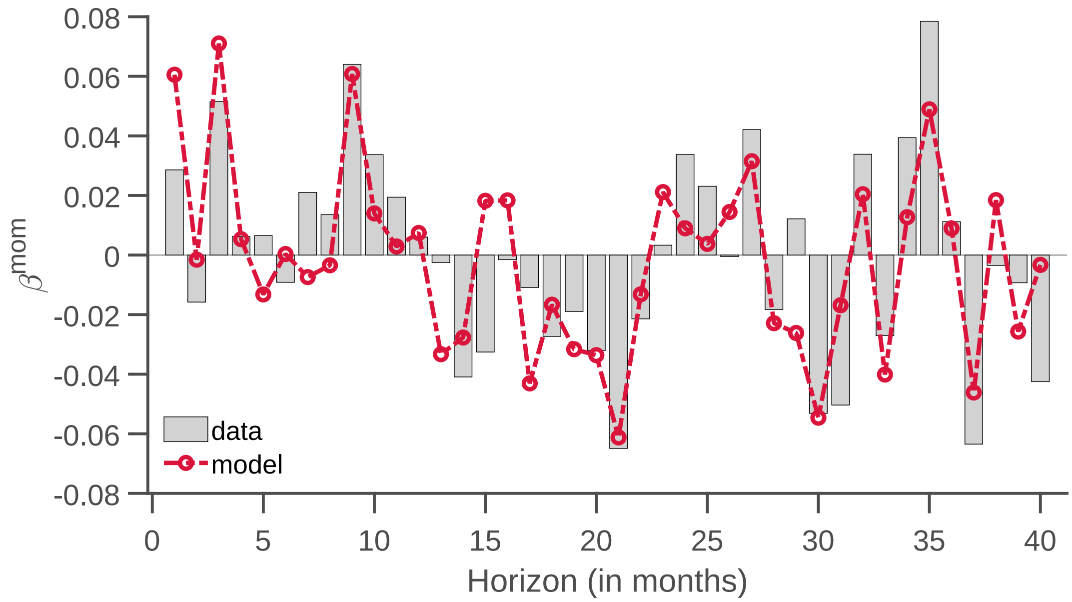

Figure 6. Momentum and Reversal: β^{mom}

Notes: This figure plots the slope coefficients of the pooled regression $\text{rx}^s_{t+h} = \alpha^{\text{mom}}_h + \beta^{\text{mom}}_h \text{rx}^s_t + \epsilon^s_{t,h}$ for $h = 1, \ldots, 40$ months and rx^s_t is the excess return of index s from the data (bars) and the model-implied regression coefficient.

return. Furthermore, equation (20) is the empirical counterpart of the momentum regression (7) in Section 2. Therefore, under the assumption that investors can only entertain single-factor models (i.e., $k = 1$), we can use the closed-form expression in Proposition 2 to calculate the model-implied slope coefficients of the momentum regression at different horizons. To this end, we calculate the autocorrelation function $\{\xi^*_\tau\}_{\tau \geq 1}$ of the dividend-price ratio for each of the equity indices in our sample and use (17) to obtain the model-implied slope coefficients of the momentum regression for each index at various horizons. Figure 6 plots the model-implied pooled slope coefficient and figure 7 plots the results against the corresponding slope coefficients estimated from the data country by country.The results clearly indicate that model-implied and empirically-estimated slope coefficients have a positive and statistically-significant relationship at short and long horizons. This relationship holds irrespective of whether the point estimates of the slope coefficients are positive or negative.[13]

We conclude by emphasizing that while the empirically-estimated slope coefficients are obtained from regressing excess returns on lagged returns, model-implied coefficients are calculated solely from the autocorrelation of dividend-price ratio (as prescribed by Proposition 2) and without using returns data.

[13]From Proposition 3, we already know that, as long as the dividend-price ratio exhibits short-run persistence, model-implied slope coefficients of the momentum regression should take opposite sings at different horizons. However, Figure 7 shows that our framework not only generates momentum and reversal, but also the resulting return predictability patterns match what is observed in the data.

