Innovation Cycles*

Jiro Kondo Dimitris Papanikolaou

Abstract

We embed frictions in the sale of ideas in a relatively standard real business cycle model. Inventors own ideas that they can either inefficiently implement on their own, or implement them efficiently in partnership with firms. However, firms can steal ideas; they are likely to do so if the idea is of sufficiently high quality. Introducing this agency problem in a dynamic general equilibrium model, we show that news about future technologies improve the firms' commitment ability not to expropriate inventors, and thus to an improvement in the pool of supplied ideas. Our model generates positive comovement of investment, consumption and labor supply in response to news about future technologies.

*Jiro Kondo is with McGill University. Dimitris Papanikolaou is with Kellogg School of Management and NBER. We thank Hengjie Ai, Larry Christiano, Simon Gilchrist, Arvind Krishnamurthy, Sergio Rebelo and seminar participants at CEPR Gerzensee, NBER, Northwestern, and UW Madison for helpful comments and discussions. Dimitris Papanikolaou thanks the Zell Center for Risk and the Jerome Kenney Fund for financial support.

Introduction

The process through which ideas are transformed into productive output is fraught with frictions. Given that innovation is an important driver of economic growth, these frictions are likely to have long-term effects. So far, most of the macroeconomic literature has focused on frictions in financing real investment. However, innovation is substantially more difficult to finance than physical capital because ideas are difficult to sell. Adding to the inherent difficulties in financing innovation, even if the quality of ideas could be verified ex-post, once ideas are disclosed they can be stolen. Arrow (1962), writes: "There is a fundamental paradox in the determination of demand for information; its value for the purchaser is not known until he has the information, but then he has acquired it without cost." In this paper, we build a real business cycle model with frictions in the sale of ideas.

The key ingredient in our model is that firms cannot commit to appropriately compensate inventors, which limits the supply of ideas ex-ante. Ideas originate with inventors; inventors are good at generating ideas but are not necessarily the best at implementing and bringing the product to market. Left to their own devices, inventors can implement projects at a low level of efficiency. By contrast, firms are good at implementing ideas – they can implement projects at a high level of efficiency – but at the same time have the ability to expropriate inventors. The deterrent to expropriation is the value of firms' reputation: if firms expropriate inventors they lose future business, that is, they can no longer attract innovators interested in partnership. The value of firms' reputations depends on the *average* value of projects; hence, firms will steal ideas only if these ideas are sufficiently higher than average quality. This lack of commitment not to expropriate inventors leads to inefficient outcomes ex-ante; inventors will choose to implement the best projects without firms, even though these projects are inefficiently implemented.

We embed this mechanism in a relatively standard real business cycle model with news about future technologies. In contrast to most existing models, good news about the likelihood

of future technological improvements lead to increases in investment, consumption, output and labor supply today. Our model features two aggregate shocks: large, but infrequent, improvements in the frontier level of technology – which determines the productivity of new capital – and changes in the arrival rate of these improvements. We refer to the latter as news about future technologies. The arrival of a positive news shock generates the following set of responses in the economy. First, good news about the future increases firms' ability to commit not to expropriate inventors. Specifically, an increase in the likelihood of improvement in the technology frontier increases the value of future projects relative to the value of existing projects. As a result, the value of firms' reputation increases relative to the immediate gains from expropriation. Inventors with projects of higher quality are willing to enter into partnership with firms, increasing both the fraction and the marginal quality of projects implemented at a high level of efficiency.

At the aggregate level, this improvement in efficiency leads to an increase in demand for new capital and thus to an increase in the equilibrium price of investment goods and the real wage. In response, both investment and labor supply increase. Most importantly, this channel is sufficiently strong to overcome the income effect in labor supply that would otherwise tend to reduce the supply of labor. Last, consumption increases in response due to an increase in the utilization rate of existing capital. Good news about the profitability of future projects lowers the value of older vintages today, leading to an increase in the rate of capacity utilization today.

Our mechanism is related to models with financial constraints. However, there are some important distinctions; even though inventors are constrained in implementing projects, firms are not. The constraint is endogenously generated by the value of firms' reputation – which depends on the rents from future projects as long as they do not expropriate inventors. Hence, the constraint lies in the supply of ideas not the supply of physical capital. This distinction is important beyond mere semantics. A relaxation in this constraint implies that the marginal

idea supplied – and thus implemented – is of *higher* quality than before. Relaxing the "idea-stealing" constraint leads to an increase not only to the fraction of projects implemented efficiently, but also to an increase in the quality of the marginal project implemented. By contrast, in standard models of financial constraints, the bottleneck lies in the supply of physical capital. Firms allocate their limited resources to the best projects available. A relaxation of financial constraints implies that the quality of the marginal project implemented has to be weakly lower, as firms implement the best projects first.

The last three decades saw the rise in prominence of an alternative form of funding for innovative companies: venture capital. By connecting the share of projects in a partnership with a firm to expectations about the arrival of future technologies, our model provides an economic foundation for the existence of venture capital cycles (Gompers and Lerner, 2006). Gompers, Kovner, Lerner, and Scharfstein (2008) document that the venture capital industry undergoes investment cycles; VCs with the most experience – a likely proxy for reputation in our model – increase their investment more during these cycles relative to firms with little experience. Moreover, even though they increase investment, their performance is not worse. Further, Nanda and Rhodes-Kropf (2011) document that, conditional on these firms going public, startups receiving their initial funding in periods with more VC funding filed more patents – and those patens receive more citations – relative to startups funded in less active years. In addition, they document that, conditional on going public, startups funded in more active years were valued higher on the IPO date. These empirical facts are consistent with our model.

Several papers study the role of frictions in entrepreneurship. A large segment of the literature focuses on credit frictions that prevent poorer but potentially highly productive entrepreneurs from entering the market (see, for instance Banerjee and Newman, 1993). Our work is closest to the work that combines RBC-style models with frictions in the sale of ideas (see, for example Silveira and Wright, 2010; Chiu, Meh, and Wright, 2011). In the models of

Silveira and Wright (2010) and Chiu et al. (2011), firms need to pay inventors the value of the idea upfront – that is, they cannot commit to pay them once the idea is implemented. Since ideas cannot be collateralized, this friction creates a demand for liquidity on the part of firms and a role for intermediation. By contrast, in our setting, firms cannot commit to pay inventors anything, either upfront or later, and paying the inventors before the idea is disclosed is not feasible due to adverse selection.

We focus on the interaction of the limited commitment friction with news about future technologies. Jaimovich and Rebelo (2009) highlight the difficulty of standard real business cycle models in generating positive responses in all of consumption, investment and labor supply; labor supply typically falls in response to good news about the future, while consumption and investment typically show opposite responses. The culprit for the former is the income effect on labor choice; Jaimovich and Rebelo (2009) propose preferences that do not have a strong income effect in the short run to generate an increase in labor supply. Closest to our paper is the work of Jermann and Quadrini (2007) and Chen and Song (2012), who show that in the presence of a standard financial friction – a collateral constraint – news about future productivity can generate an economic expansion today. Our model differs from both of these papers in several important ways. First, the friction that we focus on is on the supply of ideas, rather than the supply of capital. Second, our specification of preferences over labor/leisure choice allows for income effects. Last, the type of news shocks we consider are news about the arrival of newer vintages of technology; in contrast, Chen and Song (2012) consider news about future TFP. In a model where financial frictions arise due to collateral constraints, news about future technological vintages are unlikely to relax the constraint today.[1] Further, our model is consistent with the findings of Comin and Gertler (2006), who document the existence of substantial medium-run fluctuations in economic quantities.

Our work has implications about the measurement and identification of capital-embodied

[1] In fact, the most likely outcome is that it will lead to tighter constraints; as the arrival of future vintages becomes more likely, the value of older vintages drops.

shocks. Our model mechanism implies that news about future technologies improve real investment opportunities, and therefore affect the marginal rate of transformation between consumption and investment today. Hence, these news have a similar effect on quantities as investment-specific technology shocks (Solow, 1960; Greenwood, Hercowitz, and Huffman, 1988; Greenwood, Hercowitz, and Krusell, 1997; Fisher, 2006; Justiniano, Primiceri, and Tambalotti, 2010; Papanikolaou, 2011). However, the effect of news on the equilibrium relative price of investment goods in our model is either positive or zero, since our effect operates through the demand for new capital, similar to Christiano, Motto, and Rostagno (2013). Hence, traditional methods of identifying IST shocks based on changes in the relative price of equipment will miss this channel. Our work thus provides a micro-foundation for the marginal efficiency of investment shock in Justiniano, Primiceri, and Tambalotti (2011). Last, our paper is connected to the recent body of work that aims to disentangle whether news about future productivity or capital embodied shocks are the dominant source of business cycle fluctuations (Beaudry and Portier, 2006; Beaudry and Lucke, 2009; Fisher, 2009; Schmitt-Grohe and Uribe, 2012). In our model, news about the future improve the marginal efficiency of investment today; hence, disentangling these two types of disturbances in the data may be quite challenging.

1 The agency problem

At the heart of the model is the notion that ideas can be stolen. Idea theft is possible because intellectual property rights are not always well protected. Patents provide some measure of protection, however expropriation may still be possible. For instance, E. H. Armstrong pioneered FM radio in the 1910s and '20s. However, ,any of Armstrong's inventions were claimed by others. The regenerative circuit, which Armstrong patented in 1914 as a "wireless receiving system," was subsequently patented by Lee De Forest in 1916; De Forest then sold

the rights to his patent to AT&T. Furthermore, once disclosed, the ideas can be implemented without the innovator, who is often not crucial to the success of the venture. For example, Robert Kearns patented the intermittent windshield wiper in 1967. He tried to interest the "Big Three" auto makers in licensing the technology. They all rejected his proposal, yet began to install intermittent wipers in their cars, beginning in 1969. Kearns ultimately won the patent lawsuit against Ford in 1978 and Chrysler in 1982. The possibility of expropriation affects the choice of whether the new invention is implemented inside an established firm (or venture capitalist) or by the innovator himself.

This process can take several forms. First, the innovator can obtain an idea about a new project while working for an existing firm. He then has the choice of disclosing his idea to his superiors in the hope of obtaining a fraction of the surplus or leave the firm and implement the idea himself. His choice depends on the quality of the project and the firms' ability to commit to provide a fraction of the surplus. According to Bhide (1999), 71% of the founders of firms in the Inc 500 list of fast growing technology firms report that they replicated or modified ideas encountered through previous employment. For example, in the late 80s, software maker Peoplesoft and its founder David Duffield were sued by Integral Systems which claimed that its software was based on computer code that was stolen from the company while Mr. Duffield worked there.[2] More recently, the social networking site ConnectU sued competitor Facebook and its founder Mark Zuckerberg for allegedly copying the company's idea to build a networking website while employed at the company.[3]

Second, the inventor can approach a financier with the idea for the new project. Unfortunately, financiers that can assess the quality of the inventor's idea often have the ability to implement the idea themselves. Depending on their ability to commit – the value of their reputation – they will expropriate the innovator only if the idea is sufficiently prof-

[2]Peoplesoft was eventually acquired by Oracle in a deal valued at over $10 billion.

[3]Mr Zuckerberg was also accused of delaying the completion of ConnectU's software while bringing his product to market.

itable. Intellectual property theft is not rare in the venture capital context. For instance, EP Technologies (EPT), a medical equipment startup, sued Sierra Ventures, accusing it of using EPT's confidential information to help start CardioRhythm which offered an identical product. This information had been disclosed by EPT during funding discussions with Sierra, who had also hired consultants to evaluate EPT's technology and patents.[4] Similarly, in the joint venture context, Microdomain, a microdevice startup, sued Quinta, an already financed startup, for trade-secret theft.[5] Further, outright idea theft is not the only way that financiers can expropriate innovators; financiers can also appropriate significant rents by diluting the innovator's stake in the venture. Often, this happens after the founder has left the company or been terminated. This is often possible due to contractual features of the VC arrangement (see, for instance Kaplan and Stromberg, 2004). For instance, as described in Atanasov, Ivanov, and Litvak (2012), the founder of Pogo.com, an e-gaming company, sued the VCs on the board for issuing complicated derivative securities, effectively reducing his stake from 13% to 0.1%, and then refusing to redeem his stock in violation of prior agreement. Similarly, VCs are at times better informed than the innovators and this permits other opportunities to expropriate. For example, the founders of Epinions, a consumer product review website, sued three VC funds for fraudulently withholding information that caused them multimillion dollar losses.[6]

[4]CardioRhythm was founded by Harry Robbins, who was introduced by Sierra as a potential CEO of EPT. Sierra was also accused to have used its contacts in the VC community to limit EPT's ability to compete with Cardiorythm. According to court documents, when EPT asked another venture capital firm for financing, Sierra tried to persuade the firm to invest in CardioRhythm instead. In addition, when EPT discussed a merger with medical manufacturer Medtronic Inc., Sierra intervened and persuaded Medtronic to buy CardioRhythm. Source: Reynolds Holding and William Carlsen, San Francisco Chronicle.

