Essays in Entrepreneurial Finance

Citation

Bernstein, Shai. 2012. Essays in Entrepreneurial Finance. Doctoral dissertation, Harvard University.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:10085978

Terms of Use

Share Your Story

Accessibility

Dissertation Advisor: Andrei Shleifer   Shai Bernstein

Essays in Entrepreneurial Finance

ABSTRACT

In the first essay, I show that the transition to public equity markets have important implications to firms' innovative process. To establish a causal effect of the IPO, I compare the long-run innovation of firms that completed their filing and went public with that of firms that withdrew their filing and remained private. I use NASDAQ fluctuations during the book-building period as a source of exogenous variation that affects IPO completion but is unlikely to affect long-run innovation. Using this approach, I find that the quality of internal innovation declines by 50 percent relative to firms that remained private. The decline in innovation is driven by both an exodus of skilled inventors and a decline in productivity among remaining inventors. However, going public allows firms to attract new human capital and purchase externally generated innovations through mergers and acquisitions.

In the second essay, we explore the effects of private equity investments on the industries they invest at. This analysis looks across nations and industries to assess the impact of private equity on industry performance. Industries where PE funds have invested in the past five years have grown more quickly in terms of productivity and employment. It is hard to find support for claims that economic activity in industries with private equity backing is more exposed to aggregate

shocks. The results using lagged private equity investments suggest that the results are not driven by reverse causality.

Finally, in the third essay we model situations in which a principal offers a set of contracts to a group of agents to participate in a project. Agents' benefits from participation depend on the identity of other participating agents. We show that when assuming multilateral externalities, the optimal contracts' payoff relies on a ranking of the agents, which can be described as arising from a tournament among the agents. Rather than simply ranking agents according to a measure of popularity, the optimal contracting scheme makes use of a more refined two-way comparison between the agents. We derive results on the principal's revenue extraction and the role of the level of externalities' asymmetry.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

This is it. A five-year adventure comes to an end. Five years of excitement, frustration, inspiration, confusion, joy, and overall, hmmm...fun. I was fortunate to have gone through this journey, and I am deeply thankful to many people that made it such an experience.

First and foremost I am grateful to the members of my dissertation committee: Fritz Foley, Josh Lerner, Jeremy Stein and Andrei Shleifer. Josh, I have known you since my first year in the program, and thanks to you I was exposed to the fascinating world of entrepreneurial finance. I learned tremendously from you, either as your student or as your collaborator - always positive, generous, and amazingly productive. Andrei, your advice and direction, also in non-research areas, shaped my decisions throughout the program. Thinking about the big picture, and seeking new angles and perspectives are just few of my takeaways from our numerous interactions. Jeremy, thank you for pushing me relentlessly to think carefully about my research and leave no stone unturned. I am truly a better researcher thanks to you. Fritz, our long and thoughtful conversations helped improving not only my research but also my well-being.

I owe special thanks also to Sam Hanson, Robin Greenwood, and Adi Sunderam who provided invaluable guidance and advice (and a great mock interviews practice). Many many thanks also to Malcolm Baker, Effi Benmelech, Lauren Cohen, Ben Esty, Paul Gompers, Victoria Ivashina, Bill Kerr, Ramana Nanda, David Scharfstein, and Antoinette Schoar for their help and advice.

Of course, this Ph.D. would never have been such a great experience if not for

1. DOES GOING PUBLIC AFFECT INNOVATION?

1.1 Introduction

Does the transition to public equity markets affect innovation? Although a large body of research examines the performance of firms around their initial public offerings (IPOs), little is known about the effects of going public on innovation. This question is particularly relevant given the reliance of young and entrepreneurial firms on public equity issuances to fund their R&D investments,[1] and the critical role of innovation in promoting economic growth (Solow 1957). This paper studies the effects of going public on three important dimensions of firms' innovative activity: internally generated innovation, productivity and mobility of individual inventors, and acquisition of external innovation.

Theoretically, in frictionless financial markets selling equity publicly should have no bearing on subsequent innovative activity. However, two broad views suggest that going public should in fact matter.

The "financing" view suggests that going public may enhance innovation by overcoming financing frictions and easing access to capital. As argued by Arrow (1962) and demonstrated empirically,[2] R&D is likely to be more sensitive to financ-

---

[1] See Brown, Fazzari and Petersen (2009). In fact, Brown and Petersen (2009) demonstrate that young firms' dependence on public equity markets to finance R&D expenditure has even increased over past decades.

[2] See, for example, Brown, Fazzari and Petersen (2009), Himmelberg and Petersen (1994), and

1

ing constraints than other forms of investments. For instance, debt financing of R&D may be limited due to associated information problems, skewed and uncertain returns, and the potentially scant collateral value of intangible assets. Equity financing, on the other hand, allows investors to share upside returns and can ease the financing of R&D investments by transferring idiosyncratic innovation risk to diversified investors through public equity markets. Therefore, the financing view suggests that going public may enhance internally generated innovation and may even facilitate technology acquisitions.

