Textual Analysis in Macro-Finance: A Research Agenda

Charles W. Calomiris and Harry Mamaysky

April 2, 2019

1. Overview

Much of the early work on the application of natural language processing tools to financial economics focused on the short-term responses of individual stocks to news (for example, Das and Chen 2007 and Tetlock et al. 2008). Recently, there has been considerable interest in understanding how textual data further our understanding of macro phenomena. For example, Baker et. al (2016) show that increasing economic policy uncertainty, measured from a collection of newspaper articles, forecasts drops in industrial production and employment six to twelve months ahead. Calomiris and Mamaysky (2019a) show that country-level measures derived from news articles forecast twelve-month ahead stock market returns and drawdowns in developed and emerging markets. Furthermore, Calomiris and Mamaysky (2019b) show that forward-looking measures of monetary policy derived from central bank communications forecast twelve-month ahead currency returns, most likely via a risk premium channel.

Our current work builds on Calomiris and Mamaysky (2019a, 2019b). We have preliminary results showing that central bank textual tone (reflected in central bankers' official communications) can itself be forecast up to one-year ahead by country-level news measures. This suggests at least two alternative causal mechanisms for how country-level news flow is related to stock and currency returns: news may forecast the state of the economy, which in turn is reflected in countercyclical central bank policy responses and market risk premia, or news flow may capture variation in stochastic discount factors, which produce changes in stock market returns. If central banks are targeting the VIX or some other aspect of stock market performance, central bank policy reactions may have more to do with variation in risk premia than with variation in economic growth. Indeed, in Calomiris and Mamaysky (2019b), we found that the monetary policies of both the Fed and the ECB have reacted to changes in the VIX.

To better understand the causal links among central bank policy, stock market returns, and the business cycle, we are building our own measures of global central bank tone (the ones we have been using thus far are from a third-part data provider) which will allow us to measure which particular aspects of language best characterize central bank policy. We will then study how central bank tone reacts to country-level and global news flow. We believe this will be useful for identifying causal links among monetary policy, financial markets, and the real economy.

Furthermore, we have begun a systematic analysis of earnings call transcripts, which distinguishes among topical categories of news that is relevant for individual corporations. We are finding that corporate decision makers are keenly focused on political, regulatory, and

monetary-policy developments. We are studying how this information predicts corporate hiring, profitability, and cost of capital, as well as the macroeconomic consequences of these effects.

Our overall research agenda combines textual data from several sources – news stories, corporate earnings calls, and central bank communications – to understand how "soft" information affects risk premia across multiple markets, and what effect this has on macroeconomic conditions.

2. Research Challenges

A particular challenge in working on macro-finance questions using textual data is the sheer volume of available information. While the short-term impact of news stories on individual stock prices is relatively straightforward to measure, understanding the medium- and long-term impact of textual information on macro markets requires dimensionality reduction techniques (such as topic modeling and elastic net regressions) to obtain a parsimonious representation of the data. Combining information from multiple data sources (news, earnings calls, central banks) presents its own set of challenges, some of which can be addressed via ad hoc rules, while others (such as document-industry classification) require machine learning approaches.

3. Data and Computational Requirements

Our news data set (71 million articles, 188 gigabytes) from 1996 to today is obtained from Thomson-Reuters. Our earnings call transcripts data, from SP Global, take up 153 gigabytes. The central bank data is considerably smaller. The overall storage requirement of this project exceeds one terabyte. We analyze these data sets on Columbia Business School's Debian Linux cluster, using the Sun Grid Engine to manage parallel computation tasks. Even when running dozens of jobs at a time, many of our calculations require days to complete.

4. References

Baker, S., N. Bloom, and S. Davis, 2016, "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, 131 (4), 1593–1636.

Calomiris, C. and H. Mamaysky, 2019a, "How news and its context drive risks and returns around the world," *Journal of Financial Economics*, forthcoming.

Calomiris, C. and H. Mamaysky, 2019b, "Monetary policy and exchange rates: Time-varying risk regimes," *NBER working paper no. 25714*.

Das, S. and M, Chen, 2007, "Yahoo! for Amazon: Sentiment extraction from small talk on the web," *Management Science*, 53 (9), 1375–1388.

Tetlock, P., M. Saar-Tsechansky, and S. Macskassy, 2008, "More than words: Quantifying language to measure firms' fundamentals," *Journal of Finance*, 63 (3), 1437–1467.