

The Imperfect Knowledge Forecasting Framework and Premia on Foreign Exchange:

A New Approach to Asset Markets*

Roman Frydman** and Michael D. Goldberg***

First Draft: 25 April 2003

This Draft: 20 February 2004****

JEL: F31

Keywords: Imperfect Knowledge, Forecasting Process, Individual
Rationality, Prospect Theory, Exchange Rates, Risk Premium

* The authors are grateful to Charles Engel, Bruce Elmslie, Halina Frydman, Jan Gross, Katarina Juselius, Daniel Kahneman, Andreas Park, Ned Phelps, Andrzej Rapaczynski, Paul Romer, Stephan Schulmeister, George Soros and Larry Summers, James Wible and Mike Woodford for insightful comments on and reactions to an earlier draft of this paper. The authors also thank Marcella Frydman and Ken Murphy for editorial advice, as well as the C.V. Starr Center for Applied Economics at New York University and the Reginald F. Atkins Chair at the University of New Hampshire is for their support.

**Department of Economics, New York University, e-mail: roman.frydman@nyu.edu.

***Department of Economics, University of New Hampshire, e-mail: michael.goldberg@unh.edu

****This paper is a substantially revised and expanded version of the earlier draft which was entitled, "Imperfect Knowledge and Asset Price Dynamics: Modeling the Forecasting of Rational Agents, Dynamic Prospect Theory and Uncertainty Premia on Foreign Exchange", Institute of Economics, University of Copenhagen Discussion Paper #03-31, April, 2003.

Please do not circulate without permission from the authors.

Our capacity to predict will be confined to such general characteristics of the events to be expected and not include the capacity for predicting particular individual events....Compared with the precise predictions we have learnt to expect in the physical sciences, this sort of predictions is a second best...[However,] I am anxious to repeat, we will still achieve predictions which can be falsified and which therefore are of empirical significance....Yet the danger of which I want to warn is precisely the belief that in order to be accepted as scientific it is necessary to achieve more. This way lies charlatanism and more. I confess that I prefer true but imperfect knowledge...to a pretence of exact knowledge that is likely to be false. (Excerpts from the Nobel lecture of Friedrich Hayek, 1978, p.33).

1 Introduction

Capitalist economies have developed powerful incentives and institutional arrangements to motivate individuals to engage in profit-seeking activities. But in order to assess costs and benefits of alternative deployments of resources, economic agents have to come up with assessments of future payoffs. So, when an outside investigator – an economist – sets out to explain economic behavior, or the consequences of a government policy intervention, she has to model the forecasts of the participants in the economy.

In this paper, we propose a new framework for modeling individual forecasts in a world in which economic agents do not have access to *the* one invariant forecasting procedure that would deliver "superior" forecasts at all times. In such a world of *imperfect knowledge*, profit-seeking agents engage in a creative process that involves switches among the extant models and methods, as well as a discovery of new forecasting procedures. The approach we propose, which we call the *Imperfect Knowledge Forecasting* framework (IKF), aims to replace the *Rational Expectations Hypothesis* (REH) approach as a general framework for modeling the forecasting process in capitalist market economies.¹ We argue that REH presumes irrationality on the part of agents in a world of imperfect knowledge. In contrast, IKF is consistent with individual rationality.²

¹For a more general discussion of the related question of modeling profit-seeking behavior in capitalist market economies, see Frydman and Goldberg (2004a).

²For early analysis, see Frydman (1982) and Frydman and Phelps (1983). We restate these arguments using the formal framework developed in this paper in sections 3 and 4.

The distinctive feature of IKF is that it relies on *qualitative* restrictions to characterize the myriad of models and methods that agents could use or invent to form forecasts of payoff-relevant variables. The qualitative restrictions that comprise an IKF model are formalizations of behavioral insights. The IKF’s qualitative representations of agents’ forecasts replace the usual approaches, REH and non-REH alike, that characterize individual expectations with particular parametric models on which quantitative restrictions are imposed. The reliance on qualitative restrictions allows an economist to represent a large diversity of forecasting models and methods that individual agents could use or invent. Yet, contrary to the usual presumption — that quantitative restrictions are needed to impose “discipline” on the analysis — we show that the qualitative IKF framework leads to a rich array of testable implications.

We use the IKF framework to construct a model of the equilibrium premium in the foreign exchange market. Our specification of equilibrium makes use of our work in Frydman and Goldberg (2003a,b, 2004b). This research augments the original Kahneman and Tversky’s (1979, 1992) formulation of prospect theory with an assumption that makes the degree of loss aversion a function of the size of an agent’s speculative position. This *Endogenous Prospect Theory* implies that all agents require a premium to compensate them for their extra sensitivity to losses. As agents revise their forecasts, this premium varies over time.

In order to examine the behavior of the equilibrium premium over time, we follow the IKF approach and construct two qualitative models of forecast revisions. Our main model, which we call *the IKF Gap Model*, builds on an idea put forth by Keynes(1936). Drawing on his considerable experience in financial markets, Keynes observed that the gap between agents’ forecasts of an asset price and its perceived historical benchmark plays a key role in how agents’ revise their forecasts of this asset price and their assessments of the potential capital losses. The IKF gap model of individual forecasting behavior incorporates this insight in a formal way as a qualitative restriction imposed on our IKF representation. We show that these individual gap models lead to a positive relationship between the equilibrium uncertainty premium and an aggregate measure of the gap.

We also construct another model of individual forecasting behavior based on a different set of qualitative restrictions. In contrast to the gap effect model, the second model, which we call *the IKF Expectations Model*, is not motivated on behavioral grounds. The expectations model is comprised of a

set of qualitative restrictions that subsume, as a special case, a qualitative analog of REH . We show that these individual expectations models do not imply a systematic relationship between the (aggregate) equilibrium premium and the aggregate gap.

Thus, the IKF gap and expectations models generate competing implications for the relationship between the equilibrium premium and the aggregate gap. This shows that IKF models can generate testable implications. Moreover, we test the IKF analog of REH. against the behaviorally-based IKF gap model.

In particular, we test the implication that there is no relationship between the premium and the gap, which is implied the expectations model, against the positive relationship, which is implied by the gap model. We use survey data on exchange rate expectations from Money Market Services International (MMSI). The evidence strongly rejects the hypothesis that there is no relationship between the premium and the gap in favor of the alternative that the relationship is positive. Therefore, we conclude that the observed time path of the *ex ante* premium on foreign exchange is consistent with our IKF gap model.

We conjecture that this new empirical finding – that the gap from a historical benchmark plays an important role in understanding the dynamics of the premium on foreign exchange– may also be important in understanding the behavior of premia in other asset markets. We also show that our IKF-based model provides a simple explanation of the sign reversals that have been observed in the foreign exchange market.³ These positive findings stand in sharp contrast to the widely known difficulties of standard REH models of the risk premium in explaining the time path, volatility and sign reversals of foreign exchange premia over the modern period of floating.^{4,5}

³For example, see Mark and Wu (1998).

⁴For important review articles, see Lewis (1995) and Engel (1996).

⁵Moreover, we show in Frydman and Goldberg (2004c) that exchange rate models under IKF provide a unified explanation of other key features of exchange rate dynamics during the floating rate period that the REH approach has deemed anomalous.

1.1 IKF and Extant Approaches to the Modeling of Expectations: A Historical Sketch

Until about five decades ago, economists represented agents' forecasts by appealing, in an informal way, to psychological or sociological factors. Keynes (1936), for example, emphasized the role played by the subjective guesses of the average opinion in forming individual forecasts of asset prices. As macroeconomic models became more quantitative in the 1950s and 1960s, psychological and sociological factors were formalized by specific, parametric models for agents' expectations. For example, "adaptive expectations" were meant to capture a behavioral observation that agents tend to adjust their expectations slowly. However, as is clear from the seminal contribution of Phelps (1968), the macroeconomic models using adaptive expectations were only applied to particular episodes.⁶ Moreover, such analysis left open the question as to what other models might characterize agent's expectations. Formally, such rules for expectations were used to close macroeconomic models. However, pre-REH models remained de facto open because the implications of these models depended on an economist's choice of the exogenous rule for expectations.

The Rational Expectations revolution set out to ban from macroeconomics the apparent arbitrariness and particularity created by an analyst's choice of exogenous models for expectations. REH represents agents' forecasts to be identical to the mathematical expectation of the forecasted variable implied by the model an economist constructs. Thus, once an REH theorist chooses a model, expectations are completely determined.

This REH practice represents an agent's forecasting behavior at all times by an invariant parametric model. While REH instructs an economist to impose quantitative restrictions according to her model, REH provides no guidance as to which model out the many potential models an economist should construct.⁷ Therefore, REH alone provides no guidance as to how an economist should represent an agent's forecasting behavior.⁸

⁶The authors thank Ned Phelps for clarifying this point.

⁷These problems with REH also pose considerable, hitherto unnoticed, difficulties for the interpretation of the empirical tests of REH-based models. We take up these issues in Frydman and Goldberg (2004c).

⁸Moreover, once an economist chooses one macroeconomic model from the extant set or invents a new one, which is what economists do, she in effect replaces the episodic particularity, inherent in the pre-REH approaches, with the particularity stemming from

An REH theorist constructs her model under the assumption that economic agents are rational and then ties agents' expectations rigidly to her model. It might appear, therefore, that REH, itself, is a consequence of an assumption that agents act rationally in seeking to maximize their objective functions. But what seems to have been generally overlooked, is the fact that the rigid connection between agents' expectations and the particular model –from among the many possible – chosen by an economist does not follow from the basic postulate of rationality in economics. Rather, this rigid connection is merely assumed by fiat under REH.

The difficulties in sustaining the claim that REH follows from the postulate of individual rationality are indicative of a broader problem of establishing an objective standard of rational forecasting behavior in a world of imperfect knowledge. As Hayek argued in his trenchant critique of the possibility of socialist planning, there is a sharp distinction between individual rationality and “a problem...of the utilization of knowledge which is not given to anyone in its totality (Hayek 1948, p. 78).” In a world of imperfect knowledge, there is no mechanistic and unchanging standard of rationality of expectations that either the agents or economists can invoke.

In the absence of an objective standard of rationality, representations of forecasting behavior have to be based on observations of behavior in real-world markets. This is, indeed, the approach followed by the behavioral finance literature. However, the way in which this literature has incorporated behavioral findings into its models of expectations, shares an important feature with the REH approach. In order to formally represent empirical findings, behavioral finance theorists have modeled an agent's forecasting behavior with invariant parametric models.⁹

The premise of the IKF approach is that attempts to model forecasting behavior with fixed rules, even if based on well-established behavioral insights, are bound to fail empirically. While behaviorally-based, invariant parametric models for expectations do shed light on particular episodes of the empirical record, they do not offer a general approach that can replace the REH as a model of the forecasting process.¹⁰

an economist's choice of her model.

⁹The invariant parametric model may include a number of sub-models with a fixed (stochastic or deterministic) rule characterizing how agents switch between models. For further discussion, see sections 2 and 3 below.

¹⁰This lack of generality may explain why many economists are reluctant to abandon the REH.

The problems with the particular way in which behavioral insights have been represented in the behavioral finance literature, however, do not diminish the importance of these insights for modeling forecasting behavior. Rationality does not offer specific guidance to an economist on which model to construct. Thus, the role of behavioral findings in guiding this choice is all the more important.

The IKF approach formalizes behavioral findings in a new and rigorous way. Its premise is that at *different* points in time individual agents can, in general, form forecasts on the basis of different models and methods, formal or informal. Therefore, the qualitative restrictions that make up an IKF model represent an economist's characterization of the collection of an agent's forecasting procedures, rather than any particular model or method chosen by an agent at a specific point in time.

To implement the IKF approach, therefore, an economist must find qualitative restrictions that can characterize the collection of an agent's forecasting procedures. The IKF framework formalizes behavioral insights as such qualitative restrictions. After all, behavioral regularities do not pertain to just one particular forecasting model or method, but by their very nature, describe the qualitative properties that agents' forecasting procedures have in common. Agents in real-world markets, in general, form their forecasts on the basis of different models and methods. Thus, in our view, the best an economists can hope for is that the empirical regularities uncovered by behavioral researchers are sufficient to characterize the collections of forecasting procedures used by agents.

1.2 An Overview of the Paper

The remainder of the paper is structured as follows.

Section 2

In a companion paper, we model the speculative decision in asset markets solely on the basis of prospect theory and derive a new specification of equilibrium in the foreign exchange market (Frydman and Goldberg, 2004b). In this section, we sketch the logic behind this new momentary equilibrium condition.

Section 3

In this section, we develop a formal language for describing the individual forecasting process. We then use this language to define the concepts of

imperfect knowledge and rationality in a world of imperfect knowledge. We use our notion of imperfect knowledge to rationalize the heterogeneity of forecasts observed in real-world markets

Section 4

In this section, we discuss the implications of rationality for modeling agents' forecasts in a world of imperfect knowledge. After discussing the usual parametric representations, REH and non-REH alike, we provide an overview of the basic elements of the IKF framework. We formalize the way in which qualitative restrictions can characterize the properties of the entire collection of forecast models and methods used by an agent. We then define an IKF model of individual forecasting behavior to be a set of such qualitative restrictions. We also spell out the forms in which these qualitative restrictions can be formulated.

Section 5

The assumption that the degree of loss aversion increases with position size implies that all speculators require a premium on foreign exchange and that this premium depends on agent's forecasts of the potential capital loss from their open positions. In this section, we follow the IKF framework and construct two competing IKF models for these individual forecasts of the potential loss: the expectations model and the gap model. Both of these models impose only qualitative restrictions on the individual forecasting behavior. They differ, however, in the set of qualitative restrictions imposed

Section 6

In this section, we use our two IKF models for individual forecasts of the potential loss to derive two alternative models of the equilibrium premium on foreign exchange. We show that under mild aggregation conditions, our gap and expectations models imply competing implications for the relationship between the equilibrium premium and the aggregate gap.

Section 7

This section tests the expectations model of the equilibrium premium against the gap model. Our MMSI data consists of one-month forecasts of the German mark-U.S. dollar exchange rate from MMSI.

2 Endogenous Prospect Theory and Momentary Equilibrium in the Foreign Exchange Market

In a companion paper, we model the speculative decision in asset markets solely on the basis of prospect theory (Frydman and Goldberg, 2004b). To accomplish this task without violating the core assumptions of prospect theory, we augment the original formulation of Kahneman and Tversky (1992) with two assumptions: endogenous loss aversion and endogenous sensitivity. We find that our resulting formulation, which we call endogenous prospect theory, leads to a new specification for equilibrium in the foreign exchange market. In this section, we sketch the logic behind this new momentary equilibrium condition.

2.1 Basic Setup

Our model involves two countries, 1 and 2, and two types of nonmonetary assets, called A and B bonds. Bonds of type A and B are denominated in country-1 currency (to be referred to as the domestic currency) and country-2 currency, respectively. We impose the usual assumption of perfect capital mobility and assume, therefore, that agents can issue both A and B bonds costlessly and without limit. We express agent i 's real wealth at time t , W_t^i , in terms of country-1 currency, so that $W_t^i = (A_t^i + S_t B_t^i) / P_t$, where A_t^i and B_t^i denote the nominal value of agent i 's holdings of A and B bonds entering period t (i.e., at time $t - h$, where h denotes an infinitesimally small unit of time) and S_t and P_t denote the spot rate and domestic price level at time t , respectively.

We assume that each agent chooses a portfolio share of B bonds at time t , x_t^i , so as to maximize her objective function. We also assume that wealth holders evaluate gains and losses in terms of their own currency. Consequently, a country-1 wealth holder gains from an appreciation of the foreign currency when she holds some of her wealth in B bonds ($x_t^i > 0$) and loses when the value of her A bonds exceeds her wealth ($x_t^i < 0$).¹¹ A country-2 wealth holder, on the other hand, gains from an appreciation when the value of her B bonds exceeds her wealth ($x_t^i - 1 > 0$) and loses when she holds A

¹¹In this case, agent i issues B bonds in the market.

bonds ($x_t^i - 1 < 0$).