[5]Quinta allegedly violated a mutual nondisclosure agreement and filed patent applications including the disk-drive technology that Microdomain contended was theirs. One year later, Quinta was acquired by Seagate for $325 million.

[6]The founders alleged that the financiers persuaded them to give up their ownership interests after being led to believe that the value of their stake was zero. At the time, the VCs had indicated that the value the company was around $30 million, well below its $45 million liquidation preference. The founders alleged that, a year later, the implied value of the company was $300 million, partly due to a deal with Google and other financial results and projections that were not disclosed by the VCs.

2 A simple model

To illustrate the main intuition behind our mechanism, we first present a partial equilibrium version of the model. There exist a set of firms of measure one and a group of inventors of measure λ/β. Both groups are risk-neutral and strategic. Time is continuous. The interest rate is constant and equal to ρ.

Inventors have finite lives; they die each period with probability $\beta\, dt$, and are replaced with a new inventor. Upon entry into the economy, each inventor is endowed with a new idea (blueprint) for a project. Each idea can be implemented into a project by combining the blueprint with an amount of physical capital k. Once implemented, a project produces a constant output flow according to the following technology,

$$y = \theta^{1-\alpha} k^{\alpha}, \tag{1}$$

where θ indexes the quality of the investment opportunity; θ is i.i.d. over time, has c.d.f. $F(\theta)$ with support $[0, \infty)$. In the beginning of the period, θ is known only to the inventor. Projects expire with probability δ. Hence, the present value of the cashflows generated by the project is $M\, \theta^{1-\alpha} k^{\alpha}$, where $M = (\rho + \delta)^{-1}$. The innovator faces two options in implementing her idea.

First, she can choose to implement the project in a partnership with a firm. If the project is implemented in a partnership with the firm, it is implemented at a high level of efficiency; the firm can purchase capital at a unit price. In this case, she needs to disclose the details of the project to the firm, and in the process revealing θ. Equivalently, this choice can be framed as the sale of the idea to a firm. What is important is that the firm cannot commit to pay the innovator after the idea is disclosed. Paying the inventor before the idea is disclosed is not possible due to a standard market for lemons problem: if the firm offers a payment to inventors that is a function of the average quality of ideas, $E[\theta]$, only inventors with below

average quality will sell.

The second option available to the inventor is to implement the project herself, at a lower level of efficiency. We model this efficiency loss by assuming that the inventor buys physical capital at a price $R > 1$. The spread $R - 1$ captures the degree of efficiency loss and captures the idea that firms are more efficient at implementing projects than firms. We denote the partnership decision by the inventor by $P : \theta \to \{0, 1\}$. Thus, with a slight abuse of notation, we write the efficiency wedge as a function of θ

$$r(\theta) = \begin{cases} 1 & \text{if} \quad P(\theta) = 1 \\ R & \text{if} \quad P(\theta) = 0. \end{cases} \tag{2}$$

After the partnership decision has been undertaken, the firm – or the inventor – chooses the scale of the project k to maximize its net present value

$$\Pi(\theta) \equiv \max_k \left\{ M\, \theta^{1-\alpha} k^{\alpha} - r(\theta)\, k \right\}, \tag{3}$$

yielding

$$k^c(\theta) = \theta \left(\frac{\alpha\, M}{r(\theta)} \right)^{\frac{1}{1-\alpha}} \tag{4}$$

$$\Pi(\theta) = \theta\, \alpha^{\frac{\alpha}{1-\alpha}} (1-\alpha) \left(\frac{1}{r(\theta)} \right)^{\frac{\alpha}{1-\alpha}} M^{\frac{1}{1-\alpha}}. \tag{5}$$

Here, it is useful to compute the level of profits Π^* which assumes there is no efficiency wedge, and the level of profits Π^c that the inventor can achieve on her own. These two represent the efficient outcome and the inventor's outside option, respectively.

$$\Pi^c(\theta) \equiv \theta\, \alpha^{\frac{\alpha}{1-\alpha}} (1-\alpha) M^{\frac{1}{1-\alpha}} \left(\frac{1}{R} \right)^{\frac{\alpha}{1-\alpha}}, \tag{6}$$

$$\Pi^*(\theta) \equiv \theta\, \alpha^{\frac{\alpha}{1-\alpha}} (1-\alpha) M^{\frac{1}{1-\alpha}} \tag{7}$$

The firm will only enter a joint venture with the inventor if it knows the quality of the inventor's project. Most importantly, once the innovator decides to enter into a partnership with the firm, the latter can implement the project on its own and expropriate the innovator. If the firm expropriates the innovator, the latter obtains a payoff of zero. However, expropriating the innovator implies that future generations of inventors will refuse to do business with the firm. For a given level of project quality, θ, a partnership is feasible between the innovator and the financier if the former obtains a higher payoff under the partnership than her outside option

$$\Pi^E(\theta) \geq \Pi^c(\theta), \tag{8}$$

and the financier prefers the partnership to expropriating the inventor

$$\Pi^F(\theta) \geq \Pi^*(\theta) - V \tag{9}$$

and losing the relationship value V where

$$V_t \equiv \lambda \int_t^\infty \int_0^\infty e^{-\rho(s-t)}\, \Pi^F(\theta)\, dF(\theta)\, ds. \tag{10}$$

Each instant dt a measure of $\lambda/\beta \times \beta = \lambda\, dt$ inventors is born and randomly matched to firms. Hence, each firm faces a probability $\lambda\, dt$ of meeting an inventor each period.

We only consider efficient subgame perfect equilibria of the repeated game, namely we restrict attention to the sharing rules that satisfy $\Pi^E(\theta) + \Pi^F(\theta) = \Pi^*(\theta)$. Still, the model admits many equilibria, each indexed by a sharing rule $\eta(\theta)$. Without loss of generality, we restrict attention to equilibria where the sharing rule is constant $\eta(\theta) = \eta$.[7]

Proposition 1 *The equilibrium is characterized by a threshold θ^*, such that partnership occurs, $P(\theta) = 1$ if $\theta \leq \theta^*$ and the inventor implements the project alone, $P(\theta) = 0$, otherwise.*

[7]For every equilibrium with η a function of θ, there exists an equivalent one where $\eta(\theta)$ is constant.

Equilibrium payoffs are given by

$$\Pi^F(\theta) = \begin{cases} \Pi^*(\theta) - \Pi^c(\theta) - \eta\,(V - \Pi^c(\theta))\,, & \textit{if} \quad \theta \leq \theta^* \\ 0, & \textit{if} \quad \theta > \theta^* \end{cases}, \tag{11}$$

$$\Pi^E(\theta) = \begin{cases} \Pi^c(\theta) + \eta\,(V - \Pi^c(\theta))\,, & \textit{if} \quad \theta \leq \theta^* \\ \Pi^c(\theta), & \textit{if} \quad \theta > \theta^* \end{cases}. \tag{12}$$

The threshold θ^* is the solution to

$$\Pi^c(\theta^*) = V, \tag{13}$$

where V is the relationship value to the bank, which equals

$$V = \frac{\lambda}{\rho + \eta\,\lambda\,F(\theta^*)} \int_0^{\theta^*} \Big(\Pi^*(\theta) - (1-\eta)\,\Pi^c(\theta)\Big)\,dF(\theta) \tag{14}$$

Proposition 1 summarizes the intuition behind the financial friction in this paper. Financial relationships are limited in their ability to mitigate the hold-up problem between inventors and financiers. These relationships are insufficient exactly when they would be most valuable in the absence of agency conflicts: the best projects $\theta > \theta^*$ are not disclosed and are under-implemented. Since the financier cannot commit to not expropriate an inventor with a sufficiently high quality project, the latter will refuse to enter a partnership agreement with the financier.

The financier's continuation value V determines his ability to commit not to expropriate the inventor, and therefore the return to the marginal project (13). At this stage, it is informative to compute the average value of a new project,

$$\begin{aligned} E[\Pi(\theta)] &= \int_0^{\theta^*} \Pi^*(\theta) dF(\theta) + \int_{\theta^*}^{\infty} \Pi^c(\theta) dF(\theta) \\ &= \alpha^{\frac{\alpha}{1-\alpha}}\,(1-\alpha)\,M^{\frac{1}{1-\alpha}}\,h(\theta^*), \end{aligned} \tag{15}$$

where

$$h(x) \equiv \int_0^x \theta \, dF(\theta) + \left(\frac{1}{R}\right)^{\frac{\alpha}{1-\alpha}} \int_x^\infty \theta \, dF(\theta). \tag{16}$$

The fact that projects implemented outside a partnership are inefficient $R > 1$ implies that, for all values of θ, $\Pi^*(\theta) > \Pi^c(\theta)$. As result, the profitability of the average project is increasing in the quality of the marginal project being implemented in a partnership θ^*, implying $h'(x) > 0$.

The results of this section preview the main mechanism of our general equilibrium model. In the partial equilibrium model in this section, V is constant; more generally, shocks to the ratio of the value of firms' reputation V relative to the benefits of expropriation today M will affect the share of new projects that are implemented at a high level of efficiency, and thus the average return to new projects $h(\theta^*)$. Since partnership is the efficient outcome in terms of production, in general equilibrium these shocks affect the marginal efficiency of investment and thus the demand for new capital.

3 The model

Here, we embed the financial friction is the previous section into a dynamic general equilibrium model based on Kogan, Papanikolaou, and Stoffman (2013).

3.1 Firms and technology

There are three production sectors in the model: a sector producing intermediate consumption goods; a sector that aggregates these intermediate goods into the final consumption good; and a sector producing investment goods. Firms in the last two sectors make zero profits due to competition and constant returns to scale; the existence of the final good and investment sector helps with aggregation.

Intermediate goods

Production in the intermediate sector takes place in the form of projects. Projects are introduced into the economy by new cohorts of inventors, who are restricted in their ability to implement them on their own. An inventor may decide to enter into a partnership with an existing firm. We refer to these firms as financiers.

Active projects

Projects are differentiated from each other by three characteristics: a) their quality θ_j, which is known to the inventor; b) their scale, k_j, chosen irreversibly at their inception; and c) the level of frontier technology at the time s of project creation, ξ_s. A project j created at time s produces a flow of output at time $t > s$ equal to

$$y_{jt} = u_{jt}\, e^{\xi_s}\, \theta_j^{1-\alpha}\, k_j^{\alpha}, \qquad \alpha \in (0,1), \tag{17}$$

where ξ_s is the level of frontier technology at time the project is created, and u_{jt} denotes the level of capital utilization for project j. A more active rate of utilization increases the likelihood that the project expires. The probability that the project expires during the period t to $t + dt$ is a function of the rate of capital utilization $\delta(u)$, where $\delta'(u) > 0$ and $\delta''(u) > 0$. Variable capital utilization is not an essential feature of our model, but it helps generate positive comovement between investment and consumption growth.

Inventors and new projects

Project ideas originate with new generations of inventors. Upon entry into the economy each inventor is endowed with λ/β projects, each of unknown quality θ. The quality of ideas θ is distributed on $(0, \infty)$ with cdf $F(\theta)$. Inventors know the quality θ of their project idea. The implementation of a new project idea requires new capital purchased at the equilibrium market price p^I. Once a project is acquired, the owner of the project chooses its scale of

production k_j to maximize the value of the project. The choice of project scale is irreversible; existing projects cannot be liquidated to recover their original costs.

Most importantly, inventors cannot implement their ideas as efficiently as firms can. If the firm implements the idea, it needs to purchase k units of capital to implement a project of scale k; there is no efficiency loss. In contrast, if the inventor implements the project without a firm, she needs to purchase $R\,k$ units of capital to implement a project of scale k, where $R > 1$ captures the loss of efficiency. After the project is implemented, the inventor can sell the project in the financial markets without any adverse selection problems. The model nests the frictionless benchmark by setting $R = 1$. In this case, firms and innovators are equally adept at implementing projects.

