BIS Working Papers
No 1222

Artificial intelligence and big holdings data: opportunities for central banks

by Xavier Gabaix, Ralph S J Koijen, Robert Richmond and Motohiro Yogo

Monetary and Economic Department

October 2024

JEL classification: C5, G11, G12.

Keywords: Asset prices, central bank policies, artificial intelligence, embeddings.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Foreword

The 23rd BIS Annual Conference took place in Basel, Switzerland, on 28 June 2024. The event brought together a distinguished group of central bank Governors, leading academics and former public officials to exchange views on the theme "Navigating uncharted waters: opportunities and risks for central banks". The papers presented at the conference are released as BIS Working Papers, nos 1222, 1223, 1224 and 1225.

BIS Paper no 150 contains remarks from the closing panel on "Revisiting the last decade of monetary policy", by Michele Bullock (Reserve Bank of Australia), Pablo Hernández de Cos (Bank of Spain), Thomas Jordan (Swiss National Bank) and Sethaput Suthiwartnarueput (Bank of Thailand).

Contents

Artificial intelligence and big holdings data: Opportunities for central banks

Xavier Gabaix Ralph S.J. Koijen Robert J. Richmond Motohiro Yogo*

September 8, 2024

Abstract

Asset demand systems specify the demand of investors for financial assets and the supply of securities by firms. We discuss how realistic models of the asset demand system are essential to assess ex post, and predict ex ante, how central bank policy interventions impact asset prices, the distribution of wealth across households and institutions, and financial stability. Due to the improved availability of big holdings data and advances in modeling techniques, estimating asset demand systems is now a practical reality. We show how demand systems provide improved information for policy decisions (e.g., in the context of financial contagion, convenience yield or the strength of the dollar) or to design optimal policies (e.g., in the context of quantitative easing or designing climate stress tests). We discuss how recent AI methods can be used to improve models of the asset demand system by better measuring asset and investor similarity through so-called embeddings. These embeddings can for instance be used for policy making by central banks to understand the rebalancing channel of asset purchase programs and to measure crowded trades. JEL Codes: C5, G11, G12.

*xgabaix@fas.harvard.edu, ralph.koijen@chicagobooth.edu, rrichmon@stern.nyu.edu, myogo@princeton.edu. For financial support, Gabaix thanks the Ferrante Fund, and Koijen thanks the Center for Research in Security Prices at the University of Chicago and the Fama Research Fund at the University of Chicago Booth School of Business.

1 Introduction

Modern central banks use a rich set of policy tools to achieve their policy objectives. In this short nontechnical companion article to our main conference paper (Gabaix et al. (2024)), we explain how a quantitatively realistic model of the asset demand system is essential to assess ex post, and predict ex ante, how policy interventions impact asset prices, the distribution of wealth across households and institutions, and financial stability.

Asset demand systems specify the demand of investors for financial assets and the supply of securities by firms. By combining the model of the asset demand system with the market clearing condition, we obtain a model of asset prices. Due to the improved availability of big holdings data and advances in modeling techniques, estimating asset demand systems is now a practical reality.

Asset demand systems can be micro-founded by assumptions on investors' preferences, constraints, and beliefs, along with economic fundamentals (e.g., earnings and inflation), as is done in traditional asset pricing and macro-finance models. However, traditional asset pricing models imply an asset demand system that is misspecified in important ways, which limits their applicability for quantitative policy analyses. Specifically, traditional models imply that investors' demand is highly elastic, implying that investors aggressively trade in response to small price deviations. This implication is at odds with a growing empirical literature that documents that investors' demand is much more inelastic. The demand system approach to asset pricing accounts for these new empirical facts by directly estimating the asset demand system based on rich holdings data, asset prices, and economic fundamentals. An important objective for future research is to provide micro foundations that are consistent with those new facts in finance, also to mitigate concerns related to the Lucas critique in policy counterfactuals.

We first discuss in Section 2 how the demand system approach to asset pricing can be used to *provide better information for policy decisions.* One upside of modeling and estimating the asset demand system is that any movement in asset prices can be traced back to shifts in investors' demand curves or firms' supply curves. This is useful to understand the origins of financial contagion across sovereigns during the European sovereign debt crisis, as we will discuss, or to implement the new Transmission Protection Instrument (TPI) of the European Central Bank (ECB). This approach also helps to interpret fluctuations in key asset prices that feed into policy decisions, such as break-even inflation rates at various maturities and the strength of the dollar, or asset prices that enter financial condition indices.

We then discuss how demand system asset pricing can be used to *design optimal policies*, in particular policies that directly affect the (residual) supply of assets available to investors, e.g., when implementing asset purchase programs, or policies that induce shifts in the demand curves of regulated entities, e.g., when implementing environmental risk regulation for banks and insurance companies. We provide examples of models that have been implemented to answer these questions.

