Center on Capitalism and Society at Columbia University
www.earth.columbia.edu/ccs

Working Paper No. 24, May 2008

Macroeconomic Theory For a World of Imperfect Knowledge

Roman Frydman* and Michael D. Goldberg**

28 May 2008

** 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

Contents

I From Early Modern Economics To Imperfect Knowledge Economics 5

1	Modern Macroeconomics: Individual Forecasting and Aggregate Outcomes	6
2	Early Modern Narrative Accounts: Respecting the Limits to Knowledge	7
3	Contemporary Models: Fully Prespecifying Change	8
3.1	Rational Expectations Models	10
3.2	Behavioral Models	11
3.3	Internal Inconsistency and Flawed Microfoundations	12
3.3.1	Standard Probabilistic Representations	13
3.3.2	Diversity and Gross Irrationality	13
3.3.3	REH	14
3.3.4	Behavioral Models	17
3.3.5	Sharp Predictions and the Modern Research Program .	18
4	The Promise of Imperfect Knowledge Economics	18
4.1	Qualitative Models of Change	18
4.1.1	Non-Standard Use of Probabilistic Formalism	19
4.2	Avoiding Epistemological Flaws of Contemporary Models . . .	21
5	Contextual Rationality of IKE Models	22
5.1	Combining Insights from Economics and Other Disciplines . .	24
5.1.1	Preferences	24
5.1.2	Forecasting Behavior	25
6	IKE as the Boundary of What Macroeconomic Theory Can Deliver	27

II Why Macroeconomic Theory Cannot Ignore the Limits to Knowledge

27

7	Fully Prespecifying Change to Generate Sharp Predictions	28
7.1	Fully Predetermining Restrictions	29
7.2	Sharp Probabilistic Predictions of Change	31
7.2.1	History as the Future and Vice Versa	32
8	Epistemological Flaws of Contemporary Models in a World of Imperfect Knowledge	33
8.1	Internal Inconsistency When Forecasting Strategies Are Diverse	33
8.2	REH as a Representation of Grossly Irrational Forecasting Strategy	35
8.3	Behavioral Models	38
8.3.1	Internal Inconsistency	38
8.3.2	The Inessential Role of Diversity in Behavioral Models	39
9	Why Macroeconomics Must Open Its Models to Imperfect Knowledge	40
10	IKE: Recognizing Imperfect Knowledge and Diversity of Forecasting Strategies	41
10.1	Forecasting As Entrepreneurship	41
10.2	Imperfect Knowledge and Economic Analysis	42
10.2.1	Representing Imperfect Knowledge with Partially Predetermined Probability Distributions	43

III Imperfect Knowledge Economics of Market Fluctuations and Risk 44

11	The Failure of REH Models	48
11.1	Imperfect Knowledge and Deviations from the REH Solution .	49
11.2	Long Swings from Benchmark Levels in Asset Markets	50
12	An IKE View of Long Swings	51
12.1	Conservatism as a Qualitative Regularity of Uneven Duration	52
12.2	Uneven Swings in an Individual's Forecast	56
12.3	Bulls, Bears, and Long Swings	57

12.4 Long Swings and Fundamentals	60
12.5 Conditional Predictions of Change	60
12.6 Bounded Instability and the Role of Benchmarks	61
12.7 Benchmark Levels and Revisions of Forecasting Strategies . .	62
12.7.1 Self-Limiting Long Swings	65
13 The Market Premium and the Aggregate Gap	66
13.1 Diversity of Forecasting Strategies and the Market Premium .	68
14 How Recognizing the Limits to Knowledge Avoids Internal Inconsistency	69
14.1 The Gap Conditions and Market Premium	70
14.2 Conservatism and Long Swings	70
14.3 Bulls and Bears in a Long Swing	71
15 The Futile Search for Sharp Predictions	72
15.1 Lost Fundamentals in Currency Markets	72
15.2 Is the Market Really Grossly Inefficient?	73
16 Coming to Terms with Imperfect Knowledge	75

Part I

From Early Modern Economics To Imperfect Knowledge Economics

When macroeconomists confront their models with time-series data they often find gross inconsistencies. In our recent book (Frydman and Goldberg, 2007), we argue that contemporary models also suffer from insuperable epistemological flaws. We trace these empirical and theoretical difficulties to a common source: in modeling aggregate outcomes, contemporary economists fully prespecify the causal mechanism that underpins change in real-world markets. They do so because they have come to believe that, in order to be worthy of scientific status their models should generate sharp predictions of how individual decision making and market outcomes evolve over time.¹

In our book, we propose an alternative approach to economic analysis, which we call Imperfect Knowledge Economics (IKE). Although IKE builds on the methodology of contemporary macroeconomics by modeling aggregate outcomes on the basis of mathematical representations of individual decision making, it jettisons models that generate sharp predictions. In this paper, we elaborate on and extend the arguments that led us to propose IKE. We show analytically that in order to avoid the fundamental epistemological flaws inherent in extant models, economists must stop short of fully prespecifying change. We also show how acknowledging the limits of their knowledge may enable economists to shed new light on basic features of the observed time-series of market outcomes, such as fluctuations and risk in asset markets, which have confounded extant approaches for decades.

¹As we illustrate rigorously in section 7, a contemporary model is said to generate a sharp prediction of change if, conditional on the probability distribution of outcomes at some initial time, it represents outcomes at all other times, past and future, with a single probability distribution.

1 Modern Macroeconomics: Individual Forecasting and Aggregate Outcomes

Modern macroeconomics constructs models of aggregate outcomes on the basis of mathematical representations of individual decision making, with market participants' forecasting behavior lying at the heart of the interaction between the two levels of analysis. Individuals' forecasts play a key role in how they make decisions, and markets aggregate those decisions into prices. The causal mechanisms behind both individual decisions and aggregate outcomes, therefore, depend on market participants' understanding of the economy and how they use this knowledge to forecast the future.

By focusing on the central role of forecasting for understanding the connection between micro and macro outcomes, economists have achieved important insights. For example, building on the path-breaking work of Phelps (1968, 1970), Lucas (1976) sharply criticized econometric policy analysis that, based on Keynesian aggregate models, examined the effects of changes in tax rates, money supply, or other “policy” variables on market outcomes. This analysis presumed that the same structure — the set of causal variables and the parameters that relate them to those variables — would continue to represent adequately the causal mechanism after a change in policy. The main point of the “Lucas critique” was the untenability of that premise. He argued that changes in policy variables would alter the way market participants forecast the future — and hence their decision-making. In general, this change on the individual level would also alter the causal mechanism driving market outcomes.

Lucas offered a seemingly straightforward remedy to this fundamental difficulty. He presumed that the Rational Expectations Hypothesis (REH) would enable economists to model exactly how policy changes would affect market participants' forecasts and aggregate outcomes. REH postulates that the economist's aggregate model precisely represents the causal mechanism driving individuals' forecasts and their revisions. By constraining the predictions on the individual and aggregate levels to be one and the same, REH is thought to offer a “scientific” way to predict both the micro and macro effects of policy changes. In embracing REH-based models, economists have merely replaced Keynesian policy analysis with another mechanistic approach to evaluating the consequences of policy changes.

To be sure, the Lucas critique of Keynesian models does not depend on

REH: it requires only that policy changes influence forecasting strategies significantly enough to alter the causal mechanism driving market outcomes. What has been largely overlooked, however, is that Lucas's critical arguments point to a fundamental difficulty inherent in the entire modern research program in macroeconomics. After all, while policy changes undoubtedly play a role in market participants' alteration of their forecasting strategies, so do many other factors.²

In fact, even if one were to limit the analysis solely to policy changes, the solution that Lucas proposes is facile: REH supposes that individuals revise their forecasting strategies in mechanical ways that can be precisely specified in advance. However, in capitalist economies, individuals are strongly motivated to search for genuinely new ways to forecast the future and deploy their resources. The social context, including the institutions within which individuals make decisions, also changes in unforeseeable ways. But when the social context or how market participants' forecast future outcomes changes, so too does the causal mechanism underpinning market outcomes. Thus, change in capitalist economies is to a significant extent non-routine, for it does not follow pre-existing rules and procedures. The premise of IKE is that leaving macroeconomic models open to such change is crucial for understanding outcomes in real-world markets.

2 Early Modern Narrative Accounts: Respecting the Limits to Knowledge

The history of economic thought includes widely differing responses to the daunting challenge that change poses for economic analysis. Early modern economists relied on a largely narrative mode of analysis. Although imprecise by contemporary standards, narrative accounts had the important advantage of leaving economists relatively free to explore the complexity and opaqueness of the interdependence between individual rationality, the social context of decision-making, and market outcomes.³ Indeed, the giants of early modern economics uncovered remarkably powerful and durable in-

²For some early warnings concerning the fundamental flaws of the REH approach, see Frydman (1982), Phelps (1983), and Frydman and Phelps (1983).

³Of course, a narrative mode of analysis also constrains argument, but this constraint is relatively weak compared to the rigor of mathematical language.

sights, such as Hayek’s (1948) prescient prediction that socialist planning is bound *in principle* to fail; Knight’s (1921) assertion that standard probabilistic uncertainty cannot adequately characterize business decisions; and Keynes’s (1936) closely related arguments concerning the importance of radical uncertainty, the social context, and conventions for forecasting returns and risk on investment in real and financial assets. These insights point to the fundamental flaw in contemporary economists’ research program: the causal mechanism that underpins change in capitalist economies is not completely intelligible to *anyone*, including market participants, economists, policy officials, or social planners.

3 Contemporary Models: Fully Prespecifying Change

Largely ignoring early modern arguments concerning the inherent limits to economists’ knowledge, contemporary economists construct models that we call *fully predetermined*. These models represent the causal mechanism that underpins change on the individual and aggregate levels with mechanical rules. Thus, they leave no room for changes in individual decision making and aggregate outcomes that have not been fully specified in advance by an economist.

Economists represent an individual’s decision making by specifying her forecasts of future market outcomes, preferences that rank the future consequences of her decisions for her well being, and the constraints that she faces. They also specify a decision rule, such as maximization of an individual’s well being, which selects the preferred deployment of resources. In this way, economists represent the way an individual makes decisions in terms of a set of causal factors.

Alternative specifications of preferences, forecasting strategies, constraints, and the decision rule enable economists to formalize competing explanations of individual decision making.⁴ At some arbitrary “initial” point in time, these representations relate an individual’s choices to causal factors with qualitative conditions.⁵ However, although these representations of individ-

⁴For an overview of alternative specifications of these components and how they result in alternative representations of individual decision-making, see chapter 3 of our book.

⁵For example, it is common for economists to assume that an individual’s utility de-

ual decision-making are qualitative, contemporary economists fully prespecify change by imposing restrictions that relate their representations at all points in time, past and future, *exactly* to the properties of their representation at the initial point in time. Based on such microfoundations, contemporary models also fully prespecify, in terms of some set of causal factors, how market outcomes unfold over time.

Over time, individuals alter the way they to make decisions. In general, therefore, models with structures that vary over time are required to represent adequately individual decision making. Oddly, in view of the broad applicability of the Lucas critique, the vast majority of contemporary models do not allow for any changes in their structure. These time-invariant models do not accord any role to revisions in forecasting strategies in driving outcomes; instead, they force economists to rely solely on the movement of causal factors to explain time-series data. These factors are represented as random variables, referred to as “exogenous shocks.” The random variation in these shocks is supposed to account for new information that becomes available to market participants. Consequently, contemporary economists focus on information and its asymmetries as the principal factors driving market outcomes. As with the rest of the structure of contemporary models, the probability distributions of the causal variables or new information are usually assumed to be time-invariant.

Contemporary economists sometimes recognize the importance of incorporating in their models the fact that participants in real-world markets do not adhere endlessly to one forecasting strategy, or, more broadly, that they alter the way they make decisions. They also sometimes take into account the fact that the social context, particularly economic policy, changes over time, at least intermittently. However, because the resulting models determine exactly how individual decisions and the social context may change over time, they remain as mechanistic as their time-invariant counterparts.⁶ By adhering to the bogus belief that only models that promise sharp predic-

pend positively on her consumption of goods or that her forecast of a future market price depends positively on the current value of this price.

⁶A popular way to model policy changes in the literature is to use the Markov-switching framework of Hamilton (1988, 1990). We show in Frydman and Goldberg (2007, chapter 6) that because such models fully prespecify the set of possible structures that might represent outcomes and a probabilistic rule that determines the timing of change, their representations of aggregate outcomes are formally equivalent to those implied by models whose structure is time-invariant.

tion deserve scientific status, contemporary economic analysis locks in the presumption that individuals never forecast or alter their decision-making in new ways.

Of course, economists recognize that knowledge, including their own, is imperfect and that time-invariant models or those that fully prespecify change are fraught with error. To represent their own imperfection of knowledge, they include additive error terms to their models. As with new information, economists usually assume that the distributions characterizing these error terms never change, or if change is allowed, they fully prespecify it.⁷ Here, again, contemporary economics embodies an odd and counterproductive conception of imperfect knowledge. By fully predetermining how the probability distributions of the error terms in his model evolve over time, an economist, in effect, fully prespecifies how his own imperfection of knowledge unfolds between the initial time and all other time periods.

3.1 Rational Expectations Models

Virtually all contemporary economists have adopted a set of *a priori* assumptions that putatively characterize how rational individuals make decisions, with REH as the centerpiece of their standard of rationality. In REH, people’s beliefs “are not inputs,” but the outcomes of economists’ theories.⁸ Thus, REH rules out the possibility that market participants’ forecasting strategies play an autonomous role in driving individual decision making and aggregate outcomes. An REH model derives its representation of forecasting strategies from its specification of preferences, constraints, and the way that policy and other causal variables unfold over time. As such, the causal variables and parameters that make up an REH representation of forecasting behavior stem solely from variables and parameters that an economist uses to specify the other components of his model.

Proponents of REH-based models regard this “tightness” as the greatest virtue of their approach. They often point to the fact that it “disciplines” economic analysis in a way that was absent in earlier approaches. Indeed, to inculcate and enforce this discipline in the economics profession, every graduate student is warned early on to “[b]eware of theorists bearing free

⁷For a well-known formulation that fully prespecifies stochastic terms within an economist’s model, see Engel (2003).

⁸Thomas J. Sargent, in an interview with Evans and Honkapohja (2005, p. 566).

parameters [arising from autonomous representations of forecasting strategies].”⁹ Following this dictum, the vast majority of economists, whom we refer to as *conventional*, appeal to REH in specifying the microfoundations of their models.

Unsurprisingly, conventional models have experienced the most glaring empirical failures in financial and other markets in which revisions of forecasting strategies are among the key factors driving outcomes. Many of these failures were uncovered by successive generations of conventional economists themselves.¹⁰ Barred from constructing models that accord an autonomous role to market participants’ forecasts, REH theorists have engaged in an intensive effort to explain outcomes on the basis of alternative specifications of the other components of their models. For example, in attempting to explain the excess return of stocks over bonds as a risk premium, REH has led economists to search for alternative specifications of preferences.¹¹

3.2 Behavioral Models

Behavioral economists, for their part, have also uncovered many inconsistencies between the way market participants “actually” behave and standard representations of rational behavior. However, they have not interpreted their findings as evidence that the contemporary standard of rationality does not adequately represent rational decision-making. Instead, they have concluded that market participants are not “fully rational.” This view has led some behavioral economists to use parts of the contemporary standard of rationality to specify certain components of their models, while replacing others with their empirically motivated alternatives.

For example, some behavioral-finance theorists have continued to rely on REH to represent individual forecasting. Thus, they follow their conventional

⁹Attributed to Robert E. Lucas, Jr. in Sargent (2001, p.73).

¹⁰Frydman and Goldberg (2007, chapters 7 and 8) provide an extensive discussion of the empirical failures of REH models in the context of currency markets, along with more than 100 references to research documenting these failures. The list of such references would grow considerably if the conventional approach’s failures in other markets were included.

¹¹This research strategy has an important drawback. As Lucas (2003) points out, preference specifications designed to explain outcomes in one market are often inapplicable or not useful in modeling other markets. This is not surprising, because representations of changes in preferences in REH models must capture the effect of changes in market participants’ preferences and knowledge, and the latter clearly depend on the modeling context.

colleagues in searching for alternative specifications of preferences to remedy the failure of canonical REH models.¹² However, the behavioral approach has an important methodological advantage over its conventional counterpart: it recognizes that forecasting plays an autonomous role in driving markets and thus admits models that do not rely on REH. As a result, behavioral economists have looked for empirically motivated alternatives to REH. This search has yielded a number of valuable insights on how individuals form and revise their forecasts. For example, they have documented a regularity called “conservatism:” individuals tend to revise the strategies that they use to form beliefs about uncertain outcomes in ways that lead to gradual changes in those beliefs.¹³

Admitting such departures from REH into the microfoundations of economic models is an important advance in macroeconomics. However, behavioral economists have followed their conventional colleagues in insisting that, to be worthy of scientific status, their models should generate sharp predictions. Consequently, they have formalized their behavioral insights into how individuals revise their beliefs with mechanical rules that specify exactly all change in advance.

3.3 Internal Inconsistency and Flawed Microfoundations

Beyond ignoring key factors that drive market outcomes, the practice of fully prespecifying change creates insuperable epistemological problems for modern macroeconomics. These problems stem from the standard use of probability theory to represent decision-making under uncertainty.

¹²For a widely-cited behavioral model of the risk premium that relies on REH, see Barberis et al. (2001). This study, which augments risk averse preferences with the assumption of loss aversion, reports the results of calibration exercises that are supportive of the model in the market for equities. However, Frydman and Goldberg (2007, chapter 13) show that, although the Barberis, *et al* model might appear successful according to the calibration methodology in the market for equities, it is grossly inconsistent with the observed time-path of currency prices.

¹³See Edwards (1968) and Shleifer (2000). We make use of this regularity in modeling long swings in asset prices. See section 12.

3.3.1 Standard Probabilistic Representations

Once an economist characterizes causal factors of his model as random variables, his model becomes probabilistic. Such models represent the causal mechanism driving outcomes on the individual level at a point in time with a probability distribution that is conditional on the set of causal variables. A time-invariant model represents outcomes at every point in time with the same conditional distribution. However, even if a contemporary economist allows for change, he fully prespecifies when and how the conditional probability distributions implied by his model vary over time. Because all transitions across probability distributions are fully predetermined, such models in effect characterize individual decision making at every point in time with the same overarching probability distribution.¹⁴

3.3.2 Diversity and Gross Irrationality

Most contemporary economists ignore the diversity of forecasting strategies that we observe in real world markets. To represent such diversity, an economist would specify more than one conditional probability distribution on the individual level. However, if he were to do so, at least some, if not all, of those distributions would systematically differ from the single overarching probability distribution — a sharp prediction — generated by the aggregate model. Thus, any fully predetermined model that recognizes diversity in how market participants forecast the future is *necessarily* internally inconsistent.

Lucas (1995, pp. 254-255; 2001, p.13) has forcefully argued against such models: because an economist hypothesizes that his model adequately represents regularities in market outcomes, his representations of market participants' forecasting strategies should not be systematically inconsistent with the hypothesized regularities that the aggregate model implies. Indeed, in his Nobel lecture, Lucas (1995, p. 255) pointed to internal inconsistency as the key difficulty in constructing macroeconomic models based on representations of individual decision making.

The prevailing strategy for macroeconomic modeling in the early 1960's held that the individual or sectoral models arising out of this intertemporal theorizing could then simply be combined in a single model. But models of individual decisions over

¹⁴For a rigorous demonstration, see chapter 6 of our book.

time necessarily involve expected future prices. . . . However, . . . [aggregate] models assembled from such individual components implied behavior of actual prices. . . that bore no relation to, and were in general *grossly inconsistent* with, the price expectations that the theory imputed to individual agents (Lucas, 1995, pp. 254-255, emphasis added).

