From the Prevailing Paradigm to the Qualitative Expectations Hypothesis

Roman Frydman*, Søren Johansen†, Anders Rahbek†, and Morten Nyboe Tabor†

A Note Prepared for the Institute for New Economic Thinking
Plenary Conference in Edinburgh, October 21-23, 2017

October 10, 2017

*Department of Economics, New York University. †Department of Economics, University of Copenhagen.
The authors are participants in the Research Program on Imperfect Knowledge Economics at the Institute for New Economic Thinking (INET). The authors are grateful to INET for continuing support of this research. Johansen gratefully acknowledges support from CREATES - Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. Rahbek is grateful for funding by The Danish Council for Independent Research (Grants 12-124980 and DFF-7015-00028). Tabor greatly appreciates support from the Carlsberg Foundation (grant reference 2013-01-0972).
Correspondence to: Roman Frydman, Department of Economics, New York University, 19 West 4th Street, New York, NY 10003, USA;e-mail: rf3@nyu.edu;

In the paper that we present this afternoon, Søren Johansen, Anders Rahbek, Morten Tabor, and I introduce the Qualitative Expectations Hypothesis (QEH) as a new approach to modeling macroeconomic and financial outcomes.

QEH recognizes that economists and market participants alike face ambiguity about which quantitative model of the process driving outcomes is the correct one. Building on Frank Knight's distinction between risk and "true uncertainty," QEH formalizes ambiguity by opening an economic model to unforeseeable change in its coefficients. The defining feature of such change is that it cannot "by any method be [represented ex ante] with an objective, quantitatively determined probability" (Knight, 1921, p. 321).

By design, QEH models do not seek to account for quantitative regularities in time-series data. Instead, QEH proposes a way to build models that generate qualitative predictions of co-movements in these data.

Opening a model to unforeseeable change, and yet aiming to confront the model's predictions with empirical evidence, poses considerable challenges. In our presentation, we discuss how QEH proposes a way forward that might enable economists to overcome these challenges.

Regardless of whether QEH, or some other, yet-to-be invented approach, turns out to be useful in this regard, recognizing the inherent limits to what we can know about the future appears to be necessary for developing epistemologically coherent and empirically relevant macroeconomic and finance models. In this note, we provide a rigorous argument for this imperative.

1 The Consensus Conception of Economic Science

Today, economic models must account for quantitative regularities in time-series data to be considered scientific. To this end, an overwhelming majority of economists constrain the structure of their models – the set of variables that they include and their parameters — to be unchanging over time. When economists do recognize that the process underpinning outcomes undergoes change, they represent such change with a probabilistic rule – for example, Markov switching. Consensus-compliant models thus specify a complete dynamic stochastic process driving outcomes. The conditional moments of this process constitute the model's quantitative predictions of regularities in time-series data.

1.1 REH and Its Consequences

The most far-reaching application of models generating quantitative predictions has been to underpin REH's representations of the aggregate of market participants' forecast. These representations implement John Muth's path-breaking idea that an economist can specify the market's forecast by imposing consistency within his own model. As he put it,

> [participants' expectations] are essentially the same as the predictions of *the relevant economic theory* (emphasis added, Muth, 1961, 316).

Following the disciplinary consensus, an economist implements this insight by hypothesizing that a complete dynamic stochastic process driving outcomes is "the relevant economic theory." He then relies on REH to represent participants' forecasts with the model's quantitative predictions.

However, representing the market's forecasts as being consistent with the model's hypothesis about how outcomes will actually unfold over time should be viewed as a principle of coherent model-building in general. As such, Muth's principle can be applied in any model. As we discuss below, QEH relies on Muth's principle to represent participants' forecasts in models that recognize ambiguity. Although QEH models do not generate quantitative predictions of movements in time-series data, their qualitative predictions of regularities in such data can be used to represent the model-consistent market's forecasts.

Viewing Muth's insight as a principle of coherent model-building may help explain why so many economists have found REH so appealing and embraced it so quickly. Given the consensus view that models generating quantitative predictions are the relevant economic theory, non-REH representations are thought to presume that market participants are grossly irrational.

Indeed, on purely logical grounds, as Robert Lucas pointed out early on, representing the market's forecast as being inconsistent with the model's quantitative predictions would presume that participants ignore systematic, observable forecast errors in perpetuity. By design, REH excludes such gross irrationality from models that generate quantitative predictions. This has buttressed the common belief that REH models represent rational forecasting and decision-making.