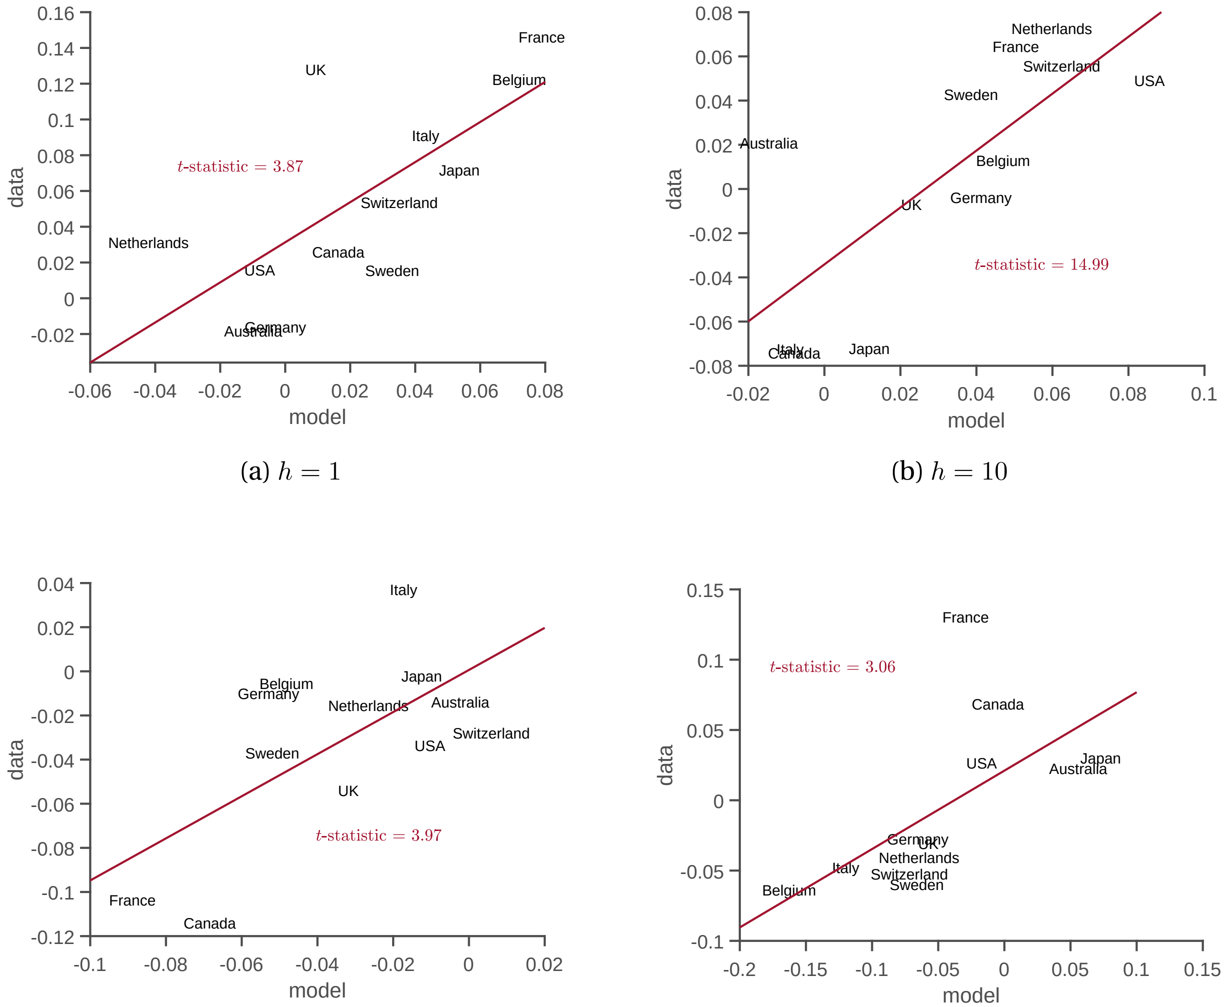

(a) $h = 1$

(b) $h = 10$

(c) $h = 20$

(d) $h = 40$

Figure 7. Momentum Coefficient Cross-Section

Notes: This figure plots the empirically-estimated momentum regression coefficients against the corresponding model-implied coefficients for twelve equity indices at horizons $h = 1, 10, 20$ and 40 months. The red line in each panel indicates the corresponding least-square fit. The country indices are Australia, Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, United Kingdom, and United States.

5 Conclusion

In this paper, we provide a framework to study the implications of model misspecification on asset prices. We develop a theoretical framework in which a sequence of exogenous fundamentals (in our applications, the interest rate differential or an asset's dividends) are generated by an n-factor model. While agents can observe the sequence of realized fundamentals, they neither observe nor know the underlying factors that drive them. As a result, they need to rely on their past observations to estimate a factor model that would allow them to make predictions about the future. As our main behavioral assumption, we posit that agents can only entertain models with at most k factors, where k may be distinct from the true number of factors. This assumption captures the idea that there is a limit to the complexity of statistical models that agents are able to consider. In this sense, the number of factors, k, indexes the degree of agents' sophistication, with a larger k corresponding to agents who can entertain a richer class of models. Importantly, whenever $k < n$, the agents will end up with a misspecified model of the world, irrespective of which model they end up using.

We then apply our framework to two applications in asset pricing: (i) the violations of the uncovered interest rate parity in foreign exchange and (ii) time-series momentum and reversal in equity returns. In both cases, we find that deviations from rational expectations generated by misspecified factor models can generate return predictability patterns consistent with the data.

A Proofs

Proof of Proposition 1

We first establish equation (8). From (6), it is immediate that the slope coefficient of the Fama regression at horizon h is given by $\beta_h^{\text{Fama}} = \mathbb{E}^*[x_t \, \text{rx}_{t+h}]/\mathbb{E}^*[x_t^2]$, where $\mathbb{E}^*[\cdot]$ denotes the expectation with respect to the true data-generating process. Furthermore, recall that excess returns satisfy (5). As a result, it is immediate that β_h^{Fama} satisfies (8).

Next, to establish (9), observe that equation (5) implies that

$$\mathbb{E}^*[\text{rx}_t \text{rx}_{t+h}] = \sum_{\tau=1}^{\infty}\sum_{s=1}^{\infty} \delta^{\tau+s} \left(\mathbb{E}_{t+1}[x_{t+\tau}] - \mathbb{E}_t[x_{t+\tau}]\right)\left(\mathbb{E}_{t+h+1}[x_{t+h+s}] - \mathbb{E}_{t+h}[x_{t+h+s}]\right).$$

Noting that the slope coefficient of the momentum regression is given by $\beta_h^{\text{mom}} = \mathbb{E}^*[\text{rx}_t \, \text{rx}_{t+h}]/\mathbb{E}^*[\text{rx}_t^2]$ then establishes (9). □

Proof of Theorem 1

To establish the existence of a constrained-rational expectations equilibrium, it is sufficient to show that $\text{argmin}_{\theta\in\Theta_k} \text{KL}(\theta^*\|\theta)$ is non-empty for all k and all θ^*. As a first observation, note that instead of optimizing over Θ_k, we can optimize over $\underline{\Theta}_1 \cup \underline{\Theta}_2 \cup \cdots \cup \underline{\Theta}_k$, where $\underline{\Theta}_r$ is the set models whose minimal realization consists of r factors. Therefore, in what follows, and without loss of generality, we assume that model $\theta = (A, b, \Sigma_\epsilon)$ is a minimal realization consisting of $r \le k$ factors.