Internally inconsistent models presume that market participants adhere to strategies that generate systematic forecast errors, and thus attribute to them gross irrationality. For Lucas, therefore, any such model of time-series regularities is “the wrong theory.”

Lucas’ argument is compelling. However, avoiding inconsistency between a model’s representations on the individual and aggregate levels is not as simple as it might appear: in a world of imperfect knowledge, economists must jettison sharp predictions — the *sine qua non* of Lucas’s own approach.

3.3.3 REH

Otherwise, an economist is left with only one way to rid macroeconomic models of internal inconsistency: he must represent market participants’ forecasting strategies with the one probability distribution generated by the aggregate model that he himself constructs. Indeed, the promise of internal consistency is precisely why Lucas and others embraced REH. As Lucas later put it, “John Muth’s [REH] focused on this inconsistency. . . and showed how it can be removed” (Lucas, 1995).

Although REH automatically removes internal inconsistency from an economist’s model, Muth understood that it should *not* be viewed as a normative hypothesis about how rational individuals should forecast the future. As he put it,

At the risk of confusing this *purely descriptive* hypothesis with a pronouncement as to what firms ought to do, we call such expectations “rational” (Muth, 1961, p. 316, emphasis added).

Of course, what applies to firms is true of other market participants as well. Unfortunately, despite Muth’s warning, REH is, in fact, commonly interpreted as a pronouncement as to what they ought to do.

Early critics pointed out a number of reasons why REH should not be relied on to represent adequately market participants’ forecasting strategies. let

alone how rational – purposeful – individuals forecast the future.¹⁵ Frydman (1982) argued that there is an inherent conflict between REH’s presumption that “people’s beliefs” can be adequately represented with the prediction of an economist’s model and the premise that market participants are motivated by self-interest: purposeful individuals would not, in general, adhere to a single forecasting strategy, let alone the strategy implied by an economist’s model. Moreover, as Phelps (1983) pointed out, economists themselves have constructed a number of alternative models of outcomes. Thus, if a particular economist’s model were somehow to represent rational forecasting, the use of REH in any other model to represent forecasting would have to presume gross irrationality. As we argue below, the only way to escape this conundrum is to jettison the belief that rational decision-making can be fully prespecified by an economist.

REH and the Mathematics of Planning REH models avoid the foregoing difficulties by ignoring the diversity among economists’ models and market participants’ forecasting strategies. Although they sometimes allow for differences in preferences and information,¹⁶ they represent forecasts with a single overarching probability distribution. Such representations are usually referred to as the “representative agent’s” forecasting strategy.

While all economic models are abstractions, the representative-agent assumption is particularly extreme: it ignores the division of knowledge, which is the key feature that distinguishes the allocation of resources by decentralized markets from an “optimal” deployment of resources by a single individual. As Hayek put it,

The economic problem of society is...not merely a problem of how to allocate "given" resources – if "given" is taken to mean given to a single mind which deliberately solves the [resource-allocation] problem....It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is *a problem of the utilization of knowledge*

¹⁵Later critics built on these arguments. Particularly striking are the arguments by Sargent (1993), who was an early advocate of REH. For further discussion and references, see chapter 2 of our book.

¹⁶See Lucas (1973) for an early example and Stiglitz (2001) for an overview of informationally driven explanations of market outcomes.

which is not given to anyone in its totality (Hayek, 1945, p. 519-520, emphasis added).

In relying on REH, economists have ignored Hayek’s arguments. Indeed, Lucas’s account of how REH led him to embrace the representative-agent construct stands in stark contrast to Hayek’s position. In discussing market outcomes for a competitive industry under perfect foresight, which is the deterministic analog of REH, Lucas (2001, p. 13) pointed out that “one can show that an industry over time will operate so as to maximize a discounted, consumer surplus integral—a problem that is mathematically no harder than the present value maximizing problem faced by a single firm.”

Lucas then asked “who, exactly, is solving this planning problem?” As Hayek did, he recognized that “Adam Smith’s ‘invisible hand,’ of course, *not any actual person*” (emphasis added). Nevertheless, in a striking leap of faith, Lucas claimed that an economist — an actual person — can adequately represent what the invisible hand of the market does by solving the value-maximizing problem faced by a single firm.

For Hayek, the division of “knowledge which is not given to anyone in its totality” was the key to his argument that central planners could not, *in principle*, substitute for markets. For Lucas, REH models, which rule out the division of knowledge and enable an economist to make use of a single-agent optimization techniques, were the right tools to comprehend market outcomes.¹⁷ As he put it,

[T]he mathematics of planning problems turned out to be just the right equipment needed to understand the decentralized interactions of a large number of producers. (Lucas, 2001, p.14)

In effect, Lucas posited that Smith’s invisible hand could be made visible and intelligible, after all. To understand markets, economists need only to learn how to solve optimal allocation problems that a fictitious central

¹⁷As we pointed out above, Lucas (1973), Stiglitz (2001) and others have relied on imperfect and asymmetric information in REH models to represent heterogeneity of forecasts in decentralized markets. Frydman (1983) formalizes the arguments in the so-called “Socialist Calculation Debate” by von Mises (1920) and Hayek (1945). Frydman shows that because REH models necessarily ignore the imperfection of knowledge concerning the causal mechanism, Von Mises’ and Hayek’s arguments against planning apply wholesale to contemporary REH models that attempt to capture “the decentralized interactions of a large number of producers” solely with informational imperfections.

planner confronts. Indeed, this is what graduate students in economics are instructed to devote most of their time to doing.

Gross Irrationality of REH Representations Contemporary economists use the “mathematics of planning” to model market outcomes partly because they believe that the representative-agent assumption is just a harmless approximation. However, our foregoing arguments make clear that this abstraction is in fact a knife-edge assumption. Once one recognizes the smallest degree of diversity of forecasting strategies, REH models become internally inconsistent: even if, somehow, an REH model were to represent adequately the average of those strategies, it would presume *ipso facto* that an individual adhered to a forecasting strategy that generates systematic forecast errors *endlessly*. Thus, REH models presume that market participants are grossly irrational. By Lucas’ own argument, they are the wrong theory of market outcomes.

3.3.4 Behavioral Models

The microfoundations of behavioral models often rely on the representative-individual construct. Sometimes, however, behavioral economists attempt to capture the fact that participants in real-world markets employ different forecasting strategies. But, even if they allow for diversity, behavioral economists, too, embrace the conventional belief that economic models should generate sharp predictions and thus fully prespecify change on the individual and aggregate levels.

For an approach that aims for psychological realism, representing market participants as robots who act according to rules that are fully prespecified by an economist is odd. But what makes this approach’s reliance on fully predetermined models particularly puzzling is that it was developed after the REH revolution, the rationale for which was that fully predetermined non-REH models, including non-REH behavioral representations, involve an inherent inconsistency between their probabilistic implications on the individual and aggregate levels. As Lucas argued, internally inconsistent models are the wrong theory of time-series regularities.

3.3.5 Sharp Predictions and the Modern Research Program

The contemporary methodology raises an intractable epistemological problem: there is an inherent conflict between the objective of modeling market outcomes on the basis of explicit, plausible microfoundations and the insistence of both conventional and behavioral economists that their models generate sharp predictions. In real world markets, contemporary models are internally inconsistent, and their “microfoundations” represent grossly irrational behavior.

4 The Promise of Imperfect Knowledge Economics

We have traced the empirical failures and epistemological flaws of contemporary models to economists’ insistence on sharp predictions. In part II, we use a simple algebraic model of the market price to show that in order to escape these flaws, economists must abandon this position. Part III uses the same simple model to illustrate how we have used IKE to model fluctuations and risk in asset markets — phenomena that have confounded conventional models for decades. It also shows how IKE avoids the epistemological flaws of extant approaches.

4.1 Qualitative Models of Change

IKE continues the modern research program in macroeconomics, which was interrupted by the REH revolution. The giants of early modern economics (Knight, 1921; Keynes, 1921, 1936; Hayek, 1945, 1948) and the originators of the modern micro-based approach to macroeconomics (Phelps, 1968, 1970) emphasized the importance of forecasting for understanding market outcomes. However, they also argued that the key feature of capitalist market economies is that they engender change that cannot be prespecified with mechanical rules. At the time Phelps pioneered modern macroeconomics, it was not apparent how to leave mathematical models open to autonomous, non-fully prespecified revisions of forecasting strategies while still representing individual decision-making mathematically. We propose IKE as such an approach and compare it with extant approaches to modern macroeconomics.

Like contemporary models, IKE models consist of representations of an individual’s preferences, the constraints that she faces, her forecasts of the future outcomes that are relevant to her wellbeing, and a decision rule that selects her preferred deployment of resources. However, IKE recognizes that knowledge is inherently imperfect: no one has access to a fully predetermined model that adequately represents, as judged by whatever criteria one chooses, the causal mechanism that underpins outcomes in all time periods, past and future. Consequently, IKE does not fully prespecify which causal variables may be relevant, or when and how these variables may enter an economist’s representation of forecasting behavior. In this way, IKE models remain open to changes in the ways individuals in real-world markets forecast the future — ways that they themselves, let alone economists, cannot specify in advance.

Although IKE jettisons sharp predictions, it aims to explain aggregate outcomes on the basis of mathematical representations of individual decision making. To this end, IKE explores the possibility that revisions of forecasting strategies, though diverse and context-dependent, might exhibit qualitative regularities that can be formalized with mathematical conditions. Nevertheless, an aggregate model based on such microfoundations generates only qualitative predictions of market outcomes.

4.1.1 Non-Standard Use of Probabilistic Formalism

Contemporary models represent outcomes at each point in time, as well as how they unfold over time with a single “overarching” conditional probability distribution.¹⁸ The relationships between the moments of this distribution and the set of causal variables constitute the model’s empirical content that they can be confronted with the time-series data.

By contrast, early modern economists argued that standard probabilistic representations cannot adequately represent change. Indeed, both Knight and Keynes emphasized that economic decisions and institutional and policy changes are fraught with *radical uncertainty*; the complete set of outcomes and their associated probabilities can neither be inferred from past data nor known in advance.

Radical uncertainty is often thought of as a situation in which no economic theory is possible: neither economists (nor market participants) are able to represent mathematically any aspects of the causal mechanism underpinning

¹⁸See chapter 6 of our book.

change. IKE adopts an intermediate position between radical uncertainty and the contemporary presumption that models that fully prespecify change are not only within reach of economic analysis, but anything less is not worthy of scientific status.

Of course, if economic decisions stem only from erratic “animal spirits,” no economic theory is possible. As Phelps (2008) has recently put it, “animal spirits can’t be modelled.” Although animal spirits may play a role, IKE explores the possibility that individual decision-making displays some regularity that can be prespecified with a mathematical model.

Departing from the position of Knight and Keynes, IKE makes use of the probabilistic formalism. This facilitates the formalization of conditions that specify the microfoundations of IKE models and the mathematical derivation of their qualitative implications. However, IKE recognizes the importance of early modern arguments that market participants, let alone economists, have access to only imperfect knowledge of which causal factors may be useful for understanding outcomes and how they influence those outcomes.

Like extant approaches, IKE represents revisions of market participants’ forecasting strategies, and more broadly change in how individuals make decisions, with transitions across probability distributions. However, IKE constrains these revisions with only qualitative conditions. Consequently, it does not follow extant approaches in presuming that individual decision making and market outcomes can be adequately represented with a single overarching probability distribution. At the same time, IKE does not adopt the other extreme position that uncertainty is so radical as to preclude economists from saying anything useful and empirically relevant about how market outcomes unfold over time.

Because its restrictions on change are qualitative, IKE models represent outcomes at every point in time with myriad of probability distributions. Nevertheless, the qualitative restrictions of IKE models constrain all transitions across probability distributions to share one or more qualitative features. These common features, which are embodied in what we call *partially predetermined probability distributions*, enable economists to model mathematically some aspects of the causal mechanism that underpins individual decision making and market outcomes.¹⁹ Such probabilistic representations constitute the empirical content of IKE models.

¹⁹For some simple examples of partially predetermined probability distributions, see chapter 3 of our book.

Although IKE acknowledges the limits to knowledge, it constrains its models sufficiently to distinguish empirically among alternative explanations of aggregate outcomes. In our book, we develop several alternative IKE models and show that their qualitative predictions enable us to reject some in favor of others on the basis of time-series data. Jettisoning sharp predictions may appear to lower the “scientific standard” that economists have self-imposed on their models. However, as Hayek anticipated, replacing the “pretense of exact knowledge” with imperfect knowledge as the foundation for economic analysis is crucial for understanding markets. Remarkably, stopping short of sharp predictions is also necessary to escape the epistemological flaws of extant fully predetermined models.

4.2 Avoiding Epistemological Flaws of Contemporary Models

As we mentioned above, although some well-known behavioral models use REH, the behavioral approach admits non-REH representations of forecasting strategies, thereby according forecasts an autonomous role in driving market outcomes.

The acceptance of the behavioral approach has weakened the position of REH as *the* way to model forecasting behavior. But, it has not diminished contemporary economists’ insistence on models that generate sharp predictions. Indeed, despite their focus on psychological realism, behavioral economists fully prespecify their non-REH representations, and leading behavioral economists have emphasized that theirs “is not meant to be a separate approach [of contemporary economics] in the long run” (Camerer, Loewenstein and Rabin, 2004, p. 42).

Like the behavioral approach, IKE accords market participants’ forecasting strategies an autonomous role in driving individual decision-making and hence market outcomes. However, because it represents revisions of these strategies with qualitative conditions, IKE enables economists to avoid the internal inconsistency inherent in behavioral models. The key reason for this assertion is already apparent in our foregoing discussion of fully predetermined models.

On the aggregate level, the predictions of an IKE model are characterized with myriad distributions, whereas on the individual level, a pluralism of distributions is required to represent diversity of forecasting strategies.

Thus, by limiting itself to qualitative predictions on the aggregate level, IKE opens the microfoundations of its models to a diversity of strategies without necessarily introducing internal inconsistency.

Although jettisoning sharp predictions is necessary to avoid internal inconsistency, it is not sufficient. As we discuss more fully in section , an IKE model avoids inconsistency in different ways, depending on the predictions that it generates at the aggregate level and which aspects of the causal mechanism it constrains on the individual level. We show that an IKE model can avoid internal inconsistency, and thus the presumption of irrationality, even if it assumes that the diversity of forecasting strategies includes some that predict the market price to rise and others that predict it to fall.

5 Contextual Rationality of IKE Models

We have sketched how IKE’s approach to modeling market outcomes on the basis of mathematical microfoundations enables an economist to escape the insurmountable flaws inherent in extant approaches, including the gross irrationality of both REH and fully predetermined non-REH representations. However, the implications of jettisoning sharp predictions are broader. Acknowledging the limits to knowledge of how individuals think about the future calls for a reexamination of the very notions of rationality and purposeful decision-making. This rather large undertaking requires a separate treatment, but our foregoing discussion underscores a few key differences between IKE’s view of rationality and that of extant approaches.

Economists usually equate rationality — being able to justify one’s actions by an appeal to one’s objectives and reason²⁰ — with self-interest. Consequently, they represent rational decision making on the part of an individual with the deployment of resources that maximizes her well-being. That much is common to all approaches, including IKE. Conventional economists, however, go much further. They believe not only that self-interest is a universal human trait, which they invoke as the main “reason” behind individual decision-making, but also that they can adequately represent self-interested behavior with mechanical rules. Indeed, contemporary economists often use the same fully predetermined representations to explain individual self-interested behavior and aggregate outcomes over many decades or in different economies or markets.

²⁰Paraphrased from Meriam-Webster’s Dictionary.

Behavioral economists share both the conventional view of rationality and the belief that it can be represented with mechanical rules. This understanding leads them to diagnose the incompatibility between actual behavior and conventional representations of “rational” behavior as a symptom of market participants’ “irrationality.” They then proceed to use mechanical rules to fully prespecify the irrationality that they believe they have found.

Our discussion of IKE suggests that even if, on the margin, conventional models’ failures have something to do with market participants’ irrationality, attempting to build economic theory on exact representations of irrationality is as futile as attempting to build it on exact representations of rationality. The key premise of IKE is that no mechanical rule can adequately represent how purposeful individuals alter their decision-making and how they forecast the future consequences of their decisions. Thus, we should expect to find what behavioral economists have found: gross discrepancies between conventional representations of rational decision making and the way individuals actually behave.²¹

The key role of forecasting in rational decision making also implies that, even if an economist were able to attribute clear objectives to a market participant, he would still be unable to assess or represent exactly the participant’s rationality or irrationality.²² Rational decisions depend on forecasts of future market outcomes, which are not only a result of the actions of many individuals, but also depend on future economic policies, political developments, and institutional changes. Thus, even if individuals are presumed to be purely self-interested, how they deploy their resources depends at least as much on the social context as it does on their personal motivations. Thus, in a world of imperfect knowledge, rationality is always contextual.

²¹In chapter 13 of our book, we discuss a striking example of how the reliance on mechanical rules can lead to absurd conclusions. Economists have convinced themselves that it is possible to make money systematically by following a rule as simple as betting against the forward exchange rate. And yet individuals in these markets who are handsomely rewarded for finding such profit opportunities have somehow ignored the forward-rate rule.

²²Kay (2004) has called this fundamental difficulty “obliquity.” As he quipped, “no one will ever be buried with the epitaph ‘He maximized shareholder value,’ ...because even with hindsight there is no way of recognising whether the objective has been achieved.”

5.1 Combining Insights from Economics and Other Disciplines

IKE is compatible with, though it does not require, the presumption that market participants are contextually rational. However, we cannot disentangle the complex and changing interdependence between an individual's motivations and her forecasts, which are shaped primarily by the context within which she makes decisions. As such, the characterization of contextual rationality with only a priori assumptions is out of reach of economic analysis.

In modeling individual behavior, therefore, economists must make use of empirical findings about how individuals actually behave. This necessity undermines the common belief among economists, which is increasingly echoed by others, that contemporary economics can rigorously explain the findings of other “soft” social sciences. The contextual view of rationality implies precisely the opposite: in order to represent purposeful individual decision-making, economists must draw on the findings of other social scientists. Consequently, IKE makes use of such findings and insights in specifying the microfoundations of its models.

5.1.1 Preferences

Many studies have found that conventional representations of preferences, which usually involve expected utility theory and the assumption of risk aversion, are grossly inconsistent with the way individuals actually behave. Much of the evidence concerning how individuals make choices is based on laboratory experiments in which the structure of payoffs from various gambles is predetermined by the experimenter. This common experimental design allows the investigator to examine the nature of an individual's preferences without the confounding problem of having to represent her forecasts of the potential payoffs from gambling. The findings concerning the importance of loss aversion and the seminal formulation of prospect theory by Kahneman and Tversky (1979) and Tversky and Kahneman (1992) made use of such a setup.

Building on prospect theory, we develop a representation of preferences for modeling decision-making under imperfect knowledge that is consistent with experimental evidence. This representation, which we call *endogenous prospect theory*, supposes that an individual's preferences share certain qual-

itative features at every point in time. This utility ranking depends on her forecasts of the outcomes of her decisions regarding the allocation of her resources, in particular, on her forecast of future returns and on her forecast of the potential loss that she might incur. The representation also assumes that an individual's degree of loss aversion increases as her forecast of the size of the potential loss increases. Because we represent forecasting with qualitative conditions, the way in which an individual's degree of loss aversion changes between any two points in time is only partially predetermined in our models.