Investment- and consumption-good firms

Firms in the capital-good sector use a constant returns to scale technology employing labor L_I and intermediate goods Y_I to produce the investment goods needed to implement new projects in the intermediate-good sector

$$I_t = Y_{It}^{\beta}\, L_{It}^{1-\beta}. \tag{18}$$

Similarly, final consumption good firms produce output using labor L_C and intermediate goods Y_{Ct}

$$C_t = Y_{Ct}^{\phi}\, L_{Ct}^{1-\phi}. \tag{19}$$

Both sectors use the output of the intermediate good sector Y as a input; their total demand for intermediate goods must clear the market

$$Y_{Ct} + Y_{It} = Y_t. \tag{20}$$

Technology

The model features two sources of uncertainty: technological revolutions, captured by changes in ξ; and news about technological revolutions, captured by changes in the arrival rate μ_t.

Technological revolution

The level of frontier technology ξ at the time the project is implemented evolves according to

$$d \ln \xi_t = \chi \, dN_t, \tag{21}$$

where N_t is a poisson jump process with a state dependent arrival rate μ_t. Each time the Poisson process hits, the level of the frontier technology is increased by proportional amount χ – a technological revolution. The variable μ_t captures the likelihood of a technological increase, and therefore represents news about future embodied shocks, and evolves according to the following stationary stochastic process:

$$d\mu_t = \kappa_\mu(\bar{\mu} - \mu_t)\, dt + \sigma_\mu \mu_t^p dZ_t^\mu. \tag{22}$$

The parameters κ_μ, $\bar{\mu}$, σ_μ and p control the degree of mean reversion, the mean, the volatility, and the skewness of the stationary distribution of μ, respectively.

3.2 Households

Households are infinitely lived, and have non-separable preferences over sequences of consumption C and leisure N. Utility J is defined recursively as

$$J_t = E_t \int_t^\infty \tilde{f}(C_s, N_s, J_s) ds, \tag{23}$$

where the aggregator $f(C, J)$ takes the form

$$\tilde{f}(C, N, J) \equiv \frac{\rho}{1-\theta^{-1}} \left(\frac{\left(CN^{\psi}\right)^{1-\theta^{-1}}}{((1-\gamma)J)^{\frac{\gamma-\theta^{-1}}{1-\gamma}}} - (1-\gamma)\, J \right) \tag{24}$$

Here, ψ is the preference weight on leisure, ρ is the subjective discount rate, γ is the coefficient of relative risk aversion, and θ is the elasticity of intertemporal substitution (EIS).[8]

To focus on our main economic mechanism, we assume the existence of perfect risk-sharing across inventors. That is, even though inventors have finite lives, we assume that they belong to a large family that perfectly shares risk across generations. We consider the case of imperfect intergenerational risk sharing as an extension in the online appendix. Allowing for imperfect risk sharing across generations does not alter the qualitative predictions of our model; instead, it reinforces its quantitative predictions.

Last, households are endowed with a fixed unit of time, that can be freely allocated between leisure, producing consumption goods or producing capital

$$L_{It} + L_{Ct} + N_t = 1. \tag{27}$$

The allocation of labor between the investment-good and consumption-good sectors is the mechanism through which the economy as a whole saves or consumes.

[8] In the special case where $\gamma = \theta = 1$ our preferences reduce to the standard log case

$$J_t = E_t \int_t^{\infty} e^{-(\rho+\beta)(s-t)}\, U(C_s, N_s) ds. \tag{25}$$

where

$$U(C, N) \equiv \log C + \psi \log N \tag{26}$$

4 Competitive equilibrium

We begin our description of the competitive equilibrium by characterizing households' optimality conditions.

4.1 Dynamic evolution of the economy

The current state of the economy is characterized by the vector $Z_t = [K_t, \omega_t, \mu_t]$, where

$$\omega \equiv \xi - (1 - \alpha\beta) \ln K. \tag{28}$$

The dynamic evolution of the aggregate state Z depends on the law of motion for ξ, given by equations (21) and the evolution of the effective stock of capital,

$$\begin{aligned} dK_t = \Big(i(\omega_t, \mu_t) - \delta(\omega_t, \mu_t) \Big) K_t \, dt, \\ \text{where} \quad i(\omega_t, \mu_t) &\equiv \lambda \, e^{\xi_t} \int_0^\infty \theta^{1-\alpha} \, k_t(\theta)^\alpha \, dF(\theta) \\ &= \lambda \, e^{\omega} \left(\frac{u^\beta(\omega_t, \mu_t) \, y_I^\beta(\omega_t, \mu_t) \, L_I^{1-\beta}(\omega_t, \mu_t)}{\lambda \, g\Big(\theta^*(\omega_t, \mu_t)\Big)} \right)^\alpha h\Big(\theta^*(\omega_t, \mu_t)\Big), \\ \delta(\omega_t, \mu_t) &\equiv \hat{\delta}(u^*(\omega_t, \mu_t)). \end{aligned} \tag{29}$$

The variable ω captures transitory fluctuations along the stochastic trend. Since $i'(\omega) > 0$, an increase in ω accelerates the growth rate of the effective capital stock, and thus the long-run growth captured by χ. Last, the variable μ captures the likelihood of a technological revolution. The following variables are stationary and depend only on ω and μ: hours worked $1 - N$, the threshold θ_t^* and consequently the supply of new ideas $h(\theta^*)$; and the allocation of labor and the share of intermediate goods used in the consumption and investment sector.

The function h defined in (16) captures the dependence of aggregate output on the

measure of new projects that are implemented in a partnership, versus by the inventor alone. Examining (16), we see that as the innovators become equally adept as firms in implementing new projects, $R \to 1$, $h(x)$ becomes a constant function equal to $E[\theta]$. In this case, relationships have no value and the model reduces to the frictionless benchmark.

At the aggregate level, our model has important similarities to the neoclassical growth model. The first state variable χ is difference-stationary and captures the stochastic trend in the economy. Long-run growth χ depends on the disembodied shock A and the effective capital stock K. The effective capital stock K depreciates at the rate δ of project expiration. The effective capital K grows as the economy allocates more resources into producing capital L_I relative to the total scale of new projects $g(\theta^*)$; when the quality of new capital improves ξ or labor becomes more productive A; and when the quality of the marginal project that is financed efficiently $h(\theta^*)$ improves.

The quality of the marginal project θ^* being financed in a partnership depends on the commitment ability of firms, which itself is a function of the present value of rents from future projects, V. Consequently, a shock that leads to an increase in V – for instance positive news about *future* technology due to an increase in μ_t – will also affect the quality of the marginal project in a partnership θ^* and therefore the marginal return of investment $h(\theta^*)$ *today*.

4.2 Household optimization

Labor supply

The labor-leisure decision is intra-temporal. The first-order condition with respect to leisure is $U'(N_t)/U'(C_t) = w_t$, implying

$$N_t = \psi\left(\frac{C_t}{w_t}\right) \tag{30}$$

As equation (30) shows, the income elasticity of labor supply $d\log(1-N)/d\log C$ is equal to the opposite of the Frisch elasticity of labor supply $d\log(1-N)/d\log w$, as in **?**. Using the labor market clearing condition (27),

$$N = (1 - L_I)\left(\frac{\psi}{\psi + 1 - \phi}\right) \tag{31}$$

we can see that labor supply is increasing in L_I.

Consumption and savings

Households have access to a complete menu of state contingent securities. Let π_t be the state-price density (the state price of consumption in units of probability). The household first-order condition is

$$\frac{\pi_s}{\pi_t} = \exp\left(\int_t^s \tilde{f}_J(C_u, N_u, J_u)\, du\right) \frac{\tilde{f}_C(C_s, N_s, J_s)}{\tilde{f}_C(C_t, N_t, J_t)} \tag{32}$$

For a derivation, see Skiadas...

Welfare

The households' utility index J satisfies the Hamilton-Jacobi-Bellman equation $\tilde{f}(C, N, J) + \mathcal{D}J = 0$. We guess, and subsequently verify, that the household's utility index J is given by

$$J_t = \frac{1}{1-\gamma} K_t^{\phi\,(1-\gamma)} j(\omega_t, \mu_t) \tag{33}$$

where the function $j(\omega_t, \mu_t)$ satisfies the PDE in the Appendix.

4.3 Firm optimization

Next, we study the optimality conditions on the production side of the economy. We conjecture – and subsequently verify – that the partnership decision takes the same form as the partial equilibrium model in section 2. Specifically, the inventor and the firm enter into a partnership for projects of quality below a threshold θ_t^*. Importantly, the threshold θ^* depends on the state of the economy. With a slight abuse of notation, we index all state-dependent variables and optimal policies with a time subscript.

Market for capital

The owner of a project in the intermediate-goods sector chooses the scale of investment, k_j, in each project to maximize its net present value, which equals the market value of a new project, minus its implementation cost. The implementation cost depends on whether the project is implemented by the firm or by the inventor, and is proportional to the cost of purchasing k units of capital at the equilibrium price p^I. For now, we guess that the market value of a new project of vintage ξ_t and quality θ is equal to $M_t\, e^{\xi_t}\, \theta_j^{1-\alpha}\, k_j^{\alpha}$, where

$$M_t = \max_{u_s} E_t \int_t^{\infty} \exp\left(-\int_t^s \delta(u_u)\, du\right) \frac{\pi_s}{\pi_t}\, p_s^Y\, u_s\, ds. \tag{34}$$

We verify our conjecture and derive M_t explicitly below. Given M_t and a partnership strategy, the net present value of a project of quality θ is

$$\max_k NPV = M_t\, \xi_t\, \theta^{1-\alpha} k^{\alpha} - r_t(\theta)\, p_t^I\, k \tag{35}$$

The optimal scale of investment is a function of the ratio of the market value of a new

project to its marginal cost of implementation,

$$k_t(\theta) = \theta \left(\frac{\alpha\, e^{\xi_t}\, M_t}{p_t^I} \right)^{\frac{1}{1-\alpha}} r_t(\theta)^{-\frac{1}{1-\alpha}}, \tag{36}$$

Equation (36) bears similarities to the q-theory of investment (Hayashi, 1982). A key difference here is that the numerator involves the market value of a new project M – akin to marginal q – which is distinct from the market value of the firm – average q. Importantly, the optimal scale of projects that are implemented by the inventor alone $P = 0$ is lower due to the loss of efficiency $R > 1$.

Aggregating across new projects, the total demand for new capital equals

$$I_t = \lambda \int_0^\infty k_t(\theta)\, dF(\theta) = \lambda \left(\frac{\alpha\, \xi_t\, M_t}{p_t^I} \right)^{\frac{1}{1-\alpha}} g(\theta_t^*) \tag{37}$$

where the function g captures the dependence of the demand for capital on the measure of projects that are implemented in a partnership, versus by the inventor alone

$$g(\theta^*) \equiv \int_0^\infty \theta\, r_t(\theta)^{-\frac{1}{1-\alpha}} dF(\theta) = \int_0^{\theta_t^*} \theta\, dF(\theta) + \left(\frac{1}{R} \right)^{1/(1-\alpha)} \int_{\theta_t^*}^\infty \theta\, dF(\theta). \tag{38}$$

The equilibrium price of investment goods, p_t^I, clears the supply (18) and the total demand for new capital (37)

$$p_t^I = \alpha\, e^{\xi_t}\, M_t \left(\frac{I_t}{\lambda\, g(\theta_t^*)} \right)^{\alpha-1}. \tag{39}$$

An increase in the level of frontier technology ξ, or to the quality of the marginal project θ_t^* leads to an increase in the demand for capital, and thus to an increase in its equilibrium price p^I.

Examining equations (37) and (39), we see that the threshold θ_t affects the quantity and price of investment in a qualitatively similar manner as the embodied shock ξ. An improvement in the quality of the marginal project that is implemented in a partnership improves the

real investment opportunities in the economy in the same way as if the productivity of new capital was higher.

Market for ideas

Relative to a standard real business cycle model, the main novel new ingredient is the market for ideas. Here, we verify that the results from the partial equilibrium model in section 2 continue to hold, specifically that partnership occurs for projects only below a certain threshold $P_t(\theta) = 1 \Leftrightarrow \theta \leq \theta_t^*$. The difference here is that the threshold θ_t^* is state-dependent.

