Number 2021-02
February 24, 2021

ECONOMIC COMMENTARY

# Inflation: Drivers and Dynamics 2020 Conference Summary

*Edward S. Knotek II, Robert Rich, Raphael Schoenle, Philippe Andrade, Marco Del Negro, Colin Hottman, Christian Höynck, Matthias Meier, Giovanni Ricco, Elisa Rubbo, Daniel Villar, and Michael Weber**

To provide insights into the processes that drive inflationary dynamics, the Federal Reserve Bank of Cleveland holds an annual conference on the topic of inflation: the *Inflation: Drivers and Dynamics* series. The 2020 installment of the conference was held on May 21-22, 2020. This *Commentary* summarizes the papers at the conference, which broadly fell into four categories: (1) empirical Phillips curves, (2) networks and Phillips curves, (3) expectations formation, and (4) price-setting behavior and inflation.

Inflation is a topic of considerable interest and importance to the public, policymakers, and academics. Researchers continue to investigate new—and, sometimes, old—aspects of the inflation process in an ongoing attempt to better understand its behavior. This effort reflects in part new theoretical models, more sophisticated empirical methods, and richer data sets that not only allow for the exploration of new questions, but also a reevaluation of conclusions drawn from previous work.

Recent research in the realm of inflation has focused on several key issues. One issue receiving particular attention is the stability of the Phillips curve, which suggests that higher inflation is linked to stronger economic activity. Because the observed relationship between inflation and economic activity at the aggregate level can change for a variety of reasons, it is critical to identify the sources of any change in order to, for example, determine if these changes are likely to persist or to assess the implications of the changes for the transmission of monetary policy. Another object of interest are inflation expectations, as it is widely understood that if individuals are forward-looking in their decision-making, then expectations of where inflation is likely to be in the future can affect inflation rates and economic activity in the present. Last, it is instructive to remember that, while inflation is an economy-wide phenomenon, it is ultimately driven by the prices that are actually set for individual goods and services. As a result, closely studying businesses' price-setting behavior may yield insights into this aggregate, economy-wide phenomenon.

*Edward S. Knotek II, Robert Rich, and Raphael Schoenle are at the Federal Reserve Bank of Cleveland. Philippe Andrade is at the Federal Reserve Bank of Boston. Marco Del Negro is at the Federal Reserve Bank of New York. Colin Hottman is at the Board of Governors of the Federal Reserve System. Christian Höynck is at the Bank of Italy. Matthias Meier is at the University of Mannheim. Giovanni Ricco is at the University of Warwick. Elisa Rubbo is at Harvard University. Daniel Villar is at the Board of Governors of the Federal Reserve System. Michael Weber is at the University of Chicago's Booth School of Business. 

*Economic Commentary* is published by the Research Department of the Federal Reserve Bank of Cleveland and is available on the Cleveland Fed's website at www.clevelandfed.org/research. To receive an email when a new *Economic Commentary* is posted, subscribe at www.clevelandfed.org/subscribe-EC.

ISSN 2163-3738 DOI: 10.26509/frbc-ec-202102

To provide insights into the processes that drive inflationary dynamics, the Federal Reserve Bank of Cleveland holds an annual conference on the topic of inflation: the *Inflation: Drivers and Dynamics* series. The 2020 installment of the conference was organized jointly by the Cleveland Fed's Center for Inflation Research and, for the first time, the European Central Bank. This *Commentary* summarizes the papers presented at the latest virtual conference held May 21–22, 2020.

The papers of this conference broadly fell into four categories:

- Empirical Phillips Curves
- Networks and Phillips Curves
- Expectations Formation
- Price Setting Behavior and Inflation

**Empirical Phillips Curves**

The Phillips curve is a central building block of models used by central banks and economists around the world to analyze and forecast inflation. Although the Phillips curve has undergone various incarnations, its central tenet—that periods of lower unemployment are associated with higher levels of inflation and periods of higher unemployment are associated with lower levels of inflation—has not changed. During the past 30 years, however, US inflation has remained relatively stable even though real activity and unemployment have continued to cycle up and down. For example, the Great Recession witnessed the unemployment rate rising to 10 percent but inflation barely dipping below 1 percent, while the recent expansion saw the unemployment rate staying below 5 percent for almost four years but inflation persistently running below 2 percent.