Although laboratory experiments have been the key to uncovering new ways to model preferences, their typical design effectively limits the economist's view of an individual's decision-making; the economist is able to observe only the subject's responses to an experimenter's stimuli. This basic framework, which is used extensively in psychological research, sidesteps a key problem: participants in real-world markets forecast payoffs — the experimenter's "stimuli" — on the basis of imperfect knowledge. Moreover, these forecasts depend not only on the subject's creativity, her analytical abilities, and other personal characteristics, but also on the unfolding social context.²³ As a result, the basic type of model used in these psychological experiments is grossly insufficient as a foundation for representing economic behavior.

5.1.2 Forecasting Behavior

The premise that self-interested or, more broadly, purposeful behavior is to an important degree context-dependent does not preclude the usefulness of insights from psychology in modeling individual behavior. Indeed, we make use of some of these insights in representing how an individual revises her forecasting strategy.²⁴

However, the importance of the social context for an individual's decision-making implies that, in searching for empirical regularities that might be useful in modeling an individual's decisions, economists will need to look beyond laboratory experiments and insights from psychology. Other social

²³Kahneman and Tversky (1979) recognized that laboratory experiments, while useful in uncovering the properties of the utility function over single outcomes, may be much less informative about an individual's choices over gambles with two or more uncertain outcomes in real-world markets.

²⁴See section 12

scientists have knowledge and intuitions concerning the social context within which individuals make decisions that may complement economists' work in modeling individual forecasting behavior.

We make use of Keynes's (1936) insight that conventions among market participants play an important role in individual decision-making. We also draw on our understanding of the qualitative regularities that have characterized aggregate outcomes; we suppose that market participants must also be aware of these of regularities when they form their forecasts. For example, the tendency of exchange rates to undergo long swings away from historical benchmark levels and then to exhibit sustained counter-movements plays a key role in our model of the premium on foreign exchange.

The distinguishing feature of IKE models is that they require an economist to prespecify neither the potential set of causal variables that underpin change in outcomes nor the influences of these variables in his representation. This is important, as the presumption that an economist can prespecify, even partially, the set of causal variables and their influences is very bold.

Nevertheless, in addressing some problems, an economist may have to represent these aspects of the causal mechanism. For example, in order to examine whether macroeconomic fundamentals matter for exchange rate movements, an economist must prespecify, at least partially, a representation of the causal mechanism, which includes the set of potential fundamentals (potential causal factors) and how they influence the exchange rate.

To this end, we consider the possibility that the stock of extant economic models summarizes economists' insights concerning the causal factors that underpin market outcomes. Presumably, these insights are shared by market participants. This idea underlies the Theories Consistent Expectations Hypothesis (TCEH) proposed by Frydman and Phelps (1990). TCEH recognizes that a set of extant economic models at best indicates to a market participant, or to an economist attempting to represent her behavior, which causal variables may be important for forecasting market outcomes; it also suggests, in a qualitative way, how these variables may influence those outcomes.²⁵

In our book, we propose a simple procedure that enables an economist to take into account the qualitative features of more than one model in con-

²⁵In an early implementation of TCEH, Goldberg and Frydman (1996a) relied on the qualitative features of the reduced forms of several REH models to represent individual forecasting behavior. In our book, we eschew REH and propose how an economist can decipher the qualitative features of the reduced forms of a set of models under imperfect knowledge.

structuring his representation of market participants' forecasting strategies. We show that TCEH representations can rationalize many of the sign predictions implied by the REH monetary models of the exchange rate. However, although TCEH may seem to be just a qualitative analog of REH, there are two fundamental differences. In order to account for the social context within which market participants' act, TCEH recognizes that an economist cannot ignore the pluralism of models. Moreover, TCEH only partially prespecifies change.

6 IKE as the Boundary of What Macroeconomic Theory Can Deliver

In our book, we show how IKE can account for the salient features of the empirical record on exchange rates, which have confounded extant models. In part III, we sketch our analysis and show it can be applied to study price movements in other asset markets. Although these results are promising, it is much too early to claim broader usefulness for IKE in macroeconomic and policy modeling.

In contrast to the conventional approach, which seeks to understand economic decisions with universal mechanistic rules, the constraints of IKE models are qualitative and context dependent. If qualitative regularities can be established in contexts other than asset markets, IKE can show how they can be incorporated into mathematical models. However, in contexts in which revisions of forecasting strategies cannot be adequately characterized with reasonably long-lasting qualitative conditions, empirically relevant models of the observed time-series may be beyond the reach of economic analysis. In this sense, IKE provides the boundary to what modern macroeconomic theory — which aims to explain empirical regularities in aggregate outcomes with models that are based on mathematical microfoundations — can deliver.

Part II

Why Macroeconomic Theory Cannot Ignore the Limits to Knowledge

7 Fully Prespecifying Change to Generate Sharp Predictions

In this part, we make to use of a simple algebraic model of the market price to formalize the key arguments of our critique of contemporary models. We begin by showing how the pursuit of sharp predictions leads economists to fully prespecify change in their models.

Our simple algebraic example is motivated by basic supply and demand analysis. In financial markets as well as in many other markets, the quantity demanded and supplied on the individual level is often thought to depend on the forecast of the future market price and a set of causal variables, such as the money supply or rate of growth of real output. Aggregating such representations of individual decision-making, and determining the market price by equating aggregate demand and supply, typically yields the following representation, in semi-reduced form, for the equilibrium price at a given point in time:

$$P_t = a_t + b_t X_t + c_t \hat{P}_{t|t+1} \quad (1)$$

where X_t is a set of causal variables, (a_t, b_t, c_t) is a vector of parameters, and $\hat{P}_{t|t+1}$ is an aggregate of market participants' forecasts formed at t of the market price at $t + 1$.

Individual forecasts that comprise the aggregate forecast, $\hat{P}_{t|t+1}$, are formed on the basis of forecasting strategies at t . Contemporary economists represent these forecasts by relating them to a set of causal variables, which represents the information sets used by market participants. An aggregate of such representations can be written as,

$$\hat{P}_{t|t+1} = \alpha_t + \beta_t Z_t \quad (2)$$

where Z_t is a vector of causal variables that represents the union of information sets used by market participants and (α_t, β_t) is a vector of parameters.

An economist formalizes his assumptions about individual decision-making and how it translates into aggregate outcomes by placing restrictions on his

representations. At each point in time, the structure of an economist’s model is characterized by the following properties:

1. The composition of the set of causal variables appearing in the representations on the individual and aggregate levels.
2. The properties of the joint probability distribution of the causal variables.²⁶
3. A functional form that relates outcomes and the causal variables, typically including the signs of partial derivatives. In cases, such as our example, in which the functional form is explicit, economists often restrict the signs of some parameters.

In general, as time passes, individuals alter the way they make decisions. Various aspects of the social context also change. These changes influence the way aggregate outcomes move over time. Thus, to model the causal mechanism over time, an economist will need different structures — different specifications of forecasting, preferences, constraints, decision and aggregation rules, or the processes driving the causal variables — at different points in time to represent individual behavior and market outcomes.

However, as we have already mentioned, most economists construct models which use the same structure to represent individual behavior and aggregate outcomes at every point in time. Such time-invariant models presume that individuals never alter, let alone devise, new ways to forecast future market outcomes. Sometimes economists do recognize the need to allow for change in their models. However, the insistence that their models should generate sharp predictions leads both conventional and behavioral economists to impose restrictions that fully prespecify this change; that is, they relate the properties of the structure of a model at all points in time, past and future, exactly to the properties of its structure at some arbitrary “initial” point in time.

7.1 Fully Predetermining Restrictions

What is perhaps most striking about contemporary representations is that they prespecify exactly how individuals revise the way they think about the future and how the social context unfolds over time.

²⁶If the model includes additive error terms, the conditions imposed by an economist also specify the joint probability distribution between these terms and the causal variables.

In general, an economist would need to fully prespecify the timing of all revisions and the representations that he believes will adequately represent all post-change forecasting strategies. In doing so, he could modify the set of causal variables that he includes in his representation, or even pick a different functional form. Because such complications would not affect any of our conclusions, we suppose that an economist represents revisions of forecasting strategies with a parametric shift in his aggregate representation at $t + 1$ and that he assumes that these strategies will remain unchanged thereafter:

$$\hat{P}_{t+\tau|t+\tau+1} = \alpha_{t+\tau} + \beta_{t+\tau} Z_{t+\tau} \quad \text{for } \tau = 1, 2, 3... \quad (3)$$

In this example, revisions, which are set to occur only $t + 1$, are represented by two constants $A_{(t,t+1)}$ and $B_{(t,t+1)}$.²⁷

$$A_{(t,t+1)} = \alpha_{t+1} - \alpha_t \quad \text{and} \quad B_{(t,t+1)} = \beta_{t+1} - \beta_t \quad (4)$$

Contemporary economists fully prespecify revisions of forecasting strategies, which in the context of our example, would constrain $A_{(t,t+1)}$ and $B_{(t,t+1)}$ to take on particular values. The assumption of time-invariance would constrain the representation of forecasting strategies to remain unchanged at all times, past and future; that is it would set $A_{(t,t+1)} = 0$ and $B_{(t,t+1)} = 0$.

Infrequently, contemporary economists allow for revisions of forecasting strategies in their models. A simple way to fully predetermined change is to set $A_{(t,t+1)}$ and $B_{(t,t+1)}$ equal to particular values $\bar{A}$ and $\bar{B}$, respectively:

$$A_{(t,t+1)} = \alpha_{t+1} - \alpha_t = \bar{A} \quad \text{and} \quad B_{(t,t+1)} = \beta_{t+1} - \beta_t = \bar{B} \quad (5)$$

We refer to such restrictions, which relate a model's pre-and post-change structures, as *fully predetermining*. Sometimes fully predetermining restrictions are probabilistic. For example, an influential class of contemporary models makes use of a rule that fully prespecifies the timing of all changes and the functions that relate $A_{(t,t+1)}$ and $B_{(t,t+1)}$ to the values of the causal variables, conditional on the the "initial" structure.²⁸

²⁷Except for purely formal complications, our conclusions in this section apply to nonlinear representations. For example, suppose that the representation of the aggregate forecasting strategy at $t + 1$ is a nonlinear function of the causal variables. In such a case, $A_{(t,t+1)}$ and $B_{(t,t+1)}$ would be nonlinear functions of the causal variables.

²⁸For example, see Hamilton (1989, 1994). Frydman and Goldberg (2007, chapter 6) show that all of our conclusions in this section apply to models that use a probabilistic rule to fully prespecify the timing of revisions and the post-change forecasting strategies.

To complete the task of fully predetermining his model, an economist would follow the foregoing procedure and fully prespecify its other components. To simplify our example further and focus on change arising from revisions in forecasting strategies, we follow much of the literature and assume that these other components are time-invariant. This assumption implies the following fully predetermining restrictions in (1):

- The composition of the set causal variables, X_t , and the properties of their joint probability distribution, remain unchanged at all times, past and future.
- The parameters (a_t, b_t, c_t) are constants, that is, $(a_t, b_t, c_t) = (a, b, c)$ for all t .

By imposing the foregoing fully predetermining restrictions on (1) and (2), an economist presumes that the following model adequately characterizes the aggregate of forecasting strategies and the market price at $t + 1$ and beyond:

$$\hat{P}_{t+\tau|t+\tau+1} = (\alpha_t + \overline{A}) + (\beta_t + \overline{B})Z_{t+\tau} \quad \text{for } \tau = 1, 2, 3... \quad (6)$$

$$P_{t+\tau} = (a + c(\alpha_t + \overline{A})) + bX_{t+\tau} + c(\beta_t + \overline{B})Z_{t+\tau} \quad \text{for } \tau = 1, 2, 3... \quad (7)$$

7.2 Sharp Probabilistic Predictions of Change

The model of the causal mechanism driving outcomes in (6) and (7) enables an economist to generate sharp probabilistic predictions of future outcomes. In order to do so, an economist must fully prespecify the movement of the causal variables over time. To keep our example simple, we assume that all causal variables follow a random walk with constant drift,

$$X_t = \mu^X + X_{t-1} + \varepsilon_t^X \quad (8)$$

$$Z_t = \mu^Z + Z_{t-1} + \varepsilon_t^Z \quad (9)$$

where $\varepsilon_t = (\varepsilon_t^X, \varepsilon_t^Z)$ is an i.i.d. vector of random variables and $E[\varepsilon_t^X] = E[\varepsilon_t^Z] = 0$.²⁹

²⁹Such random walk specifications are popular in macroeconomic models that involve variables such as the money supply, income, interest rates, and the price level.

To illustrate the concept of sharp probabilistic predictions of change on the individual level, we shift (6) and (8) one period ahead. Using (2) implies the following representation of revisions in forecasting strategies:

$$\Delta \hat{P}(t+1, t) = \hat{P}_{t+1|t+2} - \hat{P}_{t|t+1} = [\bar{A} + (\beta_t + \bar{B})\mu^Z] + (\beta_t + \bar{B})\varepsilon_t^Z \quad (10)$$

The model also implies the following representation of change in market outcomes between t and $t+1$:

$$\Delta P(t+1, t) = P_{t+1} - P_t = [c\bar{A} + b\mu^X + c(\beta_t + \bar{B})\mu^Z] + [b\varepsilon_t^X + c(\beta_t + \bar{B})\varepsilon_t^Z] \quad (11)$$

The expressions in (10) and (11) show that fully predetermined models represent forecasting strategies and aggregate outcomes as unfolding over time around deterministic time paths — $[\bar{A} + (\beta_t + \bar{B})\mu^Z]$ and $[c\bar{A} + b\mu^X + c(\beta_t + \bar{B})\mu^Z]$ — conditional on the structure of the model at some arbitrary point in time— α_t , β_t , a , b , c , μ^X , and μ^Z . Although contemporary models allow for deviations from these paths — $(\beta_t + \bar{B})\varepsilon_t^Z$ and $[b\varepsilon_t^X + c(\beta_t + \bar{B})\varepsilon_t^Z]$ — they fully prespecify how the probability distributions of these deviations might change over time.³⁰ Thus, as different as they may appear, fully predetermined probabilistic representations of change are as restrictive as their deterministic counterparts, linear or nonlinear: they both make no allowance whatsoever for the possibility that at some point, individuals decide to think about the future in ways that could not have been foreseen in advance or that policy officials alter policy in new ways.

7.2.1 History as the Future and Vice Versa

Another way to comprehend how odd this conception of change is to note that contemporary models represent history as if it were fully reversible. To illustrate this point, consider the representation of forecast revisions in (10) and suppose that one of the causal variables increases with the passage of time from t to $t+1$. An individual is supposed to revise her forecast between t and $t+1$ in a way that is fully prespecified by an economist. Then, if the causal variable subsequently declines back to its original value between $t+1$

³⁰See Engle (2003). As is usual in the literature, our example assumes that these distributions are time-invariant.

and $t + 2$, market participants are presumed to revert to the same forecast as the one they started with at t .

Although this conclusion follows immediately from the invariance restriction, it also holds in a model in which the change in the causal variable triggers a fully prespecified revision in forecasting strategies. However, as soon as the causal variable returns to its initial value, forecasting strategies also revert to their initial representations.^{31,32} This representation of change is tantamount to presupposing that as time passes, market participants do not gain any experience or come up with any genuinely new ways of thinking about the future. Remarkably, the contemporary approach presumes that the passage of time does not play an essential role in altering individual behavior and the social context within which individuals make decisions.

8 Epistemological Flaws of Contemporary Models in a World of Imperfect Knowledge

8.1 Internal Inconsistency When Forecasting Strategies Are Diverse

Hayek (1945) posited that one cannot understand how markets allocate resources without considering the central role played by the division of knowledge, which is a key factor behind diversity of forecasting strategies. Yet, a fully predetermined model that allows for such diversity is necessarily internally inconsistent.

To illustrate this inconsistency, suppose that the model in (1) adequately represents the causal mechanism that underpins the behavior of the equilibrium price. Also suppose that in modeling the microfoundations of (1), an

³¹See Frydman and Goldberg (2007, chapter 5) for a rigorous demonstration in the context of demand and supply analysis.

³²We note that change in multiple-equilibrium models is not, in general, reversible: a return of a causal variable to its initial value could be associated with a move to an equilibrium other than the initial one. Economists sometimes construct so-called hysteresis models in which the movement of an outcome variable is path-dependent. For example, see Krugman (1987). Nevertheless, both types of models fully prespecify the equilibria to which the system does move. Consequently, they both presume that individuals have not invented any genuinely new ways of thinking about the future as history unfolds. See Frydman and Goldberg (2007, chapter 6) and references therein.

economist recognizes that market participants make use of diverse forecasting strategies, which for simplicity we represent with two strategies, $\hat{P}_{t|t+1}^{(1)}$, and $\hat{P}_{t|t+1}^{(2)}$. To generate sharp predictions, the representations of these two strategies must be fully predetermined. Following the literature, and without a loss of generality, we constrain these representations to be time-invariant:

$$\hat{P}_{t|t+1}^{(1)} = \alpha^{(1)} + \beta^{(1)} Z_t^{(1)} \quad (12)$$

and

$$\hat{P}_{t|t+1}^{(2)} = \alpha^{(2)} + \beta^{(2)} Z_t^{(2)} \quad (13)$$

where, $Z_t^{(i)}$ $i = 1, 2$ are vectors of causal variables that represent information sets used by market participants in each group, $(\alpha^{(i)}, \beta^{(i)})$, $i = 1, 2$ are vectors of parameters, and the diversity of forecasting strategies implies $\alpha^{(1)} \neq \alpha^{(2)}$ and $\beta^{(1)} \neq \beta^{(2)}$. The aggregate forecast appearing in (1) can then be represented as an average of market participants' forecasts:

$$\hat{P}_{t|t+1} = \omega \hat{P}_{t|t+1}^{(1)} + (1 - \omega) \hat{P}_{t|t+1}^{(2)} \quad (14)$$

where the weight $0 < \omega < 1$ represents the importance of type 1 individuals in influencing the behavior of the equilibrium price.

Substituting (14) into (1) yields the following representation of the causal mechanism that drives the market price:

$$P_t = a + bX_t + c [\omega \alpha^{(1)} + (1 - \omega) \alpha^{(2)}] + c [\omega \beta^{(1)} Z_t^{(1)} + (1 - \omega) \beta^{(2)} Z_t^{(2)}] \quad (15)$$

It is apparent that the model is internally inconsistent in the sense articulated by Lucas: the representations of individual forecasting strategies in (12) and (13) are systematically inconsistent with the representation of the market price (15). To see this, assume that the causal variables follow random walks with constant drifts, analogous to (8) and (9). This implies that the difference between the aggregate representation and the representation for, say, type 1 individuals can be written as:

$$P_{t+1} - \hat{P}_{t|t+1}^{(1)} = C + bX_t + (c\omega - 1) \beta^{(1)} Z_t^{(1)} + c(1 - \omega) \beta^{(2)} Z_t^{(2)} + \eta_t \quad (16)$$

where C is a non-zero constant that depends on drift terms and η_t is an i.i.d. random variable that depends on error terms from the X , Z^1 , and Z^2 processes. As we discussed in section 3.3, such representations imply that individuals systematically forego obvious profit opportunities.

8.2 REH as a Representation of Grossly Irrational Forecasting Strategy

REH models avoid internal inconsistency by employing the representative-agent construct, which disregards the diversity of forecasting strategies that drive market outcomes. Of course, self-interested, rational individuals would collectively adhere to one forecasting strategy in perpetuity only if “all agents have solved their ‘scientific problems’” (Sargent 1993, p.23). In such an imaginary world, REH would be a plausible hypothesis. All market participants and economists would have discovered an overarching causal mechanism that characterizes aggregate outcomes, as well as how the causal factors unfold over time, and thus, individual creativity and, more broadly change that does not follow pre-existing rules, would cease to be economically important. Economic decisions would become purely routine and passive, and thus capable of being captured by fully predetermined representations. In this fanciful world, contemporary representations would adequately explain individual behavior and would be completely consistent with the model of aggregate outcomes. Moreover, in such an *REH world*, the heterogeneity of forecasts among market participants would stem solely from differences in information.