The approach to understanding financial markets by modeling the asset demand system is not new and has its roots in the work by Brainard and Tobin (1968) and Friedman (1977), among others. This line of work ran into challenges related to limited or incomplete data on portfolio holdings, tractability in modeling the demand system, and identification of the asset demand system. In recent years, starting with Koijen and Yogo (2019), a new literature emerged that overcomes these challenges using modern modeling techniques of asset prices and demand systems, high quality portfolio holdings data, and improved econometric tools.[1] As a result, demand systems have become a practical reality that, as we view it, are perfectly suited to enrich the toolbox of central banks and financial market regulators to assess and predict the impact of policy interventions.

A quantitatively realistic model of the asset demand system needs to accurately capture investors' demand elasticities and substitution patterns across assets within an asset class and across asset classes. For instance, when central banks purchase government or corporate bonds from investors, a key question when analyzing the rebalancing channel is which assets investors buy instead. Likewise, when assessing the fragility of financial markets and the possibility of crowded trades[2], it is essential to understand the fraction of investors implementing similar investment strategies that can cause excess comovement in prices beyond economic fundamentals.

The guiding thread of Gabaix et al. (2024) is that modern AI models for language, vision, and audio are perfectly suited to measure the similarity across assets and investors. At the core of modern large language models (LLMs) are so-called embeddings, which are numerical representations of words designed to capture their similarity. By mapping words, sentences or entire paragraphs to embeddings, the models can be used to classify text, measure its sentiment, or generate new text in generative AI models.

In the recent research on asset demand systems, and in asset pricing and macro-finance more broadly, it is common practice to capture asset or investor similarity using observed characteristics (e.g., a firm's industry, size, growth rate, et cetera). We instead explore whether asset embeddings can be learned from data instead of specified in advance, building on the success that such a modeling approach has had in the context of AI and machine learning. This is particularly important when a new economic environment calls for new characteristics, for instance, to measure a firm's exposure to COVID-19, its use of intangibles or exposure to AI.

Specifically, we show how to estimate asset and investor embeddings using big holdings data using a model architecture and training method that directly builds on recent language model, thus providing a natural segue between AI and financial economics.

In our current research, we focus on the role of AI and machine learning in modeling and estimating the asset demand system. We also discuss how similar modeling techniques can be used

[1] See for instance Gabaix and Koijen (2024) as a new approach to estimating elasticities, which has been applied in Gabaix and Koijen (2022).

[2] See for instance Khandani and Lo (2011).

to model the demand of consumers for non-financial goods, which can be a key input to modeling and understanding inflation. At the same time, as AI is increasingly viewed as a general purpose technology, it may impact other aspects of central banking that fall outside of the scope of our current research. We will briefly discuss those in Section 4, although an in-depth discussion is well beyond the scope of this paper.

2 A demand system approach to central banking

We first discuss how asset demand systems provide improved information to policymakers in Section 2.1 and how demand systems can help to design optimal policies in Section 2.2. We provide examples based on existing research to illustrate the broader conceptual point that asset demand systems are a valuable tool to modern central banks and financial market regulators. This recent effort has just started and this overview is intended to highlight the breadth of questions that can be answered using this new approach.

2.1 Improved information for policy decisions

2.1.1 Understanding financial contagion

In the demand system approach, we model investors' asset demand (with microfoundations from a portfolio choice problem), match the asset demand to actual data on portfolio holdings, and equate the aggregate demand to the supply of assets by firms and governments to solve for asset prices by market clearing. This implies that for asset prices to have changed, one of the elements of the asset demand system must have changed. As a result, we can use the estimated demand system to trace any movement in prices back to shifts in investors' demand through changes in fundamentals and portfolios flows, and shifts in the supply by firms and governments.

Using this approach, Koijen and Yogo (2020) develop a global asset demand system of short- and long-term bond markets and equity markets across 37 developed and emerging countries. The model can be used to explain fluctuations in these asset prices and exchange rates. We revisit the 2009-10 European sovereign debt crisis to illustrate the model's use for policymakers. During this period, long-term bond yields in vulnerable European countries increased sharply relative to German long-term yields. While there was a deterioration of Greek fundamentals, this was less obvious for other vulnerable countries. This contagious behavior is typically hard to understand and, as a result, contain.