The apparent disconnect between inflation and real activity has called into question the ongoing usefulness of the Phillips curve and prompted a search for possible causes of the weakened relationship over the business cycle. In the paper "What's Up with the Phillips Curve?" Marco Del Negro, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti examine the relative merits of four classes of explanations for the weakened relationship: (i) mis-measurement of either inflation or economic slack; (ii) a flatter wage Phillips curve; (iii) a flatter price Phillips curve; and (iv) a flatter aggregate demand relationship induced by an improvement in the ability of policymakers to stabilize inflation.

The authors use a combination of macroeconometric techniques—vector autoregressions (VARs), structural VARs, and a dynamic stochastic general equilibrium model—to show that changes in the structure of the labor market are unlikely to explain the weaker inflation–real activity nexus because the responses of wages and unit-labor costs to changes in unemployment have been stable over time. They also rule out measurement problems because the pattern of the business cycle's comovement with various popular measures of economic slack—unemployment, output, and hours worked—has not changed dramatically over the past 50 years.

To investigate which of the two remaining hypotheses—a flat price Phillips curve or a flat aggregate demand curve—is most consistent with the evidence, the authors study how the macroeconomic effects of demand shocks changed before and after 1990. In their simplest form, demand shocks are forces that tend to raise output and inflation at the same time, such as a shift in households' desired spending. In the most recent sample, the authors find that demand shocks continue to have an effect on real economic activity, but they now generate a markedly lower response of inflation than in the early sample period. This reduced sensitivity of inflation to demand-induced cost pressures suggests that the slope of the Phillips curve must have fallen after 1990. The important policy implication of this finding is that, with a flatter Phillips curve, a central bank would need to recalibrate its strategy to retain its ability to stabilize inflation.

The issue of a stable Phillips curve is also explored by Giovanni Ricco, Thomas Hasenzagl, Filippo Pellegrino, and Lucrezia Reichlin in their paper "A Model of the Fed's View on Inflation." The authors develop a medium-size semistructural time series model of inflation that is consistent with the view, often expressed by central banks, that three components are important to explain inflation's dynamic behavior: a trend anchored by long-run expectations, a Phillips curve, and temporary fluctuations in energy prices.

The authors find that a stable long-term inflation trend and a well-identified, steep Phillips curve are consistent with the data. Energy prices play two roles: They affect headline inflation not only via the Phillips curve but also via an independent expectations channel. In addition, the data indicate a slowdown in the trend growth rate of output rather than a widening output gap since the start of the new millennium.

Overall, the analysis indicates that the Phillips curve—understood as a relationship connecting nominal variables with real variables and inflation expectations—is alive and well and has been fairly stable since the early 1980s. Importantly, the authors' cycle decomposition shows that the Phillips curve component is not always the dominant force affecting inflation. Large oil price fluctuations can move prices away from the real–nominal relationship both by directly impacting energy services prices and by shifting consumers' expectations—that is, "unanchoring" them—and hence inducing expectations-driven fluctuations in prices. Consequently, the inflation puzzle of the last 10 years can be explained by a high-frequency energy price cycle, which is related to global factors affecting the commodity market and has often overpowered the Phillips curve.

From a policy perspective, the authors argue that the stable inflation trend is an indication of the Fed's success in anchoring expectations. However, their results also point to the challenges that policymakers have to overcome in guiding expectations and stabilizing the economy in the presence of large energy price disturbances.

**Production Networks and the Phillips Curve**

The weaker association between inflation and real activity has attracted considerable interest and a host of possible explanations. One hypothesis for the weaker relationship is that production structures have changed. Firms use a mixture of inputs to build their final products, forming a complex web of input-output linkages in which the outputs of one sector form the inputs to production in another, and changes in this mixture over time can affect the inflation process.

In "Production Networks and the Flattening of the Phillips Curve," Christian Höynck examines how such networks affect the dynamics of inflation in a multisector New Keynesian model with sectors connected through input–output linkages. The model predicts a New Keynesian Phillips curve that is a modified version of the one predicted by the standard one-sector model. In the standard New Keynesian Phillips curve, inflation depends on expected inflation and the output gap, which is a cyclical measure of economic activity. In the production network model, the Phillips curve depends on additional endogenous components; without these components, there is a bias in the estimation of the standard Phillips curve. The size of the bias depends on the network structure. In particular, the introduction of intermediate goods production adds inertia to the inflationary process, and strategic complementarity motives in price-setting reduce the sensitivity of inflation with respect to economic activity. Moreover, the relative importance of sectors in the economy changes compared with the standard New Keynesian model because firms' production fulfills a dual role: as a final consumption good and as an input to other firms' production.

