

RESEARCH STATEMENT

Kristoffer Nimark

August 2021

This note gives an overview of my past and current research. I start by describing my main current research agenda under the heading *News Media, State-dependent Reporting and Economic Decisions*. This is followed by a description of papers falling under the header *Dynamic Models with Incomplete Information*. Finally, in *Other Research*, I summarize some work that does not fall in either of these categories, including some policy modeling conducted while working at the Reserve Bank of Australia.

News Media, State-dependent Reporting and Economic Decisions

News media monitor the world and decide which events are newsworthy enough to be reported. These reporting decisions are state-dependent, and whether a given event is reported or not depends on what else have occurred. The fact that reporting is state-dependent has at least two general implications. First, the specific notion of newsworthiness that news media use to decide what to report determines what kind of information is made available to economic agents. For instance, if negative or unusual events are considered more newsworthy, then such events will be overrepresented in the news relative to their unconditional probability of occurring. Second, the specific form of state-dependence determines what agents infer about the state of the world from a given piece of information. How these effects impact economic decisions in practice depends on the details and the context under study. In a series of papers, I have investigated both the macro and micro implications of state-dependent reporting decisions by news media.

Unusual events, media reports, and aggregate fluctuations

The title of the paper **Man-bites-dog Business Cycles (American Economic Review, 2014)** refers to the journalistic dictum that “dog-bites-man is not news, but man-bites-dog *is* news”. That is, unusual events are more likely to be considered newsworthy than events that are commonplace.

In this paper I embed an information structure in which news media is more likely to report on the economy when economic developments are unusual in an otherwise standard business cycle model. This single additional feature can explain why we observe (i) occasional large changes in aggregate variables like CPI inflation and GDP growth, without an easily identifiable change in fundamentals of comparable magnitude, (ii) persistent episodes of increased macroeconomic volatility, and (iii) a positive correlation between magnitudes of changes in macroeconomic aggregates and measures of uncertainty, as well as with measures of cross-sectional dispersion of expectations. In the model, all three features are consequences of rational Bayesian agents optimally updating their beliefs in response to man-bites-dog-type signals about aggregate

productivity. On the methodological side, the paper demonstrates how to solve and estimate models of incomplete information in which agents' information sets are time-varying.

Unrepresentative sectoral news can generate aggregate demand shocks

It has been clear from at least since the early 1990s that something akin to “demand shocks” are needed to account for observed business cycle fluctuations, yet the literature has lacked a completely satisfactory theory of what could be the fundamental causes of such shocks. In the paper **Sectoral Media Focus and Aggregate Fluctuations (conditionally accepted, American Economic Review)**, joint with Ryan Chahrour (Boston College) and my former student Stefan Pitschner (Uppsala University), we argue that time-varying sectoral news coverage can generate aggregate demand-like fluctuations of the type observed in the data, even if sectoral TFP shocks are the only source of exogenous variation.

The mechanism in the model works as follows. In a multi-sector economy, firms choose their production capacity in anticipation of demand for their products. Firms trade intermediate inputs with firms in other sectors through an input-output network. Ideally, an individual firm would like to know exactly how much firms in every other sector will produce for the firm to optimally choose its own production capacity. However, firms cannot directly observe the entire cross-section of the economy. Instead, they rely on news media to monitor the economy on their behalf. News media in turn report the most newsworthy sectoral developments. For example, in the late 1990s as well as over the last decade, the booming tech sector received extensive news coverage. Similarly, reporting on the financial sector and the car industry dominated news media during the 2008-2009 financial crisis. While clearly useful and newsworthy, even accurate reports about such sectoral developments provide only a partial picture of the economy. We show that when that partial picture is *unrepresentative* of the economy as a whole, firms over- or underestimate demand for their product. A firm that is overly optimistic about demand for its output hires too much labor relative to how much it would hire if it could observe the entire cross-section of the economy. If firms across different sectors receive the same partial information via news media, over- or under-hiring of labor will be correlated across sectors, hence generating demand shock-like fluctuations in aggregate output and employment that are orthogonal to TFP.