First, we compute the outside option of the inventor $\Pi_t^c(\theta)$ and the first-best level of profits $\Pi_t^*(\theta)$, which is also equal to the maximum amount that the bank can expropriate from the inventor. Using equations (35)-(36), we obtain

$$\begin{aligned}
\Pi_t^*(\theta) &= \theta\,(1-\alpha)\,M_t\,e^{\xi_t}\,\left(\frac{I_t}{\lambda\,g(\theta_t^*)}\right)^{\alpha}, \\
\Pi_t^c(\theta) &= \theta\,(1-\alpha)\,M_t\,e^{\xi_t}\,\left(\frac{I_t}{\lambda\,g(\theta_t^*)}\right)^{\alpha}\left(\frac{1}{R}\right)^{\alpha/(1-\alpha)}.
\end{aligned} \tag{40}$$

The outside option of the inventor Π^c is equal to the value of a project implemented at a low level of efficiency. The firm has the ability to expropriate the inventor, and extract the full value of the project, Π^*. If the firm expropriates the innovator, then the firm foregoes the value of its relationship with future innovators,

$$V_t \equiv \lambda\,E_t \int_t^{\infty} \frac{\pi_s}{\pi_t}\left(\int_0^{\theta^*} \Pi_t^F(\theta)\,dF(\theta)\right) ds. \tag{41}$$

The equilibrium payoffs to the inventor Π^E and the firm Π^F are given by proposition 1, except for the fact that they depend on the state of the economy in addition to θ. Following the same logic as equation (13), the quality θ_t^* of the marginal project that gets implemented

in a partnership satisfies $\Pi_t^c(\theta_t^*) = V_t$, implying

$$\theta_t^* \, (1-\alpha) \, M_t \, e^{\xi_t} \, \left(\frac{I_t}{\lambda \, g(\theta_t^*)} \right)^{\alpha} \left(\frac{1}{R} \right)^{\alpha/(1-\alpha)} = V_t. \tag{42}$$

Market for intermediate goods and labor

The total output of the intermediate good, Y_t, equals the sum of the output of the individual projects, $Y_t = \int y_{f,t}$, and is equal to the rate of capital utilization u – which we guess and subsequently verify that it is identical across firms – times the effective capital stock

$$Y_t = u_t \, K_t \equiv u_t \int_{j \in \mathcal{J}_t} \xi_j \, \theta_j^{1-\alpha} \, k_j^{\alpha} \, dj. \tag{43}$$

adjusted for the productivity of each vintage – captured by ξ at the time the project is created – and for decreasing returns to scale. An increase in the effective capital stock K, for instance due to a positive embodied shock, leads to a lower price of the intermediate good and to displacement for productive units of older vintages.

Consumption firms purchase the intermediate good Y at a price p_Y and hire labor L_C at a wage w to maximize their value. Their first order condition with respect to their demand for intermediate goods yields

$$\phi \, Y_{Ct}^{\phi-1} \, L_{Ct}^{1-\phi} = p_t^Y \tag{44}$$

$$(1-\phi) \, Y_{Ct}^{\phi} \, L_{Ct}^{-\phi} = w_t. \tag{45}$$

Similarly, the demand of intermediate goods and labor by investment firms implies

$$\beta \, p_t^I \, Y_{It}^{\beta-1} L_{It}^{1-\beta} = p_t^Y \tag{46}$$

$$(1-\beta) \, p_t^I \, Y_{It}^{\beta} L_{It}^{-\beta} = w_t \tag{47}$$

To obtain intuition about the equilibrium allocation of resources between the investment and consumption sector, consider the case where the factor shares are equal, $\beta = \phi$. In this case, equating the marginal products in the consumption (44)-(45) and investment sector (46)-(47) imply that in equilibrium, the relative price of investment goods is equal to one, $p_t^I = 1$ always. In this special case, clearing the capital market (39) implies

$$I_t = \lambda \left(\alpha\, e^{\xi_t}\, M_t\right)^{\frac{1}{1-\alpha}}\, g(\theta_t^*). \tag{48}$$

Ceteris paribus, the total amount of resources allocated to investment is increasing in the level of frontier technology ξ, the equilibrium price of installed capital M_t and the quality of the marginal project in partnership θ_t^*. Intuitively, shocks that tend to increase the price of capital – such as an increase in ξ_t or θ_t^* – prompt a reallocation of resources from consumption to investment. The same intuition holds in the more general case $\beta \neq \phi$, see the Appendix for more details.

The limited commitment friction plays a key role in the paper. In the absence of the limited commitment friction, $g(\theta^*)$ is a constant. In this case, the equilibrium level of investment I is increasing in the frontier level of technology ξ, as in standard models. Further, an increase in the likelihood of technological innovation μ_t lowers the value of the currently available projects M; and as a result leads to a reallocation of resources away from investment towards consumption, as well as a reduction in the labor supply. This is the standard effect that leads to a recession following good news about future productivity in a standard RBC model. In contrast, in our model θ_t^* is an endogenous function of the state of the economy, including μ_t. Good news about future technologies increase the reputation value of firms V and therefore their ability to commit not to expropriate better projects, leading to an increase in θ_t^*.

Project valuation and investment decisions

The last step in characterizing the equilibrium involves computing the market value of new projects and the relationship value of firms. Given (34), the discounted value of new projects $\tilde{M}(\Omega_t) = \pi_t M_t$ satisfies the Hamilton-Jacobi-Bellman equation

$$0 = \max_u \left\{ \tilde{f}_C(C, N, J)\, p^Y\, u + \Big(f_J(C, N, J) - \delta(u) \Big)\, \tilde{M} + \mathcal{D}\tilde{M} \right\}. \tag{49}$$

The solution to (49) is given by $\tilde{M}_t = K_t^{\phi-1}\, m(\omega_t, \mu_t)$, where the function $m(\omega_t, \mu_t)$ satisfies the PDE in the appendix. After substituting, the value of existing projects is

$$M_t = K_t^{\phi-1}\, m(\omega_t, \mu_t) \left(\left(\bar{u}_t^{\phi}\, y_{Ct}^{\phi} L_{Ct}^{1-\phi} \right)^{-\theta^{-1}} j(\omega_t, \mu_t)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N_t^{\psi(1-\theta^{-1})} \right)^{-1}.$$

Firms choose the level of capital utilization u taking the decisions of other firms u^* as given. The equilibrium level of capital utilization u^* satisfies the first-order condition,

$$p_t^Y = \delta'(u_t^*)\, M_t. \tag{50}$$

The last step involves computing the equilibrium value of a relationship to the firm-financiers, which as we see in (51), is equal to the present value of the rents $\Pi_t^F(\theta)$ they receive from the projects they enter into a partnership with $\theta \leq \theta^*$. Using the version of the Feynman-Kac theorem with discounting, the discounted relationship value $\tilde{V}(\Omega_t) = \pi_t V_t$ solves the differential equation

$$0 = \left\{ \lambda \hat{g}(\theta^*)\, \tilde{M}\, e^{\xi} \left(\frac{I}{\lambda g(\theta^*)} \right)^{\alpha} + \Big(\tilde{f}_J(C, N, J) - \lambda\, \eta\, F(\theta^*) \Big)\, \tilde{V} + \mathcal{D}\, \tilde{V} \right\} \tag{51}$$

where

$$\hat{g}(\theta^*) \equiv \left(1 - (1-\eta)\, R^{-\alpha/(1-\alpha)} \right)\, (1-\alpha) \int_0^{\theta^*} \theta\, dF(\theta). \tag{52}$$

Equation (51) shows that the firms' relationship value V is, among other things, a function of the surplus rule η. We solve (51) by conjecturing that the discounted relationship value takes the form $\tilde{V}_t = K_t^{\phi}\, v(\omega_t, \mu_t)$, where the function $v(\omega, \mu)$ solves the PDE in the appendix. After substitution, the relationship value takes the form

$$V_t = K_t^{\phi}\, v(\omega_t, \mu_t) \Big[\left(\bar{u}_t^{\phi}\, y_{Ct}^{\phi} L_{Ct}^{1-\phi} \right)^{-\theta^{-1}} j(\omega_t, \mu_t)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N_t^{\psi(1-\theta^{-1})} \Big]^{-1}. \tag{53}$$

5 Model predictions

To provide insight into the model mechanism, we explore the response of model quantities to the two fundamental disturbances in the economy: news about technological revolutions, μ, and the arrival of a technological revolution, ξ. We compare the response of aggregate quantities to a version of the model without frictions ($R = 1$).

5.1 Calibration

Whenever possible, we calibrate the model using standard parameter values or by targeting standard moments. We parameterize the project depreciation rate as a function of capital utilization rate as $\delta(u) = \delta_0 + \delta_1/2\, u^2$; we choose the parameters δ_0 and δ_1 to generate an average depreciation rate of 2% per quarter, and to match the quarterly volatility of log changes of capital utilization (equal to 0.75% using the data from Basu, Fernald, and Kimball, 2006). We calibrate the share of capital in consumption $\phi = 0.45$ and investment $\beta = 0.26$ using the point estimates of Cummins and Violante (2002). We select the share of leisure in the utility function $\psi = 2$ to generate a share of leisure approximately equal to 75%. The parameter α governing the degree of decreasing returns to scale corresponds approximately to the parameter governing adjustment costs; we thus set $\alpha = 0.4$, which is close to the quadratic adjustment cost case. We calibrate the arrival rate of new projects $\lambda = 0.25$ to

match the growth rate of the economy (0.45% per quarter). We parameterize the distribution of project quality θ to be exponential with unit scale parameter, $F(\theta) = 1 - \exp(-\theta)$, so that mean project quality is equal to one, $E(\theta) = 1$.

The remaining parameters are unique to our model and thus are difficult to calibrate using standard moments. We calibrate the parameters governing the technology process $[\xi_t, \mu_t]$ to generate realistic business cycle fluctuations. We calibrate the size of technological revolution to $\chi = 0.25$ to match the mean growth rate in the quality-adjusted price of equipment (-0.52% per quarter). We calibrate the rate of arrival of a technological revolution in the long-run state to be $\bar{\mu} = 0.025$, so that revolutions are relatively infrequent, occurring on average once every 10 years. We choose the parameters governing the volatility $\sigma_\mu = 0.1$, persistence $\kappa = 0.05$ and skewness $p = 0.25$ of the arrival process μ. These parameters lead to a quarterly volatility of output growth of 0.9%, which is in line with the data. We set the parameter governing the degree of efficiency loss when the project is implemented by the inventor to $R = 3$. Taken literally, this parametrization implies that the inventor pays three times as much for capital inputs than firms. However, this parameter captures not just the increased financing costs but also the reduction in overall efficiency of projects implemented by innovators alone. Last, we se the bargaining share between firms and workers to $\eta = 0.8$. Last, we calibrate a low value for the risk aversion coefficient, $\gamma = 2$. Risk aversion has a small impact on our results. More importantly, we calibrate the intertemporal elasticity of substitution to be on the high end of typical calibrations $\theta = 2$ to generate a smooth interest rate.

5.2 Response to news about future technological innovation

Here, we describe the response of model quantities to news about future innovations – a shock to μ – and improvements in the frontier technology – a shock to ξ. The response of the economy to news about future technological innovation sharply differentiates our model

from the frictionless benchmark. We show these responses in figures 1 and 2.

As we see in Panel A of figure 1, good news about the future leads to a *reduction* in the market value of new projects Me^{ω}, since the technology has not yet improved. Good news about future innovation reduces the incentives to invest capital today; this mechanism leads to a collapse in investment following positive news about the future in the standard model. However, as we see in Panel B, a positive shock to μ leads to an increase in the value of *future* projects relative to the value of new projects, $V/Me^{-\omega}$. This increase leads to increased ability of firms to commit to fund better projects – equation (42) – and therefore to an increase in the fraction of projects $F(\theta^*)$ that are implemented in a partnership, as we see in Panel C. A higher fraction of projects in partnership implies a reduction in the efficiency loss and thus to an increase in the average productivity of new projects $h(\theta_t^*)$, since more and better projects now operate at the optimal scale, as we see in panel D. As firms can commit to fund better projects – and these projects operate at a higher scale – this increase in the threshold θ_t^* tends to increase the demand for new capital. This is the novel mechanism in this paper. The overall demand for new capital is determined by the relative strength of the two channels in Panels s A and E. Last, in panel F we show that the equilibrium price of capital shows a weakly positive response to positive news about the future.

In Panel A of figure 2, we see that a positive shock to μ leads to an increase in investment, consumption, labor supply and capital utilization. Investment increases because the increase in the demand for capital due to the increase fraction of projects in partnership $F(\theta_t^*)$ is strong enough to overcome the reduction in the market value of existing capital (Panel D). The same mechanism leads to an increase in the price of investment goods and thus to an increase in the equilibrium wage; as a result, labor supply increases in response. Last, the reduction in the market value of existing projects (Panel D in previous figure) implies an increase in the rate of capital utilization: firms are willing to utilize old capital more and risk its destruction as its value drops. This increase in the supply of intermediate goods implies

that, even though the fraction of labor and intermediate goods allocated to consumption drops, aggregate consumption increases on impact.

