

My research falls in the general area of asset pricing. Within that area, I focus primarily on understanding the impact of information in financial markets. My secondary research theme is international finance and exchange rates.

I. THE PRODUCTION AND DIFFUSION OF INFORMATION IN FINANCIAL MARKETS

Information is a key driver of agents' economic actions. Yet, because it is not directly observable, its effects are difficult to comprehend. My research focuses primarily on understanding the factors that influence the production and diffusion of information in financial markets, and their implications. Using rational choice models and novel data, I establish a link from observable characteristics (e.g. the degree of competition in firms' product markets, their coverage in the media, investors' wealth) to observable outcomes (e.g. trading volume, the cost of capital, portfolio allocations), that so far had been missing not because it is not important, but because of modelling or data limitations.

In a series of theoretical papers, I expand the main paradigm for analysing competitive trading in stock markets, developed by Grossman and Stiglitz (1980). While this framework elegantly captures the tension between the acquisition and transmission of information, it also imposes tight constraints that limit the range of questions that can be addressed – preferences display constant absolute risk aversion (“mean-variance”) and payoffs are normally distributed. Part of my work consists of extending the Grossman-Stiglitz framework to other preference and payoff structures, notably but not only through the use of approximation techniques, in order to shed light on the forces that shape the information environment.

A. Firm characteristics

The following two papers consider the influence of two characteristics of firms on their information environment, namely the degree of competition in their product market and the size of their investor base.

In **“Product Market Competition, Insider Trading and Stock Market Efficiency”**, I link industrial organization to the informational properties of stocks. I analyse how competition in a firm's product market influences its behavior in the equity market. In particular, I examine whether product market imperfections create inefficiencies in the equity market. These questions are all the more relevant given that product markets are becoming increasingly competitive across the world as impediments to trade and barriers to entry are falling.

To answer them, I embed into the Grossman-Stiglitz framework a real sector in which firms operate under monopolistic competition – their equity in contrast trades in perfectly competitive markets, and solve the model thanks to an approximation.¹ The driving force behind my results is that firms use their monopoly power to pass on shocks to customers and insulate their profits. This encourages investors to trade on their private information, expedites its capitalization into stock prices and improves the

¹ The difficulty with market power is that it makes firms' profits and therefore stock prices nonlinear in random variables. Indeed, the Grossman-Stiglitz model is solved by guessing that equilibrium stock prices are linear functions of random variables. Assuming that utility is mean-variance generates stock demands that are linear in expected payoffs and prices, while assuming that random variables are normally distributed leads to expected payoffs linear in signals including prices, thus validating the initial guess. For the same reason, I resort to an approximation when I study the impact of wealth on investors' portfolios because I need to relax the mean-variance assumption.

allocation of capital. Hence product market imperfections, rather than spreading to equity markets, tend to mitigate stock market inefficiencies. A similar argument applies to firms with lower operational or financial leverage since they too offer payoffs that are less sensitive to shocks.

The model yields predictions consistent with existing evidence (for example on the dispersion of analyst forecasts) and implications that I test directly. In particular, I document that trading volume, including trades initiated by insiders, and the information content of stock price are higher for firms with more market power, in line with the model.

In “**The Trade-off between Risk Sharing and Information Production in Financial Markets**”, I ask whether stocks held by more investors are more closely followed. This question is important both for companies and for policy makers who often take actions to promote a broad equity ownership. The answer is clearly yes if each investor produces a fixed amount of information, independent of the number of investors in the company. But as I argue in this paper, it is no longer clear-cut if this amount declines with the number of stockholders. I show that a trade-off exists between the number of informed shareholders and their research effort: the wider the shareholder base, the smaller the risk borne by each shareholder in equilibrium and the less valuable information. So a more widely-held stock can actually be *less* closely followed. This idea echoes the familiar concept from the corporate governance literature which argues that a firm’s ownership structure affects corporate control. However, the information I consider is not used for management monitoring but for portfolio selection. Therefore, it is acquired *ex ante*, i.e. before cash flows are observed, and is revealed through prices.

To establish this result, I extend the Grossman-Stiglitz framework to non-expected utility maximizers. Specifically, I assume that traders have mean-variance preferences of the Kreps-Porteus type. Though expected and non-expected mean-variance utilities lead to the same portfolios, the demand for information under mean-variance expected utility is not influenced by the extent of risk sharing – in particular it depends neither on the supply of shares nor on the number of investors. This non-dependence turns out to be an exception rather than the rule. For example, estimating numerically the demand for information under constant relative risk aversion expected utility, I find that it is indeed a function of risk sharing.