To test directly empirically whether unrepresentative news coverage can generate aggregate fluctuations, we construct a sectoral news-weighted index of economic activity. When this index is above a corresponding unweighted aggregate reference index, the news are unrepresentatively good. When it is below the reference index, news are unrepresentatively bad. According to our theory, the unrepresentativeness index should predict how beliefs deviate from what is justified by fundamentals. We test this prediction using a sign-restricted SVAR approach that does not rely on the structure of the theoretical model. This exercise confirms the prediction of our model, i.e. that unrepresentatively good sectoral news make agents too optimistic relative to what can be justified by fundamentals, leading to an increase in aggregate output.

In this paper we also proposed a conceptually new approach to model incomplete information. Firms in our model receive accurate but partial information from news media and there are no

exogenous noise shocks. What news media report in turn depends deterministically on the cross-section of productivity shocks. Agents' beliefs are thus deterministic function of the cross-section of the real economy and the news selection function used by news media. This noise-free approach allows us to make empirically tight predictions about the relationships between beliefs, developments in the real economy, and observable patterns in news coverage.

State-dependent reporting and coordination with multiple news outlets

The papers described above are concerned with the effect that news reporting has on the fluctuations of macroeconomic aggregates. In the paper **News Media and Delegated Information Choice (Journal of Economic Theory, 2019)**, written together with Stefan Pitschner, we study some of the micro economic implications of state-dependent editorial decisions, and in particular, how state-dependent news reporting affects agents' ability to coordinate in a strategic setting. Like the papers above, it has both an empirical and a theoretical component. For the empirical part, we use a Latent Dirichlet Allocation (LDA) topic model to document three stylized facts of US news coverage. First, different news outlets generally specialize in different topics. Second, the extent of total news coverage allocated to different topics varies over time and depends on what has happened. Third, major events make news coverage more homogenous across different outlets.

The stylized facts outlined above may be unsurprising, but have interesting implications for economic decisions. We propose a theoretical model that can explain these facts and use it to prove several results. First, we show that if agents are constrained in the number of news stories they can read, it is optimal for them to delegate their information choice to news providers that make state-dependent reporting decisions. By delegating the decision of what to get information about to a news organization, agents can reduce their posterior entropy relative to any ex-ante choice of which events to observe. We thus provide a novel justification for the existence of news media and the editorial service they provide.

Second, when the information providers make state-dependent decisions about what to report, these decisions are by themselves informative about the state. For example, on a slow news day, outlets may devote most of their coverage to relatively mundane events. Because readers see no stories about more important events, they can infer that no such more important events have taken place. On the other hand, when an extraordinary event does occur and gets reported, agents cannot rule out the simultaneous occurrence of more mundane events. The more newsworthy a reported event is, the smaller is thus the set of events that can be ruled out. This mechanism can thus explain why uncertainty increases when major news stories break.

Third, we show that if multiple media outlets maximize the utility of their respective readers, agents' preferences and the distributions of events jointly determine the degree to which knowledge about an event is common among agents. This is in contrast to most of the literature in which signals are either private or public by assumption. We show that an event that is reported by all news providers is common knowledge among agents only if it is also considered maximally newsworthy by all providers. The notions of newsworthiness employed by different information

providers and how these affect the state-dependence of news reporting thus determines agents' ability to coordinate in a strategic setting.

Information market structure with attention costs

In the large rational inattention literature following the approach of Sims (2003), agents are modeled as informational yeoman farmers: They produce and consume all the information they need and do so without interacting with markets for information. However, in reality, various markets for information exist. In current work together with Tommaso Denti (Cornell University), tentatively titled **Attention Costs, Economies of Scale and Markets for Information**, we study the structure of such markets in the presence of information gathering and attention costs. If there are fixed costs of gathering information, a firm providing information to multiple agents can potentially lead to pareto improvements. However, this will only be possible if multiple agents find the same kind of information useful. In a contestable equilibrium setting, we derive conditions for when markets will allow for this kind of economies of scale to be exploited. We also show that agent heterogeneity limits the scope for exploiting economies of scale only if there are attention costs of consuming information, and that attention costs may allow information producing firms to make positive profits in equilibrium in spite of free entry for information supplying firms.

