Statement of Professional Activities

April 16, 2011

Research

My work can be broadly characterized as exploring the impact of risk on equilibrium prices and quantities. My goals are to understand what are the fundamental sources of risk in the economy, and the effect of these risks on the prices of financial assets and on investment decisions.

In my main line of research, I focus on the determinants of asset risk premia. To understand the relationship between various macroeconomic risks and asset prices, it is necessary to understand how these risks affect household investment and consumption decisions, as well as the economic mechanisms that induce heterogeneity in the exposure of individual firms to these aggregate sources of risk.

I approach the above questions using both theory and empirical work. I build structural asset pricing models where the production side of the economy is modeled explicitly. I then explore to what extent these models lead to patterns that are qualitatively and quantitatively consistent with the data. I derive new testable implications of these models and explore whether the empirical evidence supports the models' main mechanisms. My work on technology shocks and organization capital exemplifies this approach.

Relating risk to fundamental asset characteristics has important implications for asset allocation and real investment decisions. In addition, establishing links between asset prices and economic fundamentals furthers our understanding of how the real side of the economy works. Financial assets are the means by which the economy transfers capital from less to more productive sectors. If the prices of financial assets are de-linked from economic fundamentals, this reallocation is impeded. In addition, if firms respond differentially to macroeconomic shocks, the cross-section of asset returns becomes informative about the state of the economy.

In other work, I empirically investigate whether the risk factors affecting the cross-section of asset returns are linked to macroeconomic quantities. In addition, I explore how different types of risk affect investment decisions, namely the impact of liquidity risk on portfolio choice, and idiosyncratic risk on firm investment.

1 Technology shocks

1.1 Investment-shocks

In my work on investment shocks, I explore how a specific type of macroeconomic disturbance, namely changes in the cost of producing new capital goods, affects asset prices. Investment shocks capture the idea that some technological innovations affect output only to the extent that they are implemented through the formation of new capital stock. Thus, investment shocks do not affect the production of the consumption good directly, but alter the tradeoff between current and future consumption.

The recent literature on the real determinants of economic growth [Greenwood, Hercowitz and Krusell, 1997; Fisher, 2006] argues that investment-specific shocks account for a significant fraction of observed growth variability. By now, investment shocks have become a standard feature of dynamic stochastic general equilibrium models [Justiniano, Primiceri and Tambalotti, 2009]. Given that these shocks are important drivers of macroeconomic fluctuations, I explore the implications of investment shocks for asset prices and firm investment decisions in a unified framework.

There are two central ideas behind my work on investment shocks. First, a reduction in the cost of producing new capital, or equivalently an increase in the productivity of the capital-goods sector, induces a reallocation of resources from producing consumption to producing investment goods. Depending on the preference of households over smoothing consumption across time (elasticity of intertemporal substitution) and smoothing across states (risk aversion), states of the world when capital is cheap, (real investment opportunities are good) can be high or low marginal utility states. I argue that the empirical evidence is consistent with the former. The implication for asset pricing is that assets which pay off when capital is cheap are more desirable and thus command a higher price, or equivalently a lower risk premium. Second, a positive investment shock affects differently firms in the cross-section. In particular, it benefits producers of capital goods relative to consumption firms, and within each sector, benefits firms with opportunities to invest relative to those that lack growth opportunities. I develop these ideas in more detail in the following papers.

In *"Investment Shocks and Asset Prices"*, I develop a general equilibrium model that features investment-specific shocks (I-shocks) and show that they are a priced source of risk. To the best of my knowledge, this is one of the first attempts to explore the implications of investment shocks for asset prices. The sign of the risk premium associated with I-shocks depends on household preferences. Under time-separable utility, they have a negative premium or in other words the marginal value of a dollar will be high in these states of the world, because households are willing to give up consumption today in the promise of more consumption tomorrow. In the more general case in which the elasticity of intertemporal substitution (EIS) differs from the reciprocal of the coefficient of relative risk aversion, the sign of the risk premium is ambiguous.