Of course, in the real world, where the scientific problem has not been solved, there is a division of knowledge among individuals. Market participants forecast not only on the basis of different factors (their information sets), but also on the basis of different strategies (their knowledge) that map these factors into forecasts. No one knows — because no one *can know* — precisely how knowledge differs among individuals.

By design, REH models are fully predetermined. Thus, if REH models were to recognize the diversity of forecasting strategies, they would be internally inconsistent. In positing that inconsistent models are the wrong theory, Lucas (2001, p.13) argued that they represent individual market participants as *grossly irrational*, in so far as they disregard *endlessly* systematic information in their forecast error. It is easy to show that this conclusion applies

to REH models in a world of imperfect knowledge.

To this end, we write the REH representation of the representative individual's forecasting strategy as a linear function in X_t :

$$\hat{P}_{t|t+1}^{\text{RE}} = \alpha^{\text{RE}} + \beta^{\text{RE}} X_t \quad (17)$$

REH instructs an economist to determine his individual and aggregate representations *jointly*. Consequently, while an REH theorist would specify individual preferences and constraints autonomously from his aggregate model, his representation of an individual's forecasting behavior is derivative of these other components. The causal variables X_t , and the parameters $(\alpha^{\text{RE}}, \beta^{\text{RE}})$ of the forecasting strategy in (17) stem from the representation of preferences, the social context, and the constraints that an individual faces in making her decisions. In this way, REH rules out an autonomous role for market participants' knowledge in shaping market outcomes.

To derive the REH representations, an economist chooses the coefficients α^{RE} and β^{RE} to be functions of the parameters a , b , and c in (1) to ensure the required consistency between the representations on the individual and aggregate levels:

$$\hat{P}_{t|t+1}^{\text{RE}} = E[P_{t+1}^{\text{EM}} | X_t] \text{ for all realizations of } X_t \quad (18)$$

where the superscripts “RE” and “EM” denote REH representation and the representation that is implied by an economist's model, respectively, and $E[\cdot]$ is the expectation of P_{t+1}^{EM} , conditional on X_t and the constraint that the structure in (1) is time-invariant.

Imposing (18) in (1), and using (8), the model implies the following REH representation for the market price at $t+1$, and the average — representative — forecast of this price formed at t :³³

$$P_{t+1}^{\text{RE}} = \frac{a(1-c) + b\mu^X}{(1-c)^2} + \frac{b}{(1-c)} X_t + \epsilon_{t+1} \quad (19)$$

$$\hat{P}_{t|t+1}^{\text{RE}} = \frac{a(1-c) + b\mu^X}{(1-c)^2} + \frac{b}{1-c} X_t \quad (20)$$

where $\epsilon_{t+1} = \frac{b}{1-c} \varepsilon_{t+1}^X$.

³³For details of the derivation of the REH solution, see chapter 3 in Frydman and Goldberg (2007).

Now suppose, as in fact every REH theorist does, that (19) and (20) represent adequately the causal mechanism that underpins the time-path of the market price and the representative forecast, respectively. Moreover, we follow the usual interpretation of a representative individual's forecasting strategy and assume that the representation in (20) stands for an aggregate of representations of individual behavior. Without a loss of generality, we write this aggregate forecast analogously to (14),

$$\widehat{P}_{t|t+1}^{\text{RE}} = \omega_1 \widehat{P}_{t|t+1}^{(1)} + \omega_2 \widehat{P}_{t|t+1}^{(2)} \quad (21)$$

The REH representation in (21) is assumed to represent the average of forecasts across market participants. Thus, the fully predetermined representations of the individual forecasting strategies are

$$\widehat{P}_{t|t+1}^{(1)} = \alpha^{(1)} + \beta^{(1)} X_t \quad (22)$$

$$\widehat{P}_{t|t+1}^{(2)} = a^{(2)} + \beta^{(2)} X_t \quad (23)$$

where again the diversity of forecasting strategies implies $\alpha^{(1)} \neq \alpha^{(2)}$ and $\beta^{(1)} \neq \beta^{(2)}$.³⁴

This REH model also implies the following representation of the forecast errors for type 1 individuals:

$$fe_{t|t+1}^{1,\text{RE}} = P_{t+1}^{\text{RE}} - \widehat{P}_{t|t+1}^{1,\text{RE}} = C' + \left[(b - \beta^{(1)}) + c\omega(\beta^{(1)} - \beta^{(2)}) + c\beta^{(2)} \right] X_t + \epsilon'_{t+1} \quad (24)$$

where again C' depends on drift terms and ϵ'_{t+1} depends on the error term for the X_{t+1} process. Thus, in a world in which individuals have not solved their scientific problems, but where an REH model nonetheless adequately represents the behavior of market outcomes, the REH representation of an individual's forecasting strategy implies that she disregards endlessly the obvious systematic information contained in her forecast errors. To paraphrase Lucas (2001, p.13), the REH model would imply that there are profit opportunities that a market participant could see. If she does see these opportunities and she is rational, she would revise her forecasting strategies. Not doing so — ever — would be grossly irrational.

³⁴These parameters must satisfy the restriction implied by the representations in (20) through (23).

8.3 Behavioral Models

Despite its epistemological implausibility and empirical failures, some behavioral economists continue to rely on REH to represent individual forecasting.³⁵ In contrast to conventional methodology, however, the behavioral approach does not oblige an economist to use REH. This has led some behavioral economists to develop non-REH representations of market participants forecasting strategies, which are autonomous from their aggregate models. Departing from REH has also enabled behavioral economists to recognize the role of diversity in driving market outcomes.³⁶

8.3.1 Internal Inconsistency

However, behavioral economists embrace the conventional insistence on sharp predictions and thus fully predetermine their non-REH models. Even if they do not allow for diversity, fully predetermined models that jettison REH are internally inconsistent.

To illustrate this point, suppose that on the basis of some empirical considerations, a behavioral economist specifies a time-invariant version of the non-REH representation in (2):³⁷

$$\hat{P}_{t|t+1} = \alpha + \beta Z_t \quad (25)$$

Substituting (25) into the time-invariant version of (1) and shifting it one period ahead yields:

$$P_{t+1} = [a + c\alpha + b\mu^X + c\beta\mu^Z] + bX_t + c\beta Z_t + [b\varepsilon_t^X + c\beta\varepsilon_t^Z] \quad (26)$$

where we used the representations for X_{t+1} and Z_{t+1} in (8) and (9), respectively.

Comparing the representation on the individual level in (25) with that on the aggregate level in (26) immediately implies that for $(\alpha, \beta) \neq (\alpha^{\text{RE}}, \beta^{\text{RE}})$, the foregoing fully predetermined behavioral model is internally inconsistent.

³⁵For example, see a widely-cited behavioral model of the risk premium by Barberis et al. (2001).

³⁶One of the earliest examples that allowed for diversity of forecasting strategies, and that anticipated the behavioral approach, is the seminal exchange-rate model of Frankel and Froot (1987).

³⁷Our argument remains valid for behavioral models that allow for change, but fully prespecify it. See chapter 6 of our book.

As Lucas argued, such internally inconsistent models are “the wrong theory” of time-series regularities.

Moreover, the microfoundations of the model based on the fully predetermined representation in (25) presume that on average market participants are grossly irrational. This can easily be seen by subtracting (25) from (26):

$$fe_{t|t+1} = P_{t+1} - \hat{P}_{t|t+1} = D + bX_t + (c - 1)\beta Z_t + \varepsilon_{t+1} \quad (27)$$

where again D depends on drift terms and ε_{t+1} depends on the error term for the X_{t+1} and Z_{t+1} processes. The expression in (27) shows that the model attributes to a representative individual a forecasting strategy that generates systematic forecast errors *endlessly*. Thus, although non-REH behavioral models recognize that market participants’ forecasts are “inputs” to rather than outputs of economic models, their microfoundations are implausible. They retain the key feature of REH models: they presume that market participants never devise ways to forecast the future that an economist cannot specify in advance.

8.3.2 The Inessential Role of Diversity in Behavioral Models

In real world markets, the imperfection of knowledge gives rise to diversity of forecasting strategies and to change that does not follow preexisting rules. Behavioral economists have constructed non-REH models to capture the role of such factors in driving market outcomes. However, because they fully prespecify change, the fundamental implication of behavioral models for the way the economy unfolds over time is no different than that of models that do not allow for diversity, regardless of whether they are behavioral or they rely on REH.

To illustrate this point, we suppose that an economist, drawing on some behavioral observations, represents diversity with the two fully predetermined representations in (12) and (13). We also suppose that this behavioral economist follows common practice and presumes that the model in (15), with its unchanging structure, adequately represents the causal mechanism that drives outcomes. Consequently, the change in the market price between t and $t + 1$ is given by

$$\Delta P(t, t+1) = \mu + c\omega_1\beta^{(1)}\Delta Z^{(1)}(t, t+1) + c\omega_2\beta^{(2)}\Delta Z^{(2)}(t, t+1) + b\Delta X(t, t+1) \quad (28)$$

where $\Delta P(t, t+1) = P_{t+1} - P_t$, and $\mu = [c\omega_1\beta^{(1)}\mu^{Z^{(1)}} + c\omega_2\beta^{(2)}\mu^{Z^{(2)}} + b\mu^X]$.

The representation in (28) highlights the inessential role played by diversity in behavioral models. The implications of those models for the way market outcomes move over time are the same as those of behavioral and REH models that ignore diversity: except for fully prespecified random deviations, behavioral models that allow for diversity represent movements in market prices as unfolding along a fully predetermined path.

9 Why Macroeconomics Must Open Its Models to Imperfect Knowledge

We now show that to escape the insurmountable epistemological difficulties inherent in extant approaches, it is *necessary* for economic models to stop short of fully prespecifying change. This is tantamount to acknowledging that sharp predictions are outside the reach of economic analysis and that the most economic theory can be expected to deliver is qualitative predictions of market outcomes.

Consider again the model in equations (12), (13), and (15). For simplicity, we continue to impose the invariance restriction in (1), but drop this assumption for the representations of market participants' forecasting in (12) and (13). We write the representations of these strategies at time $t + 1$ in terms of their structure and the realizations of the causal variables appearing in them at time t :

$$\hat{P}_{t+1|t+2}^{(1)} = [\alpha_t^{(1)} + A^{(1)}(t, t + 1)] + [\beta_t^{(1)} + B^{(1)}(t, t + 1)]Z_{t+1}^{(1)} \quad (29)$$

and

$$\hat{P}_{t+1|t+2}^{(2)} = [\alpha_t^{(2)} + A^{(2)}(t, t + 1)] + [\beta_t^{(2)} + B^{(2)}(t, t + 1)]Z_{t+1}^{(2)} \quad (30)$$

where $A^{(i)}(t, t + 1) = \alpha_{t+1}^{(i)} - \alpha_t^{(i)}$ and $B^{(i)}(t, t + 1) = \beta_{t+1}^{(i)} - \beta_t^{(i)}$, $i = 1, 2$. Analogously to (15), the model implies the following representation of the causal mechanism for the market price at time $t + 1$:

$$\begin{aligned} P_{t+1} = & a + bX_t + \left\{ c\omega[\alpha_t^{(1)} + A^{(1)}(t, t + 1)] + (1 - \omega)[\alpha_t^{(2)} + A^{(2)}(t, t + 1)] \right\} \\ & + c \left\{ \omega[\beta_t^{(1)} + B^{(1)}(t, t + 1)]Z_t^{(1)} + (1 - \omega)[\beta_t^{(2)} + B^{(2)}(t, t + 1)]Z_t^{(2)} \right\} + \eta_t \end{aligned} \quad (31)$$

where η_t depends on the error terms for the processes representing the causal factors X_t and Z_t . It is clear that to avoid inconsistency between the predictions of any one of the diverse forecasting strategies allowed for by (29) and (30) and those of the representation of the market price, the aggregate representation in (31) must be compatible with more than one conditional probability distribution.³⁸

10 IKE: Recognizing Imperfect Knowledge and Diversity of Forecasting Strategies

10.1 Forecasting As Entrepreneurship

The contemporary approach presumes that economists can represent exactly how individual decision-making, particularly individuals' forecasting strategies, and aggregate outcomes unfold over time. If true, this would suggest that understanding the past is a relatively straightforward problem: all one needs to do is to specify an appropriate econometric model and estimate it on the basis of as much past data as are available. As the future in a contemporary model follows exactly from the past, except for "sampling error", the estimated econometric model could be used to forecast market outcomes.

Purposeful individuals, however, recognize that even when it comes to the past, everyone has only imperfect knowledge about the timing of changes in the causal mechanism driving outcomes. It is unclear therefore how to implement the econometric strategy suggested by the contemporary approach: the use of past data to estimate the model requires an assessment of how far back in time it adequately represents the economy. Even the most sophisticated statistical techniques would not automatically pinpoint when the last structural break had occurred.³⁹ Thus, to understand the past causal mechanism one cannot merely estimate some model using all the available data and presume that it adequately represents historical outcomes. Unsurprisingly, interpretations of the past vary among individuals even when they

³⁸All probability distributions of future outcomes that we refer to are conditional on the structure of the model and the realizations of the causal variables at some prior point in time, which in the present case is t .

³⁹The results of statistical tests of the model's invariance depend on the significance level of these tests, which involve an element of judgement on the part of the forecaster or an economist. See chapter 15 of our book.

use formal methods.

Moreover, even if we could adequately interpret the past with a fully pre-determined model, the future would still remain imperfectly known. Because future outcomes depend on revisions of forecasting strategies that cannot be predicted “by rational and scientific methods” (Popper, 1957, p. xii), market participants do not merely rely on pre-existing rules, whether based on formal or informal considerations, to forecast future changes in the causal mechanism. Ultimately, good forecasting is much like successful entrepreneurship: it may involve the use of quantitative models, but it also relies on one’s own personal knowledge, intuition, and a bit of luck in spotting profit opportunities.

For example, consider the problem of forecasting exchange rates. Many market participants form exact forecasts of the future exchange rate, for example, that a euro will cost \$1.5 in a week. After all, a currency trader must decide on her market position at each point in time. Nevertheless, although a market participant may base her trading on exact predictions, she does not arrive at such predictions by relying solely on quantitative models, much less the same model in every time period. In forming her forecasts, a purposeful individual often combines *her preferred* quantitative model with her own insights concerning the behavior of other market participants, the historical record on exchange rate fluctuations, and her evaluation of the impact of past and future decisions by policy officials.⁴⁰ Moreover, because market participants act on the basis of different experiences, interpretations of the past, and intuitions about the future, they adopt diverse strategies in forming and revising their exchange rate forecasts over time.

10.2 Imperfect Knowledge and Economic Analysis

Even though individuals in real world markets combine both formal and informal methods in forming their forecasts, IKE recognizes that representing their behavior in a mathematical model requires an economist to find some way to formalize it. Like any scientific theory, IKE strives to generate implications that have empirical content. As a result, it must presume that purposeful decision-making, although driven by a variety of factors that

⁴⁰In a world of imperfect knowledge, even the use of quantitative models alone involves subjectivity, because there is more than one way to represent the causal mechanism mathematically. One important reason is that, as in (1), economic models typically involve representations of market participants forecasting strategies. See chapter 15 of our book.

cannot be made fully intelligible by the individuals themselves, let alone by economists, nevertheless exhibits some regularities. Otherwise, no macro-economic theory that can be confronted with time-series evidence would be possible. IKE searches for these regularities and formalizes them with the aid of probabilistic formalism.

10.2.1 Representing Imperfect Knowledge with Partially Predetermined Probability Distributions

In the context of our simple example, the constants $A^{(i)}(t, t+1)$ and $B^{(i)}(t, t+1)$, $i = 1, 2$ in (29), (30) and (31) represent revisions of forecasting strategies and the resulting change in the causal mechanism on the aggregate level. A contemporary economist would fully prespecify change by imposing quantitative conditions on these constants, typically setting $A^{(i)}(t, t+1) = 0$ and $B^{(i)}(t, t+1) = 0$. Such fully predetermining restrictions, together with stochastic representations of the causal factors, as in (8) and (9), for example, imply that the model represents the market price at $t+1$ with a single probability distribution, conditional on the structure and the values of the causal variables at t . It also represents the consequences of individual decisions with a single probability distribution.

By contrast, as we discussed in section 4.1.1, Knight and Keynes forcefully argued that business decisions are fraught with radical uncertainty, a situation in which neither market participants nor economists are able to represent any aspects of the causal mechanism that underpins change.

In the context of (29)-(31), such extreme interpretations of radical uncertainty can be formalized by not imposing any constraints on $A^{(i)}(t, t+1)$ and $B^{(i)}(t, t+1)$, $i = 1, 2$. Without constraints, equations (29)-(31) represent change with myriad of conditional probability distributions. Within the confines of our simple linear example, such representations imply that the model in (31) is compatible with *any* relationship between the causal variables and the market price. In this extreme version, radical uncertainty inherently conflicts with economists' attempts to distinguish among alternative causal explanations of market outcomes.

Although it does not impose fully predetermining restrictions on change, IKE aims to explain outcomes with mathematical models that can be confronted with empirical evidence. To this end, IKE uses probabilistic formalism to formulate its mathematical representations of forecasting strategies and their revisions.

IKE thus adopts an intermediate position between radical uncertainty, which denies the possibility that economists might be able to formulate testable mathematical models of any features of the causal mechanism driving change, and the contemporary presumption that this mechanism can be adequately represented with a standard conditional probability distribution. Because IKE recognizes that everyone has access to only imperfect knowledge, it represents individual behavior and aggregate outcomes with many conditional probability distributions at every point in time. To model change, IKE constrains the possible transitions across these distributions to satisfy some common qualitative properties. As we show in part III, the qualitative predictions concerning how the moments of the partially predetermined probability distributions unfold over time constitute the empirical content of the model, in other words, that which can be confronted with time-series data.

Part III

Imperfect Knowledge Economics of Market Fluctuations and Risk

Market outcomes often undergo protracted swings that revolve around historical benchmark levels. For example, persistent movements in the unemployment rate away from a longer-term trend, or in stock prices away from levels consistent with historical price-earnings (PE) ratios, have occurred for years at a time. But the instability in the economy is bounded; eventually economic outcomes undergo sustained countermovements back to benchmark levels, which themselves vary over time. If outcomes happen to reach these levels, they often shoot through them. Moreover, although fluctuations are a recurrent feature of market outcomes in capitalist economies, the observed swings are uneven: the magnitude and duration of upswings and downswings vary from one episode to the next in a way that does not seem to follow any consistent pattern.

Economists have grappled for decades with trying to understand the causal mechanism driving fluctuations in capitalist economies. Some contemporary models explain market fluctuations not as movements away from the benchmark, but as movements of the benchmark itself. In these mod-

els, benchmark levels unfold along fully prespecified time paths. Thus, to explain fluctuations, economists have allowed for shifts in these time paths, stemming from changes in technology, policy, or knowledge. While these factors undoubtedly play a role in how benchmarks might vary over time, their influences cannot be fully prespecified.⁴¹

The movement of the imperfectly known benchmark is likely to play a significant role in some markets. However, it is uncontroversial to view fluctuations in asset markets as movements that revolve around benchmark levels.

Figures 1 and 2 provide just two examples of the importance of viewing fluctuations in capitalist economies as swings that revolve around benchmark levels, rather than as movements of the benchmarks themselves. In figure 1, we plot the German mark-U.S. dollar (DM/\$) exchange rate and its purchasing power parity (PPP) level,⁴² while in figure 2, we plot the price of the Standard and Poor's 500 basket of stocks relative to its underlying earnings and a 20-year moving average of this PE ratio.⁴³ Both figures

⁴¹See Phelps (2008) for an argument that movements of the natural rate cannot be fully prespecified in capitalist economies.