Alternatively, the "incentives" view suggests that ownership dilution and changes in governance may lead to a decline in the quality of innovation. Following the IPO, inventors may face weaker incentives to pursue novel projects as their claims on subsequent innovations become smaller. Increases in wealth and the ability to cash out may weaken inventors' incentives even further. In addition, since equity markets may fail to correctly evaluate innovation even when outcomes are predictable and persistent (Cohen, Diether, and Malloy, 2011), career concerns and takeover threats may pressure managers to select less novel projects that are more easily communicated to stock market investors[3] (Stein, 1989; Ferreira, Manso, and Silva, 2010). Interestingly, the benefits of accessing public markets can be tied to its costs. Managers may prefer to exploit improved access to capital to acquire ready-made technologies rather than innovating internally, as this strategy is more transparent to the stock market and potentially less prone to failure.

To shed light on these two views, I use standard patent-based metrics to study

---

Mulkay, Hall, and Mairesse (2001). For detailed surveys of the literature see Bond and Van Reenen (2007) and Hall and Lerner (2009).

[3] Minton and Kaplan (2008) demonstrate that turnover rates in publicly traded firms are high and significantly related to a firm's stock performance. Additionally, Edmans, Goldstein, and Jiang (2011) find that low stock prices strongly affect the likelihood of takeover threats.

the effects of going public on innovation. Consistent with the incentives view, the main finding of the paper illustrates that going public changes firms' strategies in pursuing innovation. Following the IPO there is a substantial decline in the novelty of internally generated innovation which is partially mitigated by the increased likelihood to acquire innovation externally.

Estimating the effects of going public on innovation is challenging due to an inherent selection bias associated with the decision to go public. A standard approach in the literature uses within-firm variation to study the dynamics of firm outcomes around the IPO.[4] But, as noted by Jain and Kini (1994), this approach is likely to reveal biased IPO effects due to the selection of firms to go public at a specific stage in their life cycle. For instance, firms may choose to go public following an innovative breakthrough, as hypothesized by Pastor, Taylor, and Veronesi (2009).[5] In that case, the post-IPO performance may be affected by reversion to the mean, reflecting life cycle effects in addition to the IPO effects.

To overcome this selection bias, I construct a novel dataset of innovative firms that filed an initial registration statement with the SEC and either completed or withdrew their filing. This sample allows me to compare the innovative activity of firms that went public with private firms at a similar stage in their life cycle, namely, firms that intended to go public at the same time but withdrew their filing. But this does not completely eliminate the selection bias as the decision to withdraw may be related to a firm's R&D policy and innovative opportunities.

I use the two-month NASDAQ fluctuations following the IPO filing date as an

---

[4] See, e.g., Degeorge and Zeckhauser (1993), Jain and Kini (1994), Mikkelson, Partch, and Shah (1997), Pagano, Panetta, and Zingales (1998), Pastor, Taylor, and Veronesi (2009), and Chemmanur, He, and Nandy (2009).

[5] Chemmanur, He, and Nandy (2009) find that firms go public following productivity improvements, and experience a decline in productivity following the IPO.

instrument for IPO completion. The instrument relies on the sensitivity of filers to stock market movements during the book-building phase (Busaba, Beneveniste, and Guo, 2000; Benveniste et al., 2003; Dunbar, 1998; Dunbar and Foerster, 2008; Edelen and Kadlec, 2005). These fluctuations provide a plausibly exogenous source of variation that affects IPO completion and is unlikely to be related to innovation.

One concern regarding the instrument might be that the exclusion restriction does not hold; i.e., that two-month NASDAQ returns may relate to innovation measures through channels other than the IPO completion (see Section 2.C for a detailed discussion). There are several reasons this may not be the case. First, the analysis compares firms that filed to go public in the same year. I find that the characteristics of filers that experienced a NASDAQ drop during the book-building phase do not differ significantly from other firms that filed to go public during the same year but did not experience such a decline.[6] Second, the analysis uses firm innovation measures that are in relative terms, scaled by the average innovation measures of all patents granted in the same year and in the same technology class.[7] Therefore, even if two-month NASDAQ returns contain information about aggregate changes in innovative opportunities, such a change should affect all firms conducting research in the same area, and is therefore unlikely to affect relative innovation measures.

Using this instrumental variables approach, I find that going public caused a substantial decline of approximately 50 percent in innovation novelty as measured by patent citations. At the same time, I find no change in the scale of innovation,

---

[6] These characteristics include: firm innovation in the three years before the IPO filing, firm financials at the time of the IPO filing, venture capital backing, age, underwriter ranking, and location within the IPO wave.