In addition, I show that an investment shock benefits investment firms relative to consumption firms, since the former benefit from the reallocation of resources into investment goods. This is useful, because it suggests the relative stock returns of investment- to consumption-goods producers (IMC) as an empirical measure of investment shocks. This new measure of I-shocks complements the commonly used reciprocal of the relative price of new equipment used by Greenwood et al. (1997) and others, is based on financial data and has the advantage of being available at high frequencies. I verify that my new empirical generates the right responses in quantity variables in the data: investment and hours worked increase, whereas both the relative price of investment and the discretionary component of consumption fall following positive returns on the IMC portfolio.

Using two measures of investment shocks, the relative price of equipment and returns to the investment versus consumption producers, I can estimate the risk premium associated with I-shocks. I find that the empirical estimates of the risk premium on the investment shocks are negative, which is consistent with the model as long as households have at least as strong preferences to smooth across time than across states (EIS lower than or equal to the reciprocal of risk aversion).

Furthermore, I show that investment shocks can generate cross-sectional dispersion in risk premia between firms with different growth opportunities. I find that, for the representative firm, the market value of its existing assets relative to the value of growth opportunities falls following the arrival of a positive investment shock. The former is defined as the value of cashflows accruing from currently installed capital and the latter as the value of future investment opportunities. The intuition is that a positive investment shock is associated with a reduction of the cost of producing new capital, which makes old capital less valuable.

This mechanism opens up the possibility that investment shocks can be behind the well-documented value premium puzzle, namely the empirical finding that value firms have higher rates of return than growth firms even though they have similar measures of systematic risk, when measured using traditional methods. If growth firms derive most of their value from the present value of their growth opportunities and value firms derive most of their value from assets in place, growth firms could have lower risk premia because their price appreciates relative to value firms exactly when real investment opportunities are high.

I explore these ideas further in a series of papers with Leonid Kogan. In *"Growth Opportunities and Technology Shocks"*, we argue that a firm's stock return sensitivity to measures of investment shocks in the data is a measure of a firm's growth opportunities. We build on the insight that firms with different opportunities to invest will have differential exposure to investment shocks. Even though a firm's investment opportunities is not observable, investment shock betas are. Thus, our procedure illustrates how information extracted from the dynamics of asset returns, in particular their comovement with aggregate sources of risk, can be informative about a firm's future investment policy.

We build a simple theoretical model which expands on some of the ideas in "Investment Shocks and Asset Prices". As before, a zero-investment portfolio long the stocks of investment-good producers and short the stocks of consumption-good producers spans the investment shock. Our model delivers endogenous firm sensitivities to investment shocks, in particular a firm's investment shock beta is a linear function of the fraction of firm value due to growth opportunities.

We take the model to the data and classify firms as those with high or low growth opportunities based on their stock return beta with respect to IMC returns. Since growth opportunities are not observable, we assess the success of our procedure indirectly. In particular, our key metric is the response of firms' investment to the I-shock. Intuitively, firms with more growth opportunities should invest relatively more in response to a favorable I-shock, since they have more potential projects to invest in. We find that the IMC-betas identify heterogeneity in firms' investment responses to the I-shocks. In particular, high-β^{IMC} firms not only invest more on average, but their investment increases more in response to a positive I-shock. The economic magnitudes are large: the difference in I-shock sensitivity between the high-beta and the low-beta firms is substantially larger than the sensitivity of the average firm.

Our findings point to a large dispersion in growth opportunities in the cross-section of firms. This heterogeneity has broader implications: when evaluating the impact of macroeconomic disturbances or economic policies on investment, focusing on the behavior of the average firm can be misleading. For instance, in preliminary work titled *"Revisiting the effect of investment tax credit on investment"*, we explore the impact of changes in government subsidies of purchases on new capital using information on the cross-section of firms. These subsidies have historically taken the form of investment tax credits. From the firm's perspective, investment subsidies are similar to an investment shock since they lower the effective cost of purchasing new capital. Given that investment subsidies affect firms with high- and low- growth opportunities differentially, we can use firms with low growth opportunities as a control group to identify the effect of changes in the investment tax credit on investment behavior. We measure high- and low- growth using IMC-beta and Tobin's Q. Our results suggest that, the elimination of the investment tax-credit in 1986-87 resulted in a drop in investment rates of high-growth relative to low-growth firms.