The conviction that REH enables one to understand how rational individuals make decisions has had far-reaching implications for how economists and non-economists alike understand financial markets, and for how central banks analyze the consequences of alternative policies. In fact, the relevance of REH in representing how rational participants forecast outcomes hinges on the validity of the assumption that models that generate quantitative predictions are the relevant economic theory.

1.1.1 The Efficient Market Hypothesis

Arguably, the most far reaching consequence of tacit acceptance of this assumption has been the belief that in leaving markets unfettered, participants' rationality can be relied upon to allocate a society's savings to alternative investments nearly perfectly. This claim – known as the Efficient Market Hypothesis (EMH) – is usually interpreted to presuppose that the market uses all available information.

In a seminal critique of EMH, Grossman and Stiglitz (1980) maintained REH. Consequently, their argument that markets' allocations deviate from the supposed perfection of EMH has relied on the observation that information is asymmetric, in the sense that some market participants have access to information that for various reasons is not available to others.

However, examining the allocative consequences of informational asymmetries in the context of an REH model, as Grossman and Stiglitz do, entails the belief that the complete stochastic process specified by the model to represent an investment's future prospects accurately characterizes how returns will unfold over time.

QEH makes no such claim. As we shall discuss shortly, Frank Knight persuasively argued that the key to understanding profit-seeking behavior in real-world markets is change that cannot be foreseen with a probabilistic rule – for example, Markov switching. By design, REH models assume away such change. This may help clarify why Alan Greenspan was surprised that, following two decades of far-reaching financial deregulation, self-interest did not save us from the 2008 crisis. The expectation that self-interest would suffice ignored the possibility that REH, and thus EMH, do not characterize what profit-seeking rational participants actually do.

1.2 Behavioral Finance's Revision of the Prevailing Paradigm

The belief that REH models represent rational decision-making has profoundly shaped our understanding of how financial markets make use of available information, and what the consequences of leaving markets unfettered might be. And the same belief has shaped the profession's interpretation of and response to these models' empirical difficulties. Robert Shiller's early path-breaking study of stock-price movements provides a case in point.

Shiller (1981) presented evidence that the REH-based present-value model is grossly inconsistent with the existence of persistent swings in stock prices. This finding provided the *raison d'être* for the development of behavioral-finance models. This approach replaced model-consistent, REH representations of the market's forecast with formalizations of empirical observations of how market participants behave. However, economists who have embraced the behavioral-finance approach have followed the prevailing conception of economic science. Thus, like their REH counterparts, they have sought to construct models that generate quantitative predictions of movements in time-series data.

Relying on these models, and consequently maintaining the belief that REH represents how rational individuals forecast companies' prospects, has led to a remarkable outcome. Behavioral-finance models interpreted Shiller's research, and subsequent findings corroborating its results, as evidence that market participants are irrational and that their expectations – and thus stock-price swings – reflect psychological and technical considerations that are largely unrelated to fundamental factors. The commonly invoked examples of such considerations are market sentiment (optimism or pessimism) and bandwagon effects (participants' mechanical extrapolation of past returns into the future).

Behavioral-finance economists' interpretation of and response to REH models' difficulties underscores the insuperable obstacles inherent in any effort to move beyond these models within the consensus view of economic science. To be sure, behavioral-finance economists have highlighted an important flaw of REH models: by design, these models assume away the role of psychological factors in driving outcomes. Behavioral economists have persuasively demonstrated the empirical relevance of these factors, so it was indeed plausible that their absence could have contributed to REH models' empirical difficulties.

However, having accepted the prevailing consensus that models should specify a complete stochastic process driving outcomes, behavioral economists were left with no option but to rely on model-inconsistent representations of psychological factors to formalize their role in driving participants' expectations. Moreover, in order for their models to generate quantitative predictions, behavioral-finance economists have had to represent psychological considerations with probabilistic rules.

Thus, the decision to remain within the confines of the prevailing paradigm has forced behavioral-finance economists to presume that market participants are not only grossly irrational, but that psychological factors drive their forecasts in unchanging and quantitatively predictable ways. Even if we recognize that participants rely on psychological considerations in forming forecasts, we would not expect that they do so according to rules that can be completely specified in advance. Unsurprisingly, behavioral models' formalizations have been found to be inconsistent with empirical evidence about how market participants behave. For example, Fama (1998) and Frydman *et al* (2015) have shown that the effects of psychological factors on asset-price movements cannot be characterized with quantitative regularities. Moreover, recalling Lucas's argument, behavioral models' premise that psychological considerations can explain why participants ignore systematic and observable forecast errors in perpetuity seems untenable.