[7] Technology classes are defined by the United States Patent and Trademark Office (USPTO), and capture thechnological essence of an invention.

4

as measured by the number of patents. These results suggest that the transition to public equity markets leads firms to reposition their R&D investments toward more conventional projects. Such findings cannot be explained by the financing view which suggests that improved access to capital may enhance innovative activities.

To uncover the channels driving the decline in innovative activity, I study the effects of going public on individual inventors' productivity and mobility over time. Consistent with the incentives view, I find that the quality of innovation produced by inventors who remained at the firm substantially declines following the IPO and key inventors are more likely to leave. These effects are partially mitigated by the ability of public firms to attract new inventors.

I also find a stark increase in the likelihood that newly public firms acquire companies in the years following an IPO, particularly privately held targets. To better understand whether these acquisitions are used for purchasing new technologies, I collect information on targets' patent portfolios. I find that public firms acquire a substantial number of patents through M&A: acquired patents constitute more than one-fifth of firms' patent portfolio in the five years following the IPO. The acquired patents are more likely to be in technologies that are only weakly related to a firm's previous patents and are of higher quality than the patents produced internally. These findings are broadly consistent with both the financing and the incentives view.

To further investigate the underlying causes, I propose two incentives-related explanations. The first explanation suggests that career concerns lead managers to select more incremental projects, while the second explanation suggests that after the IPO inventors are facing weaker incentives to pursue high-quality innovation.

5

While the two explanations are likely to co-exist, I find supportive evidence for the first explanation indicating that changing managerial incentives and public market pressures affect innovation at public firms. If managerial incentives are an important determinant of innovation, firms with more entrenched managers should be less sensitive to market pressures and therefore may invest in more novel projects. As a proxy for managerial entrenchment, I use cases in which the CEO also serves as the chairman of the board. I find that when managers are more entrenched, the negative effect of going public on innovation novelty is weaker and inventors are less likely to leave the firm.

The paper is related to several strands in the literature. First, the IPO literature documents a post-IPO decline in firm performance measures such as profitability and productivity.[8] This paper adds to the literature by demonstrating a post-IPO decline in innovation. Perhaps more importantly, the paper establishes that this decline is caused by the IPO, rather than being a symptom of a particular stage of the firm life cycle. This paper is also related to a number of papers studying withdrawn IPOs.[9] By using patent data, this study is the first to investigate the performance consequences of the decision to withdraw an IPO.

The paper reveals a complex trade-off between public and private ownership. While private firms are able to generate higher quality innovation and retain skilled inventors, public firms can acquire technologies externally and attract new human capital. In that regard, the paper is also related to a growing literature that com-

---

[8] Several papers report a post-IPO decline in profitability: Degeorge and Zeckhauser (1993), Jain and Kini (1994), Mikkelson, Partch, and Shah (1997), Pagano, Panetta, and Zingales (1998), and Pastor, Taylor, and Veronesi (2009). Chemmanur, He, and Nandy (2009) reach similar findings regarding firm productivity.

[9] For example, Benveniste et al. (2003), Busaba, Benveniste, and Guo (2000), Busaba (2006), Dunbar (1998), Dunbar and Foerster (2008), Edelen and Kadlec (2005), and Hanley (1993).

pares the behavior of public and private firms along various dimensions such as investment sensitivity, capital structure, and dividend payouts.[10] Additionally, this work contributes to the theoretical and empirical literature that explores the role of governance, capital structure, and ownership concentration on corporate innovation.[11]

The rest of the paper proceeds as follows. Section 1.2 outlines the main identification strategy. Section 1.3 explains the various data sources used to construct the sample. Section 1.4 presents the results on the effects of going public on internal innovation, inventors' mobility and productivity, and firm reliance on external technologies. Section 1.5 discusses several theoretical explanations and Section 1.6 provides a conclusion.

1.2   Empirical Strategy

In this section, I discuss the standard patent-based metrics used in the analysis to measure firm innovation. Then, I describe the empirical strategy and the instrumental variables approach used in the paper.

1.2.1   Measuring Innovative Activity

An extensive literature on the economics of technological change demonstrates that patenting activity reflects the quality and extent of firm innovation. I use

---

[10] Several aspects of firm behavior are considered in that literature. Asker, Farre-Mensa, and Ljungqvist (2010), and Sheen (2009) focus on investment sensitivity, Saunders and Steffen (2009) and Brav (2009) study debt financing and borrowing costs, Michaely and Roberts (2007) explore dividend payouts, and Gao, Lemmon, and Li (2010) focus on CEO compensation.