In *"Growth Opportunities, Technology Shocks, and Asset Prices"* we explore whether investment shocks are a plausible mechanism in generating dispersion in risk premia and comovement between high- and low-growth firms. Previous research has documented two stylized facts: high-growth firms earn lower returns on average than low-growth firms; high-growth firms comove excessively with other high-growth firms even though they have similar levels of systematic risk, as measured by market beta.

We develop a dynamic, structural model of firm investment and asset prices that features some of the same ideas as above. Our aim is two-fold: we want to explore whether the model can quantitatively reproduce the observed dispersion in premia between high and low firms; we want to derive additional testable predictions of our main mechanism. Thus, the goal is to produce a model that can help us understand the joint determination of investment behavior, asset prices and cashflow dynamics.

In our model, firms are exposed to an exogenous sequence of neutral and investment-specific shocks. In this model, a positive I-shock corresponds to a decline in the price of new capital goods. In addition, each firm is endowed with a stochastic sequence of investment opportunities which it can implement by purchasing and installing new capital. The present value of the firm's growth opportunities rises in response to a positive neutral shock, while the value of assets in place responds only to the neutral productivity shock. Thus, in the model, high-growth firms exhibit higher stock return betas with respect to IST shocks than low-growth firms. Therefore, heterogeneity in firms' growth opportunities creates cross-sectional differences in risk premia. These differences are not captured by market risk alone, as long as the two types of technology shocks are not perfectly correlated.

In addition, our model generates heterogeneity in expected returns across firms with different book-to-market ratios. This happens because firms with higher growth opportunities tend to have lower book-to-market ratios. Since high-growth firms load more on the IST shocks, our model predicts that stocks with lower book-to-market ratios should have higher IST exposure. Thus, our model generates co-movement of firms with similar book-to-market ratios, giving rise to a value factor. Since, in the model, the value factor is another measure of investment shocks, a direct test of the model is that returns to the value factor should predict investment behavior. Complementing the evidence in "Growth Opportunities and Technology Shocks", we find that the investment rate high- and low- growth firms responds differentially to returns of the value factor.

In work in progress, we further explore the importance of investment shocks in understanding the cross-section of asset returns. In *"Investment shocks, firm characteristics and the cross-section of expected returns"* we extend our investigation into two other well-documented regularities in the cross-section: i) firms with high investment rates (IK) have low average returns, and comove excessively with other high-IK firms ii) firms with higher return on assets (ROA) have higher stock returns on average, and also comove excessively with other high-ROA firms. We argue that both patterns can arise due to dispersion in growth opportunities, which leads to differential sensitivities to investment shocks. Firms with high investment rates have, all else equal, higher growth opportunities and thus should earn lower risk premia. Similarly, high-ROA are firms where assets in place are highly profitable, and thus existing assets constitute a large fraction of firm value. The market value of these firms drops relative to low-ROA firms following a positive investment shock, and thus command higher risk premia.

We extend the model in "Growth Opportunities, Technology Shocks, and Asset Prices" to incorporate uncertainty about a firm's growth opportunities. By observing the arrival of investment opportunities, the firm, and the market, updates it's beliefs about the firm's growth opportunities. As a result, a firm's historical investment behavior is an informative signal about its future investment opportunities, strengthening the link between investment behavior and growth opportunities. Our calibrated model can generate a spread in expected returns between firms between high and low- investment rates and return on assets that matches its empirical counterpart. In addition, we provide empirical evidence that differential exposure to investment shocks drive these differences in risk premia among firms.