To study the evolution of the Phillips curve over time, Höynck combines the model with historical data on input–output linkages for the US economy. The author finds that the weakened relationship between inflation and economic activity coincides with changes in those interlinkages in the middle of the 1980s. The production network model can account for a reduction in the sensitivity of inflation to economic activity of up to 15 percent. While the share of production of aggregate intermediate goods has stayed relatively constant over time, the main channel for this reduced sensitivity has been an increase in the importance of industries that have more rigid prices, especially in the service sector. This trend increases the aggregate degree of price rigidity in the economy and weakens the relationship between inflation and economic activity. If instead one considers only changes in the importance of sectors for final consumption and ignores the changes in the production network, then one would miss half of this flattening in the Phillips curve.

The role of production networks for the dynamics of the economy and monetary policy is further explored in "Networks, Phillips Curves and Monetary Policy" by Elisa Rubbo. In the conventional one-sector New Keynesian framework, the central bank can fully stabilize productivity shocks by targeting zero inflation. This phenomenon provides the theoretical foundation behind inflation targeting and is often referred to as the "divine coincidence," but it is a knife-edge result. Rubbo notes that the "divine coincidence" does not hold with more realistic production structures. The workhorse one-sector model therefore leaves economists with two important open questions. First, what is the relevant measure of aggregate inflation for monetary policy to target when there are multiple sectors—should the target involve consumer prices, producer prices, or some other measure? Second, how should central banks trade off stabilizing inflation in different sectors, given that full stabilization is impossible in a multisector economy?

This paper revisits the positive and normative implications of the New Keynesian framework in an economy with multiple sectors and a general input–output structure. It derives the Phillips curve in this setting and the welfare function based on key parameters of the production process, such as sectoral consumption, labor and input shares, sectoral probabilities of price adjustment, and elasticities of substitution in production and consumption.

The analysis of the Phillips curve addresses the question of which inflation measure to target. Different measures of aggregate inflation, which correspond to different weightings of sectoral inflation rates, yield different Phillips curves. For example, the consumer price Phillips curve is much flatter than the one-sector model would predict, its slope decreases with the size of intermediate input flows, and its intercept varies over time due to productivity fluctuations. These results indicate that consumer price inflation is not a good indicator of the output gap because the relationship between the two is unstable over time. This issue is common to all measures of aggregate inflation except for one, which Rubbo refers to as the "divine coincidence" inflation index. The divine coincidence index weights sectors according to sales shares (which, as opposed to consumption shares, reflect their full role in the production process), and discounts those with more flexible prices (because inflation in these sectors responds more strongly to any given shock). This index is stabilized if output is stabilized. To validate the theoretical results, Rubbo runs Phillips curve regressions with both consumer prices and the divine coincidence index, finding that the latter provides a much better fit.

Regarding which sector to stabilize, the presence of multiple sectors presents the central bank with a tradeoff between stabilizing aggregate output and implementing the correct relative output across firms and sectors. Unlike in the one-sector case, monetary policy cannot replicate the first-best equilibrium that would prevail under flexible prices. Nonetheless, the constrained optimal policy can still be

implemented by stabilizing an appropriate aggregate inflation index that is distinct from both consumer prices and the divine coincidence index. Calibrating the model to the US economy shows that, in practice, targeting the output gap is close to the constrained optimal policy, while stabilizing consumer prices generates a substantial welfare loss.

**Expectations Formation and Inflation**

For policymakers, the management of inflation expectations plays a critical role in their efforts to stabilize the economy. The importance of this consideration takes on even greater significance when short-term interest rates are at the zero lower bound, as central banks want to avoid a deflationary environment that would increase real interest rates and place additional restraint on aggregate demand. To prevent this occurrence, policymakers have turned their attention to other tools that can be used to further stimulate economic activity, with a particular focus on the household sector because consumers do the majority of spending in the US economy.