[11] See Aghion, Van Reenen, and Zingales (2009), Atanassov, Nanda, and Seru (2007), Belenzon, Berkovitz, and Bolton (2009), Bhattacharya and Guriev (2006), Chemmanur and Jiao (2007), Fulgheieri and Sevilir (2009), Fang, Tian, and Tice (2010), Lerner, Sorensen, and Stromberg (2010), and Tian and Wang (2010).

widely accepted patent-based metrics to measure firm innovative activity (Hall, Jaffe, and Trajtenberg, 2001; Lanjouw, Pakes, and Putnam, 1998). These measures not only capture firms' technological contribution but are also economically meaningful and have been shown to translate into firm market value (see, e.g., Trajtenberg, 1990; Hall, Jaffe, and Trajtenberg, 2005).

The most basic measure of innovative output is a simple count of the number of patents granted. However, patent counts cannot distinguish between breakthrough innovation and incremental discoveries (see, e.g., Griliches, 1990). The second metric, therefore, reflects the importance or novelty of a patent by counting the number of citations a patent receives following its approval.[12] Hall, Jaffe, and Trajtenberg (2005) illustrate that citations are a good measure of innovative quality and economic importance.[13]

Both citation rates and patent filing propensity vary over time and across technologies. Variations may stem from changes in the importance of technologies over time or from changes in the patent system. Therefore, a comparison of raw patents and citations is only partially informative. To adjust for these variations, I follow Hall, Jaffe, and Trajtenberg (2001) and scale each patent citation count by the average citations of matched patents. Matched patents are defined as patents that are granted in the same year and in the same technology class.[14] Specifically, let $Cites_{itk}$

---

[12] I count citations in the year of patent approval and three subsequent calendar years. I discuss the choice of citation horizon window in Section 2.B.

[13] Specifically, they find that extra citation in firm's average citations per patent increases market value by 3%.

[14] A technological class is a detailed classification defined by the U.S. Patenting and Trademark Office (USPTO) that captures the essence of an invention. Technological classes are often much more refined than industry classifications, consisting of about 400 main (3-digit) patent classes, and over 120,000 patent subclasses. For example, under the "Communications" category one can find numerous sub-categories such as wave transmission lines and networks, electrical communications, directive radio wave systems and devices, radio wave antennas, multiplex communications, optical wave

8

be the number of citations of patent i that was granted in year t and classified in technology class k. The *scaled citations* of patent i, $SCites_{itk}$, is $Cites_{itk}$ divided by $\overline{Cites}_{-itk}$, the average number of citations of all patents granted in the same year and in the same technology class excluding patent i, that is,

$$SCites_{itk} = \frac{Cites_{itk}}{\overline{Cites}_{-itk}}$$

Similarly, to adjust for variations in patent-filing likelihood, each patent is scaled by the average number of patents generated by firms in the same year and in the same technology class. Hence, patents that were granted in technologies in which firms issue more patents receive less weight. The *scaled patent count* per year is a simple sum of the scaled patents.

The final measures, *originality* and *generality*, use the distribution of citations to capture the fundamental nature of research (Trajtenberg, Jaffe, and Henderson, 1997). A patent that cites a broader array of technology classes is viewed as having greater originality. A patent that is being cited by a more technologically varied array of patents is viewed as having greater generality.[15] Similarly to patents and citations, *scaled originality* and *scaled generality* are scaled by the corresponding average originality or generality of all patents granted in the same year and technology class.

---

guides, pulse or digital communications, etc.

[15] The originality (generality) measure is the Herfindahl index of the cited (citing) patents, used to capture dispersion across technology classes. I use the bias correction of the Herfindahl measures, described in Jaffe and Trajtenberg (2002) to account for cases with a small number of patents within technological categories.

1.2.2 Empirical Design

Identifying the effects of going public on innovation or firm outcomes in general is challenging due to inherent selection issues that arise from the decision of firms to go public. A common approach in the literature[16] uses "within-firm" estimator that compares the performance of the same firm before and after the IPO. This method is attractive as it provides an estimate of the impact of IPOs on innovation that is not affected by a firm's time-invariant characteristics. At the same time, however, this method fails to control for the selection of firms to go public at a specific stage in its life cycle. If firms are more likely to go public following a positive innovative shock,[17] as argued by Pastor, Taylor, and Veronesi (2009), regressions designed to capture the effect of going public may be biased by life cycle effects and reversion to the mean.

To overcome the selection bias associated with firms' decision to go public, I construct a dataset that includes innovative firms that submitted the initial registration statement to the SEC in an attempt to go public. Following the IPO filing, firms engage in marketing the equity issuance to investors during the book-building phase and have the option to withdraw the IPO filing. I compare the long-run innovation of firms that went public (henceforth 'IPO firms') with firms that filed to go public at the same year, but ultimately withdrew their filing and remained private (henceforth 'withdrawn firms'). This setup is attractive as it allows the comparison of the post-IPO performance of firms that went public with

---

[16] Degeorge and Zeckhauser (1993), Jain and Kini (1994), Mikkelson, Partch, and Shah (1997), Pagano, Panetta, and Zingales (1998), Pastor, Taylor, and Veronesi (2009), and Chemmanur, He, and Nandy (2009).

