NBER Panel on the Future of Asset Pricing: Models of Beliefs *

Stefan Nagel†
University of Chicago,
NBER, CEPR, and CESIfo

November 2019

*Remarks at the NBER Panel on the Future of Asset Pricing, November 8, 2019 at Stanford University.

†University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, e-mail: stefan.nagel@chicagobooth.edu

Beliefs central to asset pricing. Asset prices are forward-looking and essentially any asset-pricing model boils down to some version of the basic asset-pricing equation

$$P = \mathbb{E}_t[M_{t+1}X_{t+1}] \tag{1}$$

according to which investors price assets based on their beliefs about the joint distribution of some stochastic discount factor (SDF) M_{t+1} and payoffs X_{t+1}. Given this central role of beliefs in asset pricing, an observer outside the field of asset pricing might guess that a major part of the research efforts in asset pricing are devoted to understanding how investors form beliefs. This, at least so far, not the case.

While some work on investor belief formation exists, the vast majority of theoretical and empirical work in asset pricing is based on the rational expectations (RE) paradigm. Under RE, as in Lucas (1978), investors are assumed to know the economy's underlying model, the values of model parameters, and, given this knowledge, to forecast rationally. Within the RE paradigm, there is no role for the study of beliefs: in asset-pricing models, beliefs are implied by the model; empirically, an econometrician can recover investor beliefs from the large-sample empirical distribution of M_{t+1} and X_{t+1}. For example, anytime researchers assume that a sample average of returns approximates investors' ex-ante expectations of returns, they are working within the RE paradigm, even they don't explicitly say so.

RE-based approaches have been useful to establish a benchmark, but asset prices have not yielded easily to RE-based explanations. Seeing the increasingly complicated dynamics in preferences and endowment processes researchers use to reverse-engineer better-fitting RE models, it is fair to wonder whether the self-imposed refusal to treat belief dynamics as an object of theoretical and empirical study is an obstacle to progress. The assumption that investors price assets as if they forecast rationally has some a priori merit in competitive markets, but the assumption that they price as if they know the true model of the world does not.

In these remarks, I sketch the outlines of an asset-pricing research program in which belief dynamics take a central place and that is organized around the following three principles:

- Research focus should be on motivating, building, calibrating, and estimating models with non-RE beliefs rather than on merely rejecting RE models. By now, plenty of evidence exists that beliefs data is not easily reconciled with RE models. To make further progress, we need structural models of belief dynamics that can compete with RE models in explaining asset prices and empirically observed beliefs.

- Deviating from RE does not necessarily imply assuming irrationality. For example, models of Bayesian learning relax the RE assumption that agents know the model of the world and its parameter values, while retaining the rational forecasting assumption. Exploration of cognitive limitations, bounded rationality, and heuristics that relax the rational forecasting assumption may have promising insights to offer as well, but even models of rational learning can produce asset-price properties that are quite different from those in an RE setting.[1]

- Belief dynamics should be disciplined with data on beliefs and micro decisions. While reverse-engineering of preferences and technology to fit asset prices is common in the literature, I would argue that we should not follow this approach with beliefs. Taking beliefs seriously as an object of scientific study also means that the belief dynamics in an asset pricing model should be confronted with data on beliefs and, possibly, also microdata on investor decisions.

In the remainder of these remarks, I discuss a number of open questions that seem promising areas for future research. There is of course already some existing research looking into these questions—and I will mention a few examples—but the point is that much more work is needed.

[1] See, e.g., Timmermann (1993), Lewellen and Shanken (2002), Collin-Dufresne, Johannes, and Lochstoer (2016)

1 Belief measurement

The availability of beliefs data has improved substantially in recent years, but for beliefs data to become a standard ingredient of asset-pricing research, further progress on belief data collection is necessary.

One very basic need is for more data on investor expectations of firm cash flows. Existing work on expectations in asset pricing has often focused on return expectations, but return expectations do not provide a complete picture of how beliefs dynamics induce variation in price levels. For example, one can easily imagine a model in which time-variation in subjective expectations of cash flows lead to big price movements, but subjective risk premia are constant. In this case, return expectations do not reveal the source of price level variation. Forecasts of earnings and dividends by analysts, used in Myers and De La O (2019), and by CFOs in the Graham-Harvey survey are a start, but there is more to be done.