Table 1 reproduces Table 8 from Koijen and Yogo (2020), which estimates a variance decomposition of changes in long-term yields in Greece, Italy, and Portugal into its components: portfolio flows, macro variables, and shifts in demand beyond macro variables (which we call latent demand), which could stand for shifts in beliefs or tastes. Each of the columns sums to one. In the case of

Greece, almost half the variation in yields can be traced back to fundamentals, particularly related to the rating downgrade and spike in equity volatility. However, this is not the case for Italy and Portugal, pointing to an important role for latent demand. The demand system allows for a further breakdown of latent demand by investor geography in the second part of the table. This analysis reveals that European investors played an outsized role in spreading the turmoil from Greece to other peripheral countries.

This analysis is relevant going forward as the ECB recently introduced the Transmission Protection Instrument (TPI) to "counter unwarranted, disorderly market dynamics that pose a serious threat to the transmission of monetary policy across the euro area." A real-time demand system can provide additional information to policymakers to assess the nature of the shocks that cause key spreads to widen.

2.1.2 Understanding key asset prices and implications for financial condition indices

Building on the application to the European sovereign debt crisis, demand systems can also be used to understand the determinants of asset prices that are of relevance to policymakers, such as equity prices, corporate bond prices, break-even inflation (Bahaj et al., 2023), implied volatilities, exchange rates (Koijen and Yogo, 2020), including specifically the strength of the dollar (Jiang et al., 2024b), and even crypto currencies (Benetton and Compiani, 2024).[3]

In the aftermath of the global financial crisis, various regulators, including the Office of Financial Research (2013), expressed concerns that large institutions could amplify volatility in bad times. The policy discussion at the time was whether large asset managers such as Blackrock and Fidelity should be designated and regulated as systemically important financial institutions. We can use demand system asset pricing to answer these questions, by decomposing the relative contributions of institutions and households in explaining the stock market volatility during the global financial crisis.

Figure 2 reproduces Table 4 from Koijen and Yogo (2020), which estimates a variance decomposition of cross-sectional stock returns in 2008. Barclays Bank (now part of Blackrock) was the largest institution in 2007:4, managing \$699 billion. Its assets fell by 41% from 2007:4 to 2008:4. During this period, its contribution to the cross-sectional variance of stock returns was 0.3%. Fidelity contributed an additional 0.9% to the cross-sectional variance of stock returns. Aggregating across the 30 largest institutions, they explain only 4.4% of the cross-sectional variance of stock returns. Smaller institutions explain 40.7%, and the households explain 46.9% of the cross-sectional variance of stock returns. The largest institutions explain a relatively small share of stock market volatility because they tend to be diversified buy-and-hold investors that hold more liquid stocks with a smaller price impact. Based on these results, reducing cross-sectional volatility in bad times

[3]Demand systems can also be used to study imbalances in countries' net foreign asset positions and the returns on these foreign asset positions(Jiang et al., 2024a).

is not a simple matter of restricting trading by the largest institutions.

Many central banks develop and monitor financial condition indices, e.g., the Chicago Fed's National Financial Conditions Index (NFCI) or the Composite Indicator of Systemic Stress (CISS) of the ECB. Among other things, these indices take asset prices as inputs. A demand system can then be used to understand why key asset prices fluctuate, and whether this calls for policy interventions.

2.2 Designing optimal policies

2.2.1 The impact of unconventional monetary policy on asset prices and financial stability

Asset demand systems can also be used to model the impact of asset purchase programs on asset prices and the distribution of risks in financial markets. Koijen et al. (2021) estimate a demand system for sovereign debt markets in the euro area to assess the impact of the ECB's asset purchase program. A key finding is that foreign investors are quite price elastic and sell a disproportionate amount of their holdings. Long-term investors, such as insurance companies and pension funds, instead amplify the purchases of the central banks, consistent with the model in Domanski et al. (2017).

Gabaix and Koijen (2022) develop a dynamic demand system for the aggregate stock market, with extensions to multiple asset classes, which can be used to understand the longer-term implications of purchase programs that operate in multiple markets. Modeling investors' expectations in response to policy interventions is an important and active area of research (Haddad et al., 2023), also in the context of modeling asset demand systems.

Asset demand systems are particularly valuable when multiple central banks implement policy interventions simultaneously, as was the case during the 2008 global financial crisis or the COVID-19 turmoil. While carefully identified event studies may provide guidance when a single country implements a purchase program, such estimates are unlikely to be reliable when multiple central banks act simultaneously. It would be useful to test, out of sample, whether demand systems can be of use in such circumstances.

2.2.2 Currency reserve management

Central banks hold a significant share of developed market debt in foreign exchange reserves. In 2020, total central bank holdings amounted to \$1.025 trillion in foreign short-debt and \$4.952 trillion in foreign long-term debt (Koijen and Yogo, 2020, Table 2). These large portfolios presumably have important effects on exchange rates and bond yields around the world.