[17] As illustrated in panel D of Table 1, I find that in the three years prior to the IPO, firms produce substantially more novel patents than comparable patents within the same year and technology class.

that of private firms at a similar stage in their life cycle. My baseline specification of interest is

$$Y_i^{post} = \alpha_1 + \beta_1 IPO_i + \gamma_1 Y_i^{pre} + X_i' \delta_1 + \nu_k + \mu_t + \varepsilon_{1i} \qquad (1.1)$$

Y_i^{post} is the average innovative performance in the five years following the IPO filing: average scaled citations, average scaled originality/generality and average scaled number of patents per year. Y_i^{pre} is the equivalent measure in the three years prior to the IPO filing.[18] IPO_i is the dummy variable of interest, indicating whether a filer went public or remained private. Under the null hypothesis that going public has no effect on innovation, β_1 should not be statistically different from zero. This model includes industry (ν_k) and IPO filing year (μ_t) fixed effects.

If the decision to withdraw an IPO filing is related to unobserved firm innovation policy or innovative opportunities (captured in the error term), the β_1 estimate may be biased. Therefore, I instrument for the IPO completion choice. Specifically, I use the two-month NASDAQ returns as an instrument, calculated from the IPO filing date (i.e., the first two months of the book-building phase). The figure below illustrates the time line of the IPO filing and the NASDAQ fluctuations during the book-building phase. Firms either choose to complete the IPO or to withdraw their filing. On average, ownership choices are accepted within four months following the IPO filing. The firm-level innovation is measured over the five-year horizon after the IPO filing:[19]

---

[18] Adding a constraint of $\gamma_1 = 1$ in the model specified in equation (1) implies that the dependent variable is equivalent to innovative performance difference before and after the IPO filing. However, absent of this constraint, the above specification is more flexible and capable of capturing potential reversion to the mean that may arise following the IPO filing.

[19] The results of the analysis remain unchanged if innovation measures are calculated from the

NASDAQ fluctuations provide a plausibly exogenous source of variation that leads some firms to remain private in spite of their IPO filing. To implement the instrumental variables approach, I estimate the following first-stage regression:

$$IPO_i = \alpha_2 + \beta_2 NSDQ_i + \gamma_2 Y_i^{pre} + X_i' \delta_2 + \nu_k + \mu_t + \varepsilon_{2i}$$

where $NSDQ_i$ is the instrumental variable. The second-stage equation estimates the impact of IPO on firm innovative activity:

$$Y_i^{post} = \alpha_3 + \beta_3 \widehat{IPO}_i + \gamma_3 Y_i^{pre} + X_i' \delta_3 + \nu_k + \mu_t + \varepsilon_{1i}$$

where $\widehat{IPO}_i$ are the predicted values from (2). If the conditions for a valid instrumental variable are met, β_3 captures the causal effect of an IPO on innovation outcomes. I implement the instrumental variable estimator using two-stage least squares. I also use a quasi-maximum likelihood (QML) Poisson model to estimate the IV specification (Blundell and Powell, 2004). This model, which I describe in the Appendix, is the standard estimation method used in the innovation literature and count data analysis more generally.

To illustrate the advantage of using this instrumental variables approach in this setting consider a simple example. Assume that firm innovation following to the IPO filing is the sum of future innovation opportunities (which are unobserved at the time of the IPO filing) and the effect of ownership structure (being public or private). Specifically, the post-IPO innovative performance can be written as $Q + c \cdot IPO$, where Q stands for the unobserved quality of the issuer's future innovative

_____

ownership choice date rather than IPO filing date, as patent filings during the book-building period are not common.

12

projects, and *IPO* is a dummy that indicates whether the issuer completed the IPO filing ($IPO = 1$) or remained private ($IPO = 0$). The goal is to estimate c, the effect of public ownership on firm innovation.

Suppose that the unobserved quality of future projects is heterogeneous and affects the likelihood of completing the IPO filing. Specifically, there are three types of firms: *Sure Thing* firms, with highest-quality future innovative projects ($Q = q_H$), will complete the IPO irrespective of market conditions; *Sensitive* firms, with medium-quality innovative projects ($Q = q_M$), will not complete the IPO filing if NASDAQ drops, but will go public otherwise; and *Long Shot* firms, with the poorest innovative prospects ($Q = q_L$), will withdraw irrespective of the NASDAQ change.[20] For simplicity, assume that NASDAQ can be either *high* or *low* each with probability of 1/2, and firm types are equally likely.