2 Rethinking Economic Models

The apparent difficulty of responding to REH's shortcomings within the limits of the prevailing paradigm motivates our attempt to move beyond it. In doing so, we build on Knight's distinction between risk and "true uncertainty." This "Knightian uncertainty" arises from "our imperfect knowledge of the future" – the consequence of unforeseeable change, which cannot "by any method be reduced to an objective, quantitatively determined probability" (Knight, 1921, p. 321).

Knight argued that recognizing uncertainty arising from unforeseeable change is the key to understanding business decisions in capitalist economies. The reason is simple:

> if all changes... could be foreseen for an indefinite period in advance of their occurrence, . . . profit or loss would not arise (Knight 1921, p. 198).

If Knight is correct, we should not be surprised that the prevailing paradigm's models – which reduce uncertainty to probabilistic measures of risk – encounter empirical difficulties in accounting for outcomes in real-world markets. On the contrary, we should expect it.

In line with Knight's arguments, Frydman and Goldberg (2007, 2011) trace the epistemological and empirical shortcomings of prevailing approaches to the premise that unforeseeable change is unimportant for understanding asset prices and risk. In his Nobel lecture, Lars Hansen (2013, p. 399) conjectured that empirical puzzles that REH models have been unable to resolve for decades arise from these models' narrow representation of uncertainty as "risk conditioned on a model." He pointed out that REH representations "miss something essential: uncertainty [arising from] ambiguity about which is the correct model."

Such arguments point to the fundamental flaw in the prevailing consensus in macroeconomics and finance theory: Any model seeking to account for quantitative regularities in time-series data must ignore the inherent limits of what we can know about these regularities. But in order to examine whether recognizing ambiguity is indeed essential to understanding outcomes in real-world markets, we need a rigorous framework to formalize the limits of what we can know about the future. QEH models suggest a way to do so.

3 A Way Forward

QEH models explicitly formalize the limits of what we can know about the future. In doing so, they recognize that unforeseeable change puts the discovery of quantitative regularities in time-series data out of reach for economic analysis. Simply put, even if such regularities can be found to characterize some period in the past, unforeseeable change in the process driving outcomes will sooner or later render them inconsistent with time-series data.

QEH's approach to model building is based on the inherent connection between the ambiguity that we face when making decisions involving future outcomes and Knightian uncertainty. According to Knight, representing this uncertainty requires an economist to jettison models that represent outcomes with a complete stochastic process.

QEH formalizes ambiguity by opening an economic model to unforeseeable change in its coefficients. Once left open to such change, the model ceases to generate quantitative predictions of outcomes. However, a model that is completely open to unforeseeable change does not generate even qualitative predictions of movements in time-series data. In order to formalize a QEH model's hypothesis about qualitative regularities in these data, we must restrict unforeseeable change in the model's parameters.

To this end, we hypothesize that, although the quantitative impact of the model's variables on outcomes varies over time, they have the same qualitative impact at every point in time. Operationalizing such a hypothesis involves, for example, restricting the sign of the model's parameter to be positive.

We also restrict the unforeseeable change in a QEH model's coefficients to lie within stochastic intervals driven by the evolution of the fundamental variables in the model. As a result, the model assumes that outcomes lie within stochastic intervals.

Because a QEH model, by design, does not specify a complete dynamic stochastic process for outcomes, we cannot rely on the standard mathematical expectation to define the model's predictions. Instead, we formalize these predictions as a conditional qualitative expectation (QE) of future outcomes, which we define as the conditional expectation of the upper and lower bounds of these outcomes' assumed stochastic intervals. The QE measures the intervals within which future outcomes, according to the model, are expected to lie. But we leave the model open to ambiguity by not specifying a mechanism determining specific values that the outcomes can take within these intervals.

3.1 Model-Consistent Representations of the Market's Forecasts

Like REH, QEH relies on Muth's principle that coherent model-building requires representing the market's forecast in a way that is consistent with an economist's hypothesis about how outcomes actually unfold over time. QEH represents the market's forecast as lying within QE intervals. Thus, although QEH models, like their REH counterparts, are internally consistent, applying Muth's principle to them does not represent the market's forecast with a precise value at each point in time.