A country-1 (country-2) wealth holder, therefore, is long in foreign exchange when $x_t^i > 0$ ($x_t^i - 1 > 0$) and short when $x_t^i < 0$ ($x_t^i - 1 < 0$). We refer to agents with long and short positions as bulls and bears, respectively.

The decision problem faced by a speculator in the foreign exchange market involves a choice among three distinct allocations of her portfolio. An agent can choose to stay out of the foreign exchange market ($x_t^i = 0$ and $x_t^i - 1 = 0$ for a country-1 and -2 agent, respectively), take a long position, or take a short position. If an agent decides to enter, she must also decide on the size of her position. Because the decision problems for country-1 and -2 wealth holders are identical, we sketch our analysis from the point of view of country 1.

Following prospect theory, we assume that the carriers of value are gains and losses in wealth relative to a reference level. A “natural” reference level in the context of a decision to speculate is the level of wealth an agent would obtain if she stayed out of the market.¹² In the context of our model of foreign exchange speculation, this reference level is given by

$$\Gamma_t^i = W_t^i (1 + i_t^A - p_t) \quad \text{for all } i \text{ and } t \quad (1)$$

where i_t^A and p_t denote respectively the riskless nominal return on A bonds and country 1 inflation rate from time t to $t + 1$. To simplify, we assume that the rates of domestic and foreign inflation are deterministic.¹³

Using log approximations, we can write the change in wealth relative to the reference level in (1) as follows:

$$\begin{aligned} \Delta W_{t+1}^i &= W_t^i [(1 - x_t^i) (i_t^A - p_t) + x_t^i (s_{t+1} - s_t + i_t^B - p_t)] - \Gamma_t^i \\ &= W_t^i x_t^i r_{t+1} + W_t^i (1 + i_t^A - p_t) - \Gamma_t^i = x_t^i W_t^i r_{t+1} \end{aligned} \quad (2)$$

where i_t^B denotes the riskless nominal return on B bonds, $r_{t+1} = s_{t+1} - s_t + i_t^B - i_t^A$ denotes the return on B bonds in excess of the return on A bonds and s_t is the log level of S_t .

¹²Other studies that use this reference level in applying prospect theory to the study of asset prices include Barberis, Huang and Santos (2001) and Barberis and Huang (2001).

¹³This assumption of nonstochastic inflation rates in the short-run is common in the literature. See, for example, Krugman (1981) and Frankel (1982), as well as Dornbusch (1983).

Equation (2) implies that gains and losses in wealth relative to the reference level in (1) correspond to the positive and negative parts of the random variable r_{t+1} . For example, if a country 1 wealth holder is a bull and $r_{t+1} > 0$ ($r_{t+1} > 0$) then this bull will experience a gain (loss) in her wealth relative to her reference level. We can express a gain ($\Delta W_{t+1}^+ > 0$) and a loss ($\Delta W_{t+1}^- < 0$) from a long, L, and a short, S, position in foreign exchange as follows:

$$\Delta W_{t+1}^{+,L} = x_t^i W_t^i R_{t+1}^+ \quad \text{and} \quad \Delta W_{t+1}^{-,L} = x_t^i W_t^i R_{t+1}^- \quad \text{for} \quad x_t^i > 0 \quad (3)$$

$$\Delta W_{t+1}^{+,S} = x_t^i W_t^i R_{t+1}^- \quad \text{and} \quad \Delta W_{t+1}^{-,S} = x_t^i W_t^i R_{t+1}^+ \quad \text{for} \quad x_t^i < 0 \quad (4)$$

where we use a standard decomposition for R_{t+1} ,

$$R_{t+1} = R_{t+1}I(R_{t+1} > 0) + R_{t+1}I(R_{t+1} < 0) = R_{t+1}^+ + R_{t+1}^- \quad (5)$$

2.2 Endogenous Prospect Theory and Position Size at the Point in Time

Our specification of an agent's objective function builds on Kahneman and Tversky's seminal formulation of prospect theory. Kahneman and Tversky (1992) propose the following value function for a prospect $\Delta W_{k,t+1}^i$:

$$v_k = \begin{cases} (\Delta W_{k,t+1}^i)^\alpha & \text{for } \Delta W_{k,t+1}^i > 0 \\ -\lambda(-\Delta W_{k,t+1}^i)^\beta & \text{for } \Delta W_{k,t+1}^i < 0 \end{cases} \quad (6)$$

where $\Delta W_{k,t+1}^i = x_t^i W_t^i r_{k,t+1}^i$ is the k th *prospect* or k th realization of ΔW_{t+1} . One of the key assumptions of prospect theory is that "the marginal value of both gains and losses decreases with their size (Kahneman and Tversky, 1991)." This assumption of diminishing sensitivity implies that $\alpha < 1$ and $\beta < 1$ in (6). Although Kahneman and Tversky's formulation in (6) allows for the curvatures in the regions of gains and losses to be different, they set $\alpha = \beta$. Another key assumption of prospect theory is that the disutility from losses exceeds the utility from gains of the same size. This assumption of loss aversion implies a value of λ greater than 1 if $\alpha = \beta$.

We follow the cumulative prospect theory of Kahneman and Tversky (1992) and assume that there is a finite set of prospects and that an agent makes a decision on the basis of an aggregate of her prospects, dubbed *prospective utility*. Kahneman and Tversky assume that each agent uses

a set of decision weights, $\pi_{k,t}^i$, to aggregate her prospects. They assume that these decision weights increase monotonically with an agent's assessments of the likelihood of each prospect, $\phi_{k,t+1}^i$. We denote the set of $\phi_{k,t+1}^i$'s as \oplus_{t+1}^i . As in Kahneman and Tversky (1979), the $\pi_{k,t}^i$'s, which may not constitute a well-defined probability measure, do not equal to the $\phi_{k,t+1}^i$'s.

To obtain an expression for agent i 's prospective utility on a long position ($x_t^i > 0$), we substitute the definition of a gain and a loss from a long position in (3) into the utility function in (6), yielding:

$$\begin{aligned} V_t^{i,L} &= (x_t^i W_t^i)^\alpha \sum_k^{K^+} \pi_{k,t}^i (r_{k,t+1}^{+,i})^\alpha - \lambda^i (x_t^i W_t^i)^\beta \sum_k^{K^-} \pi_{k,t}^i (-r_{k,t+1}^{-,i})^\beta \\ &= (x_t^i W_t^i)^\alpha (P_t^i[\mathcal{G}^L] + \Lambda^i(x_t^i) P_t^i[\mathcal{L}^L]) \quad (7) \end{aligned}$$

where the function $\Lambda^i(x_t^i) = \lambda^i (x_t^i W_t^i)^{\beta-\alpha}$ is the degree of loss aversion and \mathcal{G}^L and \mathcal{L}^L are the sets of potential gains and potential losses from a unit long position in foreign exchange ($x_t^i W_t^i = 1$) to which agent i attaches nonzero weights at time t . (For ease of notation we do not index the sets \mathcal{G}^L and \mathcal{L}^L by i and $t + 1$). Because $V_t^{i,L}$ is the prospective utility from taking a long position of size $x_t^i W_t^i$, the sets \mathcal{G}^L and \mathcal{L}^L consist of the positive and negative potential realizations of the excess return on foreign exchange, respectively, i.e., the $r_{t+1,k}^{+,i}$'s and $r_{t+1,k}^{-,i}$'s. We denote the complete set of $r_{t+1,k}^{+,i}$'s and $r_{t+1,k}^{-,i}$'s as \mathcal{R}_{t+1}^i . We also refer to $P_t^i[\mathcal{G}^L]$ and $P_t^i[\mathcal{L}^L]$ as agent i 's *prospective potential gain* and *prospective potential loss* on a unit position, respectively.

In contrast to a long position, the set of potential gains, \mathcal{G}^S , consists of the *negative* realizations of the excess return (i.e, the $r_{t+1,k}^{-,i}$'s), whereas the set of potential losses, \mathcal{L}^S , consists of the *positive* realizations of the excess return (i.e., the $r_{t+1,k}^{+,i}$'s). Thus, the prospective utility from a short position in foreign exchange ($x_t^i < 0$) is as follows:

$$\begin{aligned} V_t^{i,S} &= (-x_t^i W_t^i)^\alpha \sum_k^{K^-} \pi_{k,t}^i (-r_{k,t+1}^{-,i})^\alpha - \lambda^i (-x_t^i W_t^i)^\beta \sum_k^{K^+} \pi_{k,t}^i (r_{k,t+1}^{+,i})^\beta \\ &= (-x_t^i W_t^i)^\alpha (P_t^i[\mathcal{G}^S] + \Lambda^i(-x_t^i) P_t^i[\mathcal{L}^S]) \quad (8) \end{aligned}$$

The literature on behavioral finance has followed Kahneman and Tversky and set $\alpha = \beta$ in (6). In Frydman and Goldberg (2004b), we show that if $\alpha = \beta$, then once an agent estimates that she can raise her prospective utility

from taking an open position, she will want to hold a position of unlimited size. Without limits to speculation, equilibrium is not well defined.

We show in Frydman and Goldberg (2004b), however, that if $\beta > \alpha$, an agent would limit the size of her speculative position. This result follows from the fact that in this case, the degree of loss aversion implicit in (6) is *not* constant, but is an increasing function of position size. Equations (7) and (8) reveal this result. We refer to the assumption that the degree of loss aversion is a function of a position size as *endogenous loss aversion*.¹⁴

Although Kahneman and Tversky's value function in (6) with $\beta > \alpha$ implies endogenous loss aversion and limits to speculation, it also implies that an agent is *loss-loving* rather than loss-averse for small position sizes, i.e., $\lambda^i (x_t^i W_t^i)^{\beta-\alpha} < 1$. In order to retain limits to speculation without violating the core assumption of loss aversion, we propose a simple reparameterization of the degree of loss aversion:

$$\Lambda^i [|x_t^i|] = \lambda_1^i + \lambda_2^i \frac{\alpha^i}{\alpha^i + 1} |x_t^i| \frac{1}{(-P_t^i[\mathcal{L}^L])} \quad (9)$$

where $\lambda_1^i > 1$ and $\lambda_2^i > 0$ are constants.¹⁵

The decision problem facing each agent at time t is to choose the portfolio share, x_t^i , that maximizes her prospective utility. We note that the prospective utilities, $V_t^{i,L}$ and $V_t^{i,S}$, are defined only for $x_t^i > 0$ and $x_t^i < 0$, respectively. Thus, agent i 's decision problem involves solving two constrained maximization problems, one for long positions using (7) and one for short positions using (8), assuming that the degree of loss aversion is given in (9). Agent i chooses, then, a long or short position depending on which strategy delivers a greater prospective utility.

The solution to this decision problem yields the following optimal port-

¹⁴The assumption of endogenous loss aversion formalizes anecdotal evidence that agents become more "nervous" about potential losses as they take larger and larger speculative positions. Myron Scholes, for example, emphasized this point in his presentation delivered at a conference on "Derivatives 2003: Reports from the Frontiers," Stern School of Business, January 2003. We are unaware of any experimental evidence demonstrating this intuitively appealing observation.

¹⁵We also note that our formulation of $\Lambda^i [|x_t^i|]$ in (9) implies that the marginal value of losses is endogenous (see Frydman and Goldberg, 2004b). We call this assumption *endogenous sensitivity*.

folio shares for a country 1 bull and bear:

$$x_t^{i,L} = \frac{P_t^{i,L}[\mathcal{R}_{t+1}] - (1 - \lambda_1^i)P_t^i[\mathcal{L}^L]}{\lambda_2^i} \quad \text{and} \quad x_t^{i,S} = \frac{P_t^{i,S}[\mathcal{R}_{t+1}] + (1 - \lambda_1^i)P_t^i[\mathcal{L}^S]}{\lambda_2^i} \quad (10)$$

where $P_t^{i,L}[\mathcal{R}_{t+1}] = P_t^i[\mathcal{G}^L] + P_t^i[\mathcal{L}^L]$ and $P_t^{i,S}[\mathcal{R}_{t+1}] = -(P_t^i[\mathcal{G}^S] + P_t^i[\mathcal{L}^S])$ denote the *prospective return* on a long and short position, respectively. An agent's prospective return under prospect theory plays the same role as her forecasted return does when standard preferences are used. Under standard preferences, agents take short positions when they forecast a negative excess return. We thus define the prospective return on a short position to be negative.

The optimal portfolio shares in (10) show that under endogenous prospect theory, speculators require a prospective return in excess of some minimum positive value in order to take open positions. These expressions also show that the minimum returns required by an agent for taking a long or short position, denoted by \widetilde{up}_t^L and \widetilde{up}_t^S , respectively, are simple functions of an agent's prospective loss and her loss-aversion parameter $\lambda_1^i > 1$:

$$\widetilde{up}_t^{i,L} = (1 - \lambda_1^i) P_t^i[\mathcal{L}^L] > 0 \quad \text{and} \quad \widetilde{up}_t^{i,S} = (1 - \lambda_1^i) P_t^i[\mathcal{L}^S] > 0 \quad (11)$$

Thus, all endogenously loss-averse speculators need to be compensated for their extra sensitivity to losses. Because $\widetilde{up}_t^{i,L}$ and $\widetilde{up}_t^{i,S}$ arise from agents' uncertainty concerning the magnitudes of the potential losses, we refer to these premia as *individual uncertainty premia* on a long and short position, respectively.

2.3 Momentary Equilibrium in the Foreign Exchange Market: Uncertainty Adjusted Uncovered Interest Parity (UAUIP)

The optimal portfolio shares in (10) can now be used to determine agent i 's demand to hold B bonds at time t , i.e., $B_{dt}^i = x_t^i \frac{S_t W_t^i}{P_t^i}$.¹⁶ Following common practice, we assume that domestic (foreign) money is held only by domestic (foreign) wealth holders. This assumption implies that whenever agent i

¹⁶These solutions can be used for a country-2 bull and bear because the former holds a long position in foreign exchange when the value of her B bonds exceeds her wealth ($x_t^i - 1 > 0$) and a short position when she holds A bonds ($x_t^i - 1 < 0$).

would like to increase her holdings of B bonds at time t ($B_{d_t}^i - B_t^i > 0$), she would have to purchase foreign exchange. Analogously, whenever an agent desires to reduce her holdings of B bonds at time t ($B_{d_t}^i - B_t^i < 0$), she would have to sell foreign exchange. Thus, a $B_{d_t}^i - B_t^i > 0$ ($B_{d_t}^i - B_t^i < 0$) is tantamount to a simultaneous demand for (supply of) foreign exchange of the same size. In case an agent does not want to alter her holdings of B bonds at time t (i.e. $B_{d_t}^i = B_t^i$), she neither demands nor supplies foreign exchange.

Equilibrium in the foreign exchange market can now be written as follows:

$$\sum_i (B_{d_t}^i - B_t^i) = B_{d_t} - B_t = 0 \quad (12)$$

We assume that domestic and foreign interest rates are exogenous to the spot rate process and determined by conditions of monetary equilibrium. This assumption, which is common in the literature (e.g., see Branson and Henderson, 1985), allows us to interpret (12) as a condition for the equilibrium spot rate.

This equilibrium condition for the spot rate, after plugging in the solutions for the $B_{d_t}^i$'s, can be written in terms of an aggregate of individual prospective returns, $P_t[\mathcal{R}_{t+1}]$, and an aggregate of individual prospective potential losses, $P_t[\mathcal{L}]$. In Frydman and Goldberg (2004b), we provide a way to relate prospective and forecasted values without abandoning core features of prospect theory.¹⁷ We assume, on behavioral grounds, that when an agent revises her \mathcal{R}_{t+1}^i , \oplus_{t+1}^i and her $\pi_{k,t}^i$'s, she does so in a way that leads to movements of forecasted values in the same direction as prospective values.