The OLS estimate simply compares firms that completed the IPO filing (the upper triangle) and firms that withdrew the IPO filing (the bottom triangle) and reflects the sum of the IPO effect as well as a selection bias:

$$\gamma_{OLS} = E\left[Y|IPO = 1\right] - E\left[Y|IPO = 0\right] = c + \frac{2}{3}(q_H - q_L) > c$$

Thus OLS will overestimate the effect of going public in this example because better firms are more likely to complete the IPO filing.[21]

The instrumental variables approach uses the variation in the NASDAQ – which affects the decision to complete or withdraw the IPO filing – to estimate the effects

---

[20] The decision to withdraw or complete the IPO filing is complicated and driven by many observed and unobserved factors. For simplicity in this example I assume that the decision depends only on one factor.

[21] Note that if one assumes that lower quality firms are more likely to complete the IPO filing then the sign of the bias reverses.

of an IPO on innovative outcomes. Intuitively, this is equivalent to calculating the difference in performance across columns. Specifically, simply comparing outcomes based on the NASDAQ returns generates the "reduced-form" regression:

$$E\left[Y|NSDQ = High\right] - E\left[Y|NSDQ = Low\right] = \frac{1}{3}c$$

The "first-stage" regression captures the likelihood to complete the IPO as a function of the NASDAQ variation:

$$E\left[IPO|NSDQ = High\right] - E\left[IPO|NSDQ = Low\right] = \frac{1}{3}$$

Scaling the reduced-form result by the first-stage regression coefficient generates the desired outcome:

$$\gamma_{IV} = \frac{E\left[Y|NSDQ = High\right] - E\left[Y|NSDQ = Low\right]}{E\left[IPO|NSDQ = High\right] - E\left[IPO|NSDQ = Low\right]} = c$$

The example illustrates that the IV estimator uses only the *Sensitive* firms whose IPO completion depends on NASDAQ conditions. In fact, any instrumental variables estimator use only the information of the group of firms that respond to the instrument (Imbens and Angrist, 1994).

In the example I assumed for the sake of simplicity that NASDAQ returns can take two values. Clearly, NASDAQ returns vary considerably. When the instrument is multi-valued the IV estimate is a weighted average of the sensitive subpopulation estimates along the support of the instrument (Angrist and Imbens, 1995).[22]

---

[22] Different firms have different thresholds of NASDAQ changes for which they complete the IPO filing. Roughly speaking, the IV estimate is an average of the estimates of sensitive firms along

So far, I made two important assumptions. First, I assumed that NASDAQ conditions are not correlated with firm characteristics, and second that NASDAQ returns do not affect future innovative performance. These assumptions determine the validity of the instrument. In the next section I discuss these assumptions.

1.2.3 NASDAQ Fluctuations and the Exclusion Restriction

For the instrument to be valid, it must strongly affect IPO completion choices. Additionally, it must not affect the scaled innovation measures through any channel other than the decision to complete the IPO filing. Formally, this means that the two-month NASDAQ returns must be uncorrelated with the residual in equation (1). This residual reflects unobservable characteristics that may influence firm innovation. The latter requirement, the "exclusion restriction", is the focus of the discussion below.

I start by exploring whether firms that experience a NASDAQ drop are significantly different from other firms that filed to go public within the same year. A priori this seems unlikely since it would require high-frequency compositional shifts in IPO filers, in contrast to the evidence of clustering of similar firms in IPO markets that attempt to exploit information spillovers (Beneveniste et al. 2003).

I find no significant differences in observables between firms that experienced a NASDAQ drop and other firms that filed in the same year. As illustrated in Section 3.D, these observables include characteristics at the time of the IPO filing such as firm financial information, age, venture capital backing, IPO filing characteristics, and importantly, innovation performance in the three years before the IPO

different values of NASDAQ returns. The average is weighted by the impact of NASDAQ returns on completing the IPO filing, and by the likelihood of observing the NASDAQ returns.

filing. This suggests that the two-month NASDAQ fluctuations are not correlated with firm observables and are plausibly exogenous within a year. To further address concerns about within-year compositional shifts, I control also for the three-month NASDAQ returns before the IPO filing, and for firms' location within the IPO wave.

The two-month NASDAQ fluctuations may reflect either a change in investor sentiment or in future innovative opportunities. If NASDAQ fluctuations reflect changes in future innovative opportunities, this would raise concerns regarding the exclusion restriction. However, since R&D expenditure is a slow-moving process, firms that file within the same year are likely to respond to similar changes in innovation opportunities (Hall, Griliches, and Hausman 1986; Lach and Schankerman, 1989).

Additionally, since my innovation measures are scaled by average outcomes and therefore expressed in relative terms within the same year and technology class, changes in aggregate innovative opportunities reflected by the two-month NASDAQ returns should affect all firms conducting research in the same technology. Such changes are not likely to affect the relative innovative performance. For example, consider a firm that submitted an IPO filing in 1995 and was awarded a patent three years later in 1998 in the optical communications technology. The novelty of the patent is scaled by the average novelty of all patents granted in 1998 in the optical communications technology. If the two-month NASDAQ returns following the IPO filing in 1995 reflected a change in innovative opportunities in optical communications in coming years, and thus affected patent novelty, this change should affect the novelty of all patents within this technology class. Hence, the relative patent novelty is unlikely to be affected.