Dynamic Models with Incomplete Information

A large part of the uncertainty facing economic agents is about what actions other agents will take. For instance, a firm choosing the price of its product may be uncertain about the price its competitors will set, or a firm in an oligopolistic market may be uncertain about how much productive capacity his competitors will invest in. Introducing private information in models allow us to study this form of strategic uncertainty, but also tend to make solving the models more involved, and particularly so in a dynamic setting. In a series of papers, I have studied strategic uncertainty due to private information in both a macro price setting context and in financial markets. In these papers, I use a solution method that I proposed in a separate paper. Here, I first discuss the substantive contributions of these papers, followed by a short discussion of the methodological paper. The methodological paper is as of yet unpublished. It is also incomplete, in that one of the proofs contains an error that limits what classes of models the results apply to, and I discuss this in more detail below. I end this section by discussing a paper in which I study what central banks can learn from the yield curve for the purpose of optimal monetary policy.

Forward looking price setting, private information, and inflation inertia

In standard New-Keynesian models firms set prices to equal a mark-up over expected marginal cost. The real marginal cost is determined by both exogenous and endogenous factors, where the exogenous factors are assumed to be common (and common knowledge) among all firms. In the paper **Dynamic pricing and imperfect common knowledge** (*Journal of Monetary Economics*,

2008) I relax the assumption of common knowledge of the exogenous factors, by introducing an idiosyncratic component in firms' marginal costs. This helps reconcile two apparently conflicting stylized facts that the standard model cannot account for: Aggregate inflation responds gradually and with inertia to shocks at the same time as price changes of individual goods are quite large. A common but ad hoc fix for the inability of the baseline New-Keynesian model to match the inertia of inflation involved introducing mechanical indexation to past prices. In contrast, inertial behavior of inflation in this paper is driven by optimizing price setters.

The optimal price of an individual good depends positively on a firm's own marginal cost and the price chosen by other firms. However, individual firms cannot observe the marginal cost of other firms and therefore do not know the current price chosen by other firms with certainty. One methodological contribution that was at the time novel, was to show how models with privately informed agents that optimize intertemporally can be solved with an explicit role for higher order expectations.

Bond yields, private information and speculation

A fundamental question in financial economics is what the forces are that can explain the observed variation in asset prices. That expectations about future asset prices need to be an important component in any such explanation is uncontroversial. However, in spite of strong empirical evidence pointing towards economic agents having heterogeneous expectations about future returns, most empirical asset pricing models abstract away from this heterogeneity. One reason for this state of affairs may be that the literature has not yet developed methods to credibly quantify the importance of expectations heterogeneity for asset prices. In two papers, me and my co-author Francisco Barillas (University of New South Wales) aim to help fill this gap.

Many bonds, and U.S. treasury bonds in particular, are traded in very liquid secondary markets. In such a market, the price an individual trader will pay for a long maturity bond depends on how much he or she thinks other traders will pay for the same bond in the future. If traders have access to different information, this price may differ from what an individual trader would be willing to pay for the bond if he or she had to hold it until maturity. The possibility of reselling a bond then changes its equilibrium price as traders take speculative positions in order to exploit their private information. In **Speculation and the Term Structure of Interest Rates (Review of Financial Studies, 2017)** we study the importance of this kind of speculative behavior for bond yields using a simple equilibrium model.

In the model bond prices are pinned down by the portfolio decisions of profit maximizing traders. Because individual traders have access to private information, they can identify bonds that, conditional on their own information sets, have a positive expected excess return. In the absence of arbitrage, expected returns in excess of the risk-free rate must be compensation for risk. Traders will hold more of the bonds with a higher expected return in their portfolios. In equilibrium, the increased riskiness of a less balanced portfolio is exactly offset by the higher expected return. We show formally that heterogeneous information introduces a source of time-varying expected excess

returns that, unlike the excess returns documented by for instance Fama and Bliss (1987) and Campbell and Shiller (1991), cannot be predicted conditional on past bond yields.

When aggregated, the speculative behavior of individual traders introduces new dynamics to bond prices. We demonstrate that when traders have heterogeneous information sets, bond yields are partly determined by a speculative component that reflects traders' expectations about the error in the average, or market, expectations of future risk-free interest rates. Since it is not possible for individual traders to predict the errors that other traders make based on information available to everybody, the speculative component in bond prices must be orthogonal to publicly available information. This feature also makes the speculative component statistically distinct from traditional risk premia, which can be predicted conditional on publicly available information such as bond prices. The speculative term that we identify in the data is thus not simply a relabeling of the classical terms explained by expectations about future risk-free interest rates and risk-premia.