B. Investors’ wealth

The following two papers focus on the role of wealth for investors’ portfolio decisions – by far the strongest determinant of households’ portfolios in the data. I link agents’ wealth to their portfolio decisions through their demand for information.

A pattern observed around the world is that the fraction of wealth households invest in stocks increases with their wealth. A common explanation is that relative risk aversion decreases with wealth. However, abstracting from portfolio data, there is not much evidence in favour of decreasing relative risk aversion. On the contrary, several studies reject this hypothesis using data on attitudes toward risk such as farm, survey or experimental data. In “**Wealth, Information Acquisition and Portfolio Choice**”, I suggest an alternative explanation that relies on investors’ demand for information and only requires *absolute* risk aversion to be decreasing with wealth, an assumption that is supported by all empirical studies.

I solve (with an approximation) a Grossman-Stiglitz economy under general expected utility, thus allowing for wealth effects that the traditional constant absolute risk aversion (mean-variance) utility rules out. Because the value of information increases with the amount to be invested, whereas its cost

does not, wealthier agents acquire more information, perceive stocks as being safer and consequently invest a larger fraction of their wealth in stocks, even though they do not have lower relative risk aversion. I show further how to distinguish empirically the risk aversion story from the information acquisition explanation using data on investors' portfolio returns.

In **“Information vs. Entry Costs: What Explains U.S. Stock Market Evolution”**, I focus on households' decision to participate (or not) in the equity market. The stock market participation rate has increased remarkably over the second half of the century in many countries. In the U.S. for example, it was as low of 6% in 1952 and accelerated throughout the 80's and 90's to reach about half of households today. These observations raise two questions. First, why is stock market participation so low? Second, why does it increase over time? These questions are important not only for our understanding of financial markets, but also for the design of fiscal policies and social security systems.

In this paper, I investigate whether changes in the costs of participating in the stock market can explain the increasing trend in the number of U.S. stockholders. I separate these costs into two components: an information cost (the cost of collecting information about the market) and an entry cost (all other costs, including commissions and fees). Building on my previous work, I disentangle their general equilibrium implications in a Grossman-Stiglitz economy with Kreps-Porteus mean-variance preferences. While both components seem to have fallen in the last decades, I find that the equilibrium implications of their decline are radically different. In particular, a falling information cost cannot explain the observed increase in stock market participation, unlike a falling entry cost. In addition, a falling entry cost accounts for several other features of the U.S. economy, (i) the falling equity premium, (ii) rising return variances and (iii) the boom in passive investing relative to active investing.

C. Implications for long-run economic growth

The following two papers consider the long-run real effects of information. Specifically, they investigate the influence that financial institutions (stock markets and financial intermediaries) exert on economic growth through their information processing role. There is indeed extensive empirical evidence in support of such an influence.

In **“Learning from Stock Prices and Economic Growth”**, I embed a Grossman-Stiglitz stock market into a neoclassical growth economy and evaluate its long-run benefit and cost. On one hand, the partial revelation of private information through stock prices enables investors to share their information truthfully – an essential quality for an information-sharing mechanism when signals are costly to acquire and privately observed. On the other hand, the very existence of informative prices undermines the incentive to collect costly information. I show that the information-sharing benefit outweighs the incentive cost: thanks to the stock market, capital is both more abundant (since less is spent on collecting information) and more efficiently deployed across firms, so that income grows faster. But the positive impact on the growth rate of income is only transitory because the diminishing returns to capital limit the production of information.

The model yields several predictions on the evolution of real and financial variables. In particular, the growth path is characterized by rising capital efficiency, total factor productivity, industrial specialization, stock trading intensity and idiosyncratic stock return volatility. These predictions are consistent with the empirical evidence.

The previous paper focuses on the allocation of capital to exogenous technologies. In “**Learning about Technologies and Technological Innovations**”, I endogenize the process of technological innovation – the outcome of entrepreneurs’ research effort – to investigate how it is influenced by the quality of investors’ information. In contrast to the previous papers, information is incomplete rather than asymmetric, i.e. agents’ information is limited but is equally shared.