This assumption, together with reasonable distributional assumptions on how agents' prospective returns respond to their forecasted returns, allow us to express equilibrium in the foreign exchange market as an equality between an aggregate (or market) forecast of the excess return, \tilde{r}_{t+1} , and an aggregate premium that depends on agents' forecasts of potential unit losses and the net foreign asset position of country 1:

$$\tilde{r}_{t+1} = \tilde{p}r_t \quad (13)$$

¹⁷An alternative way to relate prospective and forecasted values is to drop the assumption of diminishing sensitivity in (6) by setting $\alpha = \beta = 1$, and equating decision weights with likelihood weights, i.e., $\pi_{k,t}^i = \phi_{k,t+1}^i$ for all k . Studies that have pursued this approach include Barberis, Huang and Santos (2001) and Barberis and Huang (2001).

where

$$\tilde{p}r_t = \tilde{u}p_t + \frac{\lambda_2}{\gamma} NFA(S_t) \quad (14)$$

$$\tilde{u}p_t = \tilde{u}p_t^L - \tilde{u}p_t^S = \frac{1 - \lambda_1}{\gamma} \left(\sum_i^{M^L} \theta_t^i P_t^{F_i} \left(\tilde{l}_{t+1}^{i,L} \right) + \sum_i^{M^S} \theta_t^i P_t^{F_i} \left(\tilde{l}_{t+1}^{i,S} \right) \right) \quad (15)$$

$NFA(S_t) = \frac{S_t B_t^1 - A_t^2}{P_t W_t^M}$, B_t^1 and A_t^2 denote respectively the value of B and A bonds held by country 1 and 2 wealth holders at time t , M^L and M^S denote, respectively, the number of bulls and bears at time t , W_t^M is the aggregate wealth of all bulls and bears at time t , \tilde{r}_{t+1} denotes the aggregate or market forecast of r_{t+1} , $\tilde{u}p_t^L$ and $\tilde{u}p_t^S$ are, respectively, aggregates of the individual uncertainty premia of bulls and bears, λ_1 and λ_2 are aggregates of the individual λ_1^i 's and λ_2^i 's, $\tilde{l}_{t+1}^{i,L}$ and $\tilde{l}_{t+1}^{i,S}$ denote, respectively, an individual forecast of the potential loss on a long and short position, $P_t^{F_i}$ is an increasing function that maps forecasted values into prospective values and θ_t^i are aggregation weights that depend on wealth shares and the individual λ_1^i 's and λ_2^i 's. We define individual forecasts of the potential loss from speculation as follows:

$$\tilde{l}_{t+1}^{i,L} = F_t^i[\mathcal{L}^L] = \sum_k^{K^-} \phi_{k,t+1}^i r_{k,t+1}^{-,i} \quad \text{and} \quad \tilde{l}_{t+1}^{i,S} = F_t^i[\mathcal{L}^S] = - \sum_k^{K^-} \phi_{k,t+1}^i r_{k,t+1}^{+,i} \quad (16)$$

In general, the equilibrium premium on foreign exchange will be nonzero. This premium arises because all endogenously loss-averse agents require a premium to compensate them for their extra sensitivity to losses. As it stands, however, equation (13), which we call *uncertainty adjusted uncovered interest rate parity* (UAUIP), says little about the equilibrium premium or the exchange rate. A rise in the forecasted return on long positions, \tilde{r}_{t+1} , *ceteris paribus*, creates an incipient capital flow into foreign exchange and a rise in the exchange rate. But a rise in \tilde{r}_{t+1} may be associated with a rise in $\tilde{p}r_t$, which, when considered together, could lead to an excess supply rather than an excess demand for foreign exchange. Moreover, the equilibrium movement of the exchange rate depends on how \tilde{r}_{t+1} and $\tilde{p}r_t$ change as s_t changes. Thus, to convert UAUIP into a model of the equilibrium premium and exchange rate, we need to model how \tilde{r}_{t+1} and $\tilde{p}r_t$ depend on s_t and other factors that agents deem relevant.

Our focus in this paper is on modeling the behavior of the equilibrium premium, and not on developing a complete model of the equilibrium exchange

rate. Consequently, while we follow our IKF framework to develop complete models of the individual \tilde{l}_{t+1}^i 's, we model only the endogenous component of \tilde{r}_{t+1} , i.e., the way in which the \tilde{r}_{t+1}^i 's depend on s_t .¹⁸ This partial modeling of the \tilde{r}_{t+1}^i 's enables us to model the equilibrium premium on foreign exchange in terms of the exogenous component of \tilde{r}_{t+1} .

In the next section, we begin this task by developing a formal language to describe the individual forecasting process in a world of imperfect knowledge.

3 Economic Rationality and a Heterogeneity of Forecasts in a World of Imperfect Knowledge

In real-world markets, economics agents must choose from among a myriad of existing models and methods or invent new models in order to form forecasts of payoff relevant variables. The IKF framework provides a way for an economists to model this creative process of model selection and discovery. In this section, we develop a formal language for describing individual forecasting models and their revisions. We then use this language to define the concept of *imperfect knowledge*, which underpins our approach to the modeling of the forecasting process. We also adopt a standard notion of rationality. We argue that rational economic agents who must cope in a world of imperfect knowledge do not adhere to one parametric forecasting model endlessly. Moreover, rationality implies that there exists a heterogeneity of forecasting models used by agents.

3.1 Diversity of Agent's Forecasting Mappings: Knowledge versus Information

We present a formal description of an agent's forecasting process in terms of four basic components:

1. An Agent's Information Set

¹⁸We model the exogenous component of \tilde{r}_{t+1} using our IKF framework in other research. See Goldberg and Frydman (1996a) and Frydman and Goldberg (2003a), where we use our IKF framework to develop models capable of explaining long swings in the exchange rate.

We denote the variables in an agent’s information set by a vector \mathcal{X}_t^i . At each point in time an agent selects a subset of the variables in \mathcal{X}_t^i , denoted by \mathcal{X}_t^{i,j_t} which she uses to construct her forecasting model. We note that the set \mathcal{X}_t^i contains current (time t) and/or lagged variables. We denote the realizations of the set of variables available at time t by a vector x_t^i .

In general, the variables constituting \mathcal{X}_t^i include fundamental variables (e.g. money supply or GDP growth), denoted by $\mathcal{X}_t^{i,f}$, as well as non-fundamental factors (e.g., those based on technical trading, market “sentiment” or policy announcements). We represent measures of each of the non-fundamental factors by a set of variables, denoted by $\mathcal{X}_t^{i,nf}$.

We note that in the usual formulation, an agent is assumed to update the values of the variables in \mathcal{X}_t^{i,j_t} , but not its composition. In our setup, agents may also alter the composition of their information sets that they use over time, that is add variables to or drop variables from \mathcal{X}_t^{i,j_t} .

2. An Agent’s Collection of Forecast Mappings

To form her time- t assessments of the prospects from speculation, an agent maps observations on the variables in $\mathcal{X}_t^{i,j}$ into $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$. As we discuss in the next section, there are a myriad of mappings that an agent can use to estimate her prospects. We refer to the mappings that an agent deems relevant as a her *collection of forecast mappings*, which we index by the letter j .

We emphasize that an agent’s collection of forecast mappings may include formal (statistical) procedures and/or informal (intuitive) ways to map $\mathcal{X}_t^{i,j}$ into $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$. Thus, our formulation is with a rich variety of methods and models used by agents in real-world markets. As Keynes (1936) put it

We are merely reminding ourselves that human decisions affecting the future, whether personal or political or economic, cannot depend on strict mathematical expectation, since the basis for making such calculations does not exist; and... that *our rational selves* [are] choosing between alternatives as best as we are able, calculating where we can, but often falling back for our motive on whim or sentiment or chance (Keynes [1936], p.162, emphasis added).

Formally, we represent one of the forecast mappings in agent i ’s collection, $\mathcal{F}_t^{i,j}(\cdot)$, as follows

$$\mathcal{F}_t^{i,j}(\mathcal{X}_t^i, \theta_t^i) : \{\mathcal{X}_t^{i,j}, \theta_t^{i,j}\} \Rightarrow \{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\} \quad (17)$$

where $\theta_t^{i,j}$ denotes a set of parameters corresponding to mapping j and θ_t^i denotes the set of parameters contained in the collection of agent i 's forecast mappings. We also denote agent i 's collection of C_t^i forecast mapping $\mathcal{F}_t^i = \{\mathcal{F}_t^{i,1}(\cdot), \mathcal{F}_t^{i,2}(\cdot), \dots, \mathcal{F}_t^{i,C_t^i}(\cdot)\}$.¹⁹ To save on notation, we do not index the sets \mathcal{R}_{t+1}^i and Φ_{t+1}^i by j .

Definition 1 *Our characterization of a forecast mapping in (17) involves three components:*

- *the particular composition of the information set, $\mathcal{X}_t^{i,j}$*
- *specific values of parameters in $\theta_t^{i,j}$ and*
- *the particular method used by an agent to map $\mathcal{X}_t^{i,j}$ into $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$.*²⁰

3. An Agent's Forecast

To form her assessments $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$, an agent chooses one ($j = j_t$) forecast mapping, $\mathcal{F}_t^{i,j_t}(\cdot)$, from her collection, \mathcal{F}_t^i . Having generated her assessments of the $r_{k,t+1}^{+,i}$'s, $r_{k,t+1}^{-,i}$'s and $\phi_{k,t+1}^{i'}$'s, an agent computes her forecasts of the excess return and the potential unit loss. We represent this computation by a forecast operator $F[\cdot]$. As defined in (16), a forecast operator uses likelihood weights in Φ_{t+1}^i to aggregate the appropriate elements in \mathcal{R}_{t+1}^i :

$$\tilde{r}_{t+1}^{i,j_t} = F \left[\mathcal{R}_{t+1}^i, \oplus_{t+1}^i | x_t^{i,j_t}, \theta_t^{i,j_t} \right] \quad (18)$$

and

$$\tilde{l}_{t+1}^{i,j_t} = F \left[\mathcal{L}_{t+1}^i, \oplus_{t+1}^{\mathcal{L}^i} | x_t^{i,j_t}, \theta_t^{i,j_t} \right] \quad (19)$$

where, as in (7), \mathcal{L}_{t+1}^i is a subset of \mathcal{R}_{t+1}^i and $\Phi_{t+1}^{\mathcal{L}^i}$ is the corresponding subset of Φ_{t+1}^i . The conditioning indicates that the values of the elements in the

¹⁹As in Kahneman and Tversky (1979), we do not assume that the likelihood weights in Φ_{t+1}^i necessarily sum to unity.

²⁰For example, the particular method used by an agent in constructing her forecast mapping may involve a specific mix of statistical and informal procedures, such as guesses.

sets \mathcal{R}_{t+1}^i and Φ_t^i depend on a vector $x_t^{i,j}$ of observations on variables in $\mathcal{X}_t^{i,j}$ and the values of $\theta_t^{i,j}$.

Using the actual value of next period's excess return, r_{t+1} , we can define a forecast error, $\epsilon_{t+1}^{i,j}$, as follows:

$$r_{t+1} = \tilde{r}_{t+1}^{i,j} + \epsilon_{t+1}^{i,j} \quad (20)$$

To illustrate this setup, we consider a special case in which $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$ can be interpreted as a conditional distribution. In this conventional REH-like setting, the next period's excess return, R_{t+1} , can be interpreted as a random variable and individual forecasts of the excess return and potential unit loss are given by

$$\tilde{r}_{t+1}^{i,j} = E[R_{t+1}|x_t^{i,j}, \theta_t^{i,j}]$$

and

$$\tilde{l}_{t+1}^{i,j} = E[R_{t+1}I(R_{t+1} < 0)|x_t^{i,j}, \theta_t^{i,j}] \quad (21)$$

where $I(R_{t+1} < 0) = 1$ if $R_{t+1} < 0$ and zero otherwise.

4. Revisions of an Agent's Forecast

Using (18), a revision of agent i 's forecast of the excess return can be written as

$$\tilde{r}_{t+1+h}^{i,j_{t+h}} - \tilde{r}_{t+1}^{i,j_t} = F \left[\mathcal{R}_{t+1+h}^i, \Phi_{t+1+h}^i | x_{t+h}^{i,j_{t+h}}, \theta_{t+h}^{i,j_{t+h}} \right] - F \left[\mathcal{R}_{t+1}^i, \Phi_{t+1}^i | x_t^{i,j_t}, \theta_t^{i,j_t} \right] \quad (22)$$

where $\tilde{r}_{t+1+h}^{i,j_{t+h}}$ is an agent's forecast of next period's excess return formed at time $(t+h)$. Thus an agent revises her forecast of the excess return by either

- altering the mapping that she uses in forming her forecast, i.e. $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot) \neq \mathcal{F}_t^{i,j_t}(\cdot)$; this may involve revisions of one or more of the three components of a mapping; and /or
- leaving her mapping unchanged, i.e. $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot) = \mathcal{F}_t^{i,j_t}(\cdot)$; this involves updating only the values of the variables constituting her information set.

We also note that whenever an agent switches to a new forecast mapping at $t+h$,

- she can either select $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot)$ to be a mapping that although different than $\mathcal{F}_t^{i,j_t}(\cdot)$, has been in her collection at time t , i.e. $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot) \in \mathcal{F}_t^i(\cdot)$; or
- she can discover at time $t+h$ a new mapping that she was not aware of at time t , i.e., $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot) \notin \mathcal{F}_t^i(\cdot)$ and $\mathcal{F}_{t+h}^i(\cdot) \neq \mathcal{F}_t^i(\cdot)$.

Thus, our formulation captures an apparent feature of the forecasting process in real-world markets: in revising their forecasts, agents not only switch from one extant mapping to another, but, in general, search for and occasionally discover new ways to forecast the future values of payoff relevant variables. This leads us to a notion of a forecast mapping with an invariant structure, which plays an important role in our analysis.

Definition 2 *Suppose that in revising her forecast, as in (22), an agent leaves her forecast mapping unchanged, i.e. $\mathcal{F}_{t+h}^{i,j_{t+h}}(\cdot) = \mathcal{F}_t^{i,j_t}(\cdot)$. If agent i 's forecast mapping satisfies this condition for all t , then we refer to $\mathcal{F}_t^{i,j_t}(\cdot)$ as an invariant mapping.*

As we noted above, forecast revisions based on an invariant mapping arise *solely* as a result of new information. Conventional REH provides the most important example of forecasting based of an invariant mapping. However, the use of models that represent agents' forecasts as invariant mappings is not confined to the standard REH approach. For example, following Engel and Hamilton (1990), recently developed REH-based models involve one over-arching mapping, which consists of a finite number of *pre-specified* forecasting models along with a rule governing how agents switch between these models in updating their forecasts. However, because the sub-models and the switching rule between them are pre-specified, the over-arching forecast mapping is in fact invariant.²¹

Moreover, invariant mappings have also been used by theorists who aim to depart from REH. A large class of learning models assumes that all agents learn and forecast on the basis of a *common* model and a fixed updating mechanism throughout the period of learning.(e.g., Evans and Honkapohja, 2001).

²¹For a more recent example, see Hansen and Sargent (2001a,b).

In an early seminal non-REH model of the foreign exchange market, Frankel and Froot (1987) model agents' forecasting using an invariant mapping: a representative agent updates his expectations by switching – according to a fixed (Bayesian) rule – between a chartist model and a fundamental model.²² Finally, as we discuss in the next subsection, recent forecasting models developed on the basis of behavioral findings also assume that agents use an invariant mapping in forming their forecasts.