An earlier more theoretical literature offers a conceptual framework to understand how flows in foreign exchange markets affect currency values, and, indirectly, GDP and employment (Gabaix and Maggiori (2015); Itskhoki and Mukhin (2021)). But now, progress in methodology has allowed to quantify those effects (Camanho et al., 2022).We can use also use an demand system approach to quantify these effects to understand the importance of foreign exchange reserves in the central bank toolkit.

Using a variance decomposition based on a demand system, Koijen and Yogo (2020, Table 7) find that latent demand of foreign exchange reserves explain 10% of the annual movements in exchange rates. Koijen and Yogo (2020, Table 10) also find that foreign exchange reserves are important for explaining the convenience yield on the US dollar and US Treasuries. Foreign exchange reserves explain 0.94 percentage points of the 1.45 percentage point convenience yield on the US dollar (expressed in units of expected returns on a value-weighted portfolio of foreign short-term debt). In addition, foreign exchange reserves explain 0.28 percentage points of the 1.45 percentage point convenience yield on US Treasuries (expressed in units of yield).

2.2.3 Climate stress tests

Policymakers are increasingly concerned about the impact that climate risks, and the related transition risks, may have on financial markets and regulated institutions, such as banks and insurance companies. These risks can take various forms. For instance, the preferences of investors, consumers or employees may change in favor of green firms or climate-improving interventions. This leads to a repricing of assets and changes in the wealth distribution across investors. Quite mechanically, investors currently holding green assets will benefit at the expense of those holding brown assets. Alternatively, policymakers may impose capital requirements on banks or insurance companies, steering them away from brown firms. This again can lead to a repricing of assets. Koijen et al. (2022) implement such climate stress tests for the cross-section of US equities. This framework can be extended to other asset classes and in particular to fixed income.

3 Using AI to improve models of the asset demand system

3.1 Embeddings to measure similarity between assets and investors

A realistic model of the asset demand system needs to capture the similarity across assets and investors to determine which assets are close substitutes and which investors behave alike.

The traditional approach in financial economics is to use observed characteristics to measure assets' similarity. For instance, in equity markets, stocks of the same size in the same industry with similar growth rates and similar profitability are considered to be close substitutes. In bond markets, securities with similar maturities and ratings are thought to be close substitutes. It is common practice to follow a similar approach to classify investors by, for instance, institutional type (e.g., insurers, hedge funds or retail investors), size, and how active they are (e.g., as measured by turnover or by how much they deviate from holding a market portfolio).

However, these characteristics are far from perfect. As the economic environment changes due to events like the COVID pandemic, the increasing importance of intangible assets, or the emergence of artificial intelligence„ it is common practice to creatively find new characteristics that measure firms' exposures to those topics to determine which firms are similar under current circumstances. Similarly, when we try to explain observed behavior of investors with known characteristics, the explained variation tends to be quite low (Koijen et al., 2022).

To make progress on this important challenge, we take advantage of recent developments in the AI literature that studies language, vision, and audio models. For instance, in the context of language models, a key question is how to capture the semantic similarity between words. Instead of relying on a notion of pre-specified characteristics, as we do in financial economics, the AI literature learns numerical representations of words to capture their similarity. These so-called embeddings are at the core of modern LLMs.

The core idea in Gabaix et al. (2024) is to learn representations of assets and investors: asset embeddings and investor embeddings. We use the same model architecture as those that have been successful in language models. The remaining question is which data to use to estimate those embeddings. While it may be natural to use text data, the signal to noise ratio is fairly low and it may be hard to train or fine-tune the models quarter-by-quarter. Instead, we show that investors' portfolio holdings are ideal data to estimate asset and investor embeddings. Indeed, just as sentences organize words and this information can be used to learn the semantic similarity of words, investors organize stocks in their portfolios and we can use this to learn about asset similarity. By the same logic, by comparing different investors holding the same stock, we can learn about the similarity of investors.

Economically, we can think of asset embeddings as learned representations of the information investors use in forming their portfolios. In a large class of economic models, we can "invert" the asset demand system to learn investors' information set. Depending on the specific model, however, different (linear or nonlinear) techniques need to be used to optimally learn this information. We therefore use different models to estimate the embeddings, building on the model architectures that have been popular in the natural language processing (NLP) literature.[4]

We train the models using a large number of portfolios from mutual funds, exchange traded funds (ETFs), and hedge funds. Generally, more data on investors' portfolios will allow us to obtain better representations of investors and assets. The output of the estimation are asset and investor embeddings, which are useful for policymaking purposes as we discuss in Section 3.2.