3.2 Psychology and Fundamentals in Asset-Price Movements

Herein lies a possibility that cannot be explored in models that, by design, assume away unforeseeable change. By recognizing ambiguity, a QEH model can account for the role of both the fundamental factors on which REH models focus and the psychological factors underpinning behavioral-finance models. *And it can do so without abandoning model consistency.*

Because a QEH model assumes that market forecasts lie within stochastic intervals that are consistent with the process an economist assumes is driving outcomes, there are myriad possible model-consistent quantitative forecasts. In making decisions – for example, about how many stocks to buy or sell – market participants thus face inherent ambiguity. They select particular quantitative forecasts by relying on a combination of considerations, including formal (econometric) models, market sentiment, and other non-fundamental factors. A QEH model can formalize the qualitative effect of such factors on participants' model-consistent forecasts, by imposing additional restrictions on how the market, in forming these forecasts, revises the weights it attaches to fundamentals.

3.3 Modeling Stock-Price Movements Under Ambiguity

In the paper that we present this afternoon, we use a simple stock-price model to illustrate how a QEH model relates the market's forecast of dividends and prices to fundamentals. We also propose a formalization of how psychological factors might affect the market's forecasts and yet remain consistent with the model's prediction that future outcomes will lie within QE intervals. With these representations, we show how the model generates predictions about the qualitative regularities characterizing movements in time-series data.

In the mathematical Appendix to this note, we use the same model to illustrate the main points in our discussion of the REH and behavioral-finance approaches. We focus on how the prevailing consensus that models must account for quantitative regularities has forced economists to choose between these two extremes. Economists have either had to accept that REH represents how rational market participants relate outcomes to fundamental factors or presume that participants' reliance on psychological considerations is a symptom of their gross irrationality.

We also use our model to illustrate rigorously Knight's distinction between risk and "true uncertainty." Our example makes clear that the usual measures of risk that, as Hansen put it, are "conditioned on the model" being correct miss the uncertainty arising from unforeseeable change in real-world markets. This example buttresses Hansen's conjecture that to understand asset-price movements and risk in real-world markets, economists should recognize that they and market participants alike face ambiguity about which is the correct model of outcomes.

4 Concluding Remarks

This note argued that the only way to overcome the epistemological flaws and empirical shortcomings of the prevailing paradigm is to jettison its core premise that we can uncover the precise mechanism driving outcomes. In the Appendix we illustrate these arguments in the context of a simple model for stock prices. This sets the stage for the presentation of our QEH paper, where we develop our approach to building models that recognize the ambiguity confronting economists and market participants alike about which is the correct quantitative model of the process driving outcomes.

Much work remains to be done to determine whether QEH's application to modeling outcomes in equity markets can shed new light on the long-standing puzzle of what drives price movements. Moreover, whether our approach to formalizing ambiguity yields useful, empirically relevant results in other contexts can be established only by developing other applications.

However, opening economic models to unforeseeable change seems crucial for understanding how well asset markets allocate society's savings and what role the state might play in regulating them. Even our simple application of QEH to stock-price movements shows that EMH's claim that asset markets allocate resources nearly perfectly is an artifact of REH's assumption that there are no limits to what we can know about the future.

Despite its simplicity, the structure of the model for stock prices that we use here captures the key features of models that are typically used in other contexts. For example, the importance of forward-looking expectations is central to the New Keynesian approach, which underpins the DSGE models used by central banks to analyze the consequences of alternative economic policies. Because these models rest on REH, it seems plausible that reformulating the New Keynesian models to account for ambiguity can shed new light on why the DSGE models have encountered such difficulties in reliably guiding policy analysis. One area of future research is to assess whether QEH's approach to formalizing the inherent ambiguity that policymakers and market participants face could help us resolve some of these difficulties, and thereby enhance macroeconomic models' usefulness for policy analysis.

APPENDIX

In this Appendix, we illustrate our discussion of the consensus conception of economic science and how REH and behavioral-finance models adhere to it. We do so in the context of a model for the stock-price and dividend movements. As we point out in the Concluding Remarks, despite its simplicity, the structure of the model shares the key features of models that are typically used in macroeconomics and finance theory.

We use the model to show how the prevalent approaches close models completely to unforeseeable change, thereby generating predictions of quantitative regularities. We then illustrate how assuming away such change underpins the widely held belief that REH represents rational forecasting. We also show how the consensus that economic science necessitates such models has led behavioral-finance theorists to build models that presume that market participants suffer from gross irrationality.

We then open the dividend model to unforeseeable change and the Knightian uncertainty to which it gives rise. We use our example to illustrate formally the distinction between the standard probabilistic notion of risk and Knight's "true uncertainty."