Recall that, under model θ, agents believe that the fundamental is described by the process in (10), where $\omega_t \in \mathbb{R}^r$ is the vector of r hidden factors. As a result, conditional on $\{x_{t-\tau}\}_{\tau=0}^{\infty}$, agents believe that ω_{t+1} is normally distributed with mean $\hat{\omega}_t = \mathbb{E}_t[\omega_{t+1}]$ and variance $\hat{\Sigma}_\omega$, where $\hat{\Sigma}_\omega$ is the unique positive definite matrix that satisfies the algebraic Riccati equation

$$\hat{\Sigma}_\omega = A\left(\hat{\Sigma}_\omega - \frac{1}{b'\hat{\Sigma}_\omega b}\hat{\Sigma}_\omega b b'\hat{\Sigma}_\omega\right)A' + \Sigma_\epsilon, \tag{A.1}$$

$\hat{\omega}_t$ is defined recursively as $\hat{\omega}_t = (A - gb')\hat{\omega}_{t-1} + g x_t$, and $g \in \mathbb{R}^r$ is the Kalman gain given by

$$g = A\hat{\Sigma}_\omega b(b'\hat{\Sigma}_\omega b)^{-1}. \tag{A.2}$$

Conditional on $\{x_{t-\tau}\}_{\tau=0}^{\infty}$, agents believe that the fundamental x_{t+1} is normally distributed with mean $\mathbb{E}_t[x_{t+1}] = b'\hat{\omega}_t$ and variance $\hat{\sigma}_x^2 = b'\hat{\Sigma}_\omega b$. Furthermore, their s-step-ahead forecasts of the future realization of the fundamental is given by

$$\mathbb{E}_t[x_{t+s}] = b'A^{s-1}\sum_{\tau=0}^{\infty}(A - gb')^\tau g x_{t-\tau} \tag{A.3}$$

for all $s \ge 1$, where g is the Kalman gain in (A.2). The above expression implies that the KL divergence (11) of agents' model θ from the true data-generating process θ^* is given by

$$\begin{aligned}\text{KL}(\theta^*\|\theta) = &-\frac{1}{2}\log(\hat{\sigma}_x^{-2}) + \frac{1}{2}\log(2\pi) + \frac{1}{2}\hat{\sigma}_x^{-2}\Xi_0^* - \sum_{s=1}^{\infty}\hat{\sigma}_x^{-2}\Xi_s^* b'(A-gb')^{s-1}g \\ &+ \frac{1}{2}\sum_{s=1}^{\infty}\sum_{\tau=1}^{\infty}\hat{\sigma}_x^{-2} b'(A-gb')^{s-1}g\Xi_{\tau-s}^* b'(A-gb')^{\tau-1}g + \mathbb{E}^*[\log f^*(x_{t+1}|x_t,\dots)],\end{aligned} \tag{A.4}$$

where $\Xi_s^* = \mathbb{E}^*[x_t x_{t+s}]$ denotes the auto-covariance of the fundamental at lag s under the true process.

To minimize the KL divergence between agents' model and the true data-generating process over $\underline{\Theta}_r$, we normalize the model parameters via a change of variables. In particular, since $\hat{\Sigma}_\omega$ is positive definite, let

$$M = \hat{\Sigma}_\omega^{-1/2} A \hat{\Sigma}_\omega^{1/2} \qquad \text{and} \qquad u = \frac{\hat{\Sigma}_\omega^{1/2} b}{\sqrt{b' \hat{\Sigma}_\omega b}}. \tag{A.5}$$

Note that u is a r-dimensional vector of unit length and M is a $r \times r$ stable matrix. Given this change of variables, the agents' s-step ahead forecasts in (A.3) can be written as

$$\mathbb{E}_t[x_{t+s}] = u' M^{s-1} \sum_{\tau=0}^{\infty} [M(I - uu')]^\tau M u x_{t-\tau} \tag{A.6}$$

for all $s \geq 1$. Similarly, substituting for A, b, and g in terms of M and u in (A.4) implies that the KL divergence between agents' model and the true data-generating process is given by

$$\begin{aligned} \mathrm{KL}(\theta^* \| \theta) &= \mathbb{E}^*[\log f^*(x_{t+1} | x_t, \dots)] - \frac{1}{2} \log\left(\hat{\sigma}_x^{-2}\right) + \frac{1}{2} \log(2\pi) \\ &+ \frac{1}{2} \hat{\sigma}_x^{-2} \Xi_0^* - \sum_{s=1}^{\infty} \hat{\sigma}_x^{-2} \Xi_s^* u' [M(I - uu')]^{s-1} M u \\ &+ \frac{1}{2} \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \hat{\sigma}_x^{-2} u' [M(I - uu')]^{s-1} M u \Xi_{\tau-s}^* u' [M(I - uu')]^{\tau-1} M u. \end{aligned} \tag{A.7}$$

Given the one-to-one correspondence between (A, b, Σ_ϵ) and $(M, u, \hat{\sigma}_x^{-2})$, minimizing the above over $(M, u, \hat{\sigma}_x^{-2})$ is equivalent to minimizing (A.4) over (A, b, Σ_ϵ). We thus first minimize (A.7) with respect to $\hat{\sigma}_x^{-2}$. Taking the corresponding first-order conditions and plugging back the result into (A.7) implies that minimizing the KL divergence between agents' model and the true underlying model is equivalent to minimizing

$$H(M, u) = 1 - 2 \sum_{s=1}^{\infty} \phi_s \xi_s^* + \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \xi_{\tau-s}^* \phi_s \phi_\tau, \tag{A.8}$$

with respect to M and u, where $\xi_s^* = \Xi_s^* / \Xi_0^*$ denotes the true autocorrelation of the fundamental at lag s and $\phi_s = u'[M(I - uu')]^{s-1} M u$. Therefore, to establish that $\mathrm{argmin}_{\theta \in \cup_{r \leq k} \underline{\Theta}_r} \mathrm{KL}(\theta^* \| \theta)$ is non-empty for all k and all θ^*, it is sufficient to show that the minimum of $H(M, u)$ in (15) is always attained for some stable matrix M and unit vector u. But this is a simple consequence of the extreme value theorem and the fact that $H(M, u)$ is a continuous function of (M, u).