⁴²The PPP benchmark in the figure is based on the *Big Mac* PPP exchange rate reported in the April 1990 issue of *The Economist* magazine (which was 1.96) and CPI-inflation-rate differentials from the IMF's *International Financial Statistics*.

⁴³Figure 2 is based on data from Shiller (2000), which are updated on his web site: www.econ.yale/~shiller. PPP and average PE ratios have long traditions as benchmark levels in currency and stock markets, respectively.

Figure 1
DM/\$ Exchange Rate: 1973-1998

show prolonged periods in which the asset price tends to move persistently away from its benchmark. The figures also show that these price movements are uneven and ultimately bounded: the duration and magnitude away from benchmark levels show no obvious pattern, but eventually they are followed by extended periods of time in which the asset price moves back toward benchmark levels.

In the remainder of this paper, we focus our attention on fluctuations in asset markets. In these markets, participants' forecasts play a key role in driving movements in prices away from and toward benchmark levels.⁴⁴ Beyond modeling forecasting, we also sketch a new model of risk to help explain why long swings are ultimately bounded. Although we develop our IKE analysis in the context of asset markets, particularly those for currencies, our approach may prove useful in modeling fluctuations in other contexts in which forecasts are important in driving outcomes. However, as with the

⁴⁴A model in which price movements arise from fluctuations in the imperfectly known benchmark and from swings away from the benchmark is beyond the scope of this paper.

foreign exchange market, applying IKE in other markets will require that an economist search for

11 The Failure of REH Models

It is widely known that REH models are unable to account for the long-swings behavior of asset prices.⁴⁵ In currency markets, for example, flexible-price REH models represent exchange rate dynamics as movements of the benchmark real-exchange rate. In these models, the benchmark depends on preferences and technology.⁴⁶ Thus, to account for the swings that we observe in figure 1, flexible price-models must assume that movements in these factors, whether stochastic or deterministic, undergo the kind of swings we observe in the exchange rate. But, exchange rate swings reverse direction much too often and abruptly for shifts in preferences and technology to provide a plausible explanation.⁴⁷

Another influential class of REH models, which assume that goods prices and/or wages adjust only sluggishly over time to equilibrium levels, view exchange rate dynamics as deviations from the PPP benchmark level.⁴⁸ However, because REH ignores the role of autonomous revisions in forecasting strategies in driving outcomes, these sticky-price models are also unable to explain exchange rate fluctuations. By design, REH rigidly ties representations of individual forecasting behavior to the PPP benchmark rate: all market participants are assumed to predict a movement in the exchange rate back toward PPP whenever a deviation from this benchmark arises. Consequently, sticky price models predict that movement in the exchange rate back to this benchmark will, on average, follow any time period involving a deviation from PPP. Such predictions are inconsistent with the persistent

⁴⁵For early studies of the inconsistency of REH models with swings in stock and bond markets, see Shiller (1979, 1981, 1990). For an extensive discussion of the empirical failures in currency markets, see chapter 7 in Frydman and Goldberg (2007). REH models have also failed to explain excess returns in asset markets as risk premiums. In chapters 8 and 12 of our book, we discuss the inability of REH models to account for risk in currency markets.

⁴⁶See, for example, Stockman (1980) and Lucas (1982).

⁴⁷For criticism of flexible-price models of exchange rate swings along these lines, see Dornbusch (1989).

⁴⁸See, for example, Dornbusch (1976), Frankel (1979), and Obstfeld and Rogoff (1995).

swings away from PPP in figure 1.⁴⁹

Because they rule out an autonomous role for forecasts, REH models, regardless of whether they assume that prices are sticky or flexible, must ignore the possibility that it is market participants' forecasts that push the exchange rate persistently away from PPP.

11.1 Imperfect Knowledge and Deviations from the REH Solution

Fully predetermined models, let alone their extreme REH versions, are inconsistent with the revisions and diversity of forecasting strategies that characterize real world markets. Moreover, economists themselves have constructed many different models. Thus, the aggregate of market participants' forecasting strategies differs from the strategy that is implied by any particular REH model, such as the one in (1). This divergence opens up the possibility that autonomous revisions in forecasting strategies play the key role in understanding long swings in asset markets.

To explore this possibility, consider again the semi-reduced form for the equilibrium price in equation (1), which we rewrite as⁵⁰

$$P_t = P_t^{\text{RE}} + c \left(\hat{P}_{t|t+1}^{\text{IK}} - \hat{P}_{t|t+1}^{\text{RE}} \right) \quad (32)$$

where, as before, $\hat{P}_{t|t+1}^{\text{RE}}$ is the REH forecast and $\hat{P}_{t|t+1}^{\text{IK}}$ denotes an IKE representation of the aggregate of individuals' forecasts.

Equation (32) shows immediately that P_t undergoes a protracted swing away from P_t^{RE} during periods of time in which $\hat{P}_{t|t+1}^{\text{IK}}$ moves persistently away from $\hat{P}_{t|t+1}^{\text{RE}}$. These periods end when the swing in $\hat{P}_{t|t+1}^{\text{IK}}$ ends.

⁴⁹Economists sometimes drop the assumption of stability and rely on the bubble paths of REH models to account for exchange rate swings. As we show in chapter 7 of our book, beyond sharing epistemological flaws with canonical REH models, these REH-bubble models are unable to account for the kind of long swings we observe in asset markets.

⁵⁰Plugging $\hat{P}_{t|t+1}^{\text{RE}}$ into equation (1) gives $P_t = a + bX_t + c\hat{P}_{t|t+1}^{\text{RE}} + c \left(\hat{P}_{t|t+1} - \hat{P}_{t|t+1}^{\text{RE}} \right)$. Equation (??) follows from the fact that $P_t^{\text{RE}} = a + bX_t + c\hat{P}_{t|t+1}^{\text{RE}}$.

11.2 Long Swings from Benchmark Levels in Asset Markets

This simple framework shows that allowing for imperfect knowledge has the potential to account for swings.⁵¹ This, of course, presupposes that (1) provides an adequate semi-reduced-form representation of the equilibrium price.

This semi-reduced form has been used extensively to model asset prices, particularly those for stocks and currencies.⁵² Under the former interpretation, P_t is the price of a stock, X_t is the dividend paid each period on that stock, and $b = c$ is a discount factor. In the case of currencies, P_t is the logarithm of the exchange rate, X_t consists of log levels of relative (domestic minus foreign) money supply and income, and b and c depend on the interest elasticity of money demand.⁵³

Under imperfect knowledge, the model in equation (1) becomes the one in (32). However, to explain swings around benchmark levels with the model in (32), the REH solution would have to correspond to an empirically-plausible notion of the benchmark in the particular market under study.⁵⁴

As it turns out, this seems to be the case in equity and currency markets. In stock markets, the REH solution of the model typically sets the equilibrium price, P_t^{RE} , equal to a constant multiple of the current dividend.⁵⁵ Market participants and policy makers have long recognized measures of the histori-

⁵¹Our IKE model of long swings builds on some of the ideas of Schulmeister (1983, 1987) and Soros (1987, 2008).

⁵²The extent to which the simple model in equation (1), together with an IKE representation of forecasting behavior, applies in markets beyond those for equity and currency is an open question that we do not explore in this paper.

⁵³See Blanchard and Fisher (1989) and references therein for interpretations of equation (1) in terms of stock prices, as well as goods prices. See Frydman and Goldberg (2007, chapter 6) and references therein for interpretations in terms of currency prices.

⁵⁴This way of thinking about asset price swings—that they arise from autonomous movements in the aggregate forecast—also underlies behavioral models of swings. See, for example, the seminal study by Frankel and Froot (1987). However, these models fully prespecify revisions in market participants' forecasts; thus, as with other behavioral models of market outcomes, they presume irrationality on the part of individuals. An open question is whether the Frankel and Froot (1986) model and its extensions (see De Grauwe and Grimaldi, 2006, and references therein) can be re-interpreted in an IKE framework.

⁵⁵This REH solution assumes that the dividend process follows $D_t = D_{t-1}(1 + \mu^D)$, where μ^D is the constant growth rate of dividends. With this specification, and the assumption of a fixed risk-free rate, r , we have $P_t^{\text{RE}} = D_t / (r - \mu^d)$. The REH solution can be expressed more generally in terms of the mean growth rate of dividends and the

cal averages of this ratio and its close correlate, the PE ratio, as benchmarks around which stock prices revolve. As for currency markets, the REH solution of the model sets P_t^{RE} equal to the PPP exchange rate, P_t^{PPP} . Here, too, measures of the PPP exchange rate have a long history dating back to the 1600s as a useful benchmark in gauging the extent of currency fluctuations.⁵⁶ Thus, in the context of currency markets, we can write (32) as⁵⁷

$$P_t = P_t^{\text{PPP}} + c \left(\hat{P}_{t|t+1}^{\text{IK}} - \hat{P}_{t|t+1}^{\text{RE}} \right) \quad (33)$$

In the next section, we use (33) to illustrate that once REH is replaced by an IKE representation of forecasting behavior, the model is able to account for the tendency of the exchange rate to undergo long swings.

12 An IKE View of Long Swings

The modern research program in economics models aggregate outcomes by relating them to individual decision-making. The constraints that an IKE model imposes on its representations of individual behavior are not only qualitative, but context-specific. They depend on which aspects of the causal mechanism driving aggregate outcomes an economist seeks to explain.

For example, if the model aims to account for movements in the exchange rate, its qualitative predictions should be consistent with these prices' tendency to undergo swings of uneven duration and magnitude around PPP. This *uneven regularity* at the aggregate level suggests that an economist should look for and formalize qualitative regularities on individual behavior that are similarly uneven.

The microfoundations of our model of currency swings represent an individuals' forecast of the future exchange rate as

$$\hat{P}_{t|t+1}^{i,\text{IK}} = \beta_t^{i,\text{IK}} Z_t^{i,\text{IK}} \quad (34)$$

mean risk-free rate when μ^D and r are allowed to vary over time.

⁵⁶International macroeconomists trace the notion of PPP back to scholars at the University of Salamanca in the fifteenth and sixteenth centuries. See Officer (1976). For formal evidence that exchange rate swings do revolve around PPP levels, see Taylor and Taylor (2004) and references therein.

⁵⁷Although the benchmark levels used in Figures 1 and 2 seem to be relevant on historical grounds in both equity and currency markets, establishing their theoretical microfoundations in a world of imperfect knowledge is beyond the scope of this paper.

This representation shows that there are two key factors that underpin the evolution of an individual’s forecast over time: revisions of her forecasting strategy—changes in $\beta_t^{i,\text{IK}}$ and the composition of the set of causal factors $Z_t^{i,\text{IK}}$ —and movements in the causal variables. To model currency fluctuations, therefore, we need to look for uneven regularities in these two components.

In doing so, we focus much of our attention on revisions of forecasting strategies. To illustrate our model of currency swings, we begin with the assumption that the causal variables in the model follow random walks with constant drifts.⁵⁸

Economic policy, of course, does change from time to time as new policy makers take charge, economic and social conditions change, or as policy makers’ understanding of those conditions evolves. As with market participants’ forecasting, policy shifts may also display uneven qualitative regularities that might help in accounting for certain features of currency fluctuations. However, to highlight how IKE representations of forecasting strategies enable us to explain the uneven nature of swings, we maintain the assumption of a fixed policy environment for most of this section.⁵⁹

12.1 Conservatism as a Qualitative Regularity of Uneven Duration

To model revisions of an individual’s forecasting strategy, we explore the implications of a well-documented phenomenon that psychologists call “conservatism:” individuals tend to revise their beliefs about uncertain outcomes in ways that lead to gradual changes in those beliefs.⁶⁰ In the context of our model, an individual’s beliefs and their formation are represented by her

⁵⁸There is much evidence that the usual macroeconomic fundamentals, such as money supply and income levels, are well approximated as unit root processes with drift. See Juselius (2007) and references therein. Although most empirical researchers model many macroeconomic times series as I(1) processes, Juselius *et al* (2007) and Frydman *et al.* (2008) show that they are also well approximated as near-I(2) processes. Such behavior is also found in Johansen (1997), Kongsted et al. (1999), Juselius (2006), and Nielsen and Rahbek (2007).

⁵⁹In section 12.7.1, we show how the relaxation of this assumption can help explain the boundedness of fluctuations in currency markets. A more general IKE model that incorporates qualitative features of policymaking is beyond the scope of this paper. For a first step in this direction, see Frydman and Goldberg (2004).

⁶⁰See Edwards (1968) and Shleifer (2000).

forecast, $\hat{P}_{t|t+1}^{i,\text{IK}}$, and her forecasting strategy, $\beta_t^{i,\text{IK}} Z_t^{i,\text{IK}}$, respectively.

The decision to revise one's forecasting strategy depends on many factors, including prior forecasting success, economic and political developments, emotions, or, as we will suggest shortly, the size of the departure of the exchange rate from PPP. While market participants may tend to behave conservatively, conservatism is a regularity that is at best qualitative and uneven.^{61,62} Eventually, the unfolding historical record on market outcomes, changes in the social context, including policy, or the sheer creativity in thinking about the future, may lead a market participant to revise her forecasting strategy in a more substantial, non-conservative way.

On one hand, individuals tend to be conservative in revising their forecasting strategies. On the other hand, they eventually cease to be so at times that do not conform to any fully prespecified rule. Conservatism thus offers the possibility of accounting for the kind of price movements we observe in asset markets, in which up-swings are followed by down-swings in ways and at times that do not follow any fully prespecified rule.

How one would formalize conservatism depends on the context. In our model of currency swings under flexible goods prices, the qualitative conditions that we use to represent this regularity are easily illustrated.⁶³ Given the representation in (34) and the assumption that the causal factors follow random walks with deterministic drifts, the total change in an individual's forecast is

$$\hat{P}_{t|t+1}^{i,\text{IK}} - \hat{P}_{t-1|t}^{i,\text{IK}} = \mathcal{D}\hat{P}_{t|t+1}^{i,\text{IK}} + \epsilon_t^i \quad (35)$$

where

⁶¹This seems to be the case with other empirical observations that behavioral economists have uncovered. For example, behavioral-finance economists report much evidence that participants in financial markets often rely on technical trading rules in deciding when to take open positions. However, there is also much evidence that the importance they place of such strategies varies over time. See Schulmeister (2006, 2008) and references therein. By contrast, we would expect the regularities that characterize individuals' preferences—for example, the importance of fairness or loss aversion in individual decision making—to be more enduring.

⁶²Barberis, Shleifer, and Vishny (1998) also appeal to conservatism in modeling an individual's forecasting behavior. But, to generate sharp predictions, they formulate conservative behavior with a fixed rule that presumes that individuals under-react to earnings announcements in exactly the same way at every point in time.

⁶³See Frydman and Goldberg (2007, chapter 14) for a formulation in the context of exchange rate models that assume sticky goods prices.

$$\mathcal{D}\hat{P}_{t|t+1}^{i,\text{IK}} = \Delta\beta_t^{i,\text{IK}} Z_{t-1}^{i,\text{IK}} + \beta_{t-1}^{i,\text{IK}} \mu^{Z^i} \quad (36)$$

$\mathcal{D}\hat{P}_{t|t+1}^{i,\text{IK}}$ is the trend in the individual's forecast between $t-1$ and t , ϵ_t^i is a vector of mean-zero, i.i.d. error terms that stem from the Z_t^i process, and the “ $\mathcal{D}$ ” notation underscores the fact that with IKE, a change in $\hat{P}_{t|t+1}^{i,\text{IK}}$ may result from a change in the structure of this representation and not only from a mere updating due to a change in causal variables.⁶⁴ The expression in (36) shows that the trend in an individual's beliefs varies over time, depending on how she revises her forecasting strategy, $\Delta\beta_t^{i,\text{IK}}$, and on what we call the underlying drift in her forecast, that is, the change in $\hat{P}_{t|t+1}^{i,\text{IK}}$ arising solely from the deterministic drifts in the causal variables, $\beta_{t-1}^{i,\text{IK}} \mu^{Z^i}$.⁶⁵

Although the trend in an individual's forecast varies over time, its algebraic sign may remain unchanged for an extended period. Such behavior would depend on how the individual revises her forecasting strategy. To represent these revisions, we make use of two conservative restrictions. One constrains revisions of $\beta_t^{i,\text{IK}}$ so that the resulting change of the level of $\hat{P}_{t|t+1}^{i,\text{IK}}$ is gradual, that is:

$$|\Delta\beta_t^{i,\text{IK}} Z_{t-1}^{i,\text{IK}}| < \delta \quad (37)$$

where $|\cdot|$ denotes an absolute value and δ is a “small” magnitude that we will say more about shortly. The other conservative condition constrains revisions of $\beta_t^{i,\text{IK}}$ so that the resulting change of the underlying drift in $\hat{P}_{t|t+1}^{i,\text{IK}}$ is gradual, that is,

$$|\Delta\beta_t^{i,\text{IK}} \mu^{Z^i}| < \delta \quad (38)$$

The qualitative constraints in (37) and (38) formalize the idea that when an individual revises her forecasting strategy, she is reluctant to do so in ways that would result in dramatic changes in both the level of her forecast and

⁶⁴The $\mathcal{D}$ operator serves a role that is analogous to the total differential, which in general is not well defined when both the arguments and the structure of the representation change.

⁶⁵At every point in time, the vector $Z_t^{i,\text{IK}}$ represents all possible casual factors that an individual might use in forming her forecast. As such, the representation of change in (35) allows for the composition of the causal variables to change. For example, if an individual is presumed to use a particular variable Z_j to form her forecast at t but not at $t-1$, then $\beta_{j,t-1}^{i,\text{IK}} = 0$ and $\Delta\beta_{j,t}^{i,\text{IK}} = \beta_{j,t}^{i,\text{IK}}$.

its underlying drift.^{66,67} They restrict neither the causal variables that may enter the representation in (34), nor how exactly those variables may matter.

Because we formalize conservatism as a qualitative condition, our representation of forecasting behavior is consistent with myriad possible changes in $\beta_t^{i,\text{IK}}$ and the composition of $Z_{t-1}^{i,\text{IK}}$. Consequently, our IKE representation implies a myriad of conditional probability distributions for the equilibrium price at t —one for each possible value of $\beta_t^{i,\text{IK}}$ and set of causal variables—conditional on any one of the distributions at time $t - 1$.⁶⁸ In this way, the representation in (??) recognizes the role of both forecast revisions that do not follow pre-existing rules, and new information for explaining outcomes.

The finding that an individual often behaves conservatively suggests that there may be extended time periods in which the revisions of her forecasting strategy are consistent with the qualitative conditions (37) and (38). Using such conditions on the individual level enables us to account for the aggregate regularity that the exchange rate sometimes undergoes persistent swings.

It does not seem plausible, however, that a market participant would revise her forecasting strategy in conservative ways endlessly.⁶⁹ However, because the time at which conservative conditions cease to adequately represent revisions cannot be foreseen in advance, our model of the long swings can also account for their unevenness.

⁶⁶These restrictions constrain the change that results only from revisions of $\beta_t^{i,\text{IK}}$. As such, they do not constrain the total change in $\hat{P}_{t|t+1}^{i,\text{IK}}$ to be gradual. If changes in the causal variables between two points in time are large, then the change in $\hat{P}_{t|t+1}^{i,\text{IK}}$ may be large.

⁶⁷In general, the constraints in (37) and (38) could be specified with distinct bounds. Doing so, however, would only complicate our analysis without altering its results. We thus abstract from this complication.