Next, we compute the responses in a version of out model without frictions. By assuming that inventors and firms are equally adept at implementing projects, $R = 1$, we are effectively assuming away the existence of the limited commitment friction. We show the results in Panel B of figure 2. Similar to existing models in the literature, absent the limited commitment friction, good news about future technological growth leads to a recession today. An increase in the likelihood of technological innovation lowers the value of existing projects M_t. Absent the increase in capital demand due to the relaxation of the limited commitment friction, investment falls and consumption rises. Even though consumption increases on impact, this response is purely transitory; consumption decreases below the initial level as the economy accumulates less capital. Further, the increase in consumption induces a strong wealth effect, leading to a sharp drop in the supply of labor.

In summary, the limited commitment friction serves to generate increases in consumption, investment, and labor supply in response to positive news about future technological innovation. Good news about the future improves firms' ability to commit to implement better projects without expropriating inventors. This improvement in the quality of the project pool implemented in partnership leads to an investment boom and an increase in the supply of labor. For these responses, the limited commitment friction is key. Absent the limited commitment friction, our model displays similar behavior as existing models. News about future innovations induces a reduction in economic activity today. Since the marginal product of capital remains unchanged – the technology has not improved today – households cut investment in favor of higher consumption today.

5.3 Response to technological innovation

The response of the economy to the arrival of the technological revolution is qualitatively similar to models with embodied shocks. However, the limited commitment fiction serves to attenuate the economy's response to improvements in the technology frontier ξ. We show these responses in Figures 3 and 4.

An improvement in the level of frontier technology leads to an increase in the market value of new projects, as we see in Panel A of Figure 3. Even though the increase in technology is permanent, the market value of future projects does not quite increase as much as the value of projects today, as we see in Panel B. The reason is that due to the technological revolution, the economy will accumulate more capital in the future; in expectation, new projects in the future will be therefore less valuable than new projects today. As a result, firms are more tempted to steal inventors' ideas, hence fewer projects are implemented in partnership (Panel C). Since more projects are now inefficiently implemented, the average output of new projects – ignoring the increase in ξ – is lower relative to a model without the commitment friction, as we see in Panel D. Similarly, Panel E shows that the demand for new capital is lower relative to the frictionless model. This effect implies that the response of investment to improvements in the frontier technology is attenuated relative to the frictionless model. Last, Panel F shows that the price of capital p^I increases in response to improvements in frontier technology. In contrast, the quality-adjusted price of capital $p^I e^{-\xi}$ drops following an increase in ξ, illustrating the close connection between our model and models with investment-specific technical change.

In Figure 4, we compute the response of quantities to an improvement in the technology frontier (Panel A) and contrast these responses to a version of the model without frictions (Panel B). A shock to ξ improves the marginal efficiency of investment in both models; however, the limited commitment friction leads to endogenous attenuation as now *fewer* projects are implemented efficiently. Hence, the increase in investment, and labor supply,

is smaller in our model relative to the baseline case. In both models, improvements in the frontier technology lead to a reduction in the market value of *old* projects M, leading to an increase in the rate of utilization u. Part of this increase is allocated towards producing consumption goods, hence consumption also increases.

Last, we should note that this attenuation result is somewhat driven by the specifics of the model, namely that the arrival of technological improvement is not correlated with the arrival of future improvements. If shocks to ξ and μ were correlated, the relation between θ_t^* and ξ_t is ambiguous. An example of such a model would be the case where the arrival rate of innovations was unobservable, and the market made inferences about μ using public news and improvements in ξ.

5.4 Comovement

Here, we explore the extent to which our model can quantitatively generate fluctuations and comovement in quantities consistent with the data. In Table 1, we focus on business cycle frequencies (6 to 32 quarters). We see that, with the exception of hours worked, our model can generate realistic business cycle fluctuations. The volatility of output, consumption and investment is lower in the model than in the data, but the difference is not dramatic; the model estimates fall within the empirical confidence intervals. By contrast, hours worked are substantially less volatile in the model than in the data (0.41% vs 1.87%), though this is also generally true for the standard RBC model with only TFP shocks. Further, as we see in the bottom panel of Table 1, investment, consumption, output and hours worked in the model comove at business cycle frequencies, similar to the data.

Our model features two aggregate shocks: small and frequent fluctuations in the likelihood of innovations (news) and large infrequent changes in frontier technology. Though both shocks are responsible in generating comovement, the behavior of economic quantities at business cycle frequencies is dominated by the news shock. To illustrate this, panel C shows that the

benchmark model without the limited commitment friction generates negative comovement between consumption and investment and labor supply at business cycle frequencies.

Next, we focus on the behavior of quantities at lower frequencies (6 to 200 quarters). Comin and Gertler (2006), document the existence low-frequency fluctuations in quantities – medium run cycles using their terminology – that are substantially more volatile than the high-frequency fluctuations typically studied in the literature. Our model replicates this fact. As we see in Table 2, our model generates low frequency fluctuations in output, consumption, and investment that are in line with the data. As before, the exception is hours worked, which are substantially more volatile in the data than in the model. Focusing on the bottom panel of Table 2, our model closely generates comovement in economic quantities at lower frequencies that are in line with the data. Last, panel C illustrates that our friction is quantitatively important in generating comovement at lower frequencies; absent the limited commitment friction, consumption is negatively related to investment and labor supply at lower frequencies.

6 Conclusion

We introduce a limited commitment friction in a relatively standard real business cycle model with news. Firms implement projects in partnership with inventors. Firms cannot commit not to expropriate inventors whose ideas are of sufficiently high quality. This friction effectively restricts the supply of project ideas and therefore affects the equilibrium demand for capital. Good news about future technological innovations endogenously increases the supply of new ideas, and therefore affects the demand for capital today. Our mechanism allows us to generate realistic, news-driven business cycle fluctuations.

References

Arrow, Kenneth, 1962, Economic welfare and the allocation of resources for invention, in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, NBER Chapters, 609–626 (National Bureau of Economic Research, Inc).

Atanasov, Vladimir, Vladimir Ivanov, and Kate Litvak, 2012, Does reputation limit opportunistic behavior in the VC industry? Evidence from litigation against VCs, *The Journal of Finance* 67, 2215–2246.

Banerjee, Abhijit V, and Andrew F Newman, 1993, Occupational choice and the process of development, *Journal of Political Economy* 101, 274–98.

Basu, Susanto, John G. Fernald, and Miles S. Kimball, 2006, Are technology improvements contractionary?, *American Economic Review* 96, 1418–1448.

Beaudry, Paul, and Bernd Lucke, 2009, Letting different views about business cycles compete, Working Paper 14950, National Bureau of Economic Research.

Beaudry, Paul, and Franck Portier, 2006, Stock prices, news, and economic fluctuations, *American Economic Review* 96, 1293–1307.

Bhide, A.V., 1999, *The Origin and Evolution of New Businesses* (Oxford University Press, USA).

Chen, Kaiji, and Zheng Song, 2012, Financial frictions on capital allocation: A transmission mechanism of tfp fluctuations, Working paper, University of Chicago.

Chiu, Jonathan, Cesaire Meh, and Randall Wright, 2011, Innovation and growth with financial, and other, frictions, NBER Working Papers 17512, National Bureau of Economic Research, Inc.

Christiano, Lawrence, Roberto Motto, and Massimo Rostagno, 2013, Risk shocks, Working Paper 18682, National Bureau of Economic Research.

Christiano, Lawrence J., and Terry J. Fitzgerald, 1999, The band pass filter, Working Paper 9906, Federal Reserve Bank of Cleveland.

Comin, Diego, and Mark Gertler, 2006, Medium-term business cycles, *American Economic Review* 96, 523–551.

Cummins, Jason G., and Giovanni L. Violante, 2002, Investment-specific technical change in the U.S. (1947-2000): Measurement and macroeconomic consequences, *Review of Economic Dynamics* 5, 243–284.

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

Fisher, Jonas D. M., 2009, Comment on 'letting different views about business cycles compete', Working Paper Series WP-2010-01, Federal Reserve Bank of Chicago.

Gompers, Paul, Anna Kovner, Josh Lerner, and David Scharfstein, 2008, Venture capital investment cycles: The impact of public markets, *Journal of Financial Economics* 87, 1–23.

Gompers, Paul, and Josh Lerner, 2006, *The Venture Capital Cycle*, volume 1 of *MIT Press Books*, second edition (The MIT Press).

Gordon, Robert J., 1990, *The Measurement of Durable Goods Prices* (University of Chicago Press).

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman, 1988, Investment, capacity utilization, and the real business cycle, *American Economic Review* 78, 402–17.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, 1997, Long-run implications of investment-specific technological change, *American Economic Review* 87, 342–362.

Hayashi, Fumio, 1982, Tobin's marginal q and average q: A neoclassical interpretation, *Econometrica* 50, pp. 213–224.

Israelsen, Ryan D., 2010, Investment based valuation, Working paper, Indiana University.

Jaimovich, Nir, and Sergio Rebelo, 2009, Can news about the future drive the business cycle?, *American Economic Review* 99, 1097–1118.

Jermann, Urban J., and Vincenzo Quadrini, 2007, Stock market boom and the productivity gains of the 1990s, *Journal of Monetary Economics* 54, 413–432.

Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti, 2011, Investment shocks and the relative price of investment, *Review of Economic Dynamics* 14, 101–121.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, 2010, Investment shocks and business cycles, *Journal of Monetary Economics* 57, 132–145.

Kaplan, Steven N., and Per Stromberg, 2004, Characteristics, contracts, and actions: Evidence from venture capitalist analyses, *Journal of Finance* 59, 2177–2210.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2012, Technological innovation, resource allocation, and growth, Working paper 17769, National Bureau of Economic Research.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2013, Technological innovation: Winners and losers, Working Paper 18671, National Bureau of Economic Research.

Nanda, Ramana, and Matthew Rhodes-Kropf, 2011, Investment cycles and startup innovation, Harvard Business School Working Papers 12-032, Harvard Business School.

Papanikolaou, Dimitris, 2011, Investment shocks and asset prices, *Journal of Political Economy* 119, pp. 639–685.

Schmitt-Grohe, Stephanie, and Martin Uribe, 2012, What's news in business cycles .

Silveira, Rafael, and Randall Wright, 2010, Search and the market for ideas, *Journal of Economic Theory* 145, 1550–1573.

Solow, Robert M., 1960, Investment and technical progess, in Kenneth J. Arrow, Amuel Karlin, and Patrick Suppes, eds., *Mathematical Methods in the Social Sciences*, 89–104 (Stanford University Press, Stanford, CA).

Appendix A: Analytical Appendix

Given a disclosure policy $P(\theta)$, define $r(\theta) = 1$ if $P(\theta) = 1$ and $r(\theta) = R$ otherwise. We guess that the market value at time t of a project of vintage s, scale k and quality θ is $M_t \xi_s \theta^{1-\alpha} k^{\alpha}$. Given a disclosure policy, the owner of a project chooses scale k to maximize

$$NPV_t = \max_k \left\{ M_t \, e^{\xi_t} \theta^{1-\alpha} k^{\alpha} - q_t r(\theta) k \right\}$$

$$\text{yielding} \qquad k_t^*(\theta) = \theta \left(\frac{\alpha M_t \, e^{\xi_t}}{q_t} \right)^{\frac{1}{1-\alpha}} \left(\frac{1}{r(\theta)} \right)^{\frac{1}{1-\alpha}} = \theta \left(\frac{1}{r(\theta)} \right)^{\frac{1}{1-\alpha}} k_t$$

We guess that the optimal partnership policy takes the form $P(\theta) = 1$ iff $\theta < \theta_t^*$. Define

$$g(x) \equiv \int_0^x \theta \, dF(\theta) + \left(\frac{1}{R} \right)^{1/(1-\alpha)} \int_x^{\infty} \theta \, dF(\theta)$$

$$h(x) \equiv \int_0^x \theta \, dF(\theta) + \left(\frac{1}{R} \right)^{\alpha/(1-\alpha)} \int_x^{\infty} \theta \, dF(\theta)$$

Next, we solve for the NPV of new projects. The total demand for new capital equals

$$\lambda \int_0^{\infty} k_t^*(\theta) dF(\theta) = \lambda \left(\frac{\alpha M_t e^{\xi_t}}{q_t} \right)^{\frac{1}{1-\alpha}} g(\theta_t^*)$$

After clearing the market for investment goods $\lambda \int k^*(\Omega_t, \theta) dF(\theta) = I_t$, the equilibrium level of investment depends on

$$k_t \equiv \frac{L_{It}^{1-\beta} y_{It}^{\beta} Y_t^{\beta}}{\lambda g(\theta_t^*)},$$

the market clearing price of new capital goods is

$$q_t = \alpha e^{\xi_t} \, M_t \left(\frac{L_{It}^{1-\beta} y_{It}^{\beta} Y_t^{\beta}}{\lambda g(\theta_t^*)} \right)^{\alpha-1}$$

and the NPV of a new project is

$$NPV_t = \theta \, (1-\alpha) \left(\frac{1}{r(\theta)} \right)^{\alpha/(1-\alpha)} M_t \; e^{\xi_t} \, k_t^{\alpha}.$$