We next prove that constrained-rational expectations satisfy the law of iterated expectations. Recall that agents' forecasts of future realizations of fundamentals are given by (A.6). Therefore, for any $s \geq 1$,

$$\begin{aligned} \mathbb{E}_{t-1}[\mathbb{E}_t[x_{t+s}]] &= u' M^{s-1} \sum_{\tau=0}^{\infty} [M(I - uu')]^\tau M u \mathbb{E}_{t-1}[x_{t-\tau}] \\ &= u' M^{s-1} \sum_{\tau=1}^{\infty} [M(I - uu')]^\tau M u x_{t-\tau} + u' M^s u \mathbb{E}_{t-1}[x_t]. \end{aligned}$$

Using (A.6) one more time to express $\mathbb{E}_{t-1}[x_t]$ in terms of past realizations of the fundamental leads to

$$\begin{aligned}\mathbb{E}_{t-1}[\mathbb{E}_t[x_{t+s}]] &= u'M^{s-1}\sum_{\tau=0}^{\infty}[M(I-uu')]^{\tau+1}Mux_{t-1-\tau} + u'M^s uu'\sum_{\tau=0}^{\infty}[M(I-uu')]^{\tau}Mux_{t-1-\tau} \\ &= u'M^s\sum_{\tau=0}^{\infty}[M(I-uu')]^{\tau}Mux_{t-1-\tau}\end{aligned}$$

Note that the right-hand side of the above equation is simply equal to $\mathbb{E}_{t-1}[x_{t+s}]$, thus establishing that $\mathbb{E}_{t-1}[\mathbb{E}_t[x_{t+s}]] = \mathbb{E}_{t-1}[x_{t+s}]$ for all $s \geq 1$. Now, a simple inductive argument implies $\mathbb{E}_{t-r}[\mathbb{E}_t[x_{t+s}]] = \mathbb{E}_{t-r}[x_{t+s}]$ for all $r, s \geq 1$, thus establishing that constrained-rational expectations satisfy the law of iterated expectations. □

Proof of Theorem 2

Proof of part (a) Recall from the proof of Theorem 1 that minimizing the $\mathrm{KL}(\theta^*||\theta)$ for $\theta \in \Theta_k$ is equivalent to minimizing (A.8) with respect to M and u, where M is a square k-dimensional stable matrix and u is a k-dimensional vector of unit length.

Furthermore, we established that agents' forecasts of future realizations of fundamentals are given by (A.6) for all $s \geq 1$. Therefore, $\mathbb{E}_t[x_{t+1}] = \sum_{\tau=1}^{\infty}\phi_\tau x_{t+1-\tau}$, resulting in mean-squared forecast errors given by

$$\mathbb{E}^*[(x_{t+1}-\mathbb{E}_t[x_{t+1}])^2] = \mathbb{E}^*\left[\left(x_{t+1}-\sum_{\tau=1}^{\infty}\phi_\tau x_{t+1-\tau}\right)^2\right].$$

Consequently,

$$\begin{aligned}\mathbb{E}^*[(x_{t+1}-\mathbb{E}_t[x_{t+1}])^2] &= \mathbb{E}^*[x_{t+1}^2] - 2\sum_{\tau=1}^{\infty}\phi_\tau\mathbb{E}^*[x_{t+1}x_{t+1-\tau}] + \sum_{s=1}^{\infty}\sum_{\tau=1}^{\infty}\phi_s\phi_\tau\mathbb{E}^*[x_{t+1-s}x_{t+1-\tau}] \\ &= \mathbb{E}^*[x_{t+1}^2]H(M,u),\end{aligned}$$

where $H(M,u)$ is given by (A.8). Since $\mathbb{E}^*[x_{t+1}^2]$ only depends on the true data-generating process and is independent of agents' model θ, the above equation implies that minimizing the KL divergence is equivalent to minimizing the mean-squared forecast errors. □

Proof of Theorem 3

In the proof of Theorem 1, we already established that minimizing the KL divergence between agents' model and the true underlying model is equivalent to minimizing

$$H(M,u) = 1 - 2\sum_{s=1}^{\infty}\phi_s\xi_s^* + \sum_{s=1}^{\infty}\sum_{\tau=1}^{\infty}\xi_{\tau-s}^*\phi_s\phi_\tau,$$

with respect to M and u, where $\xi_s^* = \Xi_s^*/\Xi_0^*$ denotes the true autocorrelation of the fundamental at lag s and $\phi_s = u'\left[M\left(I-uu'\right)\right]^{s-1}Mu$. Therefore, minimizing the KL divergence between a k-factor model $\theta \in \Theta_k$ and the true underlying model θ^* is equivalent to minimizing (15) over M, and u.