The foregoing formal representation of individual forecasting models makes clear that there are two potential sources for the heterogeneity of forecasts:

- a diversity among the individual forecast mappings chosen by agents; and/or
- differences among agents in their estimates of the values of the variables that are common in their forecasting models.

As we discussed above, the REH approach attributes to agents an invariant mapping. We also note that under the standard REH approach, all agents form expectations based on *the common*, sometimes referred to as the “objective”, probability distribution. Thus, to rationalize a heterogeneity of expectations, conventional REH models *must* suppose that not all agents observe the same values of the variables appearing in the common forecast mapping.²³ The following example illustrates this point.

We assume that agents form their forecasts according to REH, which is a special case of a subjective probability model in (20). For simplicity, suppose that this common forecasting model is given by the following linear model with an invariant structure:

$$R_{t+1} = \mathcal{Z}_t\theta + \epsilon_{t+1} \tag{23}$$

where \mathcal{Z}_t is a set of random variables, θ denotes a vector of parameter values and $E[\epsilon_{t+1}|\mathcal{Z}_t] = 0$ for all t . Under REH, if all agents were to observe identical values of the variables in \mathcal{Z}_t , denoted by z_t their forecasts would be identical and given by

$$\tilde{r}_{t+1} = z_t\theta \tag{24}$$

²²For a related approach that also models the forecasting process using an invariant mapping, see Delong, Shleifer, Summers and Waldman (1990). This model consists of speculators who form expectations according to the REH, and feedback traders, who trade on the basis of simple rules.

²³See Lyons (2001) for such an approach in the foreign exchange market.

Thus, to rationalize a heterogeneity of forecasts, an REH theorist has to assume that individual agents differ in the way they measure or are able to observe z_t . Such differences among agents may be due to asymmetric information or some informational “misperceptions”.

To illustrate this point, suppose that, as in Lucas (1973), instead of \mathcal{Z}_t each agent observes some individual proxy \mathcal{Z}_t^i . We also assume that $E[\mathcal{Z}_t|\mathcal{Z}_t^i] = \mathcal{Z}_t^i\beta$ and $E[\epsilon_{t+1}|\mathcal{Z}_t^i] = 0$. An agent forms her REH forecast according to

$$E^i[R_{t+1}|\mathcal{Z}_t^i] = E^i[\mathcal{Z}_t|\mathcal{Z}_t^i]\theta + E^i[\epsilon_{t+1}|\mathcal{Z}_t^i] = \mathcal{Z}_t^i\beta\theta \quad (25)$$

Thus informational misperceptions can rationalize a heterogeneity of forecasts in REH-based models.²⁴

A diversity among individual forecast mappings used by agents is an intrinsic feature of our framework. Consequently, we do not need to rely on informational misperceptions or asymmetries to rationalize a heterogeneity of forecasts. As we suggest in the next subsection, a diversity among forecast mappings used by agents, and the resulting heterogeneity of forecasts, can be rationalized by appealing to the implications of the postulate of economic rationality in a world of imperfect knowledge.

3.2 Individual Rationality, Imperfect Knowledge and Heterogeneity of Forecasts

In this subsection, we explore the implications of the standard notion of economic rationality for the heterogeneity of forecasts in a world of imperfect knowledge. We first clarify what we mean by the concept “imperfect knowledge.” In this paper, we adopt a relatively narrow interpretation of this term that draws on its conventional usage in economic discourse. We also use our notion of an invariant mapping introduced in definition (2).

²⁴However, the assumption that an agent only observes a proxy \mathcal{Z}_t^i is very difficult to reconcile with the REH assumption that the conditional mean $E[\mathcal{Z}_t|\mathcal{Z}_t^i]$ is known to an agent. See, Frydman (1983). This latter assumption plays a crucial role in the class of REH models developed by Lucas, as well as REH-based models with asymmetric information. For example, Lucas (1973) assumes that agents know $E[\mathcal{Z}_t|\mathcal{Z}_t^i]$ in his seminal study of transient real effects of monetary policy. Grossman and Stiglitz (1980) use another version of the same assumption in their well-known examination of the informational efficiency of markets. This assumption plays a crucial role in their conclusion that under REH informationally efficient markets are impossible.

- **Definition 3** We refer to an invariant mapping, $\mathcal{F}^{i,j} [\cdot]$, as perfect if, at every time t , the forecast errors generated by $\mathcal{F}^{i,j} [\cdot]$ at payoff relevant time horizons are uncorrelated with information available to an agent at time t . We refer to such forecast errors as (statistically) *unsystematic*.

For example, the particular *invariant* mapping, $\mathcal{F}_t^{i,jt} [\cdot]$, would be perfect if, at every t , $\mathcal{F}_t^{i,j} [\cdot]$ were able to generate forecasts with white noise errors *solely* on the basis of the updating of values of an unchanging set of variables.

Definition 4 We refer to knowledge as imperfect if neither the market participants nor an outside investigator have access to a perfect mapping

We emphasize that the notion of imperfect knowledge does not rule out the possibility that some forecast mapping might generate forecasts with statistically unsystematic errors for some subperiod of the data. However, sooner or later, such a mapping would generate forecast errors that are correlated with the available information. To illustrate this point we consider the following example:

Suppose that an agent uses the following linear model as her forecast mapping:

$$R_{t+1} = \mathcal{X}_t^i \theta_t^i + \epsilon_{t+1}^i \quad (26)$$

where \mathcal{X}_t^i and ϵ_{t+1}^i are treated as random variables and θ_t^i denotes a vector parameter values. We also suppose that during the subperiod (t_1, t_2)

$$E[\epsilon_{t+1}^i | \mathcal{X}_t^i] = 0 \quad \text{for all } t = t_1, (t_1 + 1) \dots t_2 \quad (27)$$

where the conditioning set, \mathcal{X}_t^i , is assumed to have an unchanging composition.

In view of (27), we assume that an agent revises her forecasts *solely* as a result of an updating of the values of the variables in her information set. Consequently, her forecast mapping is invariant during the subperiod (t_1, t_2) and her forecasts are given by

$$R_{t+1} = x_t^i \theta^i + \epsilon_{t+1}^i \quad \text{and} \quad \tilde{r}_{t+1} = x_t^i \theta^i \quad (28)$$

where θ^i denotes parameter values that are assumed to be *unchanging* during the subperiod (t_1, t_2) .²⁵

²⁵We note that even if the forecast errors were uncorrelated with the variables in \mathcal{X}_t^i , profit-seeking agents would *not* necessarily continue to use an invariant forecasting model

Now suppose that the mapping in (28) is not perfect. Mapping imperfection can, of course, arise for many reasons. For example, this mapping might be imperfect because beyond time t_2 , the process generating R_{t+1} undergoes structural change, that results in a correlation between x_t^i and ϵ_{t+1}^i .²⁶ Thus, if an agent were to continue to use an invariant mapping in (28), her forecast errors would become correlated with the variables in \mathcal{X}_t^i . A profit-seeking agent would attempt to discover such an imperfection, and upon discovery, she would alter her mapping in an attempt to reduce the errors of her forecasts of payoff-relevant variables.

This example suggests the need to reexamine the implications of economic rationality for the formation of forecasts in the world of imperfect knowledge. We adopt here a standard notion of economic rationality: a rational agent chooses to deploy her resources so as to maximize some measure of her future payoff. However, in order to assess the costs and benefits of alternative deployments of resources, economic agents have to come up with assessments of future payoffs. Because agents do not have access to a perfect mapping, leaving their forecast mapping unchanged would eventually lead to forecasts with statistically systematic errors. Detecting these errors and the points of structural change is crucial for improving agent's forecasts. Consequently, profit-seeking agents engage in the creative process of testing their extant forecast mappings as well as attempt to discover new mappings. Thus, we reach the following conclusion:

Conclusion 1 *In a world of imperfect knowledge, economic rationality implies that an agent who does not pass up opportunities for gain, does not use an invariant mapping in forming her forecasts. We refer to an agent who never alters her forecast mapping as grossly irrational.*

Once we abandon the assumption that agents have access to a perfect mapping, a diversity of mappings and the resulting heterogeneity of forecasts is implied the assumption that agents *are not* grossly irrational. We also note

in (28). The reason is that (27) does not rule out the possibility that some variables that are not included in \mathcal{X}_t^i are nonetheless correlated with the forecast error ϵ_{t+1} . If an agent were to discover such variables and she were to revise the structure of her model in (23), she would decrease the variability of her forecasts. Thus, if an agent cares about forecast variability, she constantly searches for ways to reduce this volatility.

²⁶For evidence that foreign exchange models experience such episodic structural change, see Boughton (1987) and Goldberg and Frydman (1996a,b,2001). Sargent (1999) emphasizes that structural change is pervasive in macroeconomic models.

that a heterogeneity of forecasts arises even if agents were assumed to use only statistical procedures in testing and revising their forecast mappings. As is clear from the above example, to suppose that agents forecasts were homogeneous in a world of imperfect knowledge, would require not only that agents decide to change their models precisely at the same time, but that the models they switch to are all the same.²⁷ Thus we reach the following conclusion:

Conclusion 2 *In a world of imperfect knowledge, the assumption that agents are not grossly irrational implies, in general, that forecasts are heterogenous.*

- For example, in his examination of the epistemological difficulties of REH, Frydman (1982) formally shows that *rational* agents *do not*, in general, use a common model and/or rely on the same information when forming forecasts in a world of imperfect knowledge.

The foregoing conclusions suggest that the presumption that agents are not grossly irrational raises serious difficulties for the modeling of individual forecasts in a world of imperfect knowledge. In the remainder of this paper, we develop an approach to this problem.

4 Rationality and the Modeling of the Forecasting Process

An economist attempting to model profit-seeking behavior has to somehow represent the process involving revisions of the extant models and discoveries of the new ways to forecast payoff relevant variables.

REH representations of individual forecasts abstract from this creative process and assume that agents' forecasts are identical to the expectation of the forecasted variable implied by the model an economist constructs. Our discussion in the previous section suggests that in a world of imperfect knowledge, REH representations of individual forecasting behavior implicitly presume gross irrationality, rather the usually claimed rationality, on the part of agents.

²⁷The implications of this point for macroeconomic practice and notions of market efficiency are outside the scope of this paper.

Consequently, it is not surprising that behavioral economists have had such success in uncovering departures from REH. Behavioral economists usually interpret such departures as evidence that agents are irrational.²⁸ However, these findings imply irrationality only if REH is accepted as the standard of rational forecasting behavior. Nevertheless, the insistence of behavioral economists that actual behavior matters for the modeling of individual forecasts has substantially eroded REH's hitherto unrivaled position as the model for expectations.

However, the incompatibility of REH with the postulate of rationality is indicative of a broader problem of establishing an objective standard of rational forecasting behavior in a world of imperfect knowledge. In a world of imperfect knowledge, there is no mechanistic and unchanging standard of rationality for expectations that either the agents or economists can invoke.²⁹ While we argued that rationality implies that an agent *would* alter her forecast mappings, rationality does not seem to offer any specific guidance as how an economist should model agent's forecasting. In the absence of an objective standard of rationality, representations of forecasting behavior have to be based on observations of behavior in real-world markets.

This is, indeed, the approach followed by behavioral economists. However, the way in which behavioral finance literature has incorporated behavioral findings into its models of expectations shares an important feature with the REH approach. In order to formally represent their empirical findings, behavioral finance theorists have modeled an agent's forecasting behavior on the basis of invariant parametric representations.

For example, one of the important behavioral findings is that when agents revise their forecasts in asset markets, they frequently seem to "under-react" to new information. To represent such behavior, a behavioral economist would typically construct a parametric model of an agent's forecasts under the assumption that, relative to some hypothetical "true" model of the asset price, the degree of underreaction is *pre-specified and unchanging over time*.³⁰

²⁸For example, see Barberis and Thaler (2002).

²⁹Moreover, this lack of an unchanging standard of rationality for expectations raises serious questions for a closely related problem of modeling individual decisions. We attempt a broader discussion of this issue in Frydman and Goldberg (2004b).

³⁰For a recent example of such modeling of underreaction and references therein, see Gourinchas and Tornell (2001). For another example of a model of forecasting with an unchanging structure, see Hong and Stein (2002, 2003). Although Hong and Stein recognize that agents revise the structure of their models, they model these revisions as a

But this supposes that "underreaction" is an unchanging characteristic of individual forecasting behavior. As Fama (1998) argued

Apparent over-reaction to information is about as common as under-reaction, and post-event continuation of pre-event abnormal returns is about as frequent as post-event reversal

Fama interprets his findings as evidence against the behavioral approach in finance. In contrast, we interpret such findings as evidence against the assumption that agent's forecasting behavior can be represented by an invariant parametric model. This is the case, even if a behavioral-finance model is based on sound behavioral insights. Rational agents, pursuing profit opportunities, alter their forecast mappings and invent new ones. Thus, attempts to represent this creative process by an invariant mapping are bound to fail empirically. Moreover, even if under-reaction is interpreted as "irrationality", representing under-reaction by a fixed rule, in effect, entails a much stronger assumption that agents are *grossly* irrational, i.e. agents persist, in perpetuity, in the particular form of irrationality attributed to them.

This argument may help explain why, despite its numerous empirical failures, many economists are reluctant to abandon REH. Although the extant behaviorally-motivated forecasting models can shed light on particular episodes of the empirical record, they do not offer a general approach that can replace REH as a model of the forecasting process.

Nevertheless, in our view, behavioral insights should play an important role in the modeling of forecasting behavior. If it is indeed the case that the modeling of an agent's forecasting behavior with fixed rules is likely to fail empirically, we need to develop an alternative approach that can accord behavioral insights a significant role. Moreover, because rationality does not offer specific guidance on how to construct models of forecasting behavior, the role of behavioral findings is all the more important.

The key premise of our approach (dubbed the *Imperfect Knowledge Forecasting*, IKF, framework) is that economic agents – in formulating and revising their forecast mappings – are limited by imperfect knowledge. We also assume that individual agents are not grossly irrational. These two assumptions imply that an individual agent does not revise her forecasts solely because of new information sets. She also alters, at least intermittently, the mapping that she uses in forming her forecasts. Moreover, a profit-seeking

choice among pre-specified models according to a fixed rule.

agent not only switches between mappings she knows about but also searches for new ways to forecast payoff relevant variables.

The IKF framework characterizes this creative process by imposing *qualitative* restrictions on:

- the representation of the current and future collections of forecast mappings that an agent uses to forecast; and/or
- the way in which an agent switches between forecast mappings.

Although the postulate of individual rationality implies that an agent *will* switch from one forecast mapping to another, it offers no guidance as to which qualitative restrictions an economist, using the IKF framework, should impose on the forecasting process. It is to solve this problem that we make use of behavioral insights. The qualitative restrictions that we choose to impose on the forecasting process are, in fact, formal representations of the regularities uncovered by empirical research examining behavior in real-world markets. The IKF approach, therefore, incorporates behavioral insights in a new and rigorous way.

5 The IKF Framework: Qualitative Modeling of Individual Forecasts

By imposing only qualitative restrictions on the forecasting process, the IKF framework aims to provide a general approach to modeling agent's forecasts that can replace REH. In this section, we provide an overview of the basic elements of this qualitative approach. According to the UAUIP condition in (13), there are two types of forecasts, \tilde{r}_{t+1}^{i,j_t} and \tilde{l}_{t+1}^{i,j_t} , that underpin the equilibrium in the foreign exchange market. Our overview focuses on the modeling an agent's \tilde{r}_{t+1}^{i,j_t} as a way to highlight the basic elements of the IKF framework.