Collecting more data on long-run expectations would be useful. Much of the currently available expectations data is focused on short forecast horizons such as one year. Asset prices, however, depend on expectations over much longer horizons.

In addition to perceived first moments, investors' subjective perception of risk are also of obvious relevance to asset pricing. Data on beliefs about the perceived downside tails of the distribution would be particularly interesting. Since crashes and disasters are infrequent, objective historical data does little to pin down these tails, leaving lots of room for subjective judgements. There has been some progress in eliciting perceived probability distributions in surveys (Manski 2017). Even so, how to elicit beliefs about the shape of distributions is still a challenging problem. Many respondents may be intuitively more familiar from their day-to-day decisions with judgements of what will happen "on average" than with assessments of the "percent chance" of some event. But further research on belief elicitation methods may bring progress in this area, too.

Finally, in many asset-pricing applications, we are interested in dynamics of beliefs at frequencies of the business cycle, or even lower. This means that we need long time series.

The time series available in survey data have lengthened substantially. But there are still potentially big benefits from innovations that can help to extend beliefs data backwards in time with proxies constructed from textual analysis of print media or social media, possibly with machine-learning methods, somewhat akin to what Manela and Moreira (2017), for example, have done to extend the VIX index.

2 Beliefs and actions

If a respondent in a survey states a belief, this does not necessarily mean that the respondent is ready to act in accordance with this stated belief. In addition to a decision-relevant signal, stated beliefs likely contain measurement noise. For example, it is unlikely that respondents deliberate as carefully when stating beliefs as they would if they actually had to take an action. It is also possible that stated beliefs truly reflect respondents' perceptions, but they do not immediately take decisions to bring their actions in line with their beliefs because of actual or cognitive costs of taking an action. To sort out these issues, and to understand how to extract the decision-relevant signal from stated beliefs, more research on the connection of beliefs and actions is needed.[2]

For asset pricing, we also need to understand the properties of this measurement error when we aggregate across respondents. If measurement error averages out within the population, or within demographic groups, then matching asset-pricing models with beliefs moments based on such aggregates may be fine, even if at the individual level the links between stated beliefs and actions are distorted by measurement error. To the extent that more sophisticated agents play a bigger role in financial markets, how the belief-action relationship varies with sophistication is an important issue as well.[3]

Another important question to sort out is whose beliefs matter for pricing. Individual investors' beliefs can differ from those of professional investors. The relative importance

[2] Giglio, Maggiori, Stroebel, and Utkus (2019) is a recent example of work that looks at this question.
[3] See, e.g., D'Acunto, Hoang, Paloviita, and Weber (2019) for recent research of this kind.

of their beliefs in influencing asset prices likely differs by aggregation level: For allocation decisions at the asset class level, e.g. stocks vs. bonds, it seems likely that individuals exert substantial influence because the investment products they choose from often have pre-determined allocations to an asset class; at the individual stock level, fund managers have discretion; at the investment style level (e.g., small vs. large stocks) there is probably a mix of pre-determined choices by individuals and managerial discretion. Sorting this out empirically would also help in linking belief-based approaches with a demand-system analysis as in Koijen and Yogo (2019).

3 Modeling belief formation

Models of investor belief dynamics need to take a stand on the sources of information that investors rely on when they form their beliefs and how they digest this information to produce forecasts. Two questions in this regard seem particularly interesting.

The first is how to model belief formation when investors face high-dimensional environments. Existing asset-pricing models with parameter learning typically feature settings in which investors observe only a small number of predictor variables. Reality, however, looks different. For example, to value stocks, investors must forecast cash flows of stocks. They observe a vast number of potential predictor variables, but they don't know the precise functional relationship between predictors and future cash flows and need to learn it from observed data. The fast-growing literature on machine learning provides methods that are designed for such high-dimensional prediction problems. Modeling investors as machine-learners may therefore be a promising route towards models of belief-formation that capture the high-dimensional nature of investors' problem.[4]

The second concerns the role of memory. Data can shape beliefs only if it is remembered. In theory, one could imagine a Bayesian learner that takes into account "all available" history when learning about pricing-relevant stochastic processes. But when mapping these models