To further address concerns regarding the exclusion restriction I conduct two additional tests reported in Section 3.E. First, I perform a placebo test. The exclusion restriction requires that the two-month NASDAQ returns affect innovation only through the ownership choice channel. If this is the case, we should expect that two-month NASDAQ returns *after* the ownership choice was made should have no effect on long-run innovation. Indeed, I find that in contrast to the two-month NASDAQ returns immediately after the IPO filing, following the IPO completion choice, the two-month NASDAQ returns have no predictive power. This evidence suggests that the effect of the instrument on the long-run innovation of the firm goes through the IPO completion channel.

As a second test, I investigate directly whether the instrument can explain changes in innovative trends. I use all patents granted by the U.S. Patent and Trademark Office to calculate changes in innovative trends in the core technologies of firms as of the time of their IPO filing. I find no evidence that the two-month NASDAQ returns can predict changes in these innovative trends. While firms may switch to different technologies subsequent to the IPO filing, this test suggests that, whether or not such a switch occurred, it is not likely to be driven by the two-month change in the NASDAQ.

1.3   Data

The dataset is constructed from several data sources combining IPO filings, patent information, hand-collected financial information and other firm characteristics. In this section I describe the steps in constructing the dataset, and provide summary statistics comparing IPO firms and withdrawn firms at the time of the filing.

17

1.3.1   IPO Filings

To apply for an IPO, a firm is required to submit an initial registration statement to the SEC (usually the S-1 form), which contains the IPO filer's basic business and financial information. Following the submission of S-1 form, issuers engage in marketing the equity issuance to investors (the "book-building" phase) and have the option to withdraw the IPO filing by submitting RW form. The most common stated reason for withdrawing is "weak market conditions".

Filing withdrawals are common in IPO markets, as approximately 20 percent of all IPO filings are ultimately withdrawn. As noted by Busaba et al. (2001), the decision to withdraw is driven by various observed and unobserved considerations that affect the investors' willingness to pay and the issuer's reservation value.[23] As long as the investors' valuation is higher than the issuer's reservation price, the firm will complete the IPO application.

I identify all IPO filings using Thomson Financial's SDC New Issues database. The sample starts in 1985, when SDC began covering withdrawn IPOs systematically, and ends in 2003 since the analysis explores the innovative outcomes of firms in the five years after the IPO filing. Following the IPO literature, I exclude IPO filings of financial firms (SIC codes between 6000 and 6999), unit offers, closed-end funds (including REITs), ADRs, limited partnerships, special acquisition vehicles, and spin-offs. I identify 5,583 complete IPOs and 1,599 withdrawn IPO filings in the period of 1985 - 2003.

---

[23] The investors' valuation may be affected by the issuer's financials, innovative activity, sentiment, and other unobserved factors. Similarily, the issuer's reservation value is influenced by future investment opportunities, cash reserves, alternative funding options, and other unobserved elements such as entrepreneur's benefits from diversification and loss of private benefits of control.

1.3.2 Patent Data

The patent data is obtained from the National Bureau of Economic Research (NBER) patent database, which includes detailed information on more than three million patents submitted to the U.S. Patent and Trademark Office (USPTO) from 1976 to 2006 (Hall, Jaffe, and Trajtenberg, 2001).

I use the NBER bridge file to COMPUSTAT to match patents to firms that completed the IPO filing. Since withdrawn firms are not included in COMPUSTAT, I match these firms based on company name, industry, and geographic location, all of which are available in SDC and IPO registration forms. In ambiguous cases where firm names are similar but not identical, or the location of the patentee differs from the SDC records or SEC registration statements, I conduct web and FACTIVA searches to verify matches.

I restrict the sample to firms with at least one successful patent application in the three years before and five years after the IPO filing; this yields 1,488 innovative firms that went public and 323 that withdrew the IPO application.

The goal is to collect information on firms' innovative activity in the five years after the IPO filing. In some cases, firms are acquired, or withdrawn firms may go public in a second attempt.[24] I collect information on firms' patenting activity even after such firm exits, to avoid biases that may arise from truncating firm activity. After all, firm exits are yet another consequence of the IPO effect that influence firms' innovative path. Collecting patent information subsequent to firms' exits is complicated since if a firm is acquired its patents may be assigned to the acquiring firm. Nevertheless, I find that in most cases patents are still assigned to the acquired company after its acquisition. This allows me to capture the patenting

---

[24] See Panel F in Table 1 for a description of the acquisition statistics of IPO and withdrawn firms.