In summary, my work on investment shocks provides a unified framework to understand the joint behavior of firm investment decisions, profitability and expected returns as equilibrium outcomes in models where firms are faced with stochastic productivity shocks and investment opportunities. Firm characteristics such as book-to-market, investment rate and return on assets are endogenously related to the firm's growth opportunities, and thus are correlated with the firm's exposure to investment shocks and its risk premium. In addition, my work provides a theoretical justification for the success of empirical factor models. These models exploit the underlying comovement between firms with similar characteristics to 'explain' the cross-section of risk premia. However, without a theoretical understanding of exactly what risk these factors are proxying for, these explanations remain incomplete.

1.2 Technological Innovation

Investment-specific technology shocks model the degree of technological innovation in the sector which produces capital goods. In related work, I explore the role of technological innovation, not necessarily specific to a single sector, in economic growth. My goal is to understand how resources get reallocated from old to new technologies, the cost of this reallocation, and the implications for quantities and prices, both in the aggregate and in the cross-section. In work in progress with Leonid Kogan, Amit Seru, and Noah Stoffman titled *"Technological innovation and growth"*, we take a first pass at answering these questions.

We construct a new measure of innovation at the firm-level using stock returns and a new dataset on patents. Measuring innovation by a simple patent count is problematic, since not all patents are likely to be of the same value. Our idea is to weigh each patent by the positive part of firm's own stock return on the day the patent was applied or issued. The benefit of our quality adjustment is that it allows to investigate impulse responses since it does not use ex-post outcomes, such as citations by future patents. Nevertheless, our measure of quality predicts future citations received.

We find that firms that innovate according to our measure have subsequently higher productivity. In addition, these firms subsequently increase investment and labor hiring, they issue debt and equity and reduce payout to shareholders. In addition, if the firm's competitors innovate, while the firm doesn't, we find an opposite effect. Thus, within each industry, inputs get reallocated to the firms that innovate from those that do not.

At the aggregate level, the impulse response to our measure is U-shaped: higher technological innovation is associated with temporarily lower consumption and output growth. Subsequently, both consumption and output increase. We investigate the long-run effects of innovation by looking at the cross-section of industries: at low frequencies, industries that innovated in the past tend to grow faster.

We develop a tractable, general equilibrium framework to analyze the effect of innovation on asset prices. We model technological innovation shocks as changes in the arrival rate of new production possibilities. Following a positive innovation shock, the demand for capital goods increases, thus triggering a reallocation of resources away from consumption and into producing new capital goods. Thus, technological innovation shocks represent a systematic source of risk that is similar to investment shocks, in that a positive shock may be associated with high marginal utility states as the economy trades off current and future consumption. The model delivers implications consistent with our empirical findings, along with some new ones: an increase in the rate of innovation leads to a fall in aggregate Tobin's Q, because the market value of existing firms drops. Our empirical results support this prediction: our aggregate measure of innovation is negatively correlated with changes in aggregate Tobin's Q.

In *Technological innovation and the cross-section of returns* we extend our analysis to the cross-section of asset returns. A positive shock to the aggregate rate of innovation activity benefits high-growth firms relative to low-growth firms, since the former tend to innovate more. If high innovation is associated with high marginal utility, this is another mechanism that would generate cross-sectional differences in risk-premia between high- and low-growth firms. Our empirical evidence is consistent: growth firms are more likely to innovate than value firms; and growth firms have higher sensitivity than value firms to aggregate innovation shocks.

2 Organization Capital

In *"Organization Capital and the Cross-Section of Expected Returns"*, joint with Andrea Eisfeldt, I explore the risk characteristics of organization capital. I argue that organization capital is exposed to additional sources of risk than physical capital, because shareholders do not necessarily appropriate all the benefits accruing from it.