A A Simple Stock-Pricing Model

The model rests on an assumption that summarizes how the market sets the stock price at each point in time. Participants bid the price to the level that satisfies the following no-arbitrage condition:

$$p_t = \gamma \left[\mathcal{F}_t \left(d_{t+1} \right) + \mathcal{F}_t \left(p_{t+1} \right) \right] \quad \text{for } t = 1, 2, 3, ... \tag{1}$$

where p_t is the market price, d_t denotes dividends, $\mathcal{F}_t(\cdot)$ stands for the time-t values of the market's (an aggregate of its participants') forecasts of dividends and prices at time $t+1$, and γ is a discount factor, which, for simplicity, we set equal to a constant.

In order to derive testable implications of the no-arbitrage condition in (1) we must represent the values of the market's forecasts formally. To this end, we consider a simple model for the dividend process, which relates dividends to one fundamental factor, corporate earnings, which we denote by x_t:

$$d_t = b_t x_t + \varepsilon_t, \tag{2}$$

where ε_t are i.i.d.$(0, \sigma_d^2)$ innovations and b_t is the time-varying impact of earnings on dividends. We assume log-earnings follow a random walk with drift, and choose a drift so that x_t is a martingale:

$$E(x_t|x_{t-1}) = x_{t-1} \tag{3}$$

B The Consensus Models

Both REH and behavioral-finance models specify a complete stochastic process driving dividends and prices. We illustrate how they do so in the context of our simple model for the dividend process. The consensus models sometimes constrain the change in b_t with a probabilistic rule, such as Markov switching. But the overwhelming majority of these models constrain the impact of earnings on dividends to be the same – say, equal to a constant, b– at every point in time. As a result,

$$d_t = bx_t + \varepsilon_t \tag{4}$$

B.1 Quantitative Regularities

The specification in (4) formalizes an economist's hypothesis about how dividends actually unfold over time. It illustrates in a particularly simple way how the consensus models generate predictions of quantitative regularities in time-series data. For example, given the information on earnings at time t, the model predicts that dividends in any future period will, on average, be precisely equal to the following conditional expectation:

$$E\left[\left(bx_{t+i} + \varepsilon_{t+i}\right)|x_t\right] = bx_t \quad \text{for } i = 1, 2, , \tag{5}$$

B.2 REH

Once an economist hypothesizes that (3) and (4) characterize how dividends actually unfold over time, relying on Muth's principle leads him to represent the market's forecast as being consistent with this hypothesis. REH implements this principle by representing this forecast to be the same as the model's prediction of quantitative regularity in (5). Because this regularity holds in all time periods, the model represents the time$-t$ market forecast of dividends in any future period to be exactly the same:

$$\mathcal{F}_t\left(d_{t+i}\right) = bx_t \quad \text{for } i = 1, 2, , \tag{6}$$

B.3 Individual Rationality of REH

Muth proposed REH as a way to represent the market's forecast. However, Lucas pointed out early on that, once an economist follows the disciplinary consensus of building models that generate quantitative regularities, the REH representation of the market's forecast, such as in (6), is the *only* way to represent how a rational market participant forecasts outcomes.

In order to illustrate this claim, let $\mathcal{F}_t^j(d_{t+1})$ stand for the time-t values of a participant's j forecast of dividends at time $t+1$. Maintaining the consensus requirement that a model for the stock price should generate quantitative regularities, we can write a non-REH representation as follows:

$$\mathcal{F}_t^j(d_{t+1}) = b_1 x_t \tag{7}$$

where $b_1 \neq b$.

Given that the dividend model in (3) and (4) represents how dividends actually unfold over time, the model represents the error in a participant's j forecast as follows:

$$\begin{aligned} fe(d_{t+1}) &= d_{t+1} - \mathcal{F}_t^j(d_{t+1}) \\ &= (b-b_1)x_t + b(x_{t+1}-x_t) + \varepsilon_{t+1}, \end{aligned} \tag{8}$$

Because

$$E\left[fe(d_{t+1})|x_t\right] = (b-b_1)x_t$$

representing a participant's j forecast with the non-REH specification in (7) implies that his forecasts are biased and are correlated with the information available at time t. Thus, as Lucas pointed out, the non-REH representation in <7> presumes that a participant is grossly irrational, in the sense that he ignores observable, systematic fore cast errors in perpetuity. This implication, which holds on purely logical grounds in *any* consensus model, arguably has buttressed the widespread belief that in order to represent how rational market participants forecast outcomes, an economist *should* rely on REH.