⁶⁸This contrasts sharply with the fully predetermined representation of forecasting behavior in section ??, which implies just one probability distribution at t .

⁶⁹This claim can also be defended on theoretical grounds. As we discuss in section 12.7.1, as the exchange rate moves further and further away from the benchmark, either bulls or bears are more and more likely to cease to be conservative in the way they forecast the exchange rate. Without the limited duration of conservatism the model would imply unbounded swings.

12.2 Uneven Swings in an Individual's Forecast

We now sketch the implications of conservative behavior for how an individual's forecast unfolds over time.⁷⁰ We suppose that the revisions of an individual's forecasting strategy are consistent with the qualitative constraints in (37) and (38) for an extended period of time, say from $t = \tau$ to $t = \tau + T$. To fix ideas, we suppose that entering period τ , the individual's strategy and the drifts in the causal variables are such that the underlying drift in her forecast is positive, that is, $\beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i} > 0$.

At τ , the individual may decide to revise her forecasting strategy. This revision may result in a change of her forecast that reinforces the upward movement implied by the underlying drift, that is, $\Delta\beta_{\tau}^{i,\text{IK}} Z_{\tau-1}^{i,\text{IK}} > 0$. In this case, $\hat{P}_{t|t+1}^{i,\text{IK}}$ will tend to rise between $\tau - 1$ and τ . The individual may also revise her strategy at τ in a non-reinforcing manner, that is, $\Delta\beta_{\tau}^{i,\text{IK}} Z_{\tau-1}^{i,\text{IK}} < 0$, in which case $\hat{P}_{t|t+1}^{i,\text{IK}}$ may not rise. This would be the case if the negative magnitude of the revision exceeded the positive underlying drift, $\beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i}$. Thus, to ensure that $\hat{P}_{t|t+1}^{i,\text{IK}}$ will tend to rise between $\tau - 1$ and τ , the size of the underlying drift cannot be "too small," which we formalize by constraining $0 < \delta < \beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i}$. With this assumption, (37) implies that $\hat{P}_{t|t+1}^{i,\text{IK}}$ will tend to rise between $\tau - 1$ and τ .

Will the upward trend in the individual's forecast persist between τ and $\tau + 1$? Again, the relative impacts of the revision of her forecasting strategy and its underlying drift are important.

Consider first whether the underlying drift, $\beta_{\tau}^{i,\text{IK}} \mu^Z$, continues to be positive. The revision of $\beta_t^{i,\text{IK}}$ at τ influences this term, which we write as

$$\beta_{\tau}^{i,\text{IK}} \mu^Z = \beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i} + \Delta\beta_{\tau}^{i,\text{IK}} \mu^{Z^i}$$

This expression makes clear that the change, $\Delta\beta_{\tau}^{i,\text{IK}} \mu^{Z^i}$, may either reinforce or counteract the positive drift that prevailed between $\tau - 1$ and τ . But, even if this change is not reinforcing, as long the revision of $\beta_t^{i,\text{IK}}$ at τ satisfies the second conservative condition in (38) and $\delta < \beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i}$, the underlying drift will remain positive between τ and $\tau + 1$.

However, whether $\hat{P}_{t|t+1}^{i,\text{IK}}$ tends to rise between τ and $\tau + 1$ also depends on how the individual revises her strategy at $\tau + 1$. If the revision is again consistent with the conservative constraint in (37) and $\delta < \beta_{\tau}^{i,\text{IK}} \mu^{Z^i}$, the tendency for $\hat{P}_{t|t+1}^{i,\text{IK}}$ to move up will continue.

⁷⁰For a more detailed analysis, see Frydman and Goldberg (2007, chapter 14).

In similar fashion, if our individual continues to revise her forecasting strategy in ways that are conservative at each point in time until $\tau + T$ and $\delta < \beta_{\tau-1}^{i,\text{IK}} \mu^{Z^i}$, the underlying drift in her forecast will remain positive throughout. Moreover, while her revisions between adjacent points in time may have positive or negative impacts on the trend, as long as they are conservative, $\hat{P}_{t|t+1}^{i,\text{IK}}$ will undergo an upswing from τ to $\tau + T$.

Although the value of δ is assumed to be small, it is of course possible that $\beta_{t-1}^{i,\text{IK}}$ and the drifts in Z_t are such that, in a particular period, the underlying drift is even smaller, that is, $\delta > \beta_{t-1}^{i,\text{IK}} \mu^{Z^i}$. At such points in time, the drifts in the causal variables may not be sufficient for a rise in $\hat{P}_{t|t+1}^{i,\text{IK}}$ even if the individual does behave conservatively. It can be shown that protracted swings in $\hat{P}_{t|t+1}^{i,\text{IK}}$ will nonetheless arise in the model if the underlying drift, $\beta_{t-1}^{i,\text{IK}} \mu^{Z^i}$, exceeds δ , most of the time in periods characterized by conservative forecasting behavior.⁷¹ In the remainder of this paper, we ignore this complication and assume that $\delta < \beta_{t-1}^{i,\text{IK}} \mu^{Z^i}$ during all periods in which the constraints in (37) and (38) are satisfied.

Our analysis of an upward swing in $\hat{P}_{t|t+1}^{i,\text{IK}}$ makes clear that for this movement to end, the individual would have to revise her forecasting strategy in a non-conservative and non-reinforcing way at some point in time. Such a revision can be associated with a short-lived reversal in $\hat{P}_{t|t+1}^{i,\text{IK}}$ or one that lasts for many periods. Our foregoing arguments imply that a sustained reversal in $\hat{P}_{t|t+1}^{i,\text{IK}}$ will occur beyond $\tau + T$ if 1) the revision of $\beta_t^{i,\text{IK}}$ at $\tau + T + 1$ leads to a negative underlying drift; and 2) revisions are once again conservative for an extended period of time beyond $\tau + T + 1$.

12.3 Bulls, Bears, and Long Swings

The qualitative prediction that an individual's forecast undergoes a swing is conditional on the way she revises her forecasting strategy and how the causal variables move over time. This prediction on the individual level implies that the predictions on the aggregate level are also conditional on how market

⁷¹This conclusion depends on which of the two conservative restrictions might be violated from time to time. For example, if the second conservative restriction in (38) holds throughout the period, but the one in (37) is at times violated, then the swing $\hat{P}_{t|t+1}^{i,\text{IK}}$ would still occur. In this case, however, the frequency of short countermovements in $\hat{P}_{t|t+1}^{i,\text{IK}}$ might be greater.

participants' revise their forecasting strategies and how the causal variables develop over time.

If we were to follow the vast majority of our colleagues and assume that all market participants share the same forecast, then movements of an individual's forecast would be tantamount to movements of the aggregate forecast. However, imperfect knowledge in real-world markets gives rise to a diversity of forecasting strategies across market participants. In asset markets, this diversity takes on a striking form: at each point in time, some participants forecast a rise in the asset price while others forecast a fall.

Recognizing this kind of diversity is crucial for explaining outcomes in currency markets. However, as we discuss in section 14, to do so without the presumption of irrationality, an economist must not only stop short of fully prespecifying change, but the qualitative constraints that he imposes on the revisions of forecasting strategies cannot apply endlessly.

Market bulls hold long positions in foreign exchange because they expect the exchange rate to rise over the holding period, while bears hold short positions because they expect the exchange rate to fall. To represent the forecasting behavior of bulls and bears, we aggregate over these groups using wealth shares;

$$\hat{R}_{t|t+1}^{L,IK} = \hat{P}_{t|t+1}^{L,IK} - P_t > 0 \quad (39)$$

$$\hat{R}_{t|t+1}^{S,IK} = P_t - \hat{P}_{t|t+1}^{S,IK} > 0 \quad (40)$$

where $\hat{P}_{t|t+1}^{L,IK} = \beta_t^{L,IK} Z_t^{L,IK}$ and $\hat{P}_{t|t+1}^{S,IK} = \beta_t^{S,IK} Z_t^{S,IK}$ represent aggregates of the exchange rate forecasts of the bulls and bears, respectively, and $\hat{R}_{t|t+1}^{L,IK}$ and $\hat{R}_{t|t+1}^{S,IK}$ represent aggregates of bulls' and bears' expected returns on holding long and short positions, respectively.⁷² In the aggregate, we have,

$$\hat{P}_{t|t+1}^{IK} = \frac{1}{2} \left(\hat{P}_{t|t+1}^{L,IK} + \hat{P}_{t|t+1}^{S,IK} \right) = \beta_t^{IK} Z_t^{IK} \quad (41)$$

where Z_t^{IK} represents the union of causal variables that participants use in forming their forecasts and β_t^{IK} denotes weighted averages of the parameters that they attach to these variables.⁷³

⁷²Positive realizations of $P_{t+1} - P_t$ imply profits on long positions and losses on short positions. We thus define the return on a short position as $R_{t+1}^S = P_t - P_{t+1}$. The definitions of the returns on open positions abstracts from the cost of capital and thus from interest rates for convenience only.

⁷³In our analysis, we assume that the wealth shares of the groups of bulls and bears are exogenous and constant, so we set them equal to a half.

The persistent movement in one direction or the other over an extended period of time depends *not* on the constellation of participants' beliefs about whether the exchange will rise or fall, but on how those beliefs *change* over time. Exchange rate swings arise in the model because the aggregate of individuals' forecasts, $\hat{P}_{t|t+1}^{\text{IK}}$, undergoes a swing.

For example, if $\hat{P}_{t|t+1}^{\text{IK}}$ were to rise, on average, over an extended period of time, then the exchange rate would also tend to undergo an upswing even though there are bears in the market who bet on the opposite movement. Such change could come about because both bulls and bears steadily raise their exchange rate forecasts or because the forecasts of only one side of the market increase.⁷⁴ Whatever the case, however, the exchange would tend to rise as long as $\hat{P}_{t|t+1}^{\text{IK}}$ were to rise on average. If this movement was toward PPP initially, and the period of a rising $\hat{P}_{t|t+1}^{\text{IK}}$ were to endure, eventually the exchange rate would shoot through this benchmark and continue trending from the other side. It is also clear that this swing away from PPP would end if the swing in $\hat{P}_{t|t+1}^{\text{IK}}$ were to end.

It is not difficult to see from our analysis on the individual level how such a swing in $\hat{P}_{t|t+1}^{\text{IK}}$ arises in our model: the aggregate $\hat{P}_{t|t+1}^{\text{IK}}$ is just a weighted sum of the individual forecasts. If we assume that each market participant revises her forecasting strategy in conservative ways over an extended period of time, then each component of the aggregate will involve a trend that, although time-varying, does not change direction. Consequently, the trend in the aggregate forecast will also involve an unchanging sign and thus imply a tendency for $\hat{P}_{t|t+1}^{\text{IK}}$ to move in one direction over time.

The assumption that all market participants behave conservatively is strong. At each point in time we would expect that some individuals would behave conservatively while others would not. Whether the aggregate forecast tends to move in one direction over a period of time would then depend on the relative weight of the individuals in the market who behave conservatively. A swing would arise in the model during those periods in which this weight remained sufficiently high.

⁷⁴The rise in the forecasts of some individuals might even cause them to switch from being a bear to being a bull.

12.4 Long Swings and Fundamentals

Many international macroeconomists view the long-swings behavior of exchange rates as evidence that these prices do not depend on macroeconomic fundamentals, such as income levels and interest rates. Their empirical studies, which search for a fully prespecified (mostly fixed) relationship between the exchange rate and macroeconomic fundamentals, appear to confirm this view.⁷⁵

Our model of currency swings provides an alternative view. It is clear from our analysis that swings arise in the model even if the forecasting strategies attributed to market participants are assumed to depend solely on macroeconomic fundamentals.⁷⁶ Moreover, to account for currency swings that are uneven in duration and magnitude, we need to allow for revisions of market participants' forecasting strategies. However, equations (33) and (35) show that such behavior is associated with structural change in the relationship between the exchange rate and fundamentals. Thus, even if fundamentals matter for currency movements, we should not expect, as the majority of extant studies do, that a time-invariant model involving a fixed set of fundamentals, would adequately represent the causal relationship for periods of time as long as decades. In fact, formal empirical analysis reveals not only that this relationship has been unstable during the modern era of floating currencies, but that the set of fundamentals that matters has changed from one sub-period to another.⁷⁷

Our theoretical and empirical findings suggest that long swings in exchange rates away from parity occur not because market participants ignore fundamentals in forming their forecasts, but because knowledge is imperfect and fundamentals tend to move in particular directions.

12.5 Conditional Predictions of Change

As with the model's predictions on the individual level, its predictions about the equilibrium exchange rate are conditional on how individuals revise their forecasting strategies and on the movements of the causal variables. The

⁷⁵For a review of this literature, see Frankel and Rose (1995). For a discussion, see section 15.1.

⁷⁶This is true even if Z_t includes only those fundamentals that drive the REH forecast, that is, if we set $Z_t = X_t$.

⁷⁷See Goldberg and Frydman (1996) and Frydman and Goldberg (2007, chapter 15).

model implies an exchange rate swing during any period of time in which 1) revisions of strategies are conservative; and 2) drifts in the causal variables remain unchanged. The assumption that conservative forecasting behavior does not last forever is crucial for the model's empirical relevance. Indeed, as we have shown, in order to account for the exchange rate swings in figure 1, which are of limited and uneven duration, we need to assume that revisions of forecasting strategies are, at unpredictable points in time, non-conservative and non-reinforcing. A change in policy may also trigger a sustained reversal in the exchange rate (see section 12.6). However, in any case, we must assume that the constraints on change do not apply endlessly if we want to explain the uneven nature of exchange rate swings.⁷⁸

We cannot prespecify exactly when forecasting behavior will be conservative or when the drifts in the causal variables will be constant. Hence, we cannot prespecify exactly when exchange rate swings will begin or end.

12.6 Bounded Instability and the Role of Benchmarks

Beyond its conditional prediction that the exchange rate will undergo swings of uneven duration and magnitude, our IKE model implies that swings away from PPP are ultimately bounded. As the gap from PPP grows, either market participants eventually cease to be conservative, owing to the way they assess the risk of capital losses on holding open positions, or policy officials alter the course of the causal variables.

Our representation of risk replaces the usual assumptions of risk aversion and expected utility theory with what we call *endogenous prospect theory*. This alternative representation of preferences under uncertainty implies that individuals' degree of loss aversion increases as the size of the potential loss from speculation increases.⁷⁹ As such, market participants (both bulls and bears) expect to earn a positive return – a premium – for holding open positions in the market, the size of which depends on one's forecast of the potential loss from speculation. To model an individual's speculative decision,

⁷⁸In fact, with no change in structure of any kind, or with individuals who behave conservatively endlessly, the model implies, counterfactually, that the exchange will undergo an unbounded swing away from PPP. See chapter 14 of our book.

⁷⁹An individual is loss averse if her disutility from losses is greater than her utility from gains of the same magnitude. Endogenous prospect theory provides a way to represent the experimental findings of Kahneman and Tversky (1979) and others in a world of imperfect knowledge.

therefore, and thus the movement of the exchange rate, we must represent how revisions of her forecasting strategy alter both her forecast of the next period's exchange rate, $\hat{P}_{t|t+1}^{i, \text{IK}}$, and her forecast of the potential loss from holding an open position, which we denote by $\hat{l}_{t|t+1}^{i, \text{IK}}$.

To this end, we make use of an observation about the historical record on movements of the aggregate premium—expected return—on foreign exchange, which we denote by $\hat{pr}_{t|t+1}^{\text{IK}}$, and the gap between the aggregate forecast and the PPP exchange rate, which we denote by $\widehat{gap}_t^{\text{IK}} = \hat{P}_{t|t+1}^{\text{IK}} - \hat{P}_t^{\text{PPP}}$. In figure 2, we plot $\hat{pr}_{t|t+1}^{\text{IK}}$ and $\widehat{gap}_t^{\text{IK}}$ for the British pound-U.S. dollar (BP/\$) market.⁸⁰ The figure clearly suggests that $\hat{pr}_{t|t+1}^{\text{IK}}$ tends to move positively over time with $\widehat{gap}_t^{\text{IK}}$.⁸¹ The results of regression analysis also indicate a positive relationship between $\hat{pr}_{t|t+1}^{\text{IK}}$ and $\widehat{gap}_t^{\text{IK}}$, although they show that the quantitative relationship between these aggregates varies over the sample period.⁸²

12.7 Benchmark Levels and Revisions of Forecasting Strategies

Our representations of forecasting behavior imply that when an individual revises her strategy, she alters her forecasts of the future exchange rate and potential loss. She may be conservative in how she changes both of these forecasts. However, because conservatism does not involve any restrictions on the causal mechanism—the composition of the causal factors and how they

⁸⁰We use survey data on exchange rate expectations to measure $\hat{pr}_{t|t+1}^{\text{IK}}$ and the *Big Mac* PPP exchange rate as in figure 1. The survey data are from Money Market Services International (MMSI), which entail median responses from market participants concerning their four-week ahead point forecasts of the exchange rate. For more details concerning the time plots in figure 2, see Frydman and Goldberg (2007, chapter 12). Other studies that have used survey data from MMSI include Frankel and Froot (1987) and Froot and Frankel (1989).

⁸¹Time plots for the DM/\$ and Japanese yen/\$ markets show a similar pattern. The observation that $\hat{pr}_{t|t+1}^{\text{IK}}$ and $\widehat{gap}_t^{\text{IK}}$ tend to move positively over time came from our own research on currency markets. We first made use of this relationship in Frydman and Goldberg (2003) to explain the forward-discount puzzle.

⁸²See Frydman and Goldberg (2007), chapter 12. Our empirical analysis relies on recursive methods that allow us to test for structural change without prespecifying its timing or nature. Moreover, we look for—and find—some evidence that the market premium also depends on cumulative current account imbalances, where, again, this relationship is only qualitative.

influence the outcomes—it is not sufficient to account for the relationship between the premium and the gap. Thus, in formulating the microfoundations of a model that might explain the aggregate regularity between these variables, we must look for additional regularities on forecasting behavior on the individual level.

To this end, we build on Keynes (1936) and relate revisions of $\widehat{l}_{t|t+1}^{i,IK}$ to the gap between an individual’s forecast of the exchange rate and her measure of the benchmark exchange rate. Keynes argued that, while asset prices have a tendency to move persistently away from benchmark levels for protracted periods, they eventually undergo, at unpredictable moments, sustained countermovements back to these benchmark levels. He recognized that market participants are also aware of this feature of the social context and use it in their attempts to forecast future market outcomes. In discussing why an individual might hold cash rather than interest-bearing bonds, Keynes observed that “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).

A benchmark level is, of course, specific to each asset market. Every individual arrives at her own determination of the benchmark value and so, in general, these assessments will differ across individuals. How individuals come to decide on a benchmark level is an open question. Keynes suggests in his discussion that conventions and the historical record play an important role. For example, as we have already mentioned, PPP has long played a role in how market participants and policy makers assess the misalignment of exchange rates, and its role as a benchmark is reasonable in view of the empirical record in many currency markets. This is also the case with historical PD and PE ratios in equity markets.

Keynes’s discussion of the importance of benchmark levels as anchors for exchange rate fluctuations suggests that a market participant’s forecast of the potential loss from speculation depends on her evaluation of the gap between the asset price and its benchmark level. For example, we would expect that as the asset price rises further above the perceived benchmark level, bulls (bears) would become more concerned (more confident) about a reversal and thus greater (smaller) capital losses.⁸³

⁸³As Keynes puts it, “Unless reasons are believed to exist why future experience will be very different from past experience, a ...rate of interest [much lower than the safe rate],

As with conservatism, we would not expect that these regularities for bulls and bears would conform to any fully prespecified rule. We thus rely on qualitative conditions to represent them, which we call gap restrictions.⁸⁴ These representations of forecasting behavior are easily illustrated.