Forward guidance on interest rates, which provides a sense of the future path of monetary policy, offers a tool that can stimulate aggregate demand by raising inflation expectations and thereby lowering the real interest rate (Krugman, 1998; Eggertsson and Woodford, 2003; Werning, 2012). In practice, however, the way in which the inflation expectations channel actually works and the strength of this channel remain open questions (see Coibion, Gorodnichenko, Kumar, and Pedemonte, 2020, for a survey). In particular, survey data on households' expectations indicate that households seem to be poorly informed about inflation. They disagree strongly about where inflation is headed, with a significant fraction of respondents reporting expectations that are far outside the range of inflation outcomes observed over previous years (Mankiw, Reis, and Wolfers, 2003). Moreover, the evidence on how households' spending depends on their inflation expectations is mixed, with some results indicating the relationship is insignificant or even negative rather than positive, as theory would predict (Bachmann, Berg, and Sims, 2015). Taken together, there is little consensus on whether households' inflation expectations really matter for their consumption decisions and for the transmission of economic shocks and stabilization policies. In the paper "What Matters in Households' Inflation Expectations?" Philippe Andrade, Erwan Gautier, and Eric Mengus provide new evidence on how the inflation expectations channel operates.

The paper shows that households make consumption decisions based on the broad inflation regime, rather than on their specific expectations for inflation. They obtain this result by studying both qualitative and quantitative answers in surveys of inflation expectations for French, German, and US households. The analysis shows that differences in individual qualitative assessments of the inflation regime are associated with much larger differences in individual consumption choices than quantitative differences across individuals who share the same qualitative assessment. These results are consistent with recent models of rational inattention where agents discretize their choices—that is, they consider only a subset of potential alternatives when making decisions (see e.g., Caplin, Dean, and Leahy, 2018; Jung, Kim, Matějka, and Sims, 2019; as well as Mackowiak, Matějka, and Wiederholt, 2018, for a recent survey).

As the authors illustrate, the results imply that the ability to manage current aggregate demand by manipulating inflation expectations is more limited than suggested by models where consumption reacts continuously to changes in inflation expectations. In particular, boosting aggregate demand through inflation expectations requires convincing households to switch from a "stable prices" regime to a "positive inflation" regime. This finding can help explain why temporary value-added-tax (VAT) increases have a more pronounced expansionary impact on households' consumption than forward guidance on policy rates, as documented in D'Acunto, Hoang, and Weber below. Compared with the latter type of policy, the former type of policy is better understood and more closely followed by households, thereby making it more likely to be effective at changing the expected inflation regime. The paper concludes that there is an active inflation expectations channel, but it is much less potent than in the standard New Keynesian model used for monetary policy analysis. While the channel can help to stabilize aggregate demand, it may act to reduce the effectiveness of certain policies that rely on managing inflation expectations such as forward guidance.

The interaction of household expectations and monetary policy is explored further by Olivier Coibion, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber in "Forward Guidance and Household Expectations." Forward guidance about the path of future interest rates, along with quantitative easing, became one of the main tools that central banks developed during the global financial crisis to address the exceptional circumstances they were facing, because the effective lower bound on interest rates prevented them from stimulating the economy by pushing the policy rate (much) below zero. But because the mechanism of forward guidance is through expectations and these are not readily observable, it has been difficult to establish how policy announcements about the future actually affect individuals in the economy.

In this paper, the authors implement a large-scale randomized control trial (RCT) on a representative sample of about 25,000 US consumers to whom they provide different pieces of information about the evolution of future interest rates as well as about past and current interest rates and inflation. This RCT approach provides a transparent way to assess whether the exogenous provision of information about future interest rates changes households' economic expectations.

The authors document a number of new results from this large-scale experiment. First, prior to any information

treatment, they find that households' knowledge about interest rates is limited. The cross-sectional dispersion in their beliefs about the level of interest rates is as high as previously documented for their beliefs about inflation. Second, the provision of information about the current level of interest rates leads to large revisions in households' expectations about interest rates over short horizons, but much less so over longer horizons. Third, providing information about interest rates at longer horizons (such as two to three years) has relatively small effects on households' expectations over future interest rates. Fourth, households tend to revise their inflation expectations along with their interest rate expectations when provided with information about interest rates, but not in a one-for-one fashion, so that their perceived real interest rates still adjust. Fifth, treatments about recent or future inflation have even larger effects on real interest rates than do most treatments about nominal interest rates. Sixth, households' expectations about unemployment are little changed by the different treatments. Finally, these changes in beliefs about future interest rates and inflation caused by the exogenous provision of information are associated with revisions in whether households perceive that now is a good or bad time to purchase durable goods.