To establish (13) and (14), recall that excess returns satisfy the recursive equation (5). As a result,

$$\mathrm{rx}_{t+h} = \delta x_{t+h} + \delta u'(I - \delta M)^{-1} \sum_{\tau=0}^{\infty} [M(I - uu')]^{\tau} M u (\delta x_{t+h-\tau} - x_{t+h-1-\tau}),$$

where we are using the fact that agents' forecasts of future realizations of fundamentals are given by (A.6). Rearranging terms, we obtain

$$\mathrm{rx}_{t+h} = \delta u'(I - \delta M)^{-1} u \left(x_{t+h} - \sum_{\tau=1}^{\infty} \phi_{\tau} x_{t+h-\tau} \right). \tag{A.9}$$

The above equation, coupled with the fact that $\beta_h^{\text{Fama}} = \mathbb{E}^*[x_t \, \mathrm{rx}_{t+h}]/\mathbb{E}^*[x_t^2]$, implies that the slope coefficient of the Fama regression (6) satisfies (13). Similarly, noting that $\beta_h^{\text{mom}} = \mathbb{E}^*[\mathrm{rx}_t \, \mathrm{rx}_{t+h}]/\mathbb{E}^*[\mathrm{rx}_t^2]$ and using (A.9) establishes that the slope coefficient of the momentum regression is given by (14). □

Proof of Theorem 4

Proof of part (a) When $k \geq n$, the true model is within the set of models considered by the agents, i.e., $\theta^* \in \Theta_n \subseteq \Theta_k$. As a result, $\theta_k^{\text{CREE}} = \theta^*$, which in turn implies that agents' subjective expectations coincide with rational expectations. Consequently, equation (5) implies that

$$\mathbb{E}_t^*[\mathrm{rx}_{t+h}] = \sum_{\tau=1}^{\infty} \delta^{\tau} \left(\mathbb{E}_t^* \mathbb{E}_{t+1}^*[x_{t+\tau}] - \mathbb{E}_t^*[x_{t+\tau}] \right).$$

Hence, by the law of iterated expectations, $\mathbb{E}_t^*[\mathrm{rx}_{t+h}] = 0$, which guarantees that $\beta_h^{\text{Fama}} = \beta_h^{\text{mom}} = 0$ for all $h \geq 1$. □

Proof of part (b) We first show that if $\beta_h^{\text{Fama}} = 0$ for all h, then $k \geq n$. Suppose to the contrary that $k < n$ and define ϕ_s and ϕ_s^* as in Theorem 3 for the k- and n-factor models, respectively. By part (a) of the theorem, the slope coefficients of the return predictability regression arising from the optimal n-factor model are equal to zero at all horizons, while by assumption, we have $\beta_h^{\text{Fama}} = 0$ for all h under the k-factor model. Therefore, by equation (13),

$$\xi_h^* = \sum_{\tau=1}^{\infty} \phi_{\tau} \xi_{h-\tau}^* \tag{A.10}$$

$$\xi_h^* = \sum_{\tau=1}^{\infty} \phi_{\tau}^* \xi_{h-\tau}^* \tag{A.11}$$

Multiplying both sides of the first equation by ϕ_h^* and the second by ϕ_h and summing over all h, we get

$$\sum_{h=1}^{\infty} \phi_h^* \xi_h^* = \sum_{h=1}^{\infty} \phi_h \xi_h^*. \tag{A.12}$$

Next, note equations (A.10) and (A.11) also imply that the objective function (15) evaluated at the optimal solution within the set of all k and n-factor models is, respectively, equal to

$$H = 1 - \sum_{s=1}^{\infty} \phi_s \xi_s^*$$
$$H^* = 1 - \sum_{s=1}^{\infty} \phi_s^* \xi_s^*.$$

Comparing the above two equations with (A.12) implies that $H = H^*$. Consequently, the k-factor model results in the same KL divergence to the data generating process as the n-factor model, whose KL divergence to the data-generating process is equal to zero by assumption. This means that data-generating process has a representation with $k < n$ factors, which contradicts the assumption that n is the number of factors in the minimal representation of the data-generating process.

As the final step of the proof, we show that if the slope coefficient of the momentum regression (7) is zero at all horizons, then $k \geq n$. It is sufficient to show that $\beta_h^{\text{mom}} = 0$ for all h implies that $\beta_h^{\text{Fama}} = 0$ for all h, as we can then use the result for the slope coefficients of the Fama regression proved earlier to conclude that $k \geq n$. To this end, note that if $\beta_h^{\text{mom}} = 0$ for all h, then equation (14) in Theorem 3 implies that

$$\xi_h^* - \sum_{s=1}^{\infty} \phi_s(\xi_{h-s}^* + \xi_{h+s}^*) + \sum_{s,\tau=1}^{\infty} \phi_s \phi_\tau \xi_{h+s-\tau}^* = 0$$

for all $h \geq 1$. Multiplying both sides of the above equation by $\delta u'(I - \delta M)^{-1}u$ and using (13) leads to

$$\beta_h^{\text{Fama}} - \sum_{s=1}^{\infty} \phi_s \beta_{h+s}^{\text{Fama}} = 0. \tag{A.13}$$

Define the sequence $(\gamma_1, \gamma_2, \dots)$ recursively as $\gamma_h = \phi_h + \sum_{\tau=1}^{h-1} \phi_{h-\tau}\gamma_\tau$. By (A.13),

$$\sum_{h=1}^{\infty} \gamma_h \beta_{h+r}^{\text{Fama}} - \sum_{h=1}^{\infty}\sum_{s=1}^{\infty} \gamma_h \phi_s \beta_{h+r+s}^{\text{Fama}} = 0$$

for all $r \geq 1$. Using the recursive definition of γ_h, we obtain

$$\sum_{h=1}^{\infty} \phi_h \beta_{h+r}^{\text{Fama}} + \sum_{h=1}^{\infty}\sum_{\tau=1}^{h-1} \phi_{h-\tau}\gamma_\tau \beta_{h+r}^{\text{Fama}} - \sum_{h=1}^{\infty}\sum_{s=1}^{\infty} \gamma_h \phi_s \beta_{h+r+s}^{\text{Fama}} = 0.$$