To conclude, we discuss the important topic of interpretability. As we obtain numerical (vector-

[4]The models we explore include latent semantic analysis (LSA), which is analogous to principal components analysis (PCA), Word2Vec, and transformer models, specifically the BERT model.

based) representations of assets, how can we attach economic meaning to these learned representations? To this end, we train a retrieval-augmented generative (RAG) model[5] using firms' earnings calls. Put simply, we first find clusters of similar firms. For those firms, we then use a RAG model to summarize the unique risks and growth opportunities. This approach provides a narrative explanation of firms that are close in embedding space.

3.2 Applications of embeddings to central banks and financial stability regulators

3.2.1 Quantifying the rebalancing channel of unconventional monetary policy

A first natural application of asset embeddings is to study the rebalancing channel of monetary policy. By estimating what investors view as similar assets using embeddings, we can potentially improve predictions of what investors purchase instead when selling part of their portfolio to central banks. While the current implementation in Gabaix et al. (2024) is based on U.S. equities, we are currently extending the analysis to also cover corporate bond markets, which is of particular relevance to analyze asset purchase programs.

3.2.2 Identifying crowded trades

While the main focus in Gabaix et al. (2024) is on asset embeddings, we also recover investor embeddings as part of the estimation that capture the similarity of investors. Asset and investor embeddings combined provide a new perspective on crowded trades, which are a source of concern among market participants and financial market regulators alike since the 2007 quant crisis. However, concerns of crowded trades occur frequently, including among technology stocks in the late nineties and again in recent months with the breakthroughs in AI.

Yet, measuring crowded trades in real time, understanding which stocks are involved, and which investors drive the market themes and crowded trades is a nontrivial task. For instance, at the onset of the enthusiasm surrounding AI stocks, investors largely focused on technology companies that build the hardware (e.g., Nvidia) or that deploy these tools at scale (e.g., Apple, Google, Meta, and Microsoft). In the next phase, however, investors realized the huge electricity demands of running large-scale data and model training facilities, and suddenly utilities are part of the AI trade.

[5] RAG models are popular LLM applications as they can enrich pre-trained LLMs with external information such as, in our case, earnings calls. The need for this separate step is because the context windows of LLMs is too limited to feed it all earnings calls at once. Instead, we chunk earnings calls in paragraphs of a certain length (1000 tokens), embed those using a text model, and store them in a vector database alongside metadata (e.g., the firm name, quarter, industry, et cetera). For a given query, such as "what are the common risks for these 10 firms?" , we embed the query and instruct agents to retrieve the closest chunks in the vector database. Those chunks are then (after running a reranking algorithm) fed into the LLM as context to answer the query. In the future, as context windows expand, RAG applications may become less important.

Embeddings can provide a real-time measure of crowded trades. As discussed before, we cluster firms using asset embeddings and potentially provide labels using the RAG model. In a second step, we can assess which investors load on those trades and how large the cluster of similar investors is that takes the same active bets. Using an asset demand system, we can also estimate how much capital is dislocated due to crowded trades and which (groups of) investors are most vulnerable.

This analysis also raises the broader question how crowded trades emerge in the first place. By providing an investor-level breakdown, it may provide valuable inputs into more realistic model of investor beliefs. Also, we can assess how commonly-used risk models, such as those developed by BlackRock (Aladdin) and MSCI (Barra), amplify or mitigate crowded trades.

3.2.3 Designing stress tests

As a third application, we can think of using these models to develop new stress tests. A natural extension of the models developed in Gabaix et al. (2024) is to use the models to predict the next trade of investors, similar to generative AI models that predict the next sentence, paragraph or page of text. By aggregating those trades and imposing market clearing, and using the estimated elasticities of the asset demand system, we can identify stress scenarios that may lead to losses of systemically important financial institutions, if they were to exist.

3.3 Economic applications of embeddings beyond financial markets

We have focused so far on using embeddings to represent assets or investors. However, there are many other applications in economics of embeddings. One application that is of particular relevance to central banks is in modeling inflation dynamics and the dynamics of relative prices. Just like investors form portfolios of financial assets, consumers purchase baskets of goods and services. This can be used to estimate embeddings of goods and services, to understand the dynamics of relative prices, and consumer embeddings to understand which groups of consumers to track to understand the relative price series. Ruiz et al. (2020) provide an early example of such a model.

4 Further discussion

AI is a general purpose technology that is developing at lightning speed, and the risks and opportunities are not yet fully understood. In our research, we explore applications of AI for modeling the asset demand system and to interpret price fluctuations. This implies that we leave many opportunities and risks of AI for central banks unaddressed. We conclude this companion paper with a broad discussion of some key topics that deserve further attention. In the spirit of this section, this list is by no means exhaustive.