In order to model an agent's \tilde{r}_{t+1}^{i,j_t} , an economist needs to characterize the forecast mapping that an agent uses at a point in time, and the future forecast mappings to which an agent may switch as she updates her forecasts. The standard approach in modern macroeconomics, REH and non-REH alike, is for the economist to characterize \tilde{r}_{t+1}^{i,j_t} by choosing a specific parametric mapping on which quantitative restrictions are imposed. Moreover, this

parametric characterization is assumed to represent an agent's forecasting behavior endlessly.

We argued in the preceding section that no invariant parametric characterization can represent the creative forecasting process in real world markets. We thus reach the following conclusion:

Conclusion 3 *In order to formalize the key point that profit-seeking agents switch between mappings or invent new ones, an economist needs to characterize the properties of the entire collection of forecast mappings of an agent, rather than the properties of any one particular mapping.*

This leads us to model \hat{r}_{t+1}^{i,j_t} in a qualitative way. To this end, we represent the forecast formed by an agent at time t , \tilde{r}_{t+1}^{i,j_t} , by a differentiable function, \hat{r}_{t+1}^i .³¹

$$\hat{r}_{t+1}^i(x_t^i, \theta_t^i) = \tilde{r}_{t+1}^{i,j_t}(x_t^{i,j_t}, \theta_t^{i,j_t}) \quad (29)$$

where θ_t^i consists of parameters for the entire set of variables used in *all* of agent i 's forecast mappings, x_t^i , and not just those included in x_t^{i,j_t} . For example, if agent i 's mapping at time t is a linear model, then those parameters in θ_t^i that are attached to the variables excluded from x_t^{i,j_t} are set equal to zero. In this way, our formulation of \hat{r}_{t+1}^i allows us to represent the properties of each forecast mapping in an agent's current and future collection of these mappings.

We recall from section 3 that \tilde{r}_{t+1}^{i,j_t} is based on the mapping, $\mathcal{F}_t^{i,j_t}(\cdot)$, which an agent chooses from her collection, \mathcal{F}_t^i . At a point in time, \hat{r}_{t+1}^i represents agent i 's forecast based on this $\mathcal{F}_t^{i,j_t}(\cdot)$ mapping. However, when an agent revises her forecast, \tilde{r}_{t+1}^{i,j_t} , the mapping on which this forecast is based may change.

In order to characterize agent i 's entire collection of forecast mappings, we impose qualitative restrictions on the structure of \hat{r}_{t+1}^i and its changes over time. To this end, we represent revisions of agent i 's forecast, defined in (22) by a total differential of \hat{r}_{t+1}^i in (29):

$$d\hat{r}_{t+1}^i = \frac{\partial \hat{r}_{t+1}^i}{\partial x_t^i} dx_t^i + \frac{\partial \hat{r}_{t+1}^i}{\partial \theta_t^i} d\theta_t^i \quad (30)$$

³¹We assume the differentiability of \hat{r}_{t+1}^i to simplify our presentation. In general, we could replace all the differentials and derivatives with discrete differences and impose qualitative restrictions on these differences.

where the partial derivatives are with respect to the individual components of (x_t^i, θ_t^i) . This total differential captures the influence of both an alteration of the mapping agent i uses at time t , as well as the updating of forecasts due to new information.

The \hat{r}_{t+1}^i representation of agent i 's forecast in (29) and its revisions in (30) allows us to define an IKF model for \hat{r}_{t+1}^{i,j_t} and its revisions.

Definition 5 *An IKF model for an agent's forecast, \hat{r}_{t+1}^{i,j_t} , consists of a set of restrictions on the qualitative structure of \hat{r}_{t+1}^i in (29) and its revisions in (30). These qualitative restrictions can take one or more of the following three forms:*

- *assumptions on the properties of the sets $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$ generated by each of the forecast mappings in an agent's collection, \mathcal{F}_t^i , as well as the mappings yet to be invented; for example, an assumption that the $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$ generated by each forecast mapping from an agent's extant or future collection can be represented as a (conditional) probability distribution;*
- *restrictions that imply some variables or parameters in \hat{r}_{t+1}^i in (29) remain unchanged during the revision process; for example, a condition that sets $d[x_t^i] = 0$ or $d[\theta_t^i] = 0$ for some subset of variables or parameters in (x_t^i, θ_t^i) ;*
- *restrictions on the signs of all or a subset of the derivatives of \hat{r}_{t+1}^i in (30); for example, $\text{sign}(\frac{\partial \hat{r}_{t+1}^i}{\partial z_t^i}) < 0$.*

Remark 1 *We note that at different points in time, \hat{r}_{t+1}^i , represents an agent's forecast that can, in general, be based on different forecast mappings. Therefore, qualitative restrictions on \hat{r}_{t+1}^i represent an economist's characterization of the collection of an agent's forecast mappings, rather than any particular mapping chosen by an agent at a specific point in time.*

To implement the IKF approach, therefore, an economist must find qualitative restrictions that characterize the collection of an agent's forecast mappings. The IKF formalizes behavioral insights as such qualitative restrictions. After all, behavioral regularities do not pertain to just one particular forecasting model or method, but by their very nature, describe the qualitative properties that agents' forecasting procedures have in common.

For example, there is much research indicating that technical trading rules are heavily used in asset markets and that agents are often “conservative” in the way they change their models in the face of uncertain outcomes.³² This research implies that an economist may want to incorporate trend-following and conservative updating behavior into her model of the forecasting process. But this empirical research provides no clue as to which *specific, parametric* technical trading rule, from among thousands used in the marketplace, should be adopted by an economist to represent agents’ forecasting behavior. By the same token, the insight of conservatism does not indicate the precise speed and frequency with which agents switch between models.

Agents in real-world markets, in general, form their forecasts on the basis of different mappings. Thus, the best an economist can hope for is that the empirical regularities uncovered by behavioral researchers are sufficient to characterize the collections of forecast mappings used by agents.³³

In the next section, we follow our IKF approach to develop two competing models of the \tilde{l}_{t+1}^{i,j_t} ’s. In section 6, we follow the IKF approach to model the endogenous component of \hat{r}_{t+1} , and then combine this analysis with our formulations for the \tilde{l}_{t+1}^{i,j_t} ’s to construct models of the equilibrium premium on foreign exchange.

6 Two Competing IKF Models of an Individual Forecast of the Potential Unit Loss

The UAUIP condition developed in section 2 shows that an agent’s uncertainty premium, \widetilde{up}_r^i , at each point in time, depends on her forecast of the potential loss on a unit position, \tilde{l}_{t+1}^{i,j_t} . When an agent revises her assessments of \mathcal{R}_{t+1}^i and Φ_{t+1}^i , she, in general, revises her forecasts of the potential unit loss. These revisions, in turn, lead to changes in individual uncertainty premia and in the equilibrium (aggregate) premium on foreign exchange. Thus, in order to analyze the behavior of the equilibrium premium, we need to model the individual \tilde{l}_{t+1}^{i,j_t} ’s.

³²For example, see Edwards (1968), and Shliefer (2000) and references therein.

³³There are also other ways to formalize qualitative conditions. Beyond behavioral findings drawn from psychology and economics, the IKF framework can, in principle, incorporate insights from fields such as history and political science. Such insights, for example, could help an economist model changes in the collection of mappings over time.

In this section, we construct two competing IKF models for agents' forecasts of the potential unit loss. These models are based on two alternative rationalizations of the qualitative conditions used to restrict \hat{l}_{t+1}^i . The first IKF model imposes qualitative restrictions that are based on non-empirical (axiomatic) considerations. This model, which we call the *IKF expectations model*, assumes that agents' mappings can be represented by subjective probability distributions. In order to subsume, as a special case, a qualitative analog of the conventional REH approach, we impose an additional qualitative condition on \hat{l}_{t+1}^i .

Our second IKF model imposes qualitative restrictions on \hat{l}_{t+1}^i that are based on behavioral insights. This model, which we call the *IKF gap model*, leads to implications for the behavior of individual uncertainty premia that differ from those obtained from the IKF expectations model. We show in section 6 that our two IKF models for the \tilde{l}_{t+1}^{i,j_t} 's also imply differing implications for the behavior of the equilibrium premium. We test these competing implications in section 7.

The assumption that an agent values gains and losses of the same size differently is the key feature of prospect theory. In a typical gambling experiment considered by Kahneman and Tversky (1979), an agent is provided, by an investigator, with the actual values of payoffs and their probabilities. In our notation, such gambles involve situations in which an agent *does not* need to come up with assessments of the future $r_{t+1,k}^{+,i}$'s, $r_{t+1,k}^{-,i}$'s and $\phi_{k,t+1}^i$'s; at the time she makes her risky choices, she knows the values of her prospects and the associated probabilities. With such special gambles, Kahneman and Tversky's development of prospect theory did not need to consider the modeling of an agent's forecast of her potential loss.³⁴

In contrast, in a typical market setting, an agent not only faces uncertain outcomes, but she has to form forecasts of the values of her potential payoffs and their likelihood weights. The original formulation of prospect theory offers no guidance on the modeling of individual forecasts. In this section, we follow our IKF framework in constructing our two alternative models for \tilde{l}_{t+1}^{i,j_t} .

Our approach to the modeling of \tilde{l}_{t+1}^{i,j_t} exploits a connection between an

³⁴In the concluding section of the paper, Kahneman and Tversky (1979) conjecture that “the theory can be extended to the *typical* situation of choice, where the probabilities of outcomes are not explicitly given (emphasis added).” Of course, for the theory to describe agents' choices in such “typical” situations, an economist has to model the formation of individual forecasts.

agent's forecast of the potential loss and her forecast of the excess return. We rationalize this connection on both non-behavioral and behavioral grounds. In the next subsection, we construct our expectations model for \tilde{l}_{t+1}^{i,j_t} .

6.1 The IKF Expectations Model

Our first IKF model for \tilde{l}_{t+1}^{i,j_t} is comprised of two qualitative conditions. The first of these qualitative conditions excludes from an agent's collection any forecast mapping that cannot be represented by a probability distribution.

Probability Restriction *An agent's collection of forecast mappings satisfies the qualitative probability restriction if the sets $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$ generated by each of the mappings $\mathcal{F}_t^{i,j} \in \mathcal{F}_t^i$ in (17) at all t can be represented by a conditional probability distribution.*

We first note that the probability restriction is weaker than the restrictions implied by REH, which also represents an agent's forecast mappings at all times by a conditional probability distribution. But in addition to the probability restriction, an economist who follows REH, represents *all* of the mappings, current and future, used by an agent by *one* invariant parametric mapping, i.e., the one mapping that is consistent with the economist's model. Moreover, under REH, an economist imposes quantitative restrictions on this one parametric representation.

We first examine the implications of the probability restriction for the qualitative structure of our representations $\tilde{l}_{t+1}^i(x_t^i, \theta_t^i)$ and $\hat{r}_{t+1}^i(x_t^i, \theta_t^i)$. Under the probability restriction, we can define two random variables R_{t+1}^i and ϵ_{t+1}^i , such that $E(\epsilon_{t+1}^i | x_t^i, \theta_t^i) = 0$.

$$R_{t+1}^i = \hat{r}_{t+1}^i + \epsilon_{t+1}^i \quad (31)$$

where

$$\hat{r}_{t+1}^i = E[R_{t+1}^i | x_t^i, \theta_t^i] \quad (32)$$

The probability restriction alone does not restrict the qualitative structure of $\hat{r}_{t+1}^i(x_t^i, \theta_t^i)$ in a way that has implications for the behavior of this forecast. If we were to construct a complete IKF model for \tilde{l}_{t+1}^{i,j_t} , therefore, we would have to impose additional qualitative restrictions on \hat{r}_{t+1}^i . But our focus in this paper is on the equilibrium premium, and so we model only

the endogenous component of \tilde{r}_{t+1}^{i,j_t} . This involves modeling the endogenous component of \tilde{s}_{t+1}^{i,j_t} , which we postpone until the next section.

We now turn to the construction of our first IKF model for \tilde{l}_{t+1}^{i,j_t} . Without loss of generality, we develop our analysis for a bull only. The probability restriction implies that for a bull

$$\tilde{l}_{t+1}^{i,j_t} = E[R_{t+1}I(R_{t+1} < 0)|x_t^i, \theta_t^i] \quad (33)$$

Substituting (31) into (33) defines \tilde{l}_{t+1}^{i,j_t} as a function \tilde{r}_{t+1}^{i,j_t} and (x_t^i, θ_t^i) :

$$\tilde{l}_{t+1}^{i,j_t} = \tilde{r}_{t+1}^{i,j_t} E[I(\epsilon_{t+1}^i < -\tilde{r}_{t+1}^{i,j_t})|x_t^i, \theta_t^i] + E[\epsilon_{t+1}^i I(\epsilon_{t+1}^i < -\tilde{r}_{t+1}^{i,j_t})|x_t^i, \theta_t^i] \quad (34)$$

Therefore, the probability restriction implies that an agent can revise her expectation of the potential unit loss, \tilde{l}_{t+1}^{i,j_t} , in one or both of the following ways:

- she can revise \tilde{l}_{t+1}^{i,j_t} by revising the first moment of her distribution, \tilde{r}_{t+1}^{i,j_t} ; and/or
- she can revise \tilde{l}_{t+1}^{i,j_t} by revising the higher moments of her conditional distribution of ϵ_{t+1}^i , which we denote by \tilde{m}_{t+1}^{i,j_t} .

Letting \hat{m}_{t+1}^i denote our representation for a vector of higher moments, \tilde{m}_{t+1}^{i,j_t} , we can write the implication of the probability restriction as follows:

$$\hat{l}_{t+1}^i = \hat{l}_{t+1}^i[\hat{r}_{t+1}^i(x_t^i, \theta_t^i), \hat{m}_{t+1}^i(x_t^i, \theta_t^i)] \quad (35)$$

We now impose additional qualitative conditions on the structure of \hat{l}_{t+1}^i in (35). In accordance with definition 5, these additional conditions are qualitative restrictions on our representation of an agent's revisions of her \tilde{l}_{t+1}^{i,j_t} . Totally differentiating (35) yields:

$$d\hat{l}_{t+1}^i = \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_{t+1}^i} d\hat{r}_{t+1}^i + \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{m}_{t+1}^i} d\hat{m}_{t+1}^i \quad (36)$$

We consider first, the terms involving \hat{r}_{t+1}^i . The following proposition, which we prove in Frydman and Goldberg (2003b), examines the implications of the probability restriction for the sign of $\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_{t+1}^i}$.

Proposition 1 *Suppose that $\{\mathcal{R}_{t+1}^i, \Phi_{t+1}^i\}$ is represented by a conditional probability distribution. Then the partial derivatives of \hat{l}_{t+1}^i with respect to \hat{r}_{t+1}^i satisfy the following sign restrictions for a bull and a bear, respectively:*

$$\frac{\partial \hat{l}_{t+1}^{i,L}}{\partial \hat{r}_{t+1}^i} > 0 \quad \text{and} \quad \frac{\partial \hat{l}_{t+1}^{i,S}}{\partial \hat{r}_{t+1}^i} < 0 \quad (37)$$

An agent may revise her forecast, \tilde{r}_{t+1}^{i,j_t} because of exogenous and/or endogenous reasons. We postpone our discussion of $d[\tilde{r}_{t+1}^i(x_t^i, \theta_t^i)]$ in (36) until the next section.

Consider now the terms involving \hat{m}_{t+1}^i in (36). In general, as agents switch from one forecast mapping to another, the higher moments of their conditional expectation change. But to facilitate a comparison of the expectations model of this section with the REH, we assume that revisions of an agent's expectation of the potential unit loss involve only the first moment of her conditional distribution. The qualitative restriction that formalizes this property of the conventional REH approach can be defined as follows:

Invariance Restriction: *The representation of the potential loss in (35) is said to be consistent with the invariance restriction if each forecast mapping in an agent's current and future collections involve the same higher moments of the conditional distribution of R_{t+1} . Thus, a switch by an agent to any new forecast mapping implies only a revision of the conditional mean, \tilde{r}_{t+1}^{i,j_t} . We formalize this restriction by setting $d\hat{m}_{t+1}^i = 0$.*

As with the probability restriction, the invariance restriction is also weaker than the restrictions on updating imposed by REH. The invariance restriction does not rule out the possibility that an agent would alter the model generating her conditional mean; for example, she could drop or add variables, and/or revise the values of parameters. In contrast, conventional REH analysis restricts the updating of agents' forecasts to the arrival of new information on x_t^{i,j_t} .