It is straightforward to verify that the second and the third terms on the left-hand side above add up to zero, whereas equation (A.13) implies that the first term is equal to β_r^{Fama}. This therefore establishes that $\beta_r^{\text{Fama}} = 0$ for all $r \geq 1$, which completes the proof. $\square$

Proof of Theorem 5

As a first observation, note that since the underlying process that generates the fundamental is stationary, its autocorrelation function $\xi_s^* = \mathbb{E}^*[x_t x_{t+s}]/\mathbb{E}^*[x_t^2]$ decays at an exponential rate as $s \to \infty$. Next, we show that $\phi_s = u'[M(I - uu')]^{s-1}Mu$ also decays at an exponential rate, where

M and u are given by (A.5). To this end, first note that, $M(I - uu') = \hat{\Sigma}_\omega^{-1/2}(A - gb')\hat{\Sigma}_\omega^{1/2}$. Therefore, it is sufficient to show that all eigenvalues of $A - gb'$ are inside the unit circle. Rewriting the algebraic Riccati equation in (A.1), we obtain

$$(A - gb')\hat{\Sigma}_\omega(A - gb')' - \hat{\Sigma}_\omega + \Sigma_\epsilon = 0,$$

which is a discrete Lyapunov equation in $A - gb'$. Since (A, b, Σ_ϵ) is the minimal representation of the state-space model, Kalman's decomposition theorem implies that (i) $\hat{\Sigma}_\omega$ is positive definite and (ii) the pair $(\Sigma_\epsilon^{1/2}, A)$ is controllable. Therefore, by Lyapunov's theorem, all eigenvalues of $A - gb'$ are inside the unit circle, thus guaranteeing that ϕ_s decays at an exponential rate as $s \to \infty$.

With the above in hand, we next show that $\lim_{h\to\infty} \beta_h^{\text{Fama}} = 0$. Recall from Theorem 3 that the slope coefficient of Fama regression satisfies (13). Therefore, by triangle inequality,

$$|\beta_h^{\text{Fama}}| \leq \delta|u'(I - \delta M)^{-1}u| \left(|\xi_h^*| + \sum_{\tau=1}^{h} |\phi_\tau||\xi_{h-\tau}^*| + \sum_{\tau=1}^{\infty} |\phi_{\tau+h}||\xi_\tau^*| \right).$$

Since ξ_s^* and ϕ_s converges to zero at exponential rates, there are constants $c_1, c_2 > 0$ and $\rho_1, \rho_2 < 1$ such that $|\xi_s^*| \leq c_1\rho_1^s$ and $|\phi_s| \leq c_2\rho_2^s$ for all s. Consequently,

$$|\beta_h^{\text{Fama}}| \leq \delta|u'(I - \delta M)^{-1}u| \left(c_1\rho_1^h + c_1c_2\rho_1^h \sum_{\tau=1}^{h}(\rho_2/\rho_1)^\tau + c_1c_2\rho_2^h \sum_{\tau=1}^{h}(\rho_1\rho_2)^\tau \right),$$

and as a result, $|\beta_h^{\text{Fama}}| \leq c_3h\rho_3^h$ for all h for some constant $c_3 > 0$ and $\rho_3 = \max\{\rho_1, \rho_2\}$. This inequality then guarantees that $\lim_{h\to\infty} \beta_h^{\text{Fama}} = 0$.

To establish that the slope coefficients of the momentum regression also converge to zero in long horizons, note that the characterization result in equations (13) and (14) implies that

$$\beta_h^{\text{mom}} = \frac{\beta_h^{\text{Fama}} - \sum_{s=1}^\infty \phi_s\beta_{h+s}^{\text{Fama}}}{\beta_0^{\text{Fama}} - \sum_{s=1}^\infty \phi_s\beta_s^{\text{Fama}}},$$

with the convention that $\beta_0^{\text{Fama}} = 1 - \sum_{s=1}^\infty \phi_s\xi_s^*$. Therefore, $|\beta_h^{\text{Fama}}| \leq c_3h\rho_3^h$ implies that

$$|\beta_h^{\text{mom}}| \leq \frac{\rho_3^h}{\beta_0^{\text{Fama}} - \sum_{s=1}^\infty \phi_s\beta_s^{\text{Fama}}} \left(c_3h + c_1c_3 \sum_{s=1}^{\infty}(h+s)(\rho_1\rho_3)^s \right).$$

As a result, $\lim_{h\to\infty} \beta_h^{\text{mom}} = 0$. $\square$

Proof of Proposition 2

Recall from Theorem 3 that the slope coefficients of the Fama and momentum regressions are given by (13) and (14), where the sequence $(\phi_1, \phi_2, \dots)$ is obtained by minimizing (15) over the $k \times k$ stable matrix M and the unit vector $u \in \mathbb{R}^k$. Therefore, when $k = 1$, it must be the case that M is a scalar, denoted by m, satisfying $|m| < 1$ and $u \in \{-1, 1\}$. This immediately implies that $\phi_1 = m$ and $\phi_s = 0$ for all $s \geq 2$. As a result, the expressions for the Fama and momentum regressions reduce to

$$\beta_h^{\text{Fama}} = \frac{\delta}{1 - \delta m}\left(\xi_h^* - m\xi_{h-1}^*\right) \tag{A.14}$$

$$\beta_h^{\text{mom}} = \frac{\xi_h^* - m(\xi_{h-1}^* + \xi_{h+1}^*) + m^2\xi_h^*}{1 - 2m\xi_1^* + m^2}, \tag{A.15}$$

respectively, while the objective function in (15) is given by

$$H(M,u) = 1 - 2m\xi_1^* + m^2.$$

Optimizing the above over $m \in (-1,1)$ implies that $m = \xi_1^*$. Plugging in the results into (A.14) and (A.15) then completes the proof. □