We define organization capital (OC) as the accumulated stock of organization "know-how", and model it as a factor of production that is embodied in labor. Part of this accumulated knowledge is firm-specific, whereas part of it can be transferred into other firms. Consequently, the owners of organization capital receive payments that are determined not by its marginal product, but rather by its outside option. The outside option of key talent varies in a systematic way, as it depends on the productivity of organization capital deployed in new firms, a form of technological innovation. As I argue in my other work, states when technological innovation is high are high marginal valuation states. As a result, shareholders in firms with high organization capital are exposed to additional risk, and thus they demand a higher risk premium.

We explore the model's implication for risk premia and investment in organization capital. Following Lev and Radhakrishnan (2004), we measure organization capital by accumulating firms' Selling, General, and Administrative (SG&A) expenses. Firms that have higher organization capital, according to our measure, are more productive, score higher in managerial quality surveys, and are more likely to list 'loss of key personnel' as a risk factor in their 10K filings. Our model's asset pricing predictions are supported by the data: firms with more organization capital relative to their industry peers have higher risk premia than firms with less organization capital by 4.8% per year; stock returns of high-OC firms relative to their low-OC peers are negative correlated with innovations in aggregate measures of reallocation and executive compensation.

In addition, our model implies that investment in organization capital should be high when either the firm-level, or the frontier, productivity of organization capital is high. We find that, consistent with the first prediction, investment in organization capital, as measured by SG&A expenditures, is increasing in firm's Tobin's Q and measures of profitability. Regarding the second prediction, we use our model's implications and proxy for the level of the frontier technology with the accumulated return on the high minus low organization capital portfolio. We find that investment in organization capital is high when the accumulated return on the high minus low organization capital portfolio is low.

3 Other work

3.1 Sources of risk

My paper with Igor Makarov, titled *"Sources of Systematic Risk"* is an attempt to link asset returns to macroeconomic fluctuations. Financial assets are a claim to future output. If the economy is subject to multiple shocks, and not all assets load similarly to these shocks, then the cross-section of asset returns should be a valuable source of information about the state of the economy. The major problem is to separate structural shocks from noise. Factor analysis partially achieves this goal, but it suffers from a major drawback. Under the standard assumption of a constant factor covariance matrix, one cannot identify individual factors, but only their linear span. This lack of identification makes interpretation of the corresponding economic shocks difficult

Under the assumption that the underlying macroeconomic factors are heteroscedastic, one can separately identify them rather than their linear span. The idea of identification through heteroscedasticity is not new in econometrics [Sentana and Fiorentini, 2001]. Our methodological contribution is to devise a tractable, two-step method of moments estimator that can be speedily implemented even when the number of factors is large.

Using our econometric specification, we recover four factors from the cross-section of asset returns. In addition to the market portfolio, our procedure recovers a portfolio that is long the capital-goods and short the consumption-goods sector, a portfolio long the energy producers and short the rest of the market, and a portfolio that loads differently on durable and non-durable goods industries. Our main finding is that our extracted factors are robust predictors of both financial and macroeconomic variables. We perform the analysis by using data from both the United States and the United Kingdom and show that our procedure identifies essentially the same factors in the UK as those identified in the US. Since we do not a priori impose any economic structure on our factors, the fact that the extracted factors are similar across different countries suggests that these factors might be identifying structural economic shocks.

3.2 Risk and Investment

In *"Investment, Idiosyncratic Risk, and Ownership"*, joint with Vasia Panousi, I empirically explore the effect of a firm's idiosyncratic risk on its investment policy. In a frictionless capital market, where management and shareholder interests are aligned, idiosyncratic risk should have no effect on a firm's investment decisions. However, as a result of an optimal contract or incomplete markets, the agents who make investment decisions may be forced to bear some idiosyncratic risk. Consequently, since investment decisions are undertaken by managers on behalf of shareholders, managers may underinvest in high idiosyncratic risk projects because these projects expose managers to risk that they cannot diversify. From a diversified shareholder's perspective, such behavior represents a welfare loss, since the manager could be turning down potentially positive NPV projects.