In summary, we constructed in this subsection a model for \tilde{l}_{t+1}^{i,j_t} that subsumes, as a special case, a qualitative analog of the REH approach. This IKF expectations model (or E-model for short) is comprised of the probability and invariance restrictions. As we noted above, these are not plausible on behavioral grounds.³⁵ In the next subsection, we develop an alternative IKF model

³⁵While the invariance restriction is difficult to justify on behavioral grounds, Savage

for \tilde{l}_{t+1}^{i,j_t} . We motivate this alternative model with behavioral regularities in real-world markets.

6.2 The IKF Gap Effect Model of the Individual Forecast of the Unit Loss

Our alternative IKF model imposes neither the probability nor invariance restrictions on the qualitative structure of \hat{l}_{t+1}^i . An invariance condition is a stringent restriction on the revisions of forecasting models. It presumes, for example, that when a bull revises her \tilde{s}_{t+1}^{i,j_t} (and thus \tilde{r}_{t+1}^{i,j_t}) upward, she necessarily reduces the size of her forecasted potential loss. This assumption ignores the possibility that a bull may be concerned about a rising divergence between her forecast, \tilde{s}_{t+1}^{i,j_t} , and some historical benchmark value of the exchange rate. This concern, which Keynes (1936) built into his formulation of speculative demand for money, may cause a bull to revise upwards, rather than downwards, the size of her \tilde{l}_{t+1}^{i,j_t} .³⁶ In this subsection, we formalize this idea by a set of qualitative restrictions on how an agent's forecast of the potential unit loss changes due to a revision in \tilde{s}_{t+1}^{i,j_t} .

We note that the terms involving \hat{m}_{t+1}^i in (35) and (36) represent the effect of revisions of an agent's forecast mapping on the higher moments of an agent's probability model. In the general case in which the probability restriction does not hold, we reinterpret \hat{m}_t^i as a vector of factors, denoted by \hat{f}_{t+1}^i . These factors represent the effects of revisions of her forecast mapping on an agent's \tilde{l}_{t+1}^{i,j_t} that are not fully captured by revisions of \tilde{r}_{t+1}^{i,j_t} . In general, a switch to a new forecast mapping, and thus to a revised (x_t^i, θ_t^i) , leads to a revision of \hat{r}_{t+1}^i and the vector of factors, \hat{f}_{t+1}^i . Consequently, we represent \tilde{l}_{t+1}^{i,j_t} as the following function, \hat{l}_{t+1}^i :

$$\hat{l}_{t+1}^i = \tilde{l}_{t+1}^i[\hat{r}_{t+1}^i(x_t^i, \theta_t^i), \hat{f}_{t+1}^i(x_t^i, \theta_t^i)] \quad (38)$$

Following the IKF approach, we impose qualitative restrictions on the structure of \hat{l}_{t+1}^i . In contrast to the E-model developed in the previous subsection, we motivate our second IKF model for \tilde{l}_{t+1}^{i,j_t} on behavioral grounds. In particular, we rely on Keynes' experience with and insights on the behavior

(1954) and others have argued that the probability restriction is plausible on behavioral grounds. For evidence to the contrary, see Kahneman and Tversky (2000).

³⁶Tobin (1958) used this idea of Keynes in his formulation of the demand for money.

of agents in real world markets. In discussing the question of why an agent might hold cash rather than interest-bearing bonds, Keynes argued that

[the demand for cash] will not have a definitive quantitative relation to a given rate of interest of r ; what matters is *not the absolute level* of r but the degree of *its divergence* from what is considered a fairly safe level of r , having regard to those calculations of probability which are being relied on (Keynes, 1936, p.201, emphasis added).

In order to develop an IKF model based on this idea, we define an agent's forecast of the *gap* between \tilde{s}_{t+1}^{i,j_t} and her assessment of the historical benchmark, $\tilde{s}_t^{\text{HB}i,j_t}$ as follows:³⁷

$$\widetilde{gap}_{t+1}^{i,j_t} = \tilde{s}_{t+1}^{i,j_t} - \tilde{s}_t^{\text{HB}i,j_t}, \quad i = \text{L, S} \quad (39)$$

A benchmark is, of course, specific to each asset market. Each agent arrives at her own determination of the benchmark value and so, in general, individual estimates will differ across agents. However, there are a few general characteristics of a benchmark that are already explicit in Keynes' remarks cited above and that are important for our analysis.

1. Although assessments of the benchmark value may differ across individual agents at each time t , the range over which these individual assessments differ should be smaller (often substantially so) than the range over which the observed asset price varies. For example, this would be the case for individual benchmarks based on averages of historical data.
2. This notion of a benchmark will play an important role in speculative decisions if historical evidence suggests that: a) asset prices tend to move persistently away from their historical benchmarks for substantial periods of time; and b) asset prices eventually revert, at unpredictable moments in time, back to their historical benchmark values; and c)

³⁷The gap can also be defined in terms of the actual, rather than in terms of the forecast of next period's exchange rate, or some weighted average of the actual and forecasted gaps. See, for example, Frydman and Goldberg (2003a). The conclusions of our analysis are not affected by either of these specifications of the gap variable.

asset prices often shoot through benchmark values from one side and continue moving persistently away from the other side.³⁸

3. Individual speculators recognize the historical evidence of long swings in asset prices and believe that this evidence is used by other speculators operating in the markets.

Based on the foregoing considerations, we assume that an agent's forecast of the gap is one of the factors that influences her forecast of the potential unit loss. In this paper, we abstract from the influence of other potential factors and/or assume that the gap variable captures their influence on \tilde{l}_{t+1}^{i,j_t} .³⁹ Consequently, in what follows, we use the following specification for \hat{l}_{t+1}^i :

$$\hat{l}_{t+1}^i = \tilde{l}_{t+1}^i[\hat{r}_{t+1}^i, \widehat{gap}_{t+1}^i] \quad (40)$$

In order to complete the construction of our IKF model, we need to develop qualitative restrictions on the derivative of \hat{l}_{t+1}^i in (40) with respect to its arguments. We first consider the sign restriction on $\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_{t+1}^i}$. Without the probability assumption, the sign of this partial derivative, in general, may be positive or negative. However, we continue to impose the sign restrictions on $\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_{t+1}^i}$ for bulls and bears in (37) and (??). These sign restrictions seem to have a straightforward behavioral interpretation. As a bull (bear) increases the size of her forecasted return, \tilde{r}_{t+1}^{i,j_t} ($-\tilde{r}_{t+1}^{i,j_t}$), *ceteris paribus*, there is more room for her forecast to be wrong, and yet, she still earns a positive return from her open position. In effect, a higher \tilde{r}_{t+1}^{i,j_t} or $-\tilde{r}_{t+1}^{i,j_t}$ provides greater insurance against a capital loss. We formalize this reasoning with the following qualitative restriction:

The Insurance Effect: *The representation of the potential loss in (40) is said to be consistent with the insurance effect if each forecast mapping in an agent's current and future collections imply that revisions of \tilde{l}_{t+1}^{i,j_t} are associated with the partial derivatives in (37).*

³⁸For evidence that exchange rates exhibit long swings that revolve around historical benchmark levels see the references contained in footnote 46. For the bond market see Bec and Anders (2002) and for the stock market see Campbell and Shiller (1998), and references therein.

³⁹We note that depending on the modeling context, our framework can be extended to allow for factors other than an agent's forecast of the gap.

Consider now the partial derivative of \hat{l}_{t+1}^i with respect to \widehat{gap}_{t+1}^i . To provide intuition, suppose that agent i 's forecast of next period's exchange rate, $\tilde{s}_{t+1}^{i,jt}$, is above her assessment of the historical benchmark, so that $\widehat{gap}_{t+1}^{i,jt} > 0$. Also suppose agent i revises up her forecast, $\tilde{s}_{t+1}^{i,jt}$. According to Keynes's definition of a historical benchmark, this upward revision in $\tilde{s}_{t+1}^{i,jt}$ is usually associated with an upward revision in agent i 's $\widehat{gap}_{t+1}^{i,jt}$. This upward revision in the forecasted gap can occur despite a belief on the part of agent i that s_t will eventually revert back to its perceived benchmark level. Indeed, the forecasting problem agents face is that s_t may move away from its perceived benchmark over an extended time period prior to reverting back.

According to Keynes, if agent i is a bull, then a rising gap will create more fear of an eventual countermovement. In this case, as agent i revises up her $\widehat{gap}_{t+1}^{i,jt}$, she simultaneously revises up her assessment of size of the potential losses, i.e. $-\tilde{l}_{t+1}^{i,jt}$ increases. If, on the other hand, agent i is a bear, then a rising gap will create more confidence that a countermovement will occur. In this case, as agent i revises up her $\widehat{gap}_{t+1}^{i,jt}$, she simultaneously revises down her assessment of size of the potential losses, i.e., $-\tilde{l}_{t+1}^{i,jt}$ falls.

A rise in agent i 's $\tilde{s}_{t+1}^{i,jt}$ leads not only to a rise in $\widehat{gap}_{t+1}^{i,jt}$, but also to a rise in $\tilde{r}_{t+1}^{i,jt}$. Keynes suggests, however, that the divergence from a historical benchmark has a substantially greater effect on an agent's estimate of the potential unit loss than the opposite effect stemming from changes in the yield. Keynes describes the relative weakness of the insurance effect as follows:

Unless reasons are believed to exist why future experience will be very different from past experience, a ...rate of interest [much lower than a historical safe rate], leaves more to fear than to hope, and offers, at the same time, a running yield which is only sufficient to offset a *very small* measure of fear (Keynes, 1936, p.202, emphasis added).

We formalize Keynes's behavioral insights by the following set of restrictions on the qualitative structure of \hat{l}_{t+1}^i .

The Gap Effect: *The representation of the potential loss in (40) is said to be consistent with the gap effect if each forecast mapping in an agent's current and future collections imply that revisions of $\tilde{l}_{t+1}^{i,jt}$ can be represented by the following partial derivatives.*

- If an agent is a bull, then

$$\frac{\partial \hat{l}_{t+1}^{i,L}}{\partial \widehat{gap}_{t+1}^i} < 0 \quad (41)$$

and

$$\frac{\partial \hat{l}_{t+1}^{i,L}}{\partial \tilde{s}_{t+1}^i} = \frac{\partial \hat{l}_{t+1}^{i,L}}{\partial \hat{r}_{t+1}^i} + \frac{\partial \hat{l}_{t+1}^{i,L}}{\partial \widehat{gap}_{t+1}^i} < 0 \quad (42)$$

- If an agent is a bear, then

$$\frac{\partial \hat{l}_{t+1}^{i,S}}{\partial \widehat{gap}_{t+1}^i} > 0 \quad (43)$$

and

$$\frac{\partial \hat{l}_{t+1}^{i,S}}{\partial \tilde{s}_{t+1}^i} = \frac{\partial \hat{l}_{t+1}^{i,S}}{\partial \hat{r}_{t+1}^i} + \frac{\partial \hat{l}_{t+1}^{i,S}}{\partial \widehat{gap}_{t+1}^i} > 0 \quad (44)$$

In summary, we constructed in this section a model for \tilde{l}_{t+1}^{i,j_t} based on behaviorally-motivated qualitative restrictions on an agent's current and future collections of forecast mappings. This model is comprised of the gap and insurance effects. We refer to this model as the *IKF gap model* (or G-model for short).

In the next section, we examine the implications of the E- and G- models for the behavior of the equilibrium premium over time.

7 The Behavior of the Equilibrium Premium and UAUIP

In the preceding section, we developed two IKF models for agents' forecasts of the potential unit loss, one based on non-empirical considerations and one based on behavioral insights. In this section, we first characterize the endogenous component of \tilde{s}_{t+1}^i following the IKF framework. We then combine this characterization with the E- and G-models for \tilde{l}_{t+1}^{i,j_t} to derive two alternative models of the equilibrium premium. We find that the E- and G-models imply competing implications for the equilibrium premium. We conclude

this section with a discussion of the UAUIP condition. We also show that the G-model of the premium provides a simple explanation of sign reversals, which REH-based models have found so difficult to explain.

7.1 Modeling the Endogenous Component of \hat{s}_{t+1}^i

To simplify the presentation we assume a representative bull and a representative bear. We also set the aggregation weights of our bull and bear to be equal. Finally, we assume that $\lambda_1^i = \lambda_1$ $i = L, S$. These simplifications still preserve substantial heterogeneity among agents without affecting any of the results of our analysis.⁴⁰

To obtain a model of the equilibrium premium, we need to model the endogenous component of \hat{s}_{t+1}^i . To this end, we express this representation of agent i 's forecast \tilde{s}_{t+1}^{i,j_t} as a sum of an autonomous component, $\hat{s}_{t+1}^{a_i}(\mathbf{z}_t^i)$, and an endogenous component in the following way:

$$\hat{s}_{t+1}^i(x_t^i, \theta_t^i) = \hat{s}_{t+1}^{a_i} + \beta_t^i s_t, \quad i = L, S \quad (45)$$

where \mathbf{z}_t^i is the subset of \mathbf{x}_t^i that does not covary with s_t . The parameter β_t^i captures the endogenous influence of changes in s_t on our representation of agent i 's forecast \tilde{s}_{t+1}^{i,j_t} . We now follow the IKF framework and impose qualitative restrictions on the β_t^i 's.

There is much evidence in the asset-market literature indicating that trend-following (extrapolative) behavior is prevalent in financial markets. For example, the use of technical trading rules or “charts,” – many of which extrapolate past price trends – has been widely documented in all financial markets. In a study of the London foreign exchange market, Allen and Taylor (1990) and Taylor and Allen (1992) find that approximately ninety percent of respondents reported using some technical rules at short horizons when forming expectations.⁴¹ There is also much experimental evidence that

⁴⁰At the cost of considerable notational complexity, the analysis in this paper can be extended to allow for a heterogeneity within the groups of bears and bulls using wealth shares to aggregate. We consider this more complete aggregation problem in Frydman and Goldberg (2003b); in this paper we show that allowing for a greater degree of heterogeneity does not affect the conclusions of our analysis.

⁴¹For early studies on the use of technical trading systems and trend-following behavior, and their implications for asset price dynamics, see Schulmeister (1987), Schulmeister and Goldberg (1988) and Soros (1987). See also the study of the Hong Kong foreign exchange market in Lui and Mole (1998).

speculators in simulated markets react positively to price trends (e.g., see Andreassen and Kraus, 1990, and De Bondt, 1993).

There is also evidence that although trend-following behavior is a good description of forecasting behavior at the aggregate level of a representative bull and bear, these forecasts do not exhibit bandwagon behavior, i.e., $\tilde{s}_{t+1}^{i,jt}$ does not move more than proportionately with a change in the exchange rate. For example, the seminal study of Frankel and Froot (1987) examine three different data sets of survey data on exchange rate expectations. All three data sets indicate that expectations are extrapolative, but not bandwagon.

Thus, although a β_t^i that is either negative or greater than one may characterize the behavior of some speculators, the behavioral evidence suggests strongly that at the aggregate level of a representative bull and bear, agents' forecasts of next period's exchange rate are trend-following, but not bandwagon.