Proof of Proposition 3

Recall from Proposition 2 that when agents are restricted to single-factor models, the slope coefficient of the momentum regression at horizon h is given by (17). As a result,

$$\sum_{h=1}^{\infty}(\xi_1^*)^h\beta_h^{\text{mom}} = \frac{1}{1-\xi_1^{*2}}\sum_{h=1}^{\infty}\left(\xi_1^{*h}(\xi_h^* - \xi_1^*\xi_{h-1}^*) - \xi_1^{*h+1}(\xi_{h+1}^* - \xi_1^*\xi_h^*)\right).$$

Since $|\xi_1^*| < 1$, we have

$$\sum_{h=1}^{\infty}\xi_1^{*h}\beta_h^{\text{mom}} = \frac{\xi_1^*(\xi_1^* - \xi_1^*\xi_0^*)}{1-\xi_1^{*2}}.$$

The fact that $\xi_0^* = 1$ then guarantees that the right-hand side of the above equation is equal to zero. □

Proof of Proposition 4

Recall from Subsection 3.3 that excess returns in the heterogenous-agent economy are given by (18). On the other hand, since rational agents can fully construct the model used by behavioral agents, the regress of expectations in (18) satisfies

$$\bar{\mathbb{E}}_t\bar{\mathbb{E}}_{t+1}\ldots\bar{\mathbb{E}}_{t+\tau}[x_{t+\tau}] = \lambda^{\tau+1}\mathbb{E}_t^*[x_{t+\tau}] + (1-\lambda)\sum_{s=0}^{\tau}\lambda^s\mathbb{E}_t^*\mathbb{E}_{t+s}[x_{t+\tau}]$$

for all $\tau \geq 0$, where $\mathbb{E}[\cdot]$ and $\mathbb{E}^*[\cdot]$ are the subjective expectations of the behavioral and rational agents, respectively. Consequently, the excess returns in the heterogenous-agent economy are given by

$$\begin{aligned}\mathrm{rx}_t(\lambda) = (1-\lambda)\sum_{s=0}^{\infty}(\delta\lambda)^s\sum_{\tau=1}^{\infty}\delta^\tau\left(\mathbb{E}_t^*\mathbb{E}_{t+s}[x_{t+\tau+s-1}] - \mathbb{E}_{t-1}^*\mathbb{E}_{t+s-1}[x_{t+\tau+s-1}]\right)\\ + \sum_{\tau=1}^{\infty}(\delta\lambda)^\tau\left(\mathbb{E}_t^*[x_{t+\tau-1}] - \mathbb{E}_{t-1}^*[x_{t+\tau-1}]\right),\end{aligned}$$

where λ denotes the fraction of rational agents in the economy. Taking expectations from both sides of the above equation and using the expression in (5) for excess returns in the representative-agent economy therefore implies that

$$\mathbb{E}_{t-1}^*[\mathrm{rx}_t(\lambda)] = (1-\lambda)\sum_{s=0}^{\infty}(\delta\lambda)^s\mathbb{E}_{t-1}^*[\mathrm{rx}_{t+s}(0)]$$

That is, expected excess returns in the heterogenous-agent economy is the discounted sum of all future excess returns of a representative-agent economy only consisting of agents who can entertain k-factor models. Multiplying both sides of the above equation by x_{t-h} and taking expectations $\mathbb{E}^*[\cdot]$ then establishes the result. □

B Technical Appendix

B.1 Model-Implied Autocorrelation Function

In this appendix, we derive the expression for the model-implied autocorrelation function when agents are restricted to models consisting of at most k factors. Let $\theta = (A, b, \Sigma_\epsilon)$ denote the collection of parameters that represent agents' model in (10). To compute the perceived autocorrelation function as a function of the parameters of agents' models, first note that, for $s \geq 0$,

$$\Xi_s = \mathbb{E}[x_t x_{t-s}] = b'\mathbb{E}[\omega_t \omega_{t-s}']b = b'A^s\mathbb{E}[\omega_t\omega_t']b.$$

Using the change of variables from the proof of Theorem 3 implies that

$$\Xi_s = \hat{\sigma}_x^2 u' M^s \hat{\Sigma}_\omega^{-1/2}\mathbb{E}[\omega_t\omega_t']\hat{\Sigma}_\omega^{-1/2} u.$$

Therefore, to represent Ξ_s in terms of M and u, we need to find $\tilde{\Xi} = \hat{\Sigma}_\omega^{-1/2}\mathbb{E}[\omega_t\omega_t']\hat{\Sigma}_\omega^{-1/2}$. Equation (10) implies that $\mathbb{E}[\omega_t\omega_t'] = A\mathbb{E}[\omega_t\omega_t']A' + \Sigma_\epsilon$. Multiplying both sides of this equation from left and right by $\hat{\Sigma}_\omega^{-1/2}$ implies that $\tilde{\Xi} = M\tilde{\Xi}M' + \hat{\Sigma}_\omega^{-1/2}\Sigma_\epsilon\hat{\Sigma}_\omega^{-1/2}$. On the other hand, the algebraic Riccati equation in (A.1) can be written in terms of M and u as $M(I - uu')M' + \hat{\Sigma}_\omega^{-1/2}\Sigma_\epsilon\hat{\Sigma}_\omega^{-1/2} = I$. Combining the last two equations implies that $\tilde{\Xi}$ is the solution to the discrete Lyapunov equation:

$$\tilde{\Xi} = M\tilde{\Xi}M' + I - M(I - uu')M'. \tag{B.1}$$

Therefore, the model-implied autocorrelation function is given, for $s \geq 1$, by

$$\xi_s = \Xi_s/\Xi_0 = \frac{u'M^s\tilde{\Xi}u}{u'\tilde{\Xi}u}, \tag{B.2}$$

where M and u minimize the expression (15) and $\tilde{\Xi}$ is the solution to the discrete Lyapunov equation (B.1).