Taking expectations over θ

$$\int_0^{\infty} NPV_t dF(\theta) = (1-\alpha) \, h(\theta_t^*) M_t \; e^{\xi_t} \, k_t^{\alpha}$$

The next step is to solve for the state-dependent threshold θ_t^*. Suppose that the relationship value to the financier is V_t. The outside option of the inventor is denoted by Π^O and is equal to (6) with $r(\theta) = R$ always. The inventor and the firm split the rents from the project surplus according to

Nash bargaining. Suppose the rents that accrue to inventor are

$$\Pi_t^I = \Pi_t^O + \eta(V_t - \Pi_t^O)$$

Denote Π_t^* to be the first-best surplus, which is equal to (6) with $r(\theta) = 1$ always. The cooperation profits of the firm equal

$$\Pi_t^F(\theta) = \pi_t^* - \Pi_t^O - \eta(V_t - \Pi_t^O) = \theta \left(1 - (1-\eta)\, R^{\alpha/(\alpha-1)}\right)\, (1-\alpha)\, M_t\, e^{\xi_t}\, k_t^\alpha - \eta V_t$$

which in expectation (over θ) equal

$$\int_0^{\theta^*} \Pi_t^F(\theta)\, dF(\theta) = \hat{g}(\theta^*)\, M_t\, e^{\xi_t}\, k_t^\alpha - \eta F(\theta_t^*) V_t$$

where

$$\hat{g}(x) \equiv \left(1 - (1-\eta) R^{a/(a-1)}\right)\, (1-\alpha)\, \int_0^x \theta\, dF(\theta)$$

The relationship value to the financier is

$$V_t = \lambda\, E_t \int_t^\infty \frac{\pi_s}{\pi_t} \left(\hat{g}(\theta_s^*)\; M_s\, e^{\xi_s}\, k_s^\alpha - \eta F(\theta_s^*) V_s\right) ds$$

The firm can credibly commit not to expropriate as long as $\Pi_t^O \leq V_t$. Hence, the quality of the marginal project θ_t^* is the solution to

$$\theta_t^*\, (1-\alpha)\, \left(\frac{1}{R}\right)^{\alpha/(1-\alpha)}\; M_t\, e^{\xi_t}\; \left(\frac{L_{It}^{1-\beta} y_{It}^\beta Y_t^\beta}{\lambda g(\theta_t^*)}\right)^\alpha = V_t$$

Next, we solve for the evolution of the aggregate state variables. Define the effective capital stock K_t as the the total output of intermediate good sector if it was working at full capacity

$$Y_t = u_t K_t,$$

$$\text{where} \qquad K_t = \int_{-\infty}^t e^{-\int_t^s \delta_u\, du} \left(\lambda \int_0^\infty e^{\xi_s}\, \theta_j^{1-\alpha}\; (k_t^*(\theta))^\alpha\, dF(\theta)\right) ds$$

which evolves according to

$$\frac{dK_t}{K_t} = (i_t - \delta_t)\; dt$$

where $i_t \equiv \lambda \left(\frac{L_{It}^{1-\beta} y_{It}^\beta u_t^\beta}{\lambda g(\theta_t^*)}\right)^\alpha e^{\omega_t}\, h(\theta_t^*)$ and δ_t is the equilibrium depreciation rate determined below. The optimality conditions of consumption and investment good producers imply equations (44)-(47) in the text. The household's labor supply decision is intra-temporal, $U_N/U_C = w_t$,implying

$$N = \frac{\psi}{1-\phi}\, L_C$$

Combining (44)-(47), with the labor market clearing $L_I + L_C = 1 - N$, determines L_{It}, L_{Ct}, Y_{It}, and Y_{Ct} as a function of the aggregate state. Let $y_{Ct} = Y_{Ct}/Y_t$ and $y_{It} = Y_{It}/Y_t$. We conjecture – and subsequently verify – that the variables u_t, L_{Ct}, L_{It}, N_t, y_{Ct}, y_{It}, i_t, and δ_t are functions only of ω and μ. The next step involves computing M_t and V_t. The stochastic discount factor π_t is given by

$$\pi_t = \exp\left(\int_0^t \tilde{f}_J(C_s, N_t, J_s)\, ds\right) \tilde{f}_C(C_t, N_t, J_t)$$

We guess that the value function takes the form

$$J_t = \frac{1}{1-\gamma} K_t^{\phi(1-\gamma)} j(\omega_t, \mu_t)$$

The household's value function satisfies the HJB equation $0 = \tilde{f}(C, N, J) + \mathcal{D}_{\omega,K,\mu} J$. After substituting our guess for J and simplifying, we obtain the following PDE for j

$$\begin{aligned}
=&\rho \frac{1-\gamma}{1-\theta^{-1}} \left(u(\omega,\mu)^{\phi}\, y_C(\omega,\mu)^{\phi} L_C(\omega,\mu)^{1-\phi}\right)^{1-\theta^{-1}} j(\omega_t,\mu_t)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} \left(N(\omega,\mu)\right)^{\psi(1-\theta^{-1})} \\
&- \left[\rho \frac{1-\gamma}{1-\theta^{-1}} - \phi\,(1-\gamma)\Big(i(\omega,\mu) - \delta(\omega,\mu)\Big)\right] j(\omega,\mu) - (1-\alpha\beta)\Big[(i(\omega,\mu) - \delta(\omega,\mu)\Big] j_\omega(\omega,\mu) \\
&+ \mu\Big(j(\omega+\chi,\mu) - j(\omega,\mu)\Big) + \kappa_\mu(\bar{\mu}-\mu)\, j_\mu(\omega,\mu) + \frac{1}{2}\sigma_\mu^2\, \mu^{2p}\, j_{\mu\mu}(\omega,\mu)
\end{aligned}$$

The value of an existing project is $M_t \xi_s \theta^{1-\alpha} k^{\alpha}$, where

$$\begin{aligned}
M_t =& \max_u E_t \int_t^\infty e^{\int_t^s h_J(C_u,J_u) - \delta_u du} \frac{\pi_s}{\pi_t} p_{Ys}\, u_s\, ds \\
M_t =& K_t^{\phi-1}\, m(\omega_t,\mu_t) \left(\left(\bar{u}_t^{\phi}\, y_{Ct}^{\phi} L_{Ct}^{1-\phi}\right)^{-\theta^{-1}} j(\omega_t,\mu_t)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N_t^{\psi(1-\theta^{-1})}\right)^{-1}
\end{aligned}$$

When optimizing over capital utilization, the firm takes the aggregate utilization $\bar{u}$ and depreciation $\bar{\delta}$ as given. The value of a project satisfies the HJB equation. After simplifying,

$$\begin{aligned}
0 =& \max_u \Bigg\{ \phi \bar{u}_s^{-1}\, u_s\, y_{Cs}^{-1} \left(\bar{u}_s^{\phi}\, y_{Cs}^{\phi} L_{Cs}^{1-\phi}\right)^{1-\theta^{-1}} j(\omega_s,\mu_s)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N(\omega,\mu)^{\psi(1-\theta^{-1})} - (1-\alpha\beta)\Bigg[i(\omega,\mu) - \delta(\omega,\mu)\Bigg] \frac{\partial}{\partial\omega} m(\omega,\mu) \\
&+ \kappa_\mu(\bar{\mu}-\mu)\frac{\partial}{\partial\mu} m(\omega,\mu) + \mu\left(m(\omega+\chi,\mu) - m(\omega,\mu)\right) + \frac{1}{2}\sigma_\mu^2 \mu^{2p} \frac{\partial^2}{\partial\mu^2} m(\omega,\mu) \\
&- m(\omega,\mu)\Bigg[- h_J(\omega,\mu) - (\phi(1-\gamma)-1)\Big(i(\omega,\mu) - \delta(\omega,\mu)\Big) + \delta(u)\Bigg]\Bigg\}
\end{aligned}$$

where the first order condition for $u(\omega,\mu)$ is

$$\phi \bar{u}^{-1}\, y_C^{-1} \left(\bar{u}^{\phi}\, y_C^{\phi} L_C^{1-\phi}\right)^{1-\theta^{-1}} j(\omega,\mu)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N^{\psi(1-\theta^{-1})} = m(\omega,\mu)\, \delta'(u)$$

symmetry $\bar{u} = u$ implies that the function $u(\omega, \mu)$ is the solution to

$$\frac{\left(y_C(\omega,\mu)^{\phi}\, L_C(\omega,\mu)^{1-\phi}\right)^{1-\theta^{-1}}\, N(\omega,\mu)^{\psi(1-\theta^{-1})} j(\omega,\mu)^{\frac{\gamma-\theta^{-1}}{\gamma-1}}}{y_C(\omega,\mu)\, m(\omega,\mu)} = \delta'(u(\omega,\mu))u(\omega,\mu)^{1+\phi(\theta^{-1}-1)}$$

verifying our conjecture that u_t and δ_t are functions of ω and μ only. Last, we solve for the relationship values V_t of the firm. We conjecture that the relationship value takes the form

$$V_t = K_t^{\phi}\, v(\omega_t,\mu_t)\Big[\left(\bar{u}_t^{\phi}\, y_{Ct}^{\phi} L_{Ct}^{1-\phi}\right)^{-\theta^{-1}} j(\omega_t,\mu_t)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} N_t^{\psi(1-\theta^{-1})}\Big]^{-1}$$

Indeed,

$$\begin{aligned}\pi_t V_t =& E_t \int_t^{\infty} \pi_s\, \lambda \left(\hat{g}(\theta_s^*)\, M_s\, e^{\xi_s} \left(\frac{L_{Is}^{1-\beta} y_{Is}^{\beta} Y_s^{\beta}}{\lambda g(\theta_s^*)}\right)^{\alpha} - \eta F(\theta_s^*) V_s\right) ds \\ =& E_t \int_t^{\infty} \exp\left(\int_t^s h_J(C_u, J_u)\, du\right) \rho\, K_s^{\phi\,(1-\gamma)} \lambda \left(\hat{g}(\theta_s^*)\, m(\omega_t,\mu_t)\, e^{\omega_s} \left(\frac{L_{Is}^{1-\beta} y_{Is}^{\beta} u_s^{\beta}}{\lambda g(\theta_s^*)}\right)^{\alpha} - \eta F(\theta_s^*) v(\omega_t,\mu_t)\right) ds.\end{aligned}$$

Using the Feynman-Kac theorem with discounting, we get a PDE for v

$$\begin{aligned}0 =& \Bigg\{ \lambda\, \hat{g}(\theta^*(\omega,\mu))\, m(\omega,\mu)\, e^{\omega} \left(\frac{L_I^{1-\beta}(\omega,\mu) y_I^{\beta}(\omega,\mu) u^{\beta}(\omega,\mu)}{\lambda\, g(\theta^*(\omega,\mu))}\right)^{\alpha} - (1-\alpha\beta)\Big[i(\omega,\mu) - \delta i(\omega,\mu)\Big] \frac{\partial}{\partial\omega} v(\omega,\mu) \\ &+ \kappa_{\mu}(\bar{\mu}-\mu)\frac{\partial}{\partial\mu} v(\omega,\mu) + \mu\Big(v(\omega+\chi,\mu) - v(\omega,\mu)\Big) + \frac{1}{2}\sigma_{\mu}^2 \mu^{2p} \frac{\partial^2}{\partial\mu^2} v(\omega,\mu) \\ &- v(\omega,\mu)\Big(- h_J(\omega,\mu) - \phi(1-\gamma)\,(i(\omega,\mu) - \delta(\omega,\mu)) + \eta\,\lambda\, F(\theta^*(\omega,\mu))\Big)\Bigg\}\end{aligned}$$

Last, we verify our conjecture that the variables L_{Ct}, L_{It}, N_t, y_{Ct}, y_{It}, and θ_t^* are only functions of ω and μ. Combining (44)-(47), labor supply and labor market clearing, yields after simplification

$$\begin{aligned}&\frac{1-\phi}{1-\beta}(1-L_{It})^{-1}\left(\frac{1-\phi}{\psi+1-\phi}\right)^{-1} = \\ &\lambda^{1-\alpha}\, g(\theta_t^*)^{1-\alpha}\, \alpha e^{\omega_t}\, m(\omega_t,\mu_t)\, \left(y_C^{\phi} u_t^{\phi}\, L_C^{1-\phi}\right)^{\theta^{-1}-1} N_t^{-\psi(1-1/\theta)}\, j(\omega_t,\mu_t)^{\frac{\gamma-\theta^{-1}}{1-\gamma}}\, u_t^{\alpha\beta}\, y_{It}^{\alpha\beta} L_{It}^{\alpha(1-\beta)-1}\end{aligned}$$

and

$$y_I(\omega,\mu) = \frac{\beta\,(\psi+1-\phi) L_I(\omega,\mu)}{1 - L_I(\omega,\mu) + L_I(\omega,\mu)\,\frac{\beta\,(\psi+1-\phi)}{\phi\,(1-\beta)}}$$

Equation (42) characterizing the threshold θ_t^* simplifies to

$$\theta_t^* \left(\frac{1}{R}\right)^{\alpha/(1-\alpha)} (1-\alpha)\; m(\omega_t,\mu_t)\, e^{\omega_t} \left(\frac{L_{It}^{1-\beta} y_{It}^{\beta}}{\lambda g(\theta_t^*)}\right)^{\alpha} u_t^{\alpha\beta} = v(\omega_t,\mu_t)$$

Since $y_C = 1 - y_I$ and $N + L_I + L_C = 1$ our conjecture is verified.