4.1 Central bank communication

Forward guidance is a policy tool that has gained in prominence following the 2008 financial crisis. A growing literature parses the speeches and communication of central banks to best extract the information contained in those speeches. As language models improve, such exercises will be more effective. At the same time, central banks respond in their communication strategies as well. A recent literature studies the communication game of central banks, which now has become increasingly complex (Cieslak et al., 2021).

4.2 Cyber risk and regulatory challenges

The rapid developments in AI and machine learning naturally raise regulatory concerns about cyber risk. This poses particular challenges for insurance companies. Like other financial institutions, they may be more prone to cyber attacks as adversaries have more powerful tools at their disposal. It also creates a new opportunity as the demand for cyber insurance increases as cyber attacks become more prevalent. A central challenge will be that the risk distribution is hard to estimate as the technology develops rapidly. A particular concern is that the correlation risk is unknown (e.g., an attack on cloud computing services) and a large literature documents that insurance companies are averse to underwriting risks with an unknown risk distribution (Eling et al., 2024; Koijen and Yogo, 2023).

In addition to the direct cyber risks, AI poses further regulatory challenges. While large incumbent players in the financial sector internalize some of the risks associated with AI due to their large franchise value, this is not true for new entrants. They may therefore experiment with riskier technologies that may or may not be well covered by existing regulations. This concern applies not only to financial stability regulation but also to conduct regulation, where more effective marketing techniques may lead to bias and unequal access to financial services. In the context of insurance markets, for instance, a growing literature addresses the concerns of "inverse selection," defined as a setting in which the insurer has superior information relative to the insured (Abrardi et al., 2020; Brunnermeier et al., 2023).

4.3 Implications for labor markets and the broader macro economy

A fast-growing literature studies the impact of AI on firms, their valuations, labor markets, and the broader economy. This literature naturally studies a broad and important set of questions as AI is increasingly viewed as a general purpose technology that will affect many industries, directly and indirectly. Important contributions to this literature include Acemoglu (2024), Babina et al. (2024a,b), Kogan et al. (2023), and Eisfeldt et al. (2024).

5 Conclusion

In this nontechnical companion article to our main paper on Asset Embeddings, we discuss how modeling the asset demand system can be helpful to modern central banks that use a rich set of policy tools. We also discuss how modern AI tools can be applied to models of the asset demand system, providing a better understanding of investors' substitution patterns, the similarity of assets, and the similarity of investors.

We would like to add a plea that the BIS, and its fellow institutions provide much more holdings data, for the holdings of stocks, bonds, derivatives around the world. That will be very useful for the world, for regulators, central banks, and private actors trying to assess systemic vulnerabilities. The Security Holdings Statistics (SHS) in the Euro area are a good example. The more data that are available to estimate and train the models (perhaps with a limited delay to not impact market participants), the better the models will perform and the more useful those models will ultimately be in guiding policy decisions, and making the financial system, and the world economy, more resilient to shocks.

References

Abrardi, Laura, Colombo Luca, and Tedeschi Piero, "The value of ignoring risk: Competition between better informed insurers," 2020. Working paper.

Acemoglu, Daron, "The Simple Macroeconomics of AI," 2024. Working paper.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," 2024. Journal of Financial Economics, forthcoming.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," 2024. NBER Volume on Technology, Productivity, and Economic Growth, forthcoming.

Bahaj, Saleem, Robert Czech, Sitong Ding, and Ricardo Reis, "The Market for Inflation Risk," 2023. Working paper.

Benetton, Matteo and Giovanni Compiani, "Investors' Beliefs and Cryptocurrency Prices," *The Review of Asset Pricing Studies*, 2024, *14*, 197–236.

Brainard, William C. and James Tobin, "Pitfalls in Financial Model Building," *American Economic Review: Papers and Proceedings*, 1968, *58* (2), 99–122.

Brunnermeier, Markus, Rohit Lamba, and Carlos Segura-Rodriguez, "Inverse Selection," 2023. Working paper.

Camanho, Nelson, Harald Hau, and Hélene Rey, "Global portfolio rebalancing and exchange rates," *The Review of Financial Studies*, 2022, *35* (11), 5228–5274.

Cieslak, Anna, Semyon Malamud, and Andreas Schrimpf, "Optimal Transport of Information," 2021. Working paper.

Domanski, Dietrich, Hyun Song Shin, and Vladyslav Sushko, "The Hunt for Duration: Not Waving but Drowning?," *IMF Economic Review*, 2017, *65* (5), 113–153.

Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska, "The Labor Impact of Generative AI on Firm Values," 2024. Working paper.