B.4 EMH

EMH's claim that unfettered markets would deliver a nearly perfect allocation of resources rests on the belief in REH's representations of rational forecasting. Although the derivation of the REH model for the stock price is well known, we reproduce it below to highlight the key role of the requirement that the model should represent the price with a probability distribution.

Hypothesizing that the no-arbitrage condition in (1) characterizes how prices are actually set at each point in time, an economist relates the stock price to earnings by iterating this condition forward. Doing so one period ahead results in

$$p_{t+1} = \gamma \left[\mathcal{F}_{t+1} \left(d_{t+2} \right) + \mathcal{F}_{t+1} \left(p_{t+2} \right) \right] \tag{9}$$

Specifying how the stock price actually unfolds over time with a complete stochastic process leads an economist to represent the rational forecasts of prices in each period with a conditional distribution of that process. We can write this formally as follows:

$$\mathcal{F}_{t+i} \left(p_{t+i+1} \right) = E \left(p_{t+i+1} | x_{t+i} \right) \quad \text{for } i = 0, 1, 2, , ,$$

Substituting this representation and that of the market's forecast of dividends into (9) results in

$$p_{t+1} = \gamma \left[E \left(d_{t+2} | x_{t+1} \right) + E \left(p_{t+2} | x_{t+1} \right) \right]$$

Muth's principle enables an economist to represent the time-t market's forecast of p_{t+1} with the conditional expectation of this representation, that is,

$$\begin{aligned} \mathcal{F}_t \left(p_{t+1} \right) &= E \left(p_{t+1} | x_t \right) \\ &= \gamma E \left\{ \left[E \left(d_{t+2} | x_{t+1} \right) + E \left(p_{t+2} | x_{t+1} \right) \right] | x_t \right\} \end{aligned} \tag{10}$$

Again invoking the hypothesis that the model specifies a complete stochastic process for the stock price, an economist uses the law of iterated expectations to represent $\mathcal{F}_t \left(p_{t+1} \right)$ as follows:

$$\mathcal{F}_t \left(p_{t+1} \right) = \gamma \left[E \left(d_{t+2} | x_t \right) + E \left(p_{t+2} | x_t \right) \right]$$

Using this representation and (6) in the no arbitrage condition in (1) and repeating these steps to infinity yields

$$p_t = \sum_{i=1}^{\infty} \gamma^i E \left(d_{t+i} | x_t \right) = \theta x_t \;\; \text{where } \theta = b \frac{\gamma}{1 - \gamma} \tag{11}$$

where $\theta = b \frac{\gamma}{1-\gamma}$.

Because an economist hypothesizes that his specifications for earnings and dividends in (3) and (2) represent how company prospects and payouts actually unfold over time, this REH representation implies EMH.

The price set by the market differs from the perfect-foresight price, p_t^{F}, that would be set by an omniscient central planner by a mean-zero error term:

$$p_t = p_t^{\mathrm{F}} - \eta_t \tag{12}$$

where

$$p_t^{\mathrm{F}} \equiv \sum_{k=1}^{\infty} \gamma^i d_{t+i} \tag{13}$$

and $E\left(\eta_t | x_t\right) = 0$

B.5 REH's Difficulties in Accounting for Stock-Price Fluctuations

In a path-breaking paper, Shiller (1981) pointed out that *any* REH model predicts that stock prices set by the market should fluctuate less than the fundamental value, defined in (13). To see this, let ε_{dt+i} denote the forecast error implied by an REH forecast based on any probabilistic representation of dividends. By definition,

$$\varepsilon_{dt+k} = d_{t+k} - E\left(d_{t+k} | x_t\right)$$

As a result, the model predicts that the perfect foresight and market prices differ by

$$\eta_t = \sum_{k=1}^{\infty} \gamma^i \varepsilon_{dt+k}$$

which implies that $E(\eta_t | x_t) = 0$ and $Cov\left[(\eta_t, p_t) | x_t\right] = 0$. Thus, the relationship in (12) implies that

$$var(p_t^{\mathrm{F}} | x_t) = var(p_t | x_t) + var(\eta_t) > var(p_t | x_t)$$

Shiller pointed out that actual prices vary much more than the perfect-foresight price does. He concluded that the REH-based present-value model for stock-price movements is inconsistent with empirical evidence. As he famously put it: "stock prices move too much to be justified by subsequent changes in dividends."