The model in Barillas and Nimark (2017) is a so-called equilibrium model in which we explicitly solve for the optimal portfolio problem of the traders. For this approach to be feasible, we need to make simplifying assumption that imply that there is no time-varying risk premia at the aggregate level. To allow for such variation and to facilitate a comparison of our results to the large existing empirical term structure literature, in **Speculation and the Bond Market: An Empirical No-arbitrage Framework (Management Science 2019)** we propose a more flexible affine no-arbitrage framework for bond pricing. This framework allows for heterogeneous expectations while imposing only a minimum of structure on the data and nests the popular full-information model of Joslin, Singleton and Zhu (2011) as a special case.

Using the proposed model, we first show that heterogeneous information introduces a term in bond yields that is statistically distinct from the classical components of yield curve decompositions, that is, terms due to risk-premia and terms reflecting expectations about future risk-free short rates. In addition to the classic components due to risk premia and expectations about future risk-free short rates, and as in Barillas and Nimark (2017), heterogeneous information introduces a third term due to speculation that is orthogonal to real-time public information such as bond prices. We document that the speculative component in bond yields is quantitatively important, accounting for up to 125 basis points in the early 1990s and up to 100 basis point of yields in the low nominal yield environment of the last decades.

Allowing for heterogeneous information also changes the cyclical properties of risk premia, as compared to a full information model. Risk premia estimated from the model with heterogeneous information is less volatile than and imperfectly correlated with risk premia extracted using the nested full information model of Joslin, Singleton and Zhu (2011). In particular, the heterogeneous information model attributes much less of the high long maturity yields during the period of the Volcker disinflation in the early 1980s to large risk premia than standard models, e.g. Cochrane and Piazzesi (2008). Instead, our model attributes most of the high long maturity bond yields of that period to expectations about future short interest rates.

Together, these results suggest that the standard models that decompose bond yields into terms explained by a common market expectation about future short rates and risk premia may be inadequate to fully account for the variation in bond yields and may lead to biased estimates of these classical yield curve terms.

Solving dynamic rational expectations models with privately informed agents

The models in Nimark (2008), and Barillas and Nimark (2017, 2019) use a solution method developed and analyzed in **Dynamic Higher Order Expectations (unpublished working paper)**. In this paper I propose a recursive method for solving models with privately informed agents. To prove that there exists a unique equilibrium I show that the mapping defined by the filtering and prediction problem of the agents is a contraction on the space relevant for the endogenous processes. I then propose a procedure that iterates on the model's Euler equation, while recursively increasing the number of orders of expectations that are considered. The proposed procedure is of the form that is shown to be a contraction, and can thus be used to find the unique equilibrium dynamics of the model.

Unfortunately, there is an error in the main proof of the current version of the paper. As a consequence of this error, the results described above only apply to models in which all signals are exogenous. This is strong limitation on the class of models the method applies to since it excludes models in which agents observe endogenous variables such as prices. The lack of a (correct) formal proof that the method can always find a unique solution when agents can observe endogenous variables is the main reason why this paper is still unpublished. Having said this, the method appears to work well “in practice” and is still a numerically effective algorithm for solving models with endogenous signals. It has also been applied by other researchers to solve their models, e.g. Graham and Wright (2010), Melosi (2017) and Kohlhas (2019). I still believe that it is possible to find stronger conditions under which the solution can be shown to be unique and I plan to tackle this question again when time permits.

What can central banks learn from the yield curve?

In the paper **Monetary Policy with Signal Extraction from the Bond Market (Journal of Monetary Economics, 2008)** I analyze what a central bank can learn about the state of the economy from the yield curve. The key contribution of the paper is to show within a coherent macroeconomic model how a central bank can extract information from bond yields, while controlling for the fact that, to a large degree, the term structure reflects the financial market's expectations about future monetary policy.

In order to quantify the importance of yield curve information from an optimal monetary policy perspective, I estimate the model parameters on US data. The empirical results suggest that there is some information in short- to medium-term yields that can help Federal Reserve identify both demand and cost-push shocks in a timely manner, while there appears to be little information about

productivity shocks. Long term yields are found to be less informative, which is simply the dual of the fact that standard macroeconomic models are not very good at explaining the movements of long maturity yields (see for instance Gürkaynak et al., 2005).