We formalize this behavioral regularity with the following qualitative restriction on \hat{s}_{t+1}^i :

The Trend Restriction *The representation of agent i 's forecast of next period's exchange rate in (45) is said to be consistent with the trend effect if each forecast mapping in agent i 's current and future collections imply that revisions of $\tilde{s}_{t+1}^{i,jt}$ are associated with the partial derivative $0 < \frac{\partial \hat{s}_{t+1}^i}{\partial s_t} < 1$.*

In terms of \hat{s}_{t+1}^i in (45), the trend restriction implies that $0 < \beta_t^i < 1$, $i = L, S$. We note that this trend restriction, $0 < \beta_t^i < 1$, $i = L, S$, is a sufficient condition for the stability of the expectations model. Thus, the behavioral evidence that justifies the trend restriction also justifies the assumption of stability.

7.2 The IKF Expectations Model

We are now ready to derive our E-model of the equilibrium premium. Plugging the representation of $\tilde{l}_{t+1}^{i,jt}$ in (35) into the definition of the equilibrium premium in (13), and imposing the probability, invariance and trend restrictions on \hat{l}_{t+1}^i , yields:

$$\hat{p}r_t = \frac{(1 - \lambda_1)}{2} \sum_{i=L,S} I(i) \hat{l}_{t+1}^i (\hat{r}_{t+1}^i) + \lambda_2 NFA \quad (46)$$

where $I(i) = 1$ if $i = L$ and $I(i) = -1$ if $i = S$. The probability restriction implies that the \hat{l}_{t+1}^i 's represent individual conditional expectations. The invariance restriction sets $d\hat{m}_{t+1}^i = 0$ and allows us to suppress \hat{m}_{t+1}^i as an argument in the individual \hat{l}_{t+1}^i 's. The trend restriction allows us to express $\hat{r}_{t+1}^i = \hat{s}_{t+1}^{a_i} + (\beta_t^i - 1) s_t - fp_t$.

There is much evidence in the literature suggesting that the covariation between a country's net foreign asset position and either the exchange rate or the excess return on foreign exchange is negligible.⁴² Indeed, if this covariation were larger, the standard models of the risk premium would be more successful in explaining excess returns. We invoke this lack of covariation and examine changes in the equilibrium premium assuming that the corresponding change in the net foreign asset position of country 1 is zero, i.e., $dNFA = 0$. This assumption allows us to simplify our analysis and to focus exclusively on the aggregate uncertainty premium as the main determinant of the excess return in the foreign exchange market. None of the results of our analysis are affected by this assumption.

We also abstract from movements in the forward premium and set $dfp_t = 0$.⁴³ Our final simplification follows from the definition of a historical benchmark, which implies that the variation in $\tilde{s}_t^{\text{HB},j_t}$ is substantially lower than the variation in \tilde{s}_{t+1}^{i,j_t} and s_t .⁴⁴ We abstract, therefore, from variation in $\tilde{s}_t^{\text{HB},j_t}$, and set $d\hat{s}_t^{\text{HB},i} = 0$, $i = L, S$. This assumption implies that $\frac{d\hat{p}r_t}{d\hat{s}_{t+1}} = \frac{d\hat{p}r_t}{d\widehat{gap}_t}$.

We obtain an expression for the change in the equilibrium premium by first totally differentiating (46):

$$d\hat{p}r_t = \frac{(1 - \lambda_1)}{2} \sum_{i=L,S} I(i) \left[\frac{\partial \hat{l}_{t+1}^i}{\partial \tilde{r}_t^i} (d\hat{s}_{t+1}^{a_i} + (\beta_t^i - 1) ds_t) \right] \quad (47)$$

The equilibrium condition in (13), allows us to express ds_t as follows:

$$ds_t = \frac{1}{(1 - \beta_t)} (d\hat{s}_{t+1}^a - d\hat{p}r_t) \quad (48)$$

⁴²For example, see Frankel (1983), Lewis (1995) and Engel (1996).

⁴³As we show in Frydman and Goldberg (2004d), the IKF approach can also be used to examine the implications of changes in fp_t . We find that our IKF-based model of the equilibrium premium sheds new light on the forward-discount anomaly.

⁴⁴The survey data on exchange rate expectations used in the next section are consistent with this definition. The data reveal that when \tilde{s}_t^{HB} is proxied by a measure of purchasing power parity, the $\text{sign}(d\tilde{s}_{t+1}) = \text{sign}(\widehat{gap}_t)$ for 166 out of 168 observations.

where

$$\beta_t = \frac{1}{2} \sum_{i=L,S} \beta_t^i \text{ and } d\hat{s}_{t+1}^a = \frac{1}{2} \sum_{i=L,S} d\hat{s}_{t+1}^{a_i} \quad (49)$$

Substituting (48) into (47) yields the following expression for the change in the equilibrium premium under the E-model:

$$d\hat{p}r_t = \frac{\frac{(1-\lambda_1)}{2} \sum_{i=L,S} I(i) \frac{\partial \hat{i}_{t+1}^i}{\partial \hat{r}_{t+1}^i} \left[d\hat{s}_{t+1}^{a_i} - \frac{(1-\beta_t^i)}{(1-\beta_t)} d\hat{s}_{t+1}^a \right]}{1 + \frac{(\lambda_1-1)}{2(1-\beta_t)} \sum_{i=L,S} I(i) \frac{\partial \hat{i}_{t+1}^i}{\partial \hat{r}_{t+1}^i} (1-\beta_t^i)} = \frac{NUM_t}{DEN_t} \quad (50)$$

We first determine the sign of DEN_t . From proposition 2 we know that $I(i) \frac{\partial \hat{i}_{t+1}^i}{\partial \hat{r}_{t+1}^i} > 0$ for both bulls and bears. We also know that $\lambda_1 > 1$. Thus, with the trend restriction imposed for bulls and bears, the denominator in (50) is unambiguously positive. Therefore, the sign of $d\hat{p}r_t$ is the same as the sign of the NUM_t .

Using (49), we can express the numerator as follows:

$$NUM_t = \frac{(1-\lambda_1)}{4(1-\beta_t)} \left[(1-\beta_t^S) d\hat{s}_{t+1}^{a_L} - (1-\beta_t^L) d\hat{s}_{t+1}^{a_S} \right] \left[\frac{\partial \hat{l}_{t+1}^L}{\partial \hat{r}_{t+1}^L} + \frac{\partial \hat{l}_{t+1}^S}{\partial \hat{r}_{t+1}^S} \right] \quad (51)$$

Noting that $\frac{d\hat{s}_{t+1}}{d\hat{s}_{t+1}^a} = 1$, (51) implies that

$$\frac{d\hat{p}r_t}{d\hat{s}_{t+1}} = 0 \text{ if either } \frac{\partial \hat{l}_{t+1}^L}{\partial \hat{r}_{t+1}^L} = -\frac{\partial \hat{l}_{t+1}^S}{\partial \hat{r}_{t+1}^S} \text{ and/or } \beta_t^S d\hat{s}_{t+1}^{a_L} = \beta_t^L d\hat{s}_{t+1}^{a_S} \quad (52)$$

The first condition in (52) says that if bulls and bears revise their forecasts of s_{t+1} in a way that leads to similar revisions in their expected potential loss (of course with opposite signs), then the movement of the equilibrium premium will be unrelated to movements of the aggregate forecast. The second condition in (52) says that if the trend-following behavior of bulls is similar to that of bears, and their revisions of $\tilde{s}_{t+1}^{i,jt}$ are also similar, then, again, the equilibrium premium will be unrelated to the aggregate forecast. Since most speculators are sometimes bulls and sometimes bears, one would expect the parameters β_t^S and β_t^L , and $\frac{\partial \hat{l}_{t+1}^L}{\partial \hat{r}_{t+1}^L}$ and $-\frac{\partial \hat{l}_{t+1}^S}{\partial \hat{r}_{t+1}^S}$, to be comparable in magnitude. Moreover, even if the conditions in (52) did not hold, one would expect the quantitative values of β_t^S , β_t^L , $\frac{\partial \hat{l}_{t+1}^L}{\partial \hat{r}_{t+1}^L}$, $-\frac{\partial \hat{l}_{t+1}^S}{\partial \hat{r}_{t+1}^S}$ to be unstable

over time, as agents revise their forecasting models and/or update the values of the variables in their information sets.

Therefore, we reach one of the two main testable implications of our IKF-based analysis:

Conclusion 4 *The E-model for individual expectations of the potential unit loss from speculation, along with the trend restriction, implies that, either*

- *the movements of the equilibrium premium are unrelated to the movements of the aggregate forecast of the gap; or*
- *there is no systematic relationship between the equilibrium premium and the aggregate forecast of the gap: the relationship between \widehat{pr}_t and \widehat{gap}_t will change over time between one that is positive, negative and nonexistent.*

Although the foregoing E-model allows for agents' forecast revisions to entail new forecast mappings, this model is consistent conventional REH forecasts, which arise solely from new realizations of \mathcal{X}_t^i . Thus, the E-model developed in this section subsumes, as a special case, a qualitative analog of REH. Thus, this IKF analog of REH leads to the prediction that there is no systematic relationship between \widehat{pr}_t and \widehat{s}_{t+1} . This conclusion serves as the null hypothesis for our empirical tests presented in section 7.

7.3 The IKF Gap Model

We now derive the testable implication of the G-model of individual forecasts of the potential unit loss for the equilibrium premium on foreign exchange. To do so we impose the trend restriction of the preceding section on individual forecasts of next period's exchange rate. We show that in contrast to the E-model, the G-model implies a positive relationship between the equilibrium premium and the aggregate forecast of the gap, i.e., $\frac{d\widehat{pr}_t}{d\widehat{gap}_t} > 0$. As in the preceding subsection, we set $dfp_t = dNFA_t = d\widehat{s}_t^{\text{HB}i} = 0$, $i = \text{L,S}$.

Plugging in \widehat{l}_{t+1}^i from (40) into the definition of \widehat{pr}_t in (13), and totally differentiating yields:

$$d\widehat{pr}_t = \frac{(1 - \lambda_1)}{2} \sum_{i=\text{L,S}} I(i) \left[\frac{\partial \widehat{l}_{t+1}^i}{\partial \widehat{r}_t^i} d\widehat{r}_{t+1}^i + \frac{\partial \widehat{l}_{t+1}^i}{\partial \widehat{gap}_t} d\widehat{gap}_t \right] \quad (53)$$

We can write (50) as the following expression for the equilibrium change in the aggregate premium:

$$d\widehat{pr}_t = \frac{\frac{(1-\lambda_1)}{2} \sum_{i=L,S} I(i) \left[\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{s}_{t+1}^i} \left(d\hat{s}_{t+1}^{a_i} - \frac{(1-\beta_t^i)}{(1-\beta_t)} d\hat{s}_{t+1}^a \right) + \frac{\partial \hat{l}_{t+1}^i}{\partial \widehat{gap}_t} \frac{1}{(1-\beta_t)} d\hat{s}_{t+1}^a \right]}{1 + \frac{(1-\lambda_1)}{2(1-\beta_t)} \sum_{i=L,S} I(i) \left[\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{s}_{t+1}^i} \beta_t^i - \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_t^i} \right]} \quad (54)$$

where we used $\frac{\partial \hat{l}_{t+1}^i}{\partial \hat{s}_{t+1}^i} = \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_t^i} + \frac{\partial \hat{l}_{t+1}^i}{\partial \widehat{gap}_t}$. from the gap conditions in (42) and (44), as well as the definitions in (49)

As with the denominator for the E-model in (50), the denominator in (54) is unambiguously positive. This result follows from $\lambda_1 > 1$ and the gap and insurance restrictions imposed by the G-model, i.e., $I(i) \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{s}_{t+1}^i} < 0$ and $I(i) \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{r}_t^i} > 0$, $i = L, S$. Consequently, as with the E-model, the relationship between the equilibrium premium and the aggregate forecast of the gap depends on the sign of the numerator in (54).

The numerator of the expression in (54) shares two similar terms with the numerator in (50), i.e., $I(i) \frac{\partial \hat{l}_{t+1}^i}{\partial \hat{s}_{t+1}^i} \left(d\hat{s}_{t+1}^{a_i} - \frac{(1-\beta_t^i)}{(1-\beta_t)} d\hat{s}_{t+1}^a \right)$, $i = L, S$. Again, most speculators are sometimes bulls and sometimes bears, and so, one would expect the parameters β_t^S and β_t^L , and $\frac{\partial \hat{l}_{t+1}^L}{\partial \hat{s}_{t+1}^L}$ and $-\frac{\partial \hat{l}_{t+1}^S}{\partial \hat{s}_{t+1}^S}$ to be comparable in magnitude. Consequently, as with the E-model, the influence of the first two terms in the numerator of (54), combined, leads to no relationship between \widehat{pr}_t and \widehat{gap}_t

Thus, the relationship between the equilibrium premium and the aggregate forecast of the gap implied by the G-model depends on the last two terms in the numerator of (54), i.e., $I(i) \frac{\partial \hat{l}_{t+1}^i}{\partial \widehat{gap}_t} \frac{1}{(1-\beta_t)} d\hat{s}_{t+1}^a$, $i = L, S$. But according to the gap conditions in (41) and (43), the values of these two terms are both negative. This analysis leads to the following conclusion:

Conclusion 5 *The G-model of individual forecasts of the potential unit loss from speculation, along with the trend restriction, implies a positive relationship between the equilibrium premium and the aggregate forecast of the gap.*

This implication of the G-model serves as the alternative hypothesis for our empirical tests presented in the section 7. But before we present these empirical tests, we use the UAUIP condition to discuss the intuition behind conclusion 5.

7.4 The UAUIP Condition and the Equilibrium Premium Under the G-Model

To illuminate the intuition behind conclusion 5, we write the UAUIP condition in (13) as follows:

$$\hat{r}_{t+1} = \widehat{up}_t^L - \widehat{up}_t^S + NFA_t \quad (55)$$

Suppose, for example, that initially $NFA = 0$ and $\hat{r}_{t+1} = \widehat{up}_t^L - \widehat{up}_t^S < 0$. Also suppose that at time t , the group of bulls revise up their forecast of next period's exchange rate. This revision may stem from the arrival of new information and /or changes in the forecast mappings used by the bulls. In terms of our representation, \hat{s}_{t+1}^L rises, causing, *ceteris paribus*, increases in \hat{s}_{t+1} and \hat{r}_{t+1} .

The key restriction imposed by the G-model on the forecasting process is that when bulls (bears) revise their forecasts of next period's exchange rate, they simultaneously revise their forecasts of the *size* of the potential unit loss in the same (opposite) direction. In terms of our example, when the group of bulls revises up \hat{s}_{t+1}^L , this group also revises up $-\tilde{l}_{t+1}^L$. The bulls increase their assessment of the potential unit loss because the rising gap is assumed to create either: 1) greater worry about an eventual countermovement back to the perceived historical benchmark, if initially $\tilde{s}_{t+1}^L > \tilde{s}_t^{\text{HB}L}$; or 2) less confidence in their prediction of a foreign currency appreciation, if initially $\tilde{s}_{t+1}^L < \tilde{s}_t^{\text{HB}L}$. And, when bulls revise up $-\tilde{l}_{t+1}^L$, their extra sensitivity to potential losses causes them to require a greater premium for taking long positions in foreign exchange, i.e., \widehat{up}_t^L also increases.

Thus, the revision of \hat{s}_{t+1}^L leads to increases in both \hat{r}_{t+1} and \widehat{up}_t , *ceteris paribus*. In general, $\hat{r}_{t+1} - \widehat{up}_t$ may rise or fall. Consider the case of an increase, so that $\hat{r}_{t+1} - \widehat{up}_t > 0$.⁴⁵ In this case, the group of bulls forecast a higher return on long positions *in excess* of the minimum premium they require. This greater excess return causes an excess demand for foreign exchange and an incipient capital flow that leads to a rise in the exchange rate.