C Behavioral-Finance Response to REH's Difficulties

Shiller's findings were corroborated by a number of subsequent studies. The profession sought to remedy the apparent difficulties of REH models by supposing that participants' forecasts of prices are driven by psychological considerations that are largely unrelated to fundamental factors. However, behavioral-finance theorists followed their REH predecessors in building models that specify outcomes with a complete stochastic process. As a result, they had to rely on non-REH representations of market participants' forecasts, thereby presuming that participants are grossly irrational.

In order to illustrate how disciplinary consensus has constrained behavioral-finance models, we adopt the simplest formulation: a participant's forecast of p_{t+1} is driven by his forecast of sentiment, denoted by s_t. We formalize this as follows:

$$\mathcal{F}_t\left(p_{t+1}\right) = cs_t \tag{14}$$

Behavioral models represent market sentiment and other non-fundamental factors with a complete stochastic process. In order to focus on the main points, we represent sentiment with a random walk:

$$s_{t+1} = s_t + \xi_{t+1} \tag{15}$$

where $E\left[\xi_{t+1}|s_t, x_t\right] = 0$.

In order to facilitate direct comparison of the REH and behavioral approaches, we also assume that the market sets the price according to the no-arbitrage condition, and that in doing so its forecasts of dividends can be represented with REH.

Using (15) in the behavioral-finance specification of the market's price forecast in (14) yields

$$\mathcal{F}_t\left(p_{t+1}\right) = cs_t \tag{16}$$

Substituting this and the REH representation of the market's forecast of dividends in (6) into a no-arbitrage condition in (1) results in the following behavioral representation of the price set by the market at time t,

$$p_t = \gamma\left[bx_t + cs_t\right]$$

Iterating this forward and using (16) implies the following representation of the forecast error:

$$\begin{aligned} fe &= p_{t+1} - \mathcal{F}_t\left(p_{t+1}\right) = \gamma\left[bx_{t+1} + cs_{t+1}\right] - cs_t \\ &= \gamma bx_{t+1} + (\gamma - 1)cs_t + \gamma c\xi_{t+1} \end{aligned}$$

Thus, behavioral-finance theorists' decision to adhere to the disciplinary consensus has forced them to represent sentiment and other psychological factors in ways that are inconsistent with their own hypotheses about the process that underpins outcomes. As a result, these models presume that a market participant ignores systematic non-zero-mean forecast errors in perpetuity.

C.0.1 Modeling Change with Probabilistic Rules

The foregoing models have specified the dividend process to be unchanging over time. But even if one were to allow for change in the impact of earnings, the disciplinary consensus requires an economist to specify such change with a probabilistic rule ex ante. However, because such representations rule out unforeseeable change, the implications of REH models that allow for change effectively reduce to those that do not.

In order to illustrate this point, we adopt the approach proposed by James Hamilton (1988) and specify b_t to evolve according to the two-state Markov chain:

		b_t	
		$b^{(1)}$	$b^{(2)}$
b_{t+1}	$b^{(1)}$	p	$(1-p)$
	$b^{(2)}$	$(1-q)$	q

where, $b^{(1)}$ and $b^{(2)}$ denote two values that b_t might take in each time period, and p $(1-p)$ and q $(1-q)$ denote the probabilities that b_t will remain unchanged (switch) between any two adjacent periods, t and $t+1$.

Using this probabilistic representation of change in (2) generates the following quantitative regularity about the co-movement of dividends and earnings:

$$E\left[\left(b_{t+1}x_{t+i} + \varepsilon_{t+i}\right)|x_t,\, b_t = b^{(1)}\right] = \left[pb^{(1)} + (1-p)b^{(1)}\right]x_t \quad \text{for } i = 1, 2, , \tag{17}$$

Moreover, an analogous expression holds if $b_t = b^{(2)}$.

This shows that models that represent change with probabilistic rules effectively represent outcomes with a single probability distribution.

D Knight's Distinction Between Risk and "True Uncertainty"

The key implication of the hypothesis that change can be represented with a probabilistic rule is that an REH model represents uncertainty as what Hansen called "risk conditioned on the model." Where $b_t = b$, this risk is typically represented with the variance of the error term in (2), σ_d^2. However, (17) immediately implies that

$$d_{t+1} - E\left(d_{t+1}|x_t,\, b_t = b^{(1)}\right) = \varepsilon_t$$

Thus, even if one allows for change, but represents it with a probabilistic rule, the risk would be the same as that in time-invariant models. According to both models, actual dividends differ from their expected values by ε_t. Thus, as its time-invariant counterpart, the Markov switching model also represents uncertainty with σ_d^2.