[4] Martin and Nagel (2019) take some first steps in this direction.

into the real world, it is not clear what "all available" means. Some implementations of learning models set time zero to 1926, because this is where the CRSP database starts, but this is obviously not the true starting point of investors' learning process. Moreover, there are reasons to expect that memory could be limited. More research, both empirically and theoretically, is needed to better understand investors' formation of memory, including those of institutions (e.g., through maintained databases or established decision rules).[5]

4 Beyond asset pricing: Macro-finance

The drivers of stock price dynamics emphasized in asset pricing research are largely disconnected from the drivers of the business cycle that macroeconomists focus on (see, e.g., Cochrane 2017). This question should be revisited through the lens of models with non-RE belief dynamics. Shocks to beliefs are potential source of links between asset prices and macro quantities. Exactly how such links could play out is an open question.[6]

Beliefs effects could operate in ways that are quite different from time-varying preferences that macro and asset-pricing research has already explored in various ways. For example, belief effects can be specific to technologies or markets. An individual could be optimistic about the housing market, but, at the same time, pessimistic about the stock market. Beliefs data will be important to sort out the commonalities and differences between different sectors and markets.

Interactions of beliefs with frictions are potentially interesting. For example, belief heterogeneity can interact with frictions in way that amplifies shocks (Caballero and Simsek 2017). The housing market seems to be a particularly interesting area to explore these types of mechanisms as it features substantial frictions and it plays a big role in the macroeconomy.

[5] Collin-Dufresne, Johannes, and Lochstoer (2017), Nagel and Xu (2019), and Wachter and Kahana (2019) are recent examples of such research.

[6] As an example, work by Kozlowski, Veldkamp, and Venkateswaran (2019) suggests that beliefs about disasters could play a role.

5 Conclusion

Asset prices express investors' beliefs about the future. Our understanding of how investors form these beliefs, how they evolve over time, and how we can measure them is still limited. Empirically grounded research on investor beliefs promises to unlock some of the mysteries of asset pricing.

References

Caballero, Ricardo J, and Alp Simsek, 2017, "A Risk-Centric Model of Demand Recessions and Macroprudential Policy", *Quarterly Journal of Economics*, forthcoming.

Cochrane, John H, 2017, "Macro-Finance", *Review of Finance* 21, 945–985.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, 2016, "Parameter Learning in General Equilibrium: The Asset Pricing Implications", *American Economic Review* 106, 664–698.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, 2017, "Asset Pricing When 'This Time is Different'", *Review of Financial Studies* 30, 505–535.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber, 2019, "IQ, Expectations, and Choice", Working paper, National Bureau of Economic Research.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2019, "Five Facts About Beliefs and Portfolios", Working paper, NYU.

Koijen, Ralph S. J., and Motohiro Yogo, 2019, "A Demand System Approach to Asset Pricing", *Journal of Political Economy* 127, 1475–1515.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, 2019, "The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation", *Journal of Political Economy*, forthcoming.

Lewellen, Jonathan, and Jay Shanken, 2002, "Learning, Asset-Pricing Tests and Market Efficiency", *Journal of Finance* 57, 1113–1145.

Lucas, R., 1978, "Asset Prices in an Exchange Economy", *Econometrica* 46, 1429–1445.

Manela, Asaf, and Alan Moreira, 2017, "News Implied Volatility and Disaster Concerns", *Journal of Financial Economics* 123, 137–162.

Manski, Charles F, "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise", in *NBER Macroeconomics Annual 2017*, Martin S. Eichenbaum and Jonathan Parker, eds., Volume 32, 2017).

Martin, Ian, and Stefan Nagel, 2019, "Market Efficiency in the Age of Big Data", Working paper, University of Chicago.

Myers, Sean, and Ricardo De La O, 2019, "Subjective Cash Flows and Discount Rates", Working paper, Stanford University.

Nagel, Stefan, and Zhengyang Xu, 2019, "Asset Pricing with Fading Memory", Working paper, University of Chicago.

Timmermann, Allan G, 1993, "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices", *Quarterly Journal of Economics*, 1135–1145.

Wachter, Jessica A, and Michael J Kahana, 2019, "A Retrieved-Context Theory of Financial Decisions", Working paper, Unversity of Pennsylvania.