Eling, Martin, Anastasia V. Kartasheva, and Dingchen Ning, "The Supply of Cyber Risk Insurance," 2024. Working paper.

Friedman, Benjamin M., "Financial Flow Variables and the Short-Run Determination of Long-Term Interest Rates," *Journal of Political Economy*, 1977, *85* (4), 661–689.

Gabaix, Xavier and Matteo Maggiori, "International liquidity and exchange rate dynamics," *The Quarterly Journal of Economics*, 2015, *130* (3), 1369–1420.

Gabaix, Xavier and Ralph SJ Koijen, "In search of the origins of financial fluctuations: The inelastic markets hypothesis," *NBER Working Paper No. 28967*, 2022.

Gabaix, Xavier and Ralph SJ Koijen, "Granular instrumental variables," Technical Report 7 2024.

Gabaix, Xavier, Ralph S.J. Koijen, Robert J. Richmond, and Motohiro Yogo, "Asset Embeddings," 2024. Working paper.

Haddad, Valentin, Alan Moreira, and Tyler Muir, "Whatever it Takes? The Impact of Conditional Policy Promises," 2023. Working paper.

Itskhoki, Oleg and Dmitry Mukhin, "Exchange rate disconnect in general equilibrium," *Journal of Political Economy*, 2021, *129* (8), 2183–2232.

Jiang, Zhengyang, Robert J. Richmond, and Tony Zhang, "A Portfolio Approach to Global Imbalances," *Journal of Finance*, 2024, *79* (3), 2025–2076.

Jiang, Zhengyang, Robert J. Richmond, and Tony Zhang, "Understanding the Strength of the Dollar," 2024. Working paper.

Khandani, Amir E. and Andrew W. Lo, "What happened to the quants in August 2007? Evidence from factors and transactions data," *Journal of Financial Markets*, 2011, *14* (1), 1–46.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, "Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data," 2023. Working paper.

Koijen, Ralph S. J., François Koulischer, Benoît Nguyen, and Motohiro Yogo, "Inspecting the mechanism of quantitative easing in the euro area," *Journal of Financial Economics*, 2021, *140* (1), 1–20.

Koijen, Ralph SJ and Motohiro Yogo, "A demand system approach to asset pricing," *Journal of Political Economy*, 2019, *127* (4), 1475–1515.

Koijen, Ralph SJ and Motohiro Yogo, "Exchange Rates and Asset Prices in a Global Demand System," 2020. NBER Working Paper 27342.

Koijen, Ralph S.J. and Motohiro Yogo, *Financial Economics of Insurance*, Princeton University Press, 2023.

Koijen, Ralph SJ, Robert J Richmond, and Motohiro Yogo, "Which Investors Matter for Equity Valuations and Expected Returns?," *Working Paper*, 2022.

Office of Financial Research, "Asset Management and Financial Stability," 2013. Manuscript, U.S. Department of the Treasury.

Ruiz, Francisco J. R., Susan Athey, and David M. Blei, "SHOPPER: A Probabilistic Model of Consumer Choice with Substitutes and Complements," *The Annals of Applied Statistics*, 2020, *14* (1), 1–27.

Table 1: Variance Decomposition of Long-Term Yields in the Euro Area

Change in	Greece	Italy	Portugal
Portfolio flows	0.23	0.15	0.24
	(0.11)	(0.18)	(0.05)
Macro variables	0.47	-0.27	0.02
	(0.09)	(0.17)	(0.19)
Latent demand	0.31	1.12	0.73
	(0.03)	(0.31)	(0.21)
Reserves	0.00	0.02	-0.01
	(0.01)	(0.09)	(0.03)
North America	0.01	0.08	0.05
	(0.01)	(0.07)	(0.03)
Europe	0.19	0.96	0.64
	(0.01)	(0.19)	(0.16)
Pacific	0.02	0.06	0.05
	(0.00)	(0.02)	(0.01)
Emerging markets	0.07	0.00	0.00
	(0.02)	(0.01)	(0.00)
Other countries	0.02	0.01	0.01
	(0.01)	(0.00)	(0.00)
Observations	17	17	17

Note: Reproduced from Koijen and Yogo (2020, Table 8). Heteroskedasticity-robust standard errors are reported in parentheses. The annual sample period is 2003 to 2020.