While the zero lower bound limits the ability of the central bank to provide economic stimulus through reductions in short-term interest rates, other developments could further restrict the opportunities for stabilization policy. For example, inflated central-bank balance sheets could lessen the scope for quantitative easing, while high debt-to-GDP ratios could limit the ability of fiscal authorities to cut taxes or increase spending during a downturn. When these channels are no longer operative or are severely constrained, then stimulating consumption requires policies that affect households' expectations.

In "Managing Households' Expectations with Unexpected Policies," Michael Weber, Francesco D'Acunto, and Daniel Hoang investigate the conditions under which policymakers can successfully manage households' inflation expectations and thus influence their consumption, saving, and borrowing decisions directly without the need of transmission through financial intermediaries. They compare the effectiveness of unconventional fiscal policy—the preannouncement of higher future consumption taxes—to forward guidance—the explicit guidance by central banks about the future path of monetary policy rates. Theoretically and in the eyes of policymakers, both policies should stimulate aggregate demand by raising households' inflation expectations, which, based on the consumer Euler equation, should reduce incentives to save and induce households to spend. Unconventional fiscal policy communicates to households a higher path for future prices of consumption goods, leading them to trade-off whether to spend less for consumption today or face a higher cost of consumption in the future. Forward guidance promises to keep policy rates low for a longer time than a conventional policy function would prescribe, thereby stimulating economic activity and raising expected inflation in the future.

The authors use unique individual-level data on a large representative population across several European countries and propose a difference-in-differences design. Although theoretically both policies should raise households' inflation expectations and spending on impact, only unconventional fiscal policy announcements produce these outcomes. Forward guidance announcements do not appear to manage expectations or affect spending plans. In particular, the paper examines the unexpected announcement of a VAT increase in Germany in November 2005, which was implemented in January 2007, along with the response in Germany to the two forward-guidance announcements by former European Central Bank President Mario Draghi around the time of the European sovereign debt crisis. Households reacted homogeneously to the unconventional fiscal policy announcements, whereas no subpopulation reacted to the forward guidance announcement, either on impact or with a delay.

**Price-Setting Behavior and Inflation**

While inflation is an economy-wide phenomenon, it is ultimately driven by the prices that are actually set for individual goods and services. As a result, closely studying businesses' price-setting behavior may yield insights into the aggregate phenomenon of inflation.

The entry and exit of varieties pose challenges for measuring inflation. If consumers value product variety, then the cost of achieving a given level of utility (equivalently, "the cost of living") is lower in a world with a larger product assortment. However, official statistics that measure changes in the cost of living generally do not capture any effect from changes in product variety, and therefore these statistics may be upwardly biased measures of inflation. In "Variety Growth and Measurement Biases in Inflation, Output, and Productivity," Etienne Gagnon, Joseph W. Gruber, Colin J. Hottman, Timothy Park, and Robert J. Vigfusson use barcode-level information on the variety of packaged food and beverage products available to US consumers during the 2000s to study the magnitude and consequences of this potential upward bias. After constructing measures of the cost of living that reflect changes in product variety, they find that adjusting for variety growth leads to economically significant downward corrections to measured price inflation, suggesting that real output and productivity in the food and beverage industries were biased downward.

Matthias Meier and Timo Reinelt in "Monetary Policy, Markup Dispersion, and Aggregate TFP" argue that an important aspect of how monetary policy transmits is missing in previous research. Specifically, the missing aspect is related to the empirical regularity that tighter monetary policy lowers aggregate productivity. In other words, if monetary policy raises short-term interest rates, less

aggregate output (GDP) is produced from a given amount of aggregate inputs (the capital stock and hours worked). Prior explanations for this regularity have mostly been based on research and development (R&D) investment. However, lower R&D investments are unlikely to explain short-run fluctuations of aggregate productivity. The paper proposes a novel explanation why monetary policy affects aggregate productivity that is based on markup dispersion and heterogeneity in price rigidity across firms.

Meier and Reinelt use aggregate and firm-level data from the United States to document three new facts. First, monetary policy shocks increase markup dispersion across firms. Second, monetary policy shocks increase the markups of firms that adjust prices less frequently. Third, firms that adjust prices less frequently have higher markups on average. These facts suggest that different price rigidities across firms can explain why markup dispersion increases in response to monetary policy shocks. In a standard New Keynesian model, higher markup dispersion implies misallocation of inputs across firms, which lowers aggregate productivity. Thus, the paper shows that a sizable share of the aggregate productivity response to monetary policy is coming through this markup dispersion channel.