Knight argued that such probabilistic risk misses the "true uncertainty" in markets, which arises from unforeseeable change that cannot "by any method be [represented *ex ante*] with an objective, quantitatively determined probability" (Knight, 1921, p. 321).

In order to illustrate how Knightian uncertainty arises from unforeseeable change, suppose that $b_t = b$ until some time period $t = T$, and that at $T+1$ the impact of earnings on dividends undergoes unforeseeable change to some value $b_1 \neq b$. By design, an REH model assumes away such change, and thus generates the following forecast error:

$$d_{T+1} - E\left[\left(bx_{T+1} + \varepsilon_{T+1}\right)|x_T\right] = \left(b_1 - b\right)x_T + \varepsilon_{T+1}$$

In addition to ε_{T+1}, which represents risk, this error has a second component, $\left(b_1 - b\right)x_T$, which represents Knightian uncertainty that arises from change in the process driving outcomes between T and $T+1$. By definition, this change could not have been foreseen at T with a probabilistic rule.

E Aiming to Uncover Qualitative Regularities

Our example represents unforeseeable change by recognizing that at T an economist faces ambiguity about which is the correct model at $T+1$. The example also makes clear that, unless we impose some restrictions on unforeseeable structural change, that is on $(b_1 - b)$, the model does not have empirical content. It is compatible with any qualitative, let alone quantitative, regularity in the co-movements between dividends and earnings between T and $T+1$.

We propose QEH as a way to build models that formalize ambiguity. Hence, we do not impose a probabilistic structure on the change in the model's coefficients over time – and thus we do not specify a complete dynamic stochastic process driving outcomes. However, we constrain change in the model's structure sufficiently to generate predictions of qualitative regularities in time-series data.

F Rethinking the Consensus

We concur with Knight that any quantitative regularity generated by an economist's model eventually becomes inconsistent with how outcomes actually unfold over time. However, although we consider the discovery of quantitative regularities to lie beyond the reach of economic analysis, we do not jettison the scientific objective of searching for regularities that can be confronted with empirical evidence.

We suggest that a more fruitful way forward would be for macroeconomists and finance theorists to engage in the development of models that may help us uncover qualitative regularities. We propose QEH as a way to build such models. QEH's core premise is that recognizing the limits of what we can know about the future is the key step toward resolving long-standing empirical puzzles that consensus models, which assume away these limits, have for decades found difficult to resolve.

In the remainder of our presentation, Anders and Morten will discuss formally how QEH recognizes ambiguity in the context of the stock-price model that we used in this introductory note. They will also sketch the main implications of recognizing the limits to our knowledge about the future for representing the role of fundamental and psychological factors in driving these movements. Finally, they will discuss the considerable challenges for econometric methodology posed by recognizing ambiguity about which quantitative model of the process driving outcomes is the correct one.

References

[1] Fama, Eugene F. (1998), "Market Efficiency, Long-term Returns, and Behavioral Finance." *Journal of Financial Economics* 49(3), 283-306.

[2] Frydman, Roman and Michael D. Goldberg (2007), *Imperfect Knowledge Economics: Exchange Rates and Risk*, Princeton University Press.

[3] Frydman, Roman and Michael D. Goldberg (2011), *Beyond Mechanical Markets: Asset Price Swings, Risk, and the Role of the State*, Princeton University Press.

[4] Frydman, Roman, Michael D. Goldberg, and Nicholas Mangee (2015). "Knightian Uncertainty and Stock Price Movements: Why the REH Present-Value Model Failed Empirically," *Economics: E-Journal*, Vol. 9(24), 1-50.

[5] Roman Frydman, Søren Johansen, Anders Rahbek, and Morten Nyboe Tabor (2017), "The Qualitative Expectations Hypothesis: Model Ambiguity, Consistent Representations of Market Forecasts, and Sentiment," Discussion Paper No. 17-10, University of Copenhagen.

[6] Grossman, Sanford and Joseph E. Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," The American Economic Review, 70, 393-408.

[7] Hansen, Lars Peter (2013), "Uncertainty Outside and Inside Economic Models," *Nobel Lecture:* https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2013/hansen-lecture.pdf.

[8] Knight, Frank H. (1921), Risk, Uncertainty and Profit, Boston: Houghton Mifflin.

[9] Muth, John F. (1961), "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29: 315-335.

[10] Shiller, Robert J. (1981), "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends," American Economic Review, 71, 421-36.