The rise in the exchange rate works to equilibrate the system by reducing \hat{r}_{t+1} and further raising \widehat{up}_t . These effects work in the following way. First, with $\beta_t < 1$ (from the trend restriction), the rise in s_t leads to a lowering of both \hat{r}_{t+1}^L and \hat{r}_{t+1}^S , and thus, to a lowering of \hat{r}_{t+1} . Second, with $\beta_t^i > 0$,

⁴⁵If $\hat{r}_{t+1} - \widehat{up}_t$ falls, becoming negative, then the resulting excess supply will cause s_t to fall. In this case, with $\beta_t < 1$, \widehat{up}_t necessarily rises in equilibrium.

$i = L, S$ (from the trend restriction), the rise in s_t leads to increases in \hat{s}_{t+1}^L and \hat{s}_{t+1}^S . These increases lead to further gap effects that cause \widehat{up}_t^L to rise and \widehat{up}_t^S to fall. Finally, the reductions \hat{r}_{t+1}^L and \hat{r}_{t+1}^S have insurance effects that also cause \widehat{up}_t^L to rise and \widehat{up}_t^S to fall.

Thus, the initial rise in \hat{r}_{t+1}^L leads to a new equilibrium in which the equilibrium premium on foreign exchange also rises. It is clear from equation (55) that the rise in the equilibrium level of \widehat{up}_t may be large enough so that while $\widehat{up}_t < 0$ initially, the new equilibrium level implies $\widehat{up}_t > 0$. All that is required is that the dominant weight in the aggregate premium shifts from bears to bulls. Thus, the G-model of the equilibrium premium provides a simple way to explain sign reversals. We discuss this result more fully in Frydman and Goldberg (2004b).

Before we move on to formal tests of the IKF models developed in this section, we construct six-month moving averages of our survey data on exchange rate expectations and on a measure of the historical benchmark. Figure 1 plots these six-month averages. The survey data consist of monthly observations of the median one-month forecast of the German mark-U.S. dollar spot rate. One month forward rates were obtained from DRI. Our sample spans the period from January 1983 through December 1996. Our measure of \widehat{gap}_t is based on the Big Mac PPP exchange rate reported in the April 1993 edition of the *Economist* (which was 2.02) and inflation differentials using the CPI series from the IFS data bank.⁴⁶

Although not a statistical test, the time plots are rather suggestive that the relationship between the equilibrium premium and an aggregate forecast of the gap is indeed positive.

8 Some Empirical Evidence

The preceding section showed that the E- and G-models of the equilibrium premium lead to competing predictions concerning the relationship between

⁴⁶For evidence that PPP does serve as a historical benchmark for the foreign exchange market, see the cointegration studies of Juselius (1995) and Cheung and Lai (1993), among others, as well as studies on the long-horizon predictability of PPP and monetary fundamentals (e.g., Mark (1995) and Mark and Sul (2001)). See also Obstfeld [1995] and references therein. We note, however, that although it may be plausible to use PPP as a proxy for the benchmark level for many exchange rates, this *does not* imply that the PPP level is a long-run equilibrium in the sense of being the rate at which the foreign exchange market “settles”.

Figure 1
Six-Month Averages




Table 1: OLS Regression
 Uncertainty Premium and the Expected Gap

Constant	0.857 (0.954)
\widetilde{gap}_t	1.334*** (0.222)
Adjusted R ² =.394	DW Statistic=2.15
Standard errors (adjusted for heteroskedasticity) are in parentheses	
* * *denotes significance with a p-value of .01	

the equilibrium premium and an aggregate forecast of the gap. The former model predicts no systematic relationship, whereas the latter model predicts a positive relationship. Table 1 reports the results of a simple OLS regression of \widetilde{up}_t on \widetilde{gap}_t . To address the problem of serial correlation and possible stochastic trends we added to the regression lags of the dependent and independent variables.⁴⁷ The table reports long-run values for the estimates of the coefficients. Diagnostic tests revealed that two lags were sufficient. The results show that the expected gap is highly significant and positively related to the expected excess return on foreign exchange (in this case the U.S. dollar), as predicted by G-model. The table also shows that 40 percent of the variation in the premium can be explained by the gap.

The results reported in Table 1 should be viewed with caution. Although the qualitative relationship between \widetilde{up}_t on \widetilde{gap}_t may be positive throughout the sample, we would expect the quantitative relationship to be nonlinear: agents most likely place little weight on small deviations from benchmark values in forming their forecasts, whereas they place relatively large weight when deviations are large. Moreover, in a world of imperfect knowledge, we would expect rational agents to change forecast mappings from time to time. This, in general, would cause the parameters of conditional forecast functions to shift, which would, in turn cause reduced-form regressions like the one in table 1 to be temporally unstable.

⁴⁷See Juselius and Hendry (2000) on the validity on using OLS standard errors for inference with unit-root variables when lagged values of the dependent and independent variables are included in the regression.

Table 2
Contingency Table Analysis

	$\Delta up_t > 0$	$\Delta up_t < 0$
$\Delta \widetilde{gap}_t > 0$	43	25
$\Delta \widetilde{gap}_t < 0$	30	51

One way to handle the changing nature of the quantitative relationship between \widetilde{up}_t on \widetilde{gap}_t is to make use of a non-parametric procedure such as contingency-table analysis. Contingency-table analysis provides a way to test the qualitative relationship, while allowing for the exact form of the quantitative relationship to change over time.

Table 2 presents the contingency table results. The diagonal (off-diagonal) cells in the table denote the number of observations for which the changes in \widetilde{up}_t and \widetilde{gap}_t (denoted $\Delta \widetilde{up}_t$ and $\Delta \widetilde{gap}_t$, respectively) were in the same (opposite) direction. The number of observations along the diagonal cells are larger than the off-diagonal cells, which suggests a positive relationship between \widetilde{up}_t and \widetilde{gap}_t . A χ^2 statistic of 10.15 indicates that this positive relationship is significant at the .01 level.

The results reported in tables 1 and 2 lead to a clear rejection of the E-model in favor of the G-model. For the DM/\$ exchange rate, over a subperiod of floating that spans the 1980s and 1990s, the evidence indicates that the relationship between the equilibrium premium and an aggregate forecast of the gap is a positive one.

Additional research is needed to determine whether the main testable restriction of the IKF gap model – that the equilibrium premium depends positively on an aggregate forecast of the divergence between the asset price and its historical benchmark level – is more than a just a description of one episode in the foreign exchange market. Other exchange rates and other asset markets need to be examined. But clearly, despite its qualitative nature, macroeconomic analysis based on the IKF framework generates testable implications.

References

- Allen, H. and M. Taylor (1990), "Charts, Noise and Fundamentals in the London Foreign Exchange Market," *Economic Journal*, 100, 49-59.
- Andreassen, P. and S. Kraus (1990), "Judgemental Extrapolation and the Saliency of Change," *Journal of Forecasting*, 9, 347-372.
- Barberis, N., Huang, M. and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116, 1-53.
- Barberis, N. and R.H. Thaler (2002), "A Survey of Behavioral Finance," in G. Constantinides, M. Harris and R. Stulz (eds.), *Handbook of the Economics of Finance*, Amsterdam: North-Holland.
- Bec, F. and R. Rahbek (2002), "Vector Equilibrium Correction Models with Non-Linear Discontinuous Adjustments," working paper no.15, department of Statistics and Operations Research, University of Copenhagen.
- Campbell, J. and R. Shiller (1998), "Valuation Ratios and the Long-Run Stock Market Outlook", *Journal of Portfolio Management*, 24, 11–26.
- Cheung, Y. and K. Lai (1993), "Long-Run Purchasing Power Parity During the Recent Float," *Journal of International Economics*, 34, 181-192.
- De Bondt, W. (1993), "Betting on Trends Intuitive Forecasts of Financial Risk and Return," *International Journal of Forecasting*, 9, 355-371.
- De Long, B., A. Shleifer, L. Summers and R. Waldman. (1990), "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45, 375-395.
- Dornbusch, R. (1983), "Exchange Rate Risk and the Macroeconomics of Exchange Rate Determination," in R. Hawkins, R. Levich and C. Wihlborg (eds.), *The Internationalization of Financial Markets and National Economic Policy*, Greenwich, Conn.: JAI Press.
- Engel, C.A. (1996), "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance*, 3, 123-191.
- Engel, C. And J. Hamilton (1990), "Long Swings in the Exchange Rate: Are They in the Data and Do Markets Know It?," *American Economic Review*, September, 689-713.
- Evans, G. W. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press.
- Fama, E. F. (1998), "Market Efficiency, Long-Term Returns and Behavioral Finance," *Journal of Financial Economics*, 49, 283-306.

- Frankel, J. A. (1982), "In Search of the Exchange Risk Premium: A Six Currency Test Assuming Mean-Variance Optimization," *Journal of International Money and Finance*, 1, 255-74.
- Frankel, J. A. (1983), "Monetary and Portfolio Balance Models of Exchange Rate Determination," in J. Bhandari and B. Putnam (eds.), *Economic Interdependence and Flexible Exchange Rates*, Cambridge: MIT press.
- Frankel J. A. and K. A. Froot (1987), "Understanding the U.S. Dollar in the Eighties: The Expectations of Chartists and Fundamentalists," *Economic Record*, December, pp. 24-38, and reprinted in J.A. Frankel (ed.), *On Exchange Rates*, The MIT Press: Cambridge, Mass., 1995, chapter 14, 295-316.
- Frydman, R. (1982), "Towards an Understanding of Market Processes: Individual Expectations, Learning and Convergence To Rational Expectations Equilibrium," *American Economic Review*, 72, 652-668.
- Frydman, R. (1983), "Individual Rationality, Decentralization, and the Rational Expectations Hypothesis," in Frydman, R. and E.S. Phelps (eds.), *Individual Forecasting and Aggregate Outcomes: "Rational Expectations" Examined*, New York: Cambridge University Press.
- Frydman, R. and M. Goldberg (2003a), "Imperfect Knowledge Expectations, Uncertainty Adjusted UIP and Exchange Rate Dynamics," in Aghion, P., Frydman, R., Stiglitz, J. and M. Woodford (eds.), *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton, NJ: Princeton University Press.
- Frydman, R. and M. Goldberg (2003b), "Imperfect Knowledge and Asset Price Dynamics: Modeling the Forecasting of Rational Agents, Dynamic Prospect Theory and Uncertainty Premia on Foreign Exchange," University of Copenhagen Discussion Paper, #03-31.
- Frydman, R. and M. Goldberg (2004a), "Imperfect Knowledge and the Economics of Capitalism," paper to be presented at the conference on "Capitalism and Society," Columbia University, April.
- Frydman, R. and M. Goldberg (2004b), "Heterogeneity of Forecasts and Equilibrium in the Foreign Exchange Market: Imperfect Knowledge and Endogenous Prospect Theory," mimeo.
- Frydman, R. and M. Goldberg (2004c), *Exchange Rates and Risk in a World of Imperfect Knowledge: A New Approach to Modeling Asset Markets*, under contract, Princeton, NJ: Princeton University Press.
- Frydman, R. and M. Goldberg (2004d), " " mimeo.

- Frydman, R. and E.S. Phelps (1983), "Introduction" in Frydman, R. and E.S. Phelps (1983), *Individual Forecasting and Aggregate Outcomes: "Rational Expectations" Examined*, New York: Cambridge University Press.
- Goldberg, M. D. and R. Frydman (1996a), "Imperfect Knowledge and Behavior in the Foreign Exchange Market," *Economic Journal*, 106, July, 869-893.
- Goldberg, M. D. and R. Frydman (1996b), "Empirical Exchange Rate Models and Shifts in the Cointegrating Vector", *Journal of Structural Change and Economic Dynamics*, 7, 55-78.
- Goldberg, M. D. and R. Frydman (2001), "Macroeconomic Fundamentals and the DM/\$ Exchange Rate: Temporal Instability and the Monetary Model," *International Journal of Finance and Economics*, October, 421-435.
- Grossman, S. and J. Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393-408.
- Gourinchas, P.O. and A. Tornell (2001), "Exchange Rate Dynamics, Learning and Misperception," mimeo.
- Hansen, L.P. and T.J. Sargent (2001a), "Acknowledging Misspecification in Macroeconomic Theory," mimeo.
- Hansen, L.P. and T.J. Sargent (2001b), "Robust Control and Model Uncertainty," mimeo.
- Hayek, F.A. (1948), *Individualism and Economic Order*, Chicago: University of Chicago Press.
- Hayek, F.A. (1978), "The Pretence of Knowledge," *Le Prix Nobel en 1974*, in *New Studies in Philosophy, Politics, Economics and History of Ideas*, Chicago: The University of Chicago Press.
- Hong, H. and J. Stein (2002), "Differences of Opinion, Short-Sales Constraints and Market Crashes," mimeo.
- Hong, H. and J. Stein (2003), "Simple Forecasts and Paradigm Shifts," mimeo.
- Juselius, K. (1995), "Do the Purchasing Power Parity and the Uncovered Interest Rate Parity Hold in the Long Run? - An Application of likelihood Inference in a Multivariate Time Series Model." *Journal of Econometrics* Vol. 69 (1), 211-240.
- Juselius, K. and D. Hendry (2000), "Explaining Cointegration Analysis. Part 1," *The Energy Journal*, 21:1, 1-42.

- Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263-291.
- Kahneman, D. and A. Tversky (2000), *Choices, Values and Frames*, Cambridge, UK: Cambridge University Press.
- Keynes, J.M. (1936), *The General Theory of Employment, Interest and Money*, New York: Harcourt, Brace and World.
- Krugman, P. (1981), "Consumption Preferences, Asset Demands, and Distribution Effects in International Financial Markets," NBER Working Paper No. 651, March.
- Lewis, K. (1995), "Puzzles in International Financial Markets," in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Vol. III, Amsterdam: North-Holland, 1913-1917.
- Lucas, R.E., Jr. (1973), "Some International Evidence on Output-Inflation Trade-offs," *American Economic Review*, 63, 326-334.
- Lui, Y. and D. Mole (1998), "The Use of Fundamental and Technical Analysis by Foreign Exchange Dealers: Hong Kong Evidence," *Journal of International Money and Finance*, 17, 535-545.
- Lyons, R. (2001), *The Microstructure Approach to Exchange Rates*, Cambridge, Massachusetts: The MIT Press.
- Mark, N. (1995), "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review*, 85, 201-218.
- Mark, N. and D. Sul (2001), "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel," *Journal of International Economics*, 53, 29-52.
- Mark, N. and Y. Wu (1998), "Rethinking Deviations from Uncovered Interest Rate Parity: The Role of Covariance Risk and Noise," *Economic Journal*, 108, 1686-1706.
- Obstfeld, M. (1995), "International Currency Experience: New Lessons and Lessons Relearned," *Brookings Papers on Economic Activity*, 1, pp.119-220.
- Phelps, E.S. (1968), "Money Wage Dynamics and Labor Market Equilibrium," *Journal of Political Economy*, August, 678-711.
- Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics*, Oxford, UK: Oxford University Press.
- Sargent, T.J. (1999), *Conquest of American Inflation*, Princeton: Princeton University Press.

Savage, J. (1954), *The Foundations of Statistics*, New York: John Wiley and Sons.

Schulmeister, S. (1987), "Currency Speculation and Dollar Fluctuations," *Banca Nazionale del Lavoro Quarterly Review*, 167, December, 343-365.

Schulmeister, S. and M. Goldberg (1989), "Noise Trading and the Efficiency of Financial Markets," in Giacomo Luciani (ed.), *The American Financial System: Between Euphoria and Crisis*, (Quaderni della Fondazione Adriano Olivetti: Rome).

Taylor, M. and H. Allen (1992), "The Use of Technical Analysis in the Foreign Exchange Market," *Journal of International Money and Finance*, 11, 304-314.

Tversky, A. and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.