Appendix B: Data construction

Productivity is utilization-adjusted TFP from Basu et al. (2006). Populations is from the U.S. Census Bureau (http://www.census.gov/popest/national/national.html). Output and consumption are from the Bureau of Economic Analysis. Output is gross domestic product (NIPA Table Table 1.1.5) divided by the consumption price index (St Louis Fed, CPIAUCNS). Consumption is consumption of non-durables plus services, deflated by the price index of non-nudrables and services respectively (NIPA Tables 1.1.5, 2.3.4). To get hours worked, we merge series CEU0500000007 and EEU00500005 from the BLS, times total private employment(BLS, CEU0500000001) divided by population. The price of new equipment is from Israelsen (2010), who extends the quality-adjusted series constructed by Gordon (1990) and extended by Cummins and Violante (2002). The data on patents is from Kogan, Papanikolaou, Seru, and Stoffman (2012).

Table 1: Volatility and comovement, business cycle frequencies

	A. Data				B. Model				C. Frictionless Model			
	c	i	y	l	c	i	y	l	c	i	y	l
	0.83	6.52	1.71	1.87	0.71	4.01	1.06	0.41	0.95	6.27	1.05	1.09
	[0.47, 1.71]	[3.67, 13.40]	[0.96, 3.51]	[1.05, 3.84]	[0.40, 1.08]	[2.20, 5.88]	[0.58, 1.68]	[0.18, 0.72]	[0.52, 1.44]	[3.48, 9.37]	[0.49, 1.88]	[0.53, 1.85]
					Correlation							
	c	i	y		c	i	y		c	i	y	
i	73.6				82.5				-43.7			
	[60.0, 87.2]				[58.0, 94.4]				[-76.4, -9.2]			
y	85.3	87.2			96.6	93.6			35.0	61.7		
	[78.9, 91.7]	[78.6, 95.8]			[92.7, 99.0]	[82.4, 98.0]			[-1.9, 73.3]	[17.7, 86.9]		
l	82.4	76.7	87.0		74.9	98.4	88.9		-53.9	97.3	55.3	
	[74.7, 90.1]	[64.9, 88.6]	[81.1, 93.0]		[49.6, 92.0]	[96.7, 99.6]	[72.4, 96.9]		[-82.0, -20.2]	[93.9, 99.2]	[5.5, 84.0]	

Table compares moments of consumption, investment, output and hours in the data (Panel A) and in the model (B) at business cycle frequencies (6-32 quarters). All series are filtered using the bandpass filter (Christiano and Fitzgerald, 1999).

Table 2: Volatility and comovement, medium run

	A. Data				B. Model				C. Frictionless Model			
	c	i	y	l	c	i	y	l	c	i	y	l
	3.01	10.60	4.02	4.01	2.69	10.61	3.32	1.05	2.83	15.92	3.63	2.58
	[1.69, 6.19]	[5.96, 21.79]	[2.26, 8.26]	[2.26, 8.25]	[0.83, 6.19]	[4.90, 16.63]	[1.20, 6.77]	[0.40, 1.82]	[0.97, 6.33]	[7.29, 23.95]	[1.09, 7.66]	[0.96, 4.35]
					Correlation							
	c	i	y		c	i	y		c	i	y	
i	65.2				67.8				-6.1			
	[51.3, 79.1]				[38.7, 89.3]				[-67.8, 49.8]			
y	92.6	79.3			95.7	85.3			59.4	69.9		
	[88.7, 96.5]	[71.2, 87.5]			[90.9, 98.7]	[70.0, 96.1]			[16.0, 88.9]	[12.0, 90.2]		
l	63.5	63.9	74.7		49.5	96.1	71.1		-21.1	97.3	59.6	
	[52.1, 74.8]	[51.2, 76.6]	[66.4, 83.0]		[-4.1, 87.9]	[86.2, 99.4]	[25.6, 93.7]		[-77.9, 39.6]	[93.3, 99.3]	[-5.3, 86.5]	

Table compares moments of consumption, investment, output and hours in the data (Panel A) and in the model (B) at medium-run frequencies (6-200 quarters), as in Comin and Gertler (2006). All series are filtered using the bandpass filter (Christiano and Fitzgerald, 1999).