We find that an increase in a firm's idiosyncratic risk is associated with lower future investment. This effect is stronger for firms where insiders (managers) have a higher ownership stake, which is consistent with a managerial risk aversion story. Furthermore, conditional on fraction of insider ownership this effect is weaker for firms with more convex compensation schemes, suggesting that the risk-shifting incentives induced by executive options may undo some of the effects of risk aversion. Finally, we find that the impact of insider ownership on the investment-risk relationship is weaker for firms with high levels of institutional ownership. A potential explanation is that institutional investors are more effective monitors than individual shareholders, since large concentrated ownership stakes can overcome the free-rider problem.

In *"Portfolio Choice with Illiquid Assets"*, joint with Andrew Ang and Mark Westerfield, we theoretically explore the impact of illiquidity risk on the investor's portfolio choice problem. We define illiquidity as the inability to trade continuously. We extend the standard Merton model to incorporate stochastic trading opportunities for a subset of assets, which coincide with the arrival of a Poisson process. In our view, this modeling choice captures two key features of real-world markets. First, certain securities are only periodically tradeable, as opposed to always tradeable at a cost reflecting transactions fees or price impact. This illiquidity can arise as a result of a limited number of market participants, possibly due to the specialized skills or systems needed to trade these assets. Second, the waiting time before a counterparty is found is random. This represents an additional source of trading risk because investors cannot anticipate their next opportunity to trade.

Our model offers a number of insights. First, the investor's illiquid wealth can drop sharply between rebalancing times. This is economically similar to a situation where the asset price follows a jump diffusion process and naturally reduces the incentive to buy illiquid assets. Second, consumption is financed through the investor's liquid wealth. Irrespective of his total wealth, if the investor's liquid wealth drops to zero he will not be able to finance consumption until the next opportunity to trade. This reduces the incentive to buy all risky assets, liquid or illiquid, and induces the investor to hold more cash. Third, the strength of the two effects above depend on the fraction of total wealth allocated to illiquid assets. As a result, the investor's effective risk aversion varies over time, which naturally leads to time-varying investment and consumption policies even in an environment where returns are independently and identically distributed.

3.3 Financing innovation

My previous work on technological innovation takes its output as an exogenous shock. In recent work with Jiro Kondo, titled *"Financing Innovation"*, I take a closer look at the process through which ideas become new projects. In particular, we explore the financing and implementation of innovation activity, and how the division of surplus between innovators and financiers affects the incentives to innovate.

Taking our cue from Arrow's fundamental paradox of information, we model innovators as possessing ideas of uncertain quality but limited funds, and financiers as possessing funds but no ideas of their own. Financiers have the expertise to evaluate the quality of ideas, but the same expertise gives them the ability to undertake the project themselves once the idea is disclosed to them, thus denying the innovator any rents. We model this interaction as a repeated game, where the financier's concern about his reputation disciplines his behavior.

We find that, in equilibrium, very high quality ideas are not financed. When the idea is of sufficiently high quality, the financier cannot commit not to expropriate the innovator: the value of his reputation depends on the quality of the average idea, so when faced with the possibility of stealing a sufficiently above-average idea the financier will do so. The expropriation problem becomes worse as competition among financiers increases, given that reputational rents are bid away. However, since the incentives to innovate increase with the rents accruing to innovators, there is a optimal level of competition between financiers in the industry.

References

Fisher, J. D. M. (2006). The dynamic effects of neutral and investment-specific technology shocks, *Journal of Political Economy* **114**(3): 413–451.

Greenwood, J., Hercowitz, Z. and Krusell, P. (1997). Long-run implications of investment-specific technological change, *American Economic Review* **87**(3): 342–362.

Justiniano, A., Primiceri, G. E. and Tambalotti, A. (2009). Investment shocks and the relative price of investment, *Discussion Paper 7598*, Centre for Economic Policy Research.

Lev, B. and Radhakrishnan, S. (2004). The valuation of organization capital, *Working paper*, NYU Stern School of Business.

Sentana, E. and Fiorentini, G. (2001). Identification, estimation and testing of conditionally heteroskedastic factor models, *Journal of Econometrics* **102**(2): 143–164.