An individual's expected loss from speculation can be written as,

$$\widehat{l}_{t|t+1}^{i,L} = E_t^i[R_{t+1}^L < 0 | Z_t^i] < 0$$

for a bull and,

$$\widehat{l}_{t|t+1}^{i,S} = E_t^i[R_{t+1}^S < 0 | Z_t^i] < 0 \quad (42)$$

for a bear, where $R_{t+1}^L = P_{t+1} - P_t$ and $R_{t+1}^S = P_t - P_{t+1}$ are the returns on long and short positions on foreign exchange, respectively, and $E_t^i[R_{t+1}^L < 0 | Z_t^i]$ and $E_t^i[R_{t+1}^S > 0 | Z_t^i]$ denote the expected value of the realizations on R_{t+1}^L and R_{t+1}^S that imply a loss for a bull and bear, respectively, conditional on an individual's forecasting strategy and information set. We omit the superscript "IK" for notational ease.

The gap conditions for an individual bull and bear can be written as follows

$$\frac{\mathcal{D}\widehat{l}_{t|t+1}^{i,L}}{\widehat{\mathcal{D}gap}_t^{i,L}} < 0 \quad \text{and} \quad \frac{\mathcal{D}\widehat{l}_{t|t+1}^{i,S}}{\widehat{\mathcal{D}gap}_t^{i,S}} > 0 \quad (43)$$

where the $\mathcal{D}$ operator, as before, denotes a total change, $\widehat{\mathcal{D}gap}_t^{i,IK} = \widehat{P}_{t|t+1}^{i,IK} - \widehat{P}_t^{i,PPP}$, and $\widehat{P}_t^{i,PPP}$ denotes an individual's assessment at t of the PPP exchange rate.⁸⁵ These qualitative conditions constrain our representation of how an individual revises her forecasting strategy by restricting all post-change conditional probability distributions to imply a gap effect: distributions that imply a higher $\widehat{\mathcal{D}gap}_t^{i,IK}$ also imply an expectation of greater (smaller) potential losses if the individual is a bull (bear).⁸⁶

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 [of capital loss]" (Keynes, 1936, p.202).

⁸⁴We also make use of additional qualitative conditions to model revisions in $\widehat{l}_{t|t+1}^{i,IK}$ that appeal to theoretical considerations. Economic theory suggests that market participants might also consider current account imbalances in forecasting potential losses. See Frydman and Goldberg (2007), chapter 12.

⁸⁵The gap could also be defined in terms of the time- t exchange rate, rather than the forecast of the future exchange rate, or some weighted average of the two without affecting the conclusions of our analysis. See Frydman and Goldberg (2003).

⁸⁶Because $\widehat{l}_{t|t+1}^{i,L}$ and $\widehat{l}_{t|t+1}^{i,S}$ are both negative, greater (smaller) losses implies a fall (rise) in these magnitudes. Hence, the less-than and greater-than inequalities in formulating the gap conditions for a bull and bear, respectively.

12.7.1 Self-Limiting Long Swings

To see how our IKE model with the gap conditions in (43) implies that exchange rate swings away from PPP will eventually end, consider an extended period of time in which market participants revise their forecasting strategies in conservative ways. Suppose that initially, forecasting strategies and the drifts in the causal variables imply a positive underlying drift in $\hat{P}_{t|t+1}^{\text{IK}}$, that is, $\beta_{\tau-1}^{\text{IK}}\mu^Z > 0$. During this period of time, then, $\hat{P}_{t|t+1}^{\text{IK}}$ and the exchange rate, P_t , will tend to rise. To keep our example simple, we focus only on the role of bulls and assume that the rise in $\hat{P}_{t|t+1}^{\text{IK}}$ is associated with a rise in $\hat{P}_{t|t+1}^{\text{L,IK}}$.⁸⁷

While conservative revisions of strategies and movements in causal factors lead bulls to raise their forecasts and bid the exchange rate further above PPP, our IKE model of risk indicates that they simultaneously become more concerned about a sustained counter-movement — that is, about capital losses. This leads them to raise the premiums that they require to increase their long positions. According to our model, if the swing away from PPP were to continue, a threshold would eventually be reached at which bulls would become so concerned about a reversal that they would no longer revise their forecasting strategies in conservative ways. At that point, they would either reduce their long positions or abandon them altogether, which would precipitate a reversal in the exchange rate.

PPP matters because market participants have come to rely on this benchmark in their attempts to forecast potential losses and assess the riskiness of holding open positions. When, exactly, the gap from PPP, and thus the potential loss, is perceived by bulls to be too large for them to continue bidding up the exchange rate depends on many factors, including economic, political, and policy considerations that no one can fully prespecify. Thus, no one can fully prespecify when long swings away from PPP will eventually end.

Policy makers also use PPP as a benchmark level in setting economic policy and their actions also play an important role in keeping exchange rate swings bounded. We have so far assumed in our discussion of long swings that trends in macroeconomic fundamentals remain fixed. However, the empirical

⁸⁷The focus on bulls enables us to avoid distributional issues. But, of course, the bears also play a role in keeping the swing away from PPP bounded. See Frydman and Goldberg (2007).

record shows that policy officials sometimes worry about large departures in the exchange rate from PPP,⁸⁸ and alter policy to engender a reversal.⁸⁹ As with revisions of forecasting strategies, we would not expect that such a regularity would follow any fully prespecified rule.

This reasoning leads us to assume that beyond some threshold, which we do not prespecify, policy makers respond to large departures from PPP by altering policy to limit the misalignment. The fact that many of the major reversals in currency markets are proximate to major changes in policy suggests that the behavior of policy makers plays an important role in keeping long swings in the exchange rate bounded.^{90,91}

13 The Market Premium and the Aggregate Gap

Beyond its implications for keeping exchange rate swings away from PPP bounded, our IKE representation of individuals' expected losses from holding open positions implies a model of the causal mechanism for the equilibrium premium on foreign exchange, $\widehat{pr}_{t+1}^{\text{IK}}$, that is able to account for the qualitative regularity evident in figure 2.

Endogenous prospect theory implies that the equilibrium premium on

⁸⁸Large misalignments pose challenges for firms that are engaged in international business by leading to changes in competitiveness. These effects, in turn, lead to calls for protectionist measures, which may reduce the benefits from international trade for economic activity.

⁸⁹Examples of such behavior include the coordinated interventions by central banks and the changes in monetary and fiscal policy that were aimed at bringing down U.S. dollar rates in 1985 and yen rates in 1995, as well as the interventions by the U.S. Federal Reserve and the European Central Bank to stem the recent fall of the dollar.

⁹⁰For example, the major reversals in U.S. dollar rates in late 1979 and early 1985 were associated with the arrival of Paul Volcker and James Baker, respectively, both of whom quickly engineered major changes in policy. An example of a connection between policy and major reversals in other asset markets is provided by the downturn in U.S. equity markets that began in August 2000, which came on the heels of the Federal Reserve's decision to raise the federal funds rate from 4.74 percent in July 1999 to 6.5 percent in May 2000.

⁹¹Beyond the policy channel, departures in the exchange rate from PPP influence trends in macroeconomic fundamentals endogenously in ways that also keep exchange rate swings bounded. For example, swings in exchange rates eventually lead to changes in current account imbalances and economic growth that would tend to limit such swings.

foreign exchange depends on an aggregate of the returns that bulls expect to earn from holding long positions relative to the returns that bears expect from their short positions:

$$\widehat{pr}_{t|t+1}^{\text{IK}} = \frac{1}{2} (\widehat{up}_{t|t+1}^{\text{L}} - \widehat{up}_{t|t+1}^{\text{S}}) \quad (44)$$

where, echoing Knights (1921) distinction between uncertainty and risk, we refer to the expected returns of the bulls and bears, $\widehat{up}_{t|t+1}^{\text{L}} > 0$ and $\widehat{up}_{t|t+1}^{\text{S}} > 0$, respectively, as uncertainty premiums.⁹² We show that if forecasting behavior is consistent with the gap conditions in (43), the market premium will move positively with the aggregate gap, $\widehat{gap}_t^{\text{IK}}$, over time. The intuition follows from the connections between individuals' premiums and their forecasts of the potential losses that, as $\widehat{gap}_t^{\text{IK}}$ rises, foreign exchange bulls (bears) become more (less) concerned about the potential losses from holding open positions in the market; thus, in equilibrium, the uncertainty premiums of the bulls (bears) rise (fall). Both of these movements lead to an increase in the equilibrium premium.⁹³

One would not expect that the influence of the aggregate gap on the market premium would be the same in every time period. Indeed, we might expect that this effect would be small (large) when the size of the $\widehat{gap}_t^{\text{IK}}$ was small (large).⁹⁴ Nevertheless, our model implies that, although the quantitative relationship between the market premium and the gap varies over time, it does so in ways that preserve the qualitative relationship between these variables: a rise (fall) in $\widehat{gap}_t^{\text{IK}}$ is associated with) a rise (fall) in $\widehat{pr}_{t|t+1}^{\text{IK}}$.

Beyond its ability to account for the positive relationship between the market premium and the aggregate gap that we observe in the data, our IKE model is able to explain a feature of the empirical record that extant approaches have found extremely puzzling: the frequent switches in the algebraic sign of $\widehat{pr}_{t|t+1}^{\text{IK}}$ that are apparent in figure 1 and in other currency markets. The expression for the market premium in (44) shows that the algebraic sign of $\widehat{pr}_{t|t+1}^{\text{IK}}$ depends on the size of the bulls' uncertainty premium

⁹²The market premium also depends positively on the international financial position of the domestic country. See Frydman and Goldberg (2007, chapter 11). To keep the discussion simple, we ignore this feature of the model.

⁹³There are some distributional issues in deriving this result that we address in Frydman and Goldberg (2007, chapter 12).

⁹⁴We report evidence in Frydman and Goldberg (2007, chapter 12) that supports this view of a non-linear relationship.

relative to the bears'. When this relative magnitude changes sign, so, too, does the sign of the market premium.⁹⁵ One of the implications of our model is that the frequency of sign reversals in $\widehat{pr}_{t|t+1}^{\text{IK}}$ declines when the size of the aggregate gap is "large." We show that this is, indeed, the case in the BP, German mark, and Japanese yen markets using thresholds for $\widehat{gap}_t^{\text{IK}}$ of 5, 10, and 20 percent. This result makes clear that allowing for the presence of bulls and bears in the market is crucial to accounting for the time paths of the premium in currency markets.

13.1 Diversity of Forecasting Strategies and the Market Premium

The important role played by the forecasts of bulls and bears in driving the time-path of the premium makes points to the key reason why REH models have had such difficulties explaining risk in financial markets. Although one may be able to represent the behavior of bulls and bears by assuming some special informational asymmetries and a single REH forecasting strategy, in general, the distinction between the two groups' forecasts stems from differences in their forecasting strategies. Moreover, in the context of our model of the premium, the individual gap restrictions in (43) necessarily require that bulls and bears follow diverse forecasting strategies.

To see this, assume that individuals do make use of the same forecasting strategy for the next period's return R_{t+1} and write the conditional mean of the distribution representing this strategy as

$$E_t[R_{t+1}|Z_t] = \widehat{R}_{t|t+1}^+ + \widehat{R}_{t|t+1}^- \quad (45)$$

where $\widehat{R}_{t|t+1}^+ = E_t[R_{t+1} > 0|Z_t] > 0$ and $\widehat{R}_{t|t+1}^- = E_t[R_{t+1} < 0|Z_t] < 0$ denote the expected values of the positive and negative realizations of R_{t+1} , respectively. Furthermore, suppose that a bull and a bear both revise upward their forecast of the future exchange rate between $t-1$ and t with no change in their estimates of the PPP exchange rate. The increases in $\widehat{P}_{t|t+1}^{i,\text{L}}$ and

⁹⁵In the case of a non-zero international financial position, the algebraic sign of $\widehat{pr}_{t|t+1}$ still depends on the relative magnitude of $\widehat{up}_{t|t+1}^{\text{L}} - \widehat{up}_{t|t+1}^{\text{S}}$. Indeed, it is this term, which is absent in conventional models, regardless of whether they allow for bulls and bears, that enables us to account for sign reversals in the premium. See Frydman and Goldberg (2007, chapter 11).

$\widehat{P}_{t|t+1}^{i,s}$ are then tantamount to upward revisions of $\widehat{gap}_t^{i,L}$ and $\widehat{gap}_t^{i,s}$, and, because individuals take P_t as given, a rise in $\widehat{R}_{t+1}^{i,L}$ and a fall in $\widehat{R}_{t+1}^{i,s}$.⁹⁶

Consider the implications of these revisions of the forecasting strategies for the bull. The gap conditions in (43) imply that the rise in $\widehat{gap}_t^{i,L}$ is associated with a rise in $-\widehat{l}_{t|t+1}^{i,L} = -\widehat{R}_{t|t+1}^{i,-}$. However, in order for both $\widehat{R}_{t+1}^{i,L}$ and $-\widehat{R}_{t|t+1}^{i,-}$ to rise, $\widehat{R}_{t|t+1}^{i,+}$ must rise, too. Now consider the bear. The gap conditions in (43) imply that the rise in $\widehat{gap}_t^{i,s}$ is associated with a fall in $-\widehat{l}_{t|t+1}^{i,s} = \widehat{R}_{t|t+1}^{i,+}$. However, if the bull and the bear are assumed to follow the same forecasting strategy at $t-1$, then this common strategy can imply only an increase or a decrease in $\widehat{R}_{t|t+1}^{i,+}$, but not both. Thus, the gap conditions in (43), which specify part of the microfoundations of our model of the market premium, require that bulls and bears follow diverse forecasting strategies at every point in time.

14 How Recognizing the Limits to Knowledge Avoids Internal Inconsistency

We have sketched how our microfounded IKE model of the currency market accounts for three regularities on the aggregate level: long swings in exchange rates of uneven duration and magnitude, a positive relationship between the market premium and the aggregate gap, and the market premium's tendency to undergo sign reversals less frequently when the size of the aggregate gap is large. Allowing for the presence of bulls and bears, and the diversity of forecasting strategies among them, played a crucial role in our model of the premium.

In section 8.1, we showed that if a fully predetermined model allows for diversity of forecasting strategies, its multiple probability distributions representing these strategies are necessarily inconsistent with the single, overarching distribution—sharp prediction—that the model generates on the aggregate level. Following Lucas, we argued that such inconsistent models are the “wrong theory.”

We also showed that in order to avoid inconsistency, an economist must stop short of fully prespecifying the microfoundations of his model, in particular, revisions of forecasting strategies. Because IKE models do so, they imply

⁹⁶For bears, $\widehat{R}_{t+1}^{i,s} = P_t - \widehat{P}_{t|t+1}^{s,IK} > 0$, so a rise in $\widehat{P}_{t|t+1}^{s,IK}$ implies a fall in $\widehat{R}_{t+1}^{i,s}$.

myriad conditional probability distribution for P_{t+1} on the aggregate level at every point in time. But, although acknowledging the limits to knowledge is necessary to avoid internally inconsistent models, it is not sufficient.

14.1 The Gap Conditions and Market Premium

Consider the model's predictions concerning the causal mechanism underpinning the premium on foreign exchange. On the aggregate level, the model implies that $\widehat{pr}_{t|t+1}^{\text{IK}}$ and $\widehat{gap}_t^{\text{IK}}$ move positively together over time.⁹⁷ The model also generates myriad of conditional distributions for R_{t+1} , all of which imply a positive relationship between the conditional mean of R_{t+1} and the conditional mean of gap_{t+1} . Thus, to check that our model is not theoretically inconsistent, we must show that the gap conditions in (43), which constrain revisions of forecasting strategies on the individual level, are compatible with a positive relationship between the conditional mean of R_{t+1} and the conditional mean of gap_{t+1} for *both* bulls and bears.

We note that the aggregate regularity between $E_t[R_{t+1}|Z_t]$ and $E_t[gap_{t+1}|Z_t]$ implies no restrictions on how $\widehat{R}_{t|t+1}^+$ and $\widehat{R}_{t|t+1}^-$ may vary separately with $E_t[gap_{t+1}|Z_t]$. Thus, because the gap conditions in (43) constrain how bulls and bears revise their forecasts of the potential loss—which involve, separately, $\widehat{R}_{t|t+1}^-$ for the bulls and $\widehat{R}_{t|t+1}^+$ for the bears—they are compatible with the positive relationship between $E_t[R_{t+1}|Z_t]$ and $E_t[gap_{t+1}|Z_t]$ on the aggregate level.⁹⁸

14.2 Conservatism and Long Swings

As for the gap conditions, IKE's use of partially predetermined probability distributions enabled us to specify our microfoundations by constraining an

⁹⁷ $\widehat{pr}_{t|t+1}^{\text{IK}}$ and $\widehat{gap}_t^{\text{IK}}$ are aggregates of the conditional means of the distributions used to represent the forecasting of individual bulls and bears. Unless one has access to survey data, as we do in figure 2, such aggregates of individuals' expectations are, in general, not in general observable; thus, the relationship between them does not, by itself, have any implications for the qualitative properties of probability distributions representing forecasting on the individual level.

⁹⁸The model constrains a bull's revisions of her forecasting strategy in such a way that if the revised strategy involves a higher (lower) $\widehat{P}_{t|t+1}^{i,L}$ and thus $\widehat{gap}_t^{i,L}$, it will also involve a higher (lower) forecast of the potential loss from speculation, that is, a lower (higher) $\widehat{R}_{t|t+1}^-$ ($\widehat{l}_{t|t+1}^{i,L} < 0$). The converse is true for the bears.

aspect of the causal mechanism that differs from the qualitative feature of this mechanism predicted by the aggregate model. This, of course, would not be possible in the context of a fully predetermined model, because these models imply a single overarching distribution of outcomes and thus must disregard diversity—as REH models in fact do—to avoid internal inconsistency.

Conditional on conservative revisions and a fixed policy environment, the aggregate model predicts the direction of change in the exchange rate, either away from or toward the benchmark level. However, it does not have any implications for the way the causal mechanism—the relationship between the exchange rate and a set of causal factors—might change over time.

By contrast, the conservative restrictions constrain the way an individual revises her forecasts of the causal mechanism that underpins the exchange rate. However, they do not restrict in any way an individual’s prediction concerning *the direction* of change in the exchange rate. Indeed, the microfoundations of our model allow for both bulls, who forecast appreciation, and bears, who forecast depreciation. Thus, the conservative revisions of forecasting strategies on the individual level are compatible with the prediction of long-swings behavior of the exchange rate implied by the model on the aggregate level.

14.3 Bulls and Bears in a Long Swing

The model predicts that, conditional on conservative revisions of forecasting strategies and constant drifts in the causal variables, the exchange rate undergoes protracted swings that revolve around PPP. This prediction appears to conflict with the assumption that the microfoundations of our model allow for bulls and bears: during every up-swing or down-swing, there are market participants who are assumed to bet on a movement of the exchange rate in the opposite direction. The key to avoiding the inconsistency, however, is that our explanation of the swings recognizes that they are uneven in duration and magnitude: conservative behavior and constant drifts do not last forever. In the model, non-conservative and non-reinforcing revisions in strategies or changes in policy could occur at any point in time during a protracted swing either up or down. The model thus implies that a reversal in the exchange rate could occur at any point in time.

Consequently, the aggregate prediction of a protracted upswing (down-swing) is not incompatible with the presence of both bulls and bears on the individual level of our model. Because our model does not fully prespecify

when exchange rate movements might reverse direction, an individual who remains a bear during an upswing or a bull during a downswing is justified in doing.

15 The Futile Search for Sharp Predictions

Economists are trained early on to believe that models that do not generate sharp predictions are not worthy of consideration. However, the opposite is true. As John Kay (2007) put it in an article from which the title of this section comes, “the quest for exact knowledge gets in the way of useful knowledge.”