These empirical findings can be explained by a simple New Keynesian model in which firms differ in how rigid their prices are. Prices of more rigid firms naturally respond less to shocks and at the same time these firms have a precautionary motive to set higher average prices. As a result, markup dispersion increases after monetary policy shocks that lower marginal costs. Quantitatively in the model, monetary policy shocks lower aggregate productivity by about half of the empirical estimate. Misattributing endogenous productivity fluctuations to exogenous technology shocks would imply that the true effects of monetary policy are not being accurately captured.

A major focus of the literature on the effects of monetary policy on the real economy is on the role that price selection plays in the effectiveness of monetary policy. Price selection captures the extent to which the prices furthest from their desired level are most likely to change. Price selection determines the average response of prices to a monetary shock and the degree of monetary non-neutrality. Most of the work on evaluating the strength of the price selection effect has been based on micro-founded models. However, Caballero and Engel (1993) had proposed an alternative, model-free approach to analyze the relationship between price rigidities and non-neutrality. This approach features the probability of a price change as a function of the misalignment with respect to its desired level: the price adjustment hazard function.

In the paper "The Price Adjustment Hazard Function: Evidence from High Inflation Periods," Shaowen Luo and Daniel Villar use the micro data underlying the US consumer price index to estimate this hazard function for the United States. This estimate yields implications for monetary non-neutrality and insights into price rigidity without having to specify or solve a model. While the main argument of the hazard function (the price misalignment) is unobserved, the function is estimated indirectly by matching moments of the (observed) price change distribution implied by the hazard function to those observed in the data.

The paper shows that the comovement between inflation and the skewness of price changes is very informative about the slope of the hazard function, which corresponds to the degree of price selection. Hazard functions with a steep slope imply a strong negative relationship between inflation and skewness, which is not seen in the data. As a result, in the estimated hazard function the probability of price adjustment stays relatively low even for sizeable misalignments. The hazard function therefore implies a high degree of non-neutrality. While a constant hazard function, which features no price selection as in a Calvo-style model, implies the highest degree of non-neutrality in the class of hazard functions, the hazard function estimated in the paper implies about 70 percent as much non-neutrality. This is more than double what is implied by hazard functions consistent with price adjustment (menu) costs. Finally, the paper finds that the estimated hazard function is asymmetric. The probability of a price increase is higher than that of a price decrease of the same magnitude. This could reflect asymmetries in the constraints that firms face and suggest possible asymmetries in the response of inflation to monetary shocks.

**References**

Bachmann, Rüdiger, Tim O. Berg, and Eric Sims. 2015. "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence." *American Economic Journal: Economic Policy*, 7 (1): 1–35. https://doi.org/10.1257/pol.20130292.

Caplin, Andrew, Mark Dean, and John Leahy. 2019. "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice." *The Review of Economic Studies*, 86 (3): 1061–94. https://doi.org/10.1093/restud/rdy037.

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte. 2020. "Inflation Expectations as a Policy Tool?" *Journal of International Economics*, 124 (May): 103297. https://doi.org/10.1016/j.jinteco.2020.103297.

Eggertsson, Gauti B., and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 34 (1): 139–235. https://doi.org/10.1353/eca.2003.0010.

Jung, Junehyuk, Jeong Ho (John) Kim, Filip Matějka, and Christopher A. Sims. 2019. "Discrete Actions in Information-Constrained Decision Problems." *The Review of Economic Studies*, 86 (6): 2643–67. https://doi.org/10.1093/restud/rdz011.

Krugman, Paul R. 1998. "It's Baaack: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity*, 1998 (2): 137. https://doi.org/10.2307/2534694.

Mackowiak, Bartosz Adam, Filip Matějka, and Mirko Wiederholt. 2018. "Survey: Rational Inattention, a Disciplined Behavioral Model." Centre for Economic Policy Research, CEPR Discussion Papers, No. 13243. https://ideas.repec.org/p/cpr/ceprdp/13243.html.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual*, 18 (January): 209–48. https://doi.org/10.1086/ma.18.3585256.

Werning, Iván. 2012. "Managing a Liquidity Trap: Monetary and Fiscal Policy." MIT.