Table 2: Variance Decomposition of Stock Returns in 2008

AUM ranking	Institution	AUM ($ billion)	Change in AUM (%)	% of variance	
	Supply: Shares outstanding, stock characteristics & dividend yield			8.1	(1.0)
1	Barclays Bank	699	-41	0.3	(0.1)
2	Fidelity Management & Research	577	-63	0.9	(0.2)
3	State Street Corporation	547	-37	0.3	(0.0)
4	Vanguard Group	486	-41	0.4	(0.0)
5	AXA Financial	309	-70	0.3	(0.1)
6	Capital World Investors	309	-44	0.1	(0.1)
7	Wellington Management Company	272	-51	0.4	(0.1)
8	Capital Research Global Investors	270	-53	0.1	(0.1)
9	T. Rowe Price Associates	233	-44	-0.2	(0.1)
10	Goldman Sachs & Company	182	-59	0.1	(0.1)
11	Northern Trust Corporation	180	-46	0.1	(0.0)
12	Bank of America Corporation	159	-50	0.0	(0.1)
13	J.P Morgan Chase & Company	153	-51	0.1	(0.1)
14	Deutsche Bank	136	-86	0.3	(0.1)
15	Franklin Resources	135	-60	0.2	(0.1)
16	College Retire Equities	135	-55	0.0	(0.0)
17	Janus Capital Management	134	-53	0.3	(0.1)
18	MSDW & Company	133	45	0.1	(0.1)
19	Amvescap London	110	-42	0.0	(0.1)
20	Dodge & Company	93	-65	0.0	(0.0)
21	UBS Global Asset Management	90	-63	0.0	(0.1)
22	Davis Selected Advisers	87	-54	0.0	(0.0)
23	Neuberger Berman	86	-73	0.0	(0.1)
24	Blackrock Investment Management	86	-69	0.0	(0.0)
25	OppenheimerFunds	83	-64	0.2	(0.1)
26	Wells Fargo & Norwest Corporation	75	-56	0.1	(0.1)
27	MFS Investment Management	73	-44	0.0	(0.0)
28	Putnam Investment Management	73	-76	0.1	(0.1)
29	Marsico Capital Management	73	-56	0.0	(0.0)
30	Lord, Abbett & Company	72	-61	0.3	(0.1)
	Subtotal: 30 largest institutions	6,050	-48	4.4	
	Smaller institutions	6,127	-53	40.7	(2.3)
	Households	6,322	-47	46.9	(2.6)
	Total	18,499	-49	100.0	

Note: Reproduced from Koijen and Yogo (2019, Table 4). The cross-sectional variance of annual stock returns in 2008 is decomposed into supply- and demand-side effects. This table reports the total demand-side effect for each institution due to changes in assets under management (AUM), the coefficients on characteristics, and latent demand. The 30 largest institutions are ranked by AUM in 2007:4. Heteroskedasticity-robust standard errors are reported in parentheses.

Previous volumes in this series

1221 October 2024	Bank geographic diversification and funding stability	Sebastian Doerr
1220 October 2024	The road to net zero: a fund flow investigation	Louisa Chen and Koji Takahashi
1219 October 2024	Stablecoins, money market funds and monetary policy	Iñaki Aldasoro, Giulio Cornelli, Massimo Ferrari Minesso, Leonardo Gambacorta, Maurizio Michael Habib
1218 October 2024	Bank specialisation and corporate innovation	Hans Degryse, Olivier De Jonghe, Leonardo Gambacorta and Cédric Huylebroek
1217 October 2024	Global inflation, inflation expectations and central banks in emerging markets	Ana Aguilar, Rafael Guerra and Berenice Martinez
1216 October 2024	Trade credit and exchange rate risk pass-through	Bryan Hardy, Felipe Saffie, Ina Simonovska
1215 October 2024	CB-LMs: language models for central banking	Leonardo Gambacorta, Byeungchun Kwon, Taejin Park, Pietro Patelli and Sonya Zhu
1214 September 2024	The impact of financial crises on industrial growth: lessons from the last 40 years	Carlos Madeira
1213 September 2024	Chinese Banks and their EMDE Borrowers: Have Their Relationships Changed in Times of Geoeconomic Fragmentation?	Catherine Casanova, Eugenio Cerutti, and Swapan-Kumar Pradhan
1212 September 2024	House price responses to monetary policy surprises: evidence from US listings data	Denis Gorea, Oleksiy Kryvtsov and Marianna Kudlyak
1211 September 2024	Non-bank lending and the transmission of monetary policy	Dominic Cucic and Denis Gorea
1210 September 2024	Which exchange rate matters to global investors?	Kristy Jansen, Hyun Song Shin and Goetz von Peter
1209 September 2024	Latin America's non-linear response to pandemic inflation	Rafael Guerra, Steven Kamin, John Kearns, Christian Upper and Aatman Vakil
1208 September 2024	Generative AI and labour productivity: a field experiment on coding	Leonardo Gambacorta, Han Qiu, Shuo Shan and Daniel M Rees

All volumes are available on our website www.bis.org.