15.1 Lost Fundamentals in Currency Markets

The detrimental effect of the belief that only models that generate sharp predictions are worthy of scientific status is perhaps most evident in the field of financial economics. Consider, for example, how the contemporary approach has impeded economists’ thinking about whether macroeconomic fundamentals matter for currency movements.

There is much anecdotal evidence in the popular media, supported by survey research, that participants in the foreign exchange market pay close attention to fundamental variables in forming their forecasts of future exchange rates. It is obvious, for example, that market participants hang on every word that central bank officials utter, listening for hints of a change in monetary policy. Similarly, two years or so prior to the writing of this paper, market participants clearly responded to announcements of large and growing US current account deficits by selling the dollar. Because individuals’ forecasts drive their behavior in financial markets, we would expect fundamental variables to have considerable influence on exchange-rate fluctuations. Indeed, we showed how currency swings can arise in IKE-based monetary models even if all individuals rely solely on macroeconomic fundamentals in forming their forecasts.

However, in order to build models on the foundation of individual rational behavior while remaining faithful to the contemporary approach, conventional exchange-rate theorists modeled individual behavior and aggregate outcomes with fully predetermined representations. These conventional models were thought to offer the way to understand how macroeconomic funda-

mentals and rational behavior affect the exchange rate.

When they failed to find an overarching relationship between the exchange rate and macroeconomic fundamentals, conventional economists concluded that swings in exchange rates away from benchmark levels were unconnected to changes in these fundamentals. Obstfeld and Rogoff (2000) have referred to this “anomalous” finding as the “exchange-rate-disconnect puzzle.” This puzzle led many to presume not only that fundamentals do not matter, but also that some or all market participants behave “irrationally.”

However, in a dynamic world economy, we should not expect to find that a single set of economic fundamentals has mattered in exactly the same way since floating currencies became the norm in the 1970’s. In fact, the exchange-rate-disconnect puzzle disregards empirical evidence, much of it reported by conventional economists themselves, showing that, while macroeconomic fundamentals matter for exchange-rate movements, the causal mechanism that underpins these movements is temporally unstable: not only do the coefficients of empirical models change from one sub-period of floating to another, but the sets of fundamentals that seem to matter for exchange rates also change. Fully predetermined models cannot account for such structural change, the nature and timing of which depends on how market participants revise their forecasting strategies and on unforeseeable changes in the social context.

IKE acknowledges that an overarching model for currency movements is beyond the reach of economic analysis. Moreover, once we allow exchange rate models to undergo structural change at the points in time that are not prespecified, statistical analysis reveals that fundamentals do matter for exchange rate movements, after all.⁹⁹

15.2 Is the Market Really Grossly Inefficient?

Relying on invariant empirical relationships, many researchers report that future returns in currency markets co-vary negatively with the current value of the forward premium.¹⁰⁰ To explain this behavior, conventional econo-

⁹⁹In chapter 15 of our book, we show that if one allows for structural change, the forecast errors generated by fundamentals-based models are smaller, in mean square error, than the errors generated by the random walk. Thus, we reverse the Meese and Rogoff (1983) conclusion that flipping a coin outperforms structural models.

¹⁰⁰The forward premium depends on the difference between the forward and spot exchange rates.

mists have constructed exchange rate models in which risk-averse individuals require a positive return, a premium, to hold risky positions in currency markets. It is widely recognized, however, that this research effort has been unsuccessful.¹⁰¹

Unable to explain the negative co-variation between the return on foreign exchange and the forward premium that their studies report, economists have reached the startling conclusion that “one can make predictable profits by betting against the forward rate” (Obstfeld and Rogoff, 1996, p. 589). The apparent anomaly that these profits remain unexploited has become one of the major “puzzles” in the international finance literature.

There are several well-known studies in the literature that indicate that the relationship between the return on foreign exchange and the forward premium is temporally unstable. In chapter 13 of our book, we add to this evidence and show that the correlation between the return on foreign exchange and the forward premium is sometimes negative, sometimes positive, and sometimes insignificantly different from zero.

Acknowledging the importance of temporal instability goes a long way toward resolving the forward-rate “puzzle.” A returns process that gives rise to both negative and positive correlations with the forward premium implies that betting against the forward rate will be profitable during some time periods but not in others. We show that a trading rule based on betting against the forward rate does not deliver significant profits over the modern period of floating in the major currency markets.

Because the contemporary approach has led economists to construct fully predetermined, mostly invariant models of foreign exchange returns that ignore temporal instability, the “finding” of a negative correlation between returns and the forward premium has led them to conclude that there is easy money to be made in the foreign exchange market. But, since the correlation is sometimes negative and at other times positive, fully predetermined trading rules based on the forward rate do not deliver profits. As in the case of the disjunction between the exchange rate and macroeconomic fundamentals, the forward-rate “puzzle” is another artifact of the epistemological flaws inherent in the contemporary approach. The fact that there are literally hundreds of studies attempting to explain this “puzzle” provides an example par excellence of how contemporary economics’ insistence on sharp predictions has misdirected research and impeded its progress.

¹⁰¹See chapter 8 of our book for a discussion of this literature

16 Coming to Terms with Imperfect Knowledge

The research program of contemporary economics is predicated on the belief that it is possible to fully prespecify economic change over periods of time as long as decades. The premise that seems, at least implicitly, to motivate this mechanistic way of modeling market outcomes is that there exists a fully predetermined causal mechanism that underpins actual behavior on the individual and aggregate levels.

But, forecasting behavior on the part of purposeful individuals alters the causal mechanism that underpins market outcomes in ways—and at points in time—that cannot be fully prespecified. Moreover, changes in the social context, including the evolution of institutions, values, and norms, are all important in engendering temporal instability in causal relationships in real-world markets.

If change in capitalist economies is not governed by a fully predetermined causal mechanism, then attempting to explain individual behavior and aggregate outcomes with representations that presume the existence of such a mechanism is misguided. It is not surprising, then, that the contemporary approach has had great difficulties in discovering the “mechanics of economic development” (Lucas, 2002, p. 21) in many markets where profit-seeking inherently involves coping with ever-imperfect knowledge.

Beyond the contemporary research program’s empirical difficulties, its reliance on fully predetermined models precludes the possibility of building any coherent economic theory of aggregate outcomes based on plausible micro-foundations. Remarkably, what economic theory requires to escape extant approaches’ epistemological flaws is also required to overcome their empirical difficulties: economists must give up their insistence on sharp predictions.

Economics calls for a new approach that represents individual behavior and aggregate outcomes mathematically, and that, at the same time, refrains from fully prespecifying economic change. IKE offers a way to take up this task.

References

- [1] Barberis, Nicolas C., Huang, Ming. and Tano Santos (2001), “Prospect Theory and Asset Prices,” *Quarterly Journal of Economics*, pp. 1-53.
- [2] Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998), “A Model of Investor Sentiment,” *Journal of Financial Economics*, vol. 49, 307-43.
- [3] Blanchard, Olivier. J. and Stanley Fischer (1989), *Lectures on Macroeconomics*, Cambridge, MA.: The MIT Press.
- [4] Camerer, Colin F., Loewenstein, George and Matthew Rabin (2004), *Advances in Behavioral Economics*, Princeton, NJ:Princeton University Press.
- [5] De Grauwe, Paul and Marianna Grimaldi (2006), *The Exchange Rate in a Behavioural Finance Framework*, Princeton, NJ:Princeton University Press.
- [6] Dornbusch, Rudiger (1976), “Expectations and Exchange Rate Dynamics,” *Journal of Political Economy*, December, 1161-1174.
- [7] Edwards, W. (1968), “Conservatism in Human Information Processing,” in B. Kleinmuth (ed.), *Formal Representation of Human Judgement*, New York: John Wiley and Sons.
- [8] Engle, Robert F. (2003), “Risk and Volatility: Econometric Models and Financial Practice,” *Nobel Lecture*, <http://nobelprize.org/economics/laureates/2003/engle-lecture.pdf>.
- [9] Evans, George W. and Seppo Honkapohja, “An Interview with Thomas J. Sargent,” *Macroeconomic Dynamics*, 9, 2005, pp. 561-583.
- [10] Frankel, Jeffrey A. (1979), “On the Mark: A Theory of Floating Exchange Rate Based On Real Interest Differentials,” *American Economic Review*, vol. 69, September, 610-22.
- [11] Frankel, Jeffrey and Kenneth Froot (1987), “Understanding the U.S. Dollar in the Eighties: The Expectations of Chartists and Fundamentalists,” *Economic Record*, December, 24-38, and reprinted in Jeffrey A.

- Frankel (ed.), *On Exchange Rates*, The MIT Press: Cambridge, Mass., 1995, chapter 14, pp. 295-316.
- [12] Frankel, Jeffrey A. and Kenneth Froot (1987), "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review*, vol. 77, March, 133-53.
 - [13] Frankel, Jeffrey and Andrew K. Rose (1995), "Empirical Research on Nominal Exchange Rates," in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Vol. III, Amsterdam: North-Holland, 1689-1729.
 - [14] Froot, Kenneth and Jeffrey A. Frankel (1989), "Forward Discount Bias: Is it an Exchange Risk Premium?," *Quarterly Journal of Economics*, vol. 104, February, 139-161, reprinted in Jeffrey A. Frankel, ed., *On Exchange Rates*, Cambridge, MA: The MIT Press, 1995.
 - [15] Frydman, Roman (1982), "Towards an Understanding of Market Processes: Individual Expectations, Learning and Convergence To Rational Expectations Equilibrium," *American Economic Review*, 72, 652-668.
 - [16] Frydman, Roman (1983), "Individual Rationality, Decentralization and the Rational Expectations Hypothesis," in Frydman, Roman and Edmund S. Phelps (eds.), *Individual Forecasting and Aggregate Outcomes: "Rational Expectations" Examined*, New York: Cambridge University Press, pp. 97-122.
 - [17] Frydman, Roman and Michael D. Goldberg (2003), "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.
 - [18] Frydman, Roman and Michael D. Goldberg, (2004), "Limiting Exchange Rate Swings in a World of Imperfect Knowledge," in Peter Sorensen (ed.), *European Monetary Integration: Historical Perspectives and Prospects for the Future. Essays in Honour of Niels Thygesen*. Copenhagen: DJOEF Publishing.

- [19] Frydman, Roman and Michael D. Goldberg (2007), *Imperfect Knowledge Economics: Exchange Rates and Risk*, Princeton, NJ.: Princeton University Press.
- [20] Frydman, Roman, Michael D. Goldberg, Katarina Juselius, Soren Johansen (2008), "Is Imperfect Knowledge the Key to Solving the Purchasing Power Parity Puzzle?," in process.
- [21] Frydman, Roman and Edmund S. Phelps (1983), "Introduction" in Frydman, Roman and Edmund S. Phelps (eds.), *Individual Forecasting and Aggregate Outcomes: "Rational Expectations" Examined*, New York: Cambridge University Press, 1-30.
- [22] Frydman, Roman and Edmund S. Phelps (1990), "Pluralism of Theories Problems in Post-Rational-Expectations Modeling," paper presented at the 1990 Siena Summer Workshop on "Expectations and Learning", June.
- [23] Goldberg, Michael D. and Roman Frydman (1996a), "Imperfect Knowledge and Behavior in the Foreign Exchange Market," *Economic Journal*, 106, 869-893.
- [24] Goldberg, Michael D. and Roman Frydman (1996b), "Empirical Exchange Rate Models and Shifts in the Co-Integrating Vector", *Journal of Structural Change and Economic Dynamics*, 7, 55-78.
- [25] Hamilton, J.D. (1988), "Rational-Expectations Econometric Analysis of Changes in Regime: An Investigation of the Term Structure of Interest Rates," *Journal of Economics Dynamics and Control*, 12, pp. 385-423.
- [26] Hamilton, James D., (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, pp. 357-384.
- [27] Hamilton, James D., (1990), "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics*, 45, 39-70.
- [28] Hamilton, James D., (1994), *Time Series Analysis*, Princeton, NJ: Princeton University Press.

- [29] Hayek, Friedrich A. (1945), "The Use of Knowledge in Society," *American Economic Review*, 35, pp. 519-30.
- [30] Hayek, Friedrich A. (1948), *Individualism and Economic Order*, Chicago: The University of Chicago Press.
- [31] Hayek, Friedrich.A. (1978), "The Pretence of Knowledge," *Nobel Lecture*, 1974, in *New Studies in Philosophy, Politics, Economics and History of Ideas*, Chicago: The University of Chicago Press.
- [32] Johansen, Soren. (1992), "An $I(2)$ cointegration analysis of the purchasing power parity between Australia and USA," in Colin Hargreaves (ed.), *Macroeconomic modelling of the long run*, Edward Elgar.
- [33] Juselius, Katarina (2007), *The Cointegrated VAR Model: Methodology and Applications*, forthcoming, Oxford University Press.
- [34] Juselius, Katarina, Soren Johansen, Roman Frydman, and Michael D. Goldberg (2008), "Testing Hypotheses in an $I(2)$ Model with Applications to the Persistent Long Swings in the Dmk/\$ Rate," invited paper for a special issue of the *Journal of Econometrics*.
- [35] Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, pp. 263-291.
- [36] Kay, John (2004), "Obliquity," *Financial Times*, January 17th.
- [37] Kay, John (2007), "The Futile Search for Sharp Prediction," *Financial Times*, October 17th.
- [38] Keynes, John Maynard. (1936), *The General Theory of Employment, Interest and Money*, Harcourt, Brace and World.
- [39] Knight, Frank H. (1921), *Risk, Uncertainty and Profit*, Boston: Houghton Mifflin.
- [40] Kongsted, H. C., A. Rahbek, and C. Jørgensen (1999), Trend stationarity in the $I(2)$ cointegration model," *Journal of Econometrics*, 90, 265-289.
- [41] Lucas, Robert E. Jr. (1973), "Some International Evidence on Output-Inflation Trade-offs," *American Economic Review*, 63, pp. 326-334.

- [42] Lucas, Robert E. Jr. (1976), "Econometric Policy Evaluation: A Critique," in Brunner Karl and Allan H. Meltzer, *The Phillips Curve and Labor Markets*, Carnegie_Rochester Conference Series on Public Policy, Amsterdam: North-Holland Publishing Company, 19-46. Reprinted in Robert E. Lucas, Jr. (1981), *Studies in Business-Cycle Theory*, Cambridge, MA: The MIT Press, pp. 104-130.
- [43] Lucas, Robert E. Jr. (1982), "Interest Rates and Currency Prices in a Two-Country World," *Journal of Monetary Economics*, vol. 10, November, 335-59.
- [44] Lucas, Robert E. Jr. (1995), "The Monetary Neutrality," *The Nobel Lecture*, Stockholm: The Nobel Foundation.
- [45] Lucas, Robert E. Jr. (2001), "Professional Memoir," mimeo. <http://home.uchicago.edu>.
- [46] Lucas, Robert E. Jr. (2002), *Lectures on Economic Growth*, Cambridge, MA: Harvard University Press.
- [47] Lucas, Robert E. Jr. (2003), "Macroeconomic Priorities," *American Economic Review*, 93, pp. 1-14.
- [48] Meese, Richard A. and Kenneth Rogoff (1983), "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?," *Journal of International Economics*, February, vol. 14, 3-24.
- [49] Mises, Ludwig von (1920), *Economic Calculation in a Socialist Commonwealth*, Reprinted in 1990 by the Ludwig von Mises Institute.
- [50] Muth, John F. (1961), "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, pp. 315-335.
- [51] Nielsen, H. B. and A. Rahbek (2007), "The likelihood ratio test for cointegration ranks in the I(2) Model," *Econometric Theory*, 23, 615-637.
- [52] Obstfeld, Maurice and Kenneth Rogoff (1995), "Exchange Rate Dynamics Redux," *Journal of Political Economy*, vol. 103, June, 624-60.
- [53] Obstfeld, Maurice and Kenneth Rogoff (1996), *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press.

- [54] Obstfeld, Maurice and Kenneth Rogoff (2000), "The Six Major Puzzles in International Macroeconomics: Is there a Common Cause?," in Ben Bernanke and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2000*, Cambridge, MA: The MIT Press, 339-90.
- [55] Officer, L. (1982), "The Purchasing Power Parity Theory of Exchange Rates: A Review Article," *International Monetary Fund Staff Papers*, 23, 1-60.
- [56] Phelps, Edmund S. (1968), "Money Wage Dynamics and Labor Market Equilibrium," *Journal of Political Economy*, 76, 687-711.
- [57] Phelps, Edmund S. *et al.* (1970), *Microeconomic Foundations of Employment and Inflation*, New York: Norton.
- [58] Phelps, Edmund S. (1983), "The Trouble with 'Rational Expectations' and the Problem of Inflation Stabilization," in Frydman, Roman and Edmund S. Phelps (eds.), *Individual Forecasting and Aggregate Outcomes: "Rational Expectations" Examined*, New York: Cambridge University Press.
- [59] Phelps, Edmund S. (1994), *Structural Slumps: The Modern Equilibrium Theory of Unemployment, Interest and Assets*, Cambridge, MA. Harvard University Press.
- [60] Phelps, Edmund S. (2008), "Our Uncertain Economy," *The Wall Street Journal*, March 14th.
- [61] Popper, Karl (1957), *The Poverty of Historicism*, London and New York: Routledge.
- [62] Sargent, Thomas J. (1993), *Bounded Rationality in Macroeconomics*, Oxford, UK: Oxford University Press.
- [63] Sargent, Thomas J. (2001), *The Conquest of American Inflation*, Princeton, NJ: Princeton University Press.
- [64] Schulmeister, Stephan (1983), "Exchange Rates, Prices and Interest Rates: Reconsidering the Basic Relationships of Exchange Rate Determination," C.V. Starr Center For Applied Economics Working Paper No. 83-13, July.

- [65] Schulmeister, Stephan (1987), "Currency Speculation and Dollar Fluctuations," *Banca Nazionale del Lavoro Quarterly Review*, 167, pp. 343-65.
- [66] Schulmeister, Stephan (2006), "The Interaction between Technical Currency Trading and Exchange Rate Fluctuations," *Finance Research Letters*, 3, 212-33.
- [67] Schulmeister, Stephan (2008), "Components of the Profitability of Technical Currency Trading," forthcoming in *Applied Financial Economics*.
- [68] Shiller, Robert J. (1979), "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure," *Journal of Political Economy*, vol. 87, December, 1190-1219.
- [69] Shiller, Robert J. (1981), "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends," *American Economic Review*, vol. 71, no. 3, 421-36.
- [70] Shiller, Robert J. (1990), "Speculative Prices and Popular Models," *Journal of Economic Perspectives*, vol. 4, Spring, 55-65.
- [71] Shiller, Robert J. (2000), *Irrational Exuberance*, New York: Broadway Books.
- [72] Shleifer, Andrei (2000), *Inefficient Markets*, Oxford: Oxford University Press.
- [73] Soros, George (1987), *The Alchemy of Finance*, New York: Wiley.
- [74] Soros, George (2008), *The New Paradigm for Financial Markets: The Credit Crisis of 2008 and What It Means*, New York: Public Affairs.
- [75] Stiglitz, Joseph E. (2001), "Information and the Change in the Paradigm in Economics," *Nobel Lecture*.
- [76] Stockman, Alan C. (1980), "A Theory of Exchange Rate Determination," *Journal of Political Economy*, vol. 88, no. 4, 673-98.
- [77] Taylor, Alan and Mark P. Taylor (2004), "The Purchasing Power Parity Debate," forthcoming in *Journal of Economic Perspectives*.

- [78] Tversky, Amos and Daniel Kahneman (1992), “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 5, pp. 297-323.