The main implication of the model is that knowledge about technologies and technological knowledge feed on each other. That is, entrepreneurs innovate more when investors are better informed about the profitability of their projects because they expect to receive more funding should their projects be successful. Conversely, investors collect more information about projects when entrepreneurs innovate more because the opportunity cost of misinvesting, i.e. of allocating capital to unsuccessful projects, is higher. Per capita income grows and its growth is entirely driven by growth in total factor productivity, consistent with the evidence. I derive new testable predictions and find that they are supported by data on analyst following and R&D expenditures across manufacturing industries and OECD countries.

D. The role of the media

In a group of empirical papers, I study the role of the mass media in transmitting information to investors. Though they expend considerable resources, media outlets should not play any role in frictionless markets. But they could if information frictions are severe, by propagating news to a broad population of investors. In order to assess the impact of the media, I constructed a database that systematically records the newspaper coverage of a large sample of U.S. stocks.

In “**Media Coverage and the Cross-Section of Stock Returns**”, Lily Fang and I study the cross-sectional relation between media coverage and expected stock returns. We find that stocks with no media coverage earn higher returns than stocks with high media coverage, even after controlling for well-known risk factors. This effect is robust, not subsumed by a host of well-documented return anomalies and remarkably stable over time. We also report that it is most pronounced among the stocks that face the most severe information problems such as small stocks and stocks with high individual ownership, low analyst following and high idiosyncratic volatility.

Given publication delays, it is unlikely that the information contained in mass print media is genuine news. Our finding therefore suggests that the dissemination of information matters to stock returns, and that the media play a central role in that process. Our work implies further that firms can reduce their cost of capital through media-relations activities. In recent years, regulation changes in the securities industry and cuts in Wall Street research departments have left many firms without analyst coverage. Our results suggest that the media may offer a substitute or a supplement to traditional channels of corporate information.

The media’s impact may stem from drawing investors’ attention to events or news. A growing body of theoretical and empirical work argues that investors’ lack of attention explains why the market appears to underreact recurrently to corporate news. In “**Media Coverage and Investors’ Attention to Earnings Announcements**”, I use media coverage to proxy for investors’ attention and investigate whether investors’ inattention contributes to the post-earnings announcement drift. I compare announcements made by the same firm in the same year and generating the same earnings surprise, when one announcement is covered in *The Wall Street Journal* while the other is not. I find that announcements with media coverage generate a stronger price and trading volume reaction at the time of the announcement and less subsequent drift. This finding indicates that investors underreact because they are inattentive.

II. INTERNATIONAL FINANCE AND EXCHANGE RATES

My second research theme is in the area of international finance and focuses specifically on exchange rates. In **“Do Demand Curves for Currencies Slope Down? Evidence from the MSCI Global Index Change”**, Harald Hau, Massimo Massa and I use a natural experiment to provide direct empirical support for the traditional portfolio balance theory. The theory derives a downward sloping currency demand function from limited international asset substitutability but enjoyed little empirical support because causal inference is hampered by a lack of clear identification. Indeed, while flows may trigger exchange rate movements, flows may themselves be induced by investors’ trend chasing behaviour.

We examine the exchange rate effect of a major redefinition of the MSCI global equity index in 2001 and 2002. The index redefinition implied large changes in the representation of different countries in the MSCI world index and therefore produced strong exogenous equity flows by index funds. Our event study reveals that countries with a relatively increasing equity representation experienced a relative currency appreciation upon announcement of the index change. Moreover, changes to country weights also modify the permanent correlation structure of exchange rates. Individual stocks have been shown to commove more with an index upon their addition to the index. We find that the same is true of currencies. We show that upweighted (downweighted) currencies tend to comove more (less) with the other currencies in the MSCI index.

In **“Optimal Portfolios of Foreign Currencies”**, I demonstrate with Jamil Baz, Francis Breedon and Vasant Naik how to form portfolios of currencies that benefit from the forward bias and trade off risk and return optimally. We show that a mean-variance analysis applied under the assumption that exchange rates behave as random walks leads to portfolio weights that are stable over time without resorting to exogenous constraints on weights. Optimal currency portfolios invested in the German deutschemark, the Japanese yen, the British pound and the Swiss franc, with the U.S. dollar as the risk-free asset, generate an average excess return of 2.79% per year over the period 1989 through 1999. Portfolios returns have a better Sharpe ratio than Treasury indices and are uncorrelated with major fixed-income and equity indexes. Thus, this methodology provides a useful benchmark for fund managers interested in optimal currency overlays.

February 2009