Figure 1: Response to news

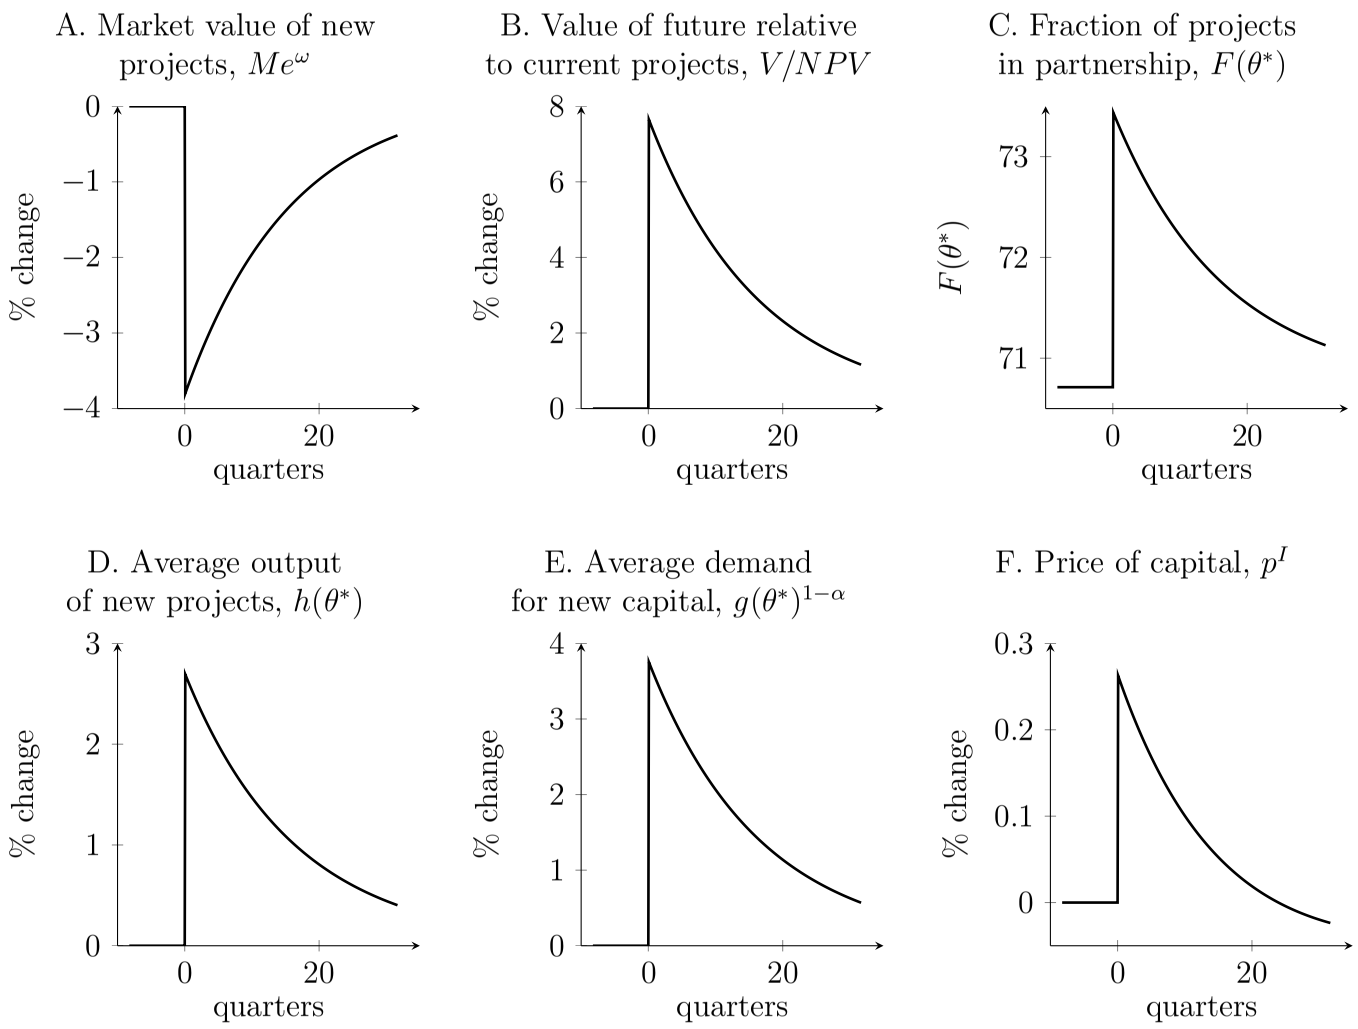

Figure 2: Response to news

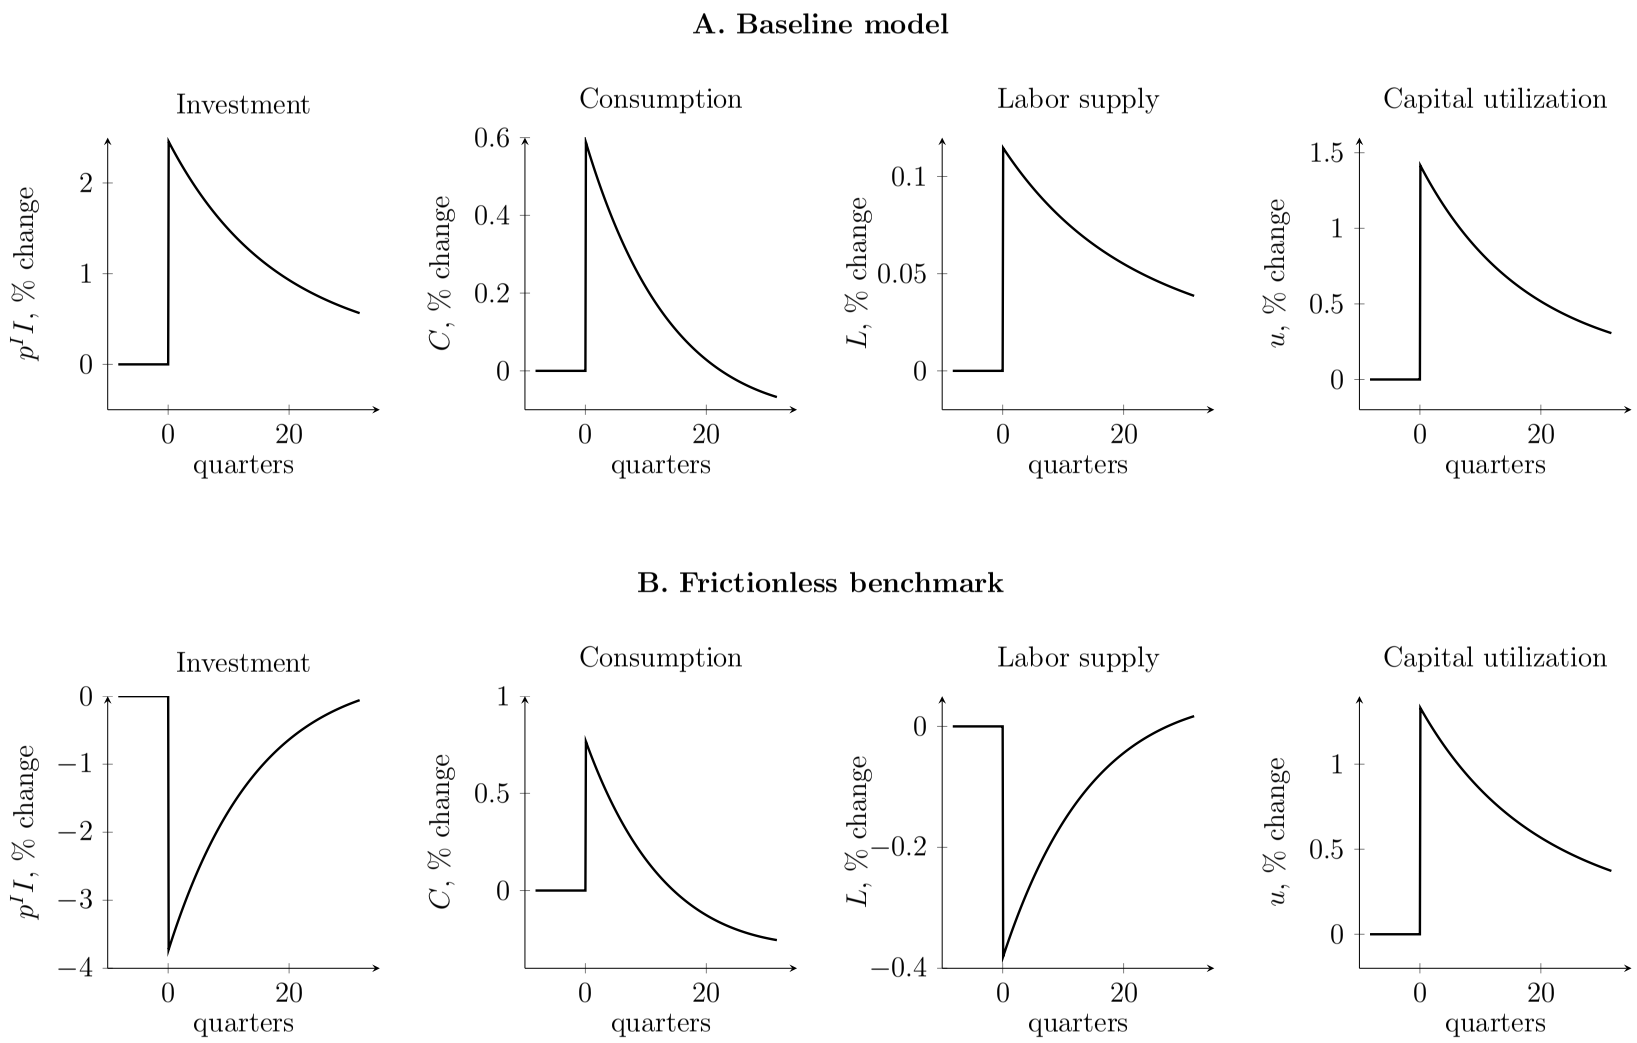

Figure 3: Response to technological innovation

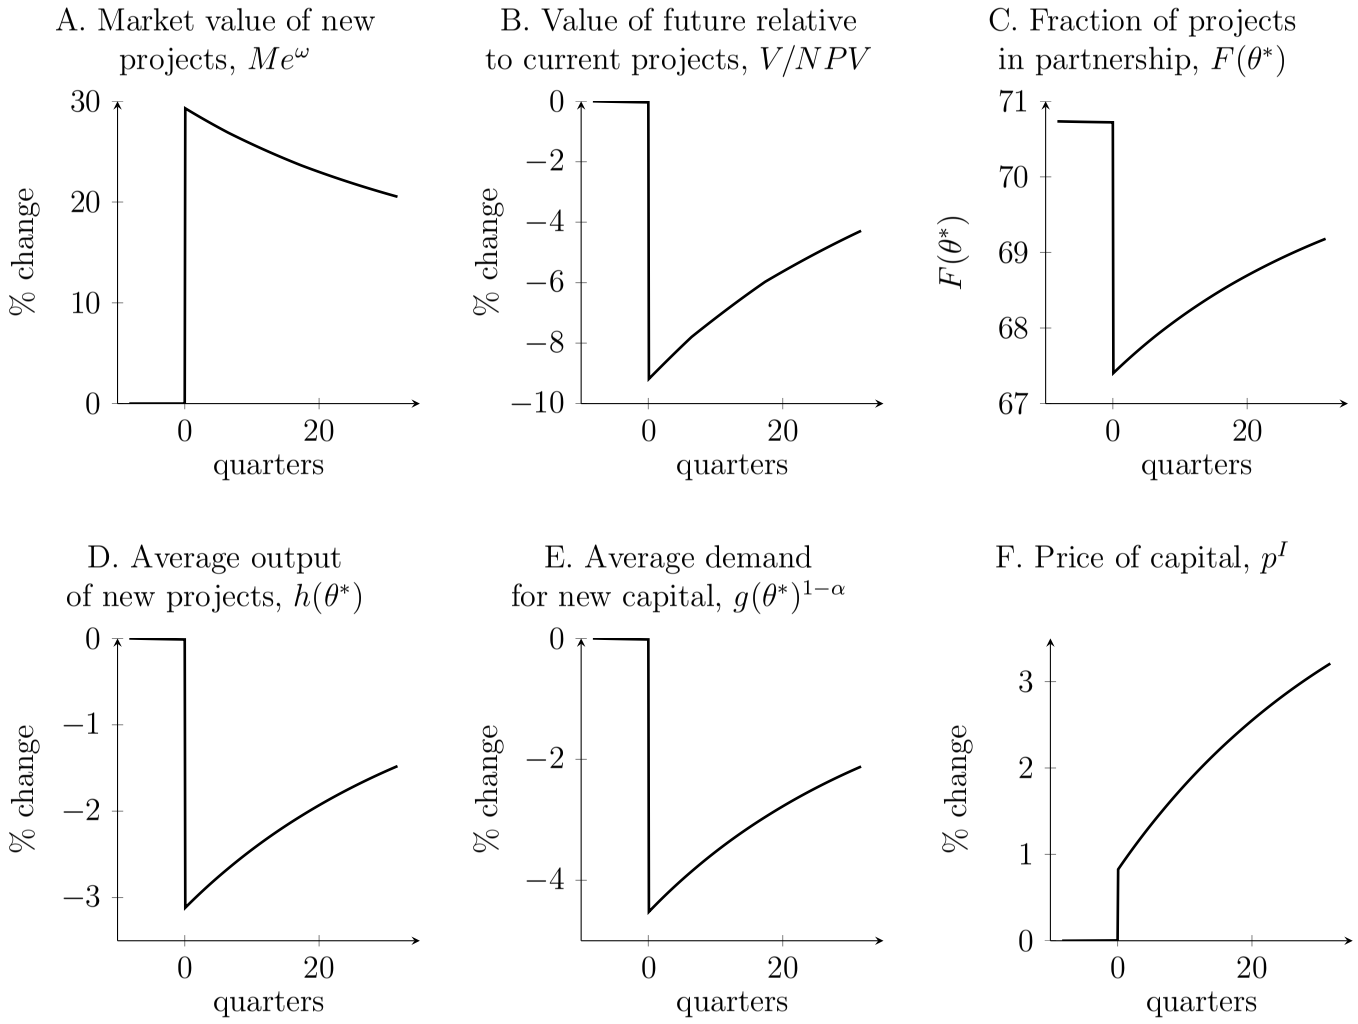
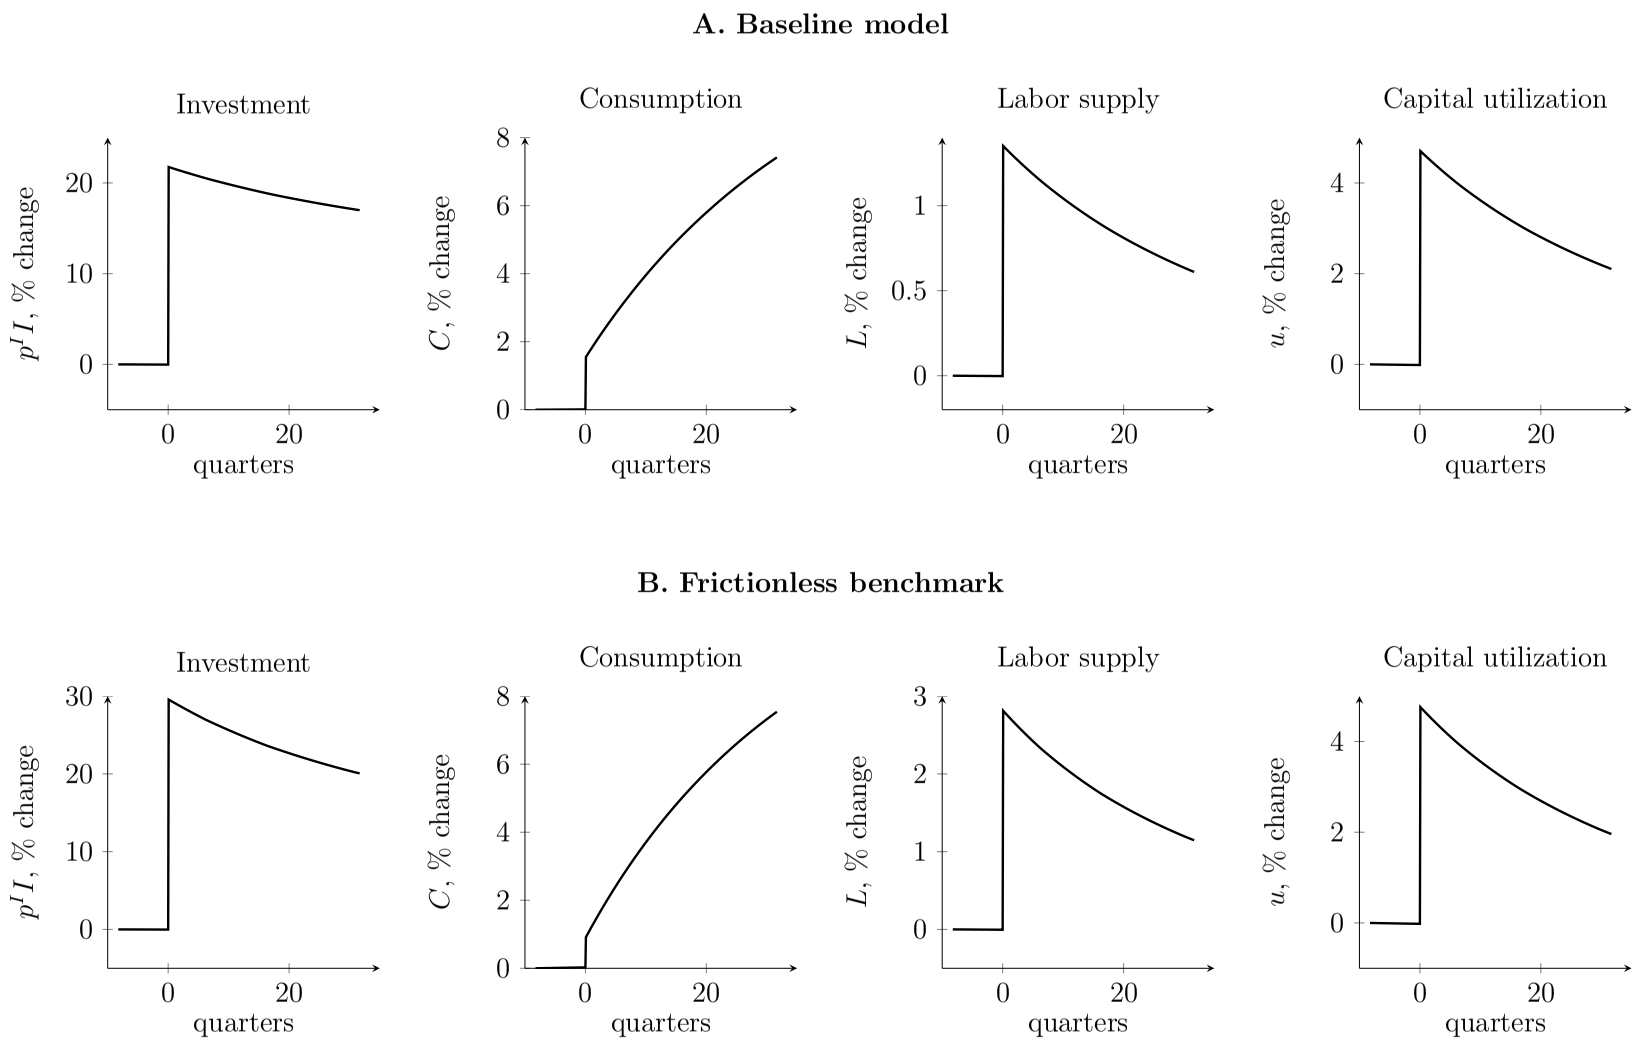
Figure 4: Response to technological innovation