References

Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini (2016), "Stock market volatility and learning." *Journal of Finance*, 71(1), 33–82.

Allen, Franklin, Stephen Morris, and Hyun Song Shin (2006), "Beauty contests and iterated expectations in asset markets." *Review of Financial Studies*, 19(3), 719–752.

Ang, Andrew and Geert Bekaert (2006), "Stock return predictability: Is it there?" *Review of Financial Studies*, 20(3), 651–707.

Angeletos, George-Marios and Zhen Huo (2019), "Myopia and anchoring." NBER Working Paper No. 24545.

Angeletos, George-Marios and Chen Lian (2018), "Forward guidance without common knowledge." *American Economic Review*, 108(9), 2477–2512.

Angeletos, George-Marios and Karthik A. Sastry (2020), "Managing expectations: Instruments vs. targets." NBER Working Paper No. 25404.

Antoulas, Athanasios C. (2005), *Approximation of Large-Scale Dynamical Systems*. Society for Industrial and Applied Mathematics.

Bacchetta, Philippe and Eric van Wincoop (2010), "Infrequent portfolio decisions: A solution to the forward discount puzzle." *American Economic Review*, 100(3), 870–904.

Bacchetta, Philippe and Eric van Wincoop (2019), "Puzzling exchange rate dynamics and delayed portfolio adjustment." NBER Working Paper, No. 26259.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2015), "X-CAPM: An extrapolative capital asset pricing model." *Journal of Financial Economics*, 115(1), 1–24.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998), "A model of investors sentiment." *Journal of Financial Economics*, 49(3), 307–343.

Barillas, Francisco and Kristoffer P. Nimark (2017), "Speculation and the term structure of interest rates." *Review of Financial Studies*, 30(11), 4003–4037.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019), "Diagnostic expectations and stock returns." *Journal of Finance*, 74(6), 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, Yueren Ma, and Andrei Shleifer (2020), "Over-reaction in macroeconomic expectations." *American Economic Review*, 110(9), 2748–2782.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), "Diagnostic expectations and credit cycles." *Journal of Finance*, 73(1), 199–227.

Campbell, John Y. (1991), "A variance decomposition for stock returns." *Economic Journal*, 101(405), 157–179.

Choi, James J. and Thomas M. Mertens (2019), "Extrapolative expectations and the equity premium." Working paper, Yale University.

Cutler, David M., James M. Poterba, and Lawrence H. Summers (1991), "Speculative dynamics." *Review of Economic Studies*, 58(3), 529–546.

da Silveira, Rava Azeredo, Yeji Sung, and Michael Woodford (2020), "Optimally imprecise memory and biased forecasts." Working paper.

Engel, Charles (2016), "Exchange rates, interest rates, and the risk premium." *American Economic Review*, 106(2), 436–474.

Esponda, Ignacio and Demian Pouzo (2016), "Berk–Nash equilibrium: A framework for modeling agents with misspecified models." *Econometrica*, 84(3), 1093–1130.

Esponda, Ignacio and Demian Pouzo (2019), "Equilibrium in misspecified Markov decision processes." Working Paper, UC Santa Barbara.

Fama, Eugene F. (1984), "Forward and spot exchange rates." *Journal of Monetary Economics*, 14(3), 319–338.

Farhi, Emmanuel and Iván Werning (2019), "Monetary policy, bounded rationality, and incomplete markets." *American Economic Review*, 109(11), 3887–3928.

Fuster, Andreas, Benjamin Hebert, and David Laibson (2012), "Natural expectations, macroeconomic dynamics, and asset pricing." *NBER Macroeconomics Annual*, 26, 1–48.

Fuster, Andreas, David Laibson, and Brock Mendel (2010), "Natural expectations and macroeconomic fluctuations." *Journal of Economic Perspectives*, 24(4), 67–84.

Gabaix, Xavier (2014), "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics*, 129(4), 1661–1710.

Gabaix, Xavier (2017), "Behavioral macroeconomics via sparse dynamic programming." NBER Working Paper No. 21848.

Guo, Hongye and Jessica A. Wachter (2019), "Superstitious investors." Working Paper, University of Pennsylvania.

Hirshleifer, David, Jun Li, and Jianfeng Yu (2015), "Asset pricing in production economies with extrapolative expectations." *Journal of Monetary Economics*, 76, 87–106.

Hong, Harrison, Jeremy C. Stein, and Jialin Yu (2007), "Simple Forecasts and Paradigm Shifts." *Journal of Finance*, 62(3), 1207–1242.

Martin, Ian and Stefan Nagel (2020), “Market efficiency in the age of big data.” Working Paper, London School of Economics.

Molavi, Pooya (2019), “Macroeconomics with learning and misspecification: A general theory and applications.” Working Paper, Northwestern University.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen (2012), “Time series momentum.” *Journal of Financial Economics*, 104(2), 228–250.

Nagel, Stefan and Zhengyang Xu (2019), “Asset pricing with fading memory.” NBER Working Paper No. 26255.

Newey, Whitney K. and Kenneth D. West (1987), “A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix.” *Econometrica*, 55(3), 703–708.

Rabin, Matthew and Dimitri Vayanos (2010), “The gambler’s and hot-hand fallacies: Theory and applications.” *Review of Economic Studies*, 77(2), 730–778.

Sandberg, Henrik (2019), “Introduction to model order reduction.” Lecture notes, http://www.diva-portal.org/smash/get/diva2:1338145/FULLTEXT01.pdf.

Valchev, Rosen (2020), “Bond convenience yields and exchange rate dynamics.” *American Economic Journal: Macroeconomics*, 12(2), 124–166.