Other research

Why do agents disagree about verifiable facts?

Many countries have experienced increases in political polarization and in disagreement about objective facts. For example, disagreement in the US about whether climate change is real and caused by human activities has increased and the views on what is essentially an empirical scientific question is well predicted by party affiliation. Political polarization and disagreement about facts clearly has many causes and politically motivated disinformation, is likely to be one of them. Stating a particular belief about a given fact may for many people also be more an expression of group belonging, rather than an expression of a sincerely held belief about the true nature of the world. However, in the paper **Inattention and Belief Polarization (Journal of Economic Theory, 2019)** joint with Savitar Sundaresan (Imperial College London), we show that even ex ante identical, rational agents may self-sort into different informational bubbles, where agents within one group permanently hold beliefs about a fact that are the opposite of the beliefs of the members of the other group.

This result is driven by two effects, both of which are consequences of agents' endogenous information choices. The first effect, which we call the *confirmation effect*, causes agents to choose to observe signals that are more precise in states they believe to be more likely. Those signals are therefore more likely to confirm their prior beliefs. If two agents initially observe different realizations of signals drawn from the same distribution, the confirmation effect then makes it less likely that agents' beliefs will converge over time. The second effect, which we call the *complacency effect*, causes agents to choose less precise channels as their uncertainty decreases. For sufficiently precise beliefs, this effect causes agents to choose completely uninformative signals. Combined, the two effects imply that the beliefs of ex ante identical agents over time will cluster in two distinct groups on opposite ends of the belief space.

We also show that these results hold when agents' information cost is specified in terms of channel capacity, but they do not necessarily hold when costs are proportional to mutual information, or entropy reduction as in the rational inattention models following Sims (2003). To understand why the two cost functions imply different behavior, consider the decision problem of an agent who wants to decide whether a given signal is worth paying attention to or not. When the precision of the agent's prior is higher, the marginal value of observing an additional signal decreases since the agent is already pretty sure about the state of the world. The expected reduction in uncertainty from observing the signal is then small. But if the cost of the signal is measured in terms of entropy reduction, a given signal also becomes cheaper as the prior precision of an agent's beliefs increases.

In the limit with perfectly precise priors, any signal, regardless of its precision, can be observed for free.

On the other hand, if the cost of paying attention to a given signal depends only on the precision of the signal, agents demand less and less precise signals as the precision of their beliefs increase, since the marginal usefulness of the signal then decreases. This implies that when an agent's beliefs are precise enough, he will choose to observe completely uninformative signals and not update his beliefs further. Importantly, an agent may stop updating his beliefs before they become degenerate. In fact, even an agent that attaches a higher probability to the incorrect state than to the correct state of the world may stop updating his beliefs. The beliefs of different agents may then cluster permanently in two distinct groups, where one group is almost certain that one state has occurred and the other group is almost certain that it has not.

How does uncertainty affect the input-output structure of an economy?

In current work, titled **Endogenous Production Networks under Uncertainty (working paper)** and joint with Mathieu Taschereau-Dumouchel (Cornell), Alexandr Kopytov (University of Hong Kong) and Bineet Mishra (Cornell), we construct a macroeconomic model of endogenous network formation to investigate how uncertainty affects firms' sourcing decisions and how, in turn, these decisions affect the macroeconomy. In the model, each firm produces a differentiated good that can be consumed by a representative household, or be used as an intermediate input by another firm. Firms can produce their goods using different techniques that imply combining inputs in different proportions. For each good, there exists an "ideal" technique that maximizes the firms total factor productivity and that the firms would choose if all inputs had the same price and in the absence of uncertainty. However, firms may choose to deviate from this ideal technique if some inputs are cheaper or has less uncertain costs than others.

Each firm is also subject to an exogenous productivity shock, and we assume that production techniques must be chosen before these shocks are realized. Beliefs about firm-level productivities can then influence the choice of technique and, thus, the structure of the production network. For instance, a firm might decide not to adopt a given high-productivity producer as a supplier if the uncertainty around that productivity is also high. While this producer sells its good at a low price on average, it is also more likely to suffer from a large negative productivity shock, in which case its price is going to blow up. Potential customers take this possibility into account and balance concerns about average efficiency and stability when choosing a production technique.

We prove that there always exists an efficient equilibrium in this environment, which allows us to characterize the equilibrium production network through the problem of a social planner that maximizes the utility of the representative household. Building on that insight, we show that the equilibrium production network seeks to strike a balance between maximizing the expected value of GDP on the one hand and minimizing its variance on the other, with the importance of the variance determined by the risk aversion of the representative household.

We then establish how the equilibrium production network reacts to changes in the distribution

of firm-level productivity shocks. Under a mild condition on the set of techniques available, we show that the importance of a firm as a supplier increases if (i) the expected value of its productivity increases, (ii) the variance of its productivity decreases, and (iii) its productivity becomes less correlated with that of its customers. These results imply that changes in perceived uncertainty can lead to changes in the structure of the production network.

One contribution of this paper is to highlight a novel mechanism through which uncertainty can lower expected output. We show that an increase in uncertainty about the firm productivity, through its impact on the structure of the production, leads to a decline in expected GDP. The mechanism is as follows: when uncertainty becomes more important, firms seek stability and, as a result, move away from the most productive (in expectation) suppliers in favor of producers that are less susceptible to risk. In the aggregate, the change in the structure of the production network triggered by this flight to safety translates into a decline in expected GDP and aggregate consumption, but an increase in the representative household's welfare.

A sometimes useful Kalman filter modification

A Low Dimensional Kalman Filter for Systems with Lagged States in the Measurement Equation (Economics Letters, 2015) is a short paper that derives a Kalman filter for systems with lagged observables. The standard way of handling lagged observables is to rewrite the system in the companion form. For most applications, this works well. However, this approach at least doubles the dimension of the state vector, which is unattractive in settings where the state vector is already of high dimension. The method proposed here is thus useful for situations where increasing the dimension of the state is costly.

Policy modeling at the Reserve Bank of Australia

While working at the Reserve Bank of Australia, I wrote two papers on structural models that were to be used in the policy process. In **A Structural Model of Australia as a Small Open Economy (Australian Economic Review, 2009)** I set up a small-scale New Keynesian model and estimate it using Australian data. In **A Medium-Scale New Keynesian Open Economy Model of Australia (Economic Record 2011)** me and my then colleague Jarkko Jaaskela set up and estimated a medium scale New Keynesian model based on the RAMSES model used by the Swedish Riksbank. While using state-of-the-art business cycle models, neither papers broke much new theoretical or methodological ground, perhaps with the exception of the inclusion of a commodities sector in the model of Jaaskela and Nimark (2011). While at the RBA, I also wrote the paper **Combining Multivariate Density Forecasts using Predictive Criteria (Reserve Bank of Australia Discussion Paper 2008-02)** with my then colleague Hugo Gerard. In that paper, we proposed to use predictive criteria, e.g. out of sample forecasting performance, to choose weights when constructing density forecasts by combining multiple models.

References

- Campbell, J. Y., and Shiller, R. J. 1991. Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies* 58:495–514.
- Cochrane, J., and Piazzesi, M. 2008. Decomposing the yield curve. *Mimeo*, University of Chicago.
- Fama, E. F., and Bliss, R. R. 1987. The information in long-maturity forward rates. *American Economic Review* 77:680–92.
- Graham, L. and Wright, S., 2010. Information, heterogeneity and market incompleteness. *Journal of Monetary Economics*, 57(2), pp.164-174.
- Gürkaynak, Refet, S., Brian Sack, and Eric Swanson. 2005. The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models. *American Economic Review*, 95 (1): 425-436.
- Joslin, S., Singleton, K. and Zhu. H. 2011. A new perspective on Gaussian DTSMs. *Review of Financial Studies* 24:926–70.
- Kohlhas, A.N., 2019. Learning by Sharing: Monetary Policy and Common Knowledge. *American Economic Journal: Macroeconomics*.
- Melosi, L., 2017. Signalling effects of monetary policy. *The Review of Economic Studies*, 84(2), pp.853-884.
- Sims, C.A., 2003. Implications of rational inattention. *Journal of Monetary Economics*, 50(3), pp.665-690.