activity in more than 90 percent of firm-year observations, irrespective of whether a firm was acquired. In the remaining firm-years, no patent was assigned to the acquired firm. This could be either because the acquired firm did not generate additional patents, or because any patents generated were assigned to the acquiring company. To identify missing patents, I use inventor identifiers and geographic location to isolate patents that were produced by the acquired rather than the acquiring company.[25]

I calculate the number of citations a patent receives in the calendar year of its approval and in the subsequent three years. This time frame is selected to fit the nature of the sample. Since many of the IPO filings in the sample occur toward the end of the 1990s, increasing the time horizon of citation counts will reduce sample size. Given that citations are concentrated in the first few years following patent's approval and the considerable serial correlation in citation rates (Akcigit and Kerr, 2011), three years is reasonably sufficient to capture the patent's importance.[26]

Since the NBER patent database ends in 2006, I supplement it with the Harvard Business School (HBS) patent database, which covers patents granted through December 2009. This enables calculating the citations of patents granted toward the end of the sample. Overall, the sample consists of 39,306 granted patents of IPO firms and 4,835 granted patents of withdrawn firms.

Panel A of Table 1 summarizes the distribution of IPO filings by year. IPO filings are concentrated in the 1990s and drop after 2000, with 95 of the 323 withdrawn filings occurring in 2000. The absence of transactions conducted before

---

[25] Specifically, I start by collecting inventor identifiers of patents produced by the acquired firm before the acquisition. These unique inventor identifiers are available through the Harvard Business School patent database. Then, I go over the patents produced by the acquiring firm in the post-acquisition years, to identify all patents produced by the same inventors.

[26] I verify that the results are not sensitive to the selected citation horizon.

1985 and after 2003 reflects the construction of the sample. Panel A also displays the patent applications and awards of IPO firms and withdrawn firms separately. Each patent is associated with an application date and grant date, reflecting the lag in patent approvals. Since the sample includes only patents granted by December 2006, the number of approved patent applications declines in 2005 and 2006.

Panels B and C detail the composition of firms and patents across industries and technology classes. The majority of the firms in the sample are concentrated in technological industries such as electronic equipment, software, drugs, and medical equipment. Similarly, most patents are concentrated in the industries that rely on intellectual property, such as computer, drugs, and electronics industries.

Panel D compares the patenting activity of withdrawn and IPO firms in the three years prior to the IPO filing. I find no significant differences across any of the patenting measures. Since a value of one in the scaled citations measure implies that a firm is producing patents of average quality, it is interesting to note that both IPO firms and withdrawn firms produce patents that are substantially more frequently cited than comparable average patents (80 percent higher for withdrawn firms and 89 percent higher for IPO firms). This evidence suggests that firms that select to go public are likely to do so following innovative breakthroughs, which may raise concerns of post-IPO reversion to the mean.

1.3.3   Financial Information and Firm Characteristics

The analysis of private firms is complicated by data limitations. While patents are useful in capturing the innovative activity of both public and private firms, no financial information is readily available for withdrawn firms when using standard financial databases. To partially overcome this constraint, I collect withdrawn

Table 1 - Summary Statistics

The table reports summary statistics of the key variables in the analysis, which are defined in Section A of the Appendix. Panel A describes the distribution of IPO filings and patents over time. Panels B and C detail the distribution of firms across industries and the distribution of patents across technology classes. The industry classification is based on Fama-French 10, and the technology classification is based on Hall, Jaffe, and Trajtenberg (2001). Panel D describes average innovative measures in the three years up to (and through) the IPO filing year. Panel E provides information on firm characteristics at the time of filing. Panel F describes firm exit characteristics in the five years after the IPO filing, where firm exits are corporate events such as acquisition, bankruptcy, or an IPO of withdrawn firms. *, **, and *** indicate that differences in means are statistically significant at the 10%, 5%, and 1% levels.

Panel A - Distribution by year

Year	IPO Filing Complete	IPO Filing Withdrawn	Patent Applications Complete	Patent Applications Withdrawn	Patent Grants Complete	Patent Grants Withdrawn
1983	N/A	N/A	4	2	0	0
1984	N/A	N/A	18	9	1	0
1985	4	2	16	8	9	8
1986	10	5	58	18	9	5
1987	11	6	111	17	39	11
1988	14	4	202	34	62	13
1989	42	6	356	74	147	27
1990	34	10	527	86	231	56
1991	120	2	715	62	321	59
1992	119	33	1169	125	525	68
1993	144	14	1457	106	797	89
1994	105	18	2152	162	1050	87
1995	140	8	3568	318	1309	94
1996	169	29	3220	262	1760	133
1997	114	25	3857	444	2298	199
1998	66	20	3672	509	3317	310
1999	169	15	4249	634	3658	388
2000	167	95	4225	586	3360	457
2001	17	13	4144	555	3448	531
2002	12	17	3082	431	3483	517
2003	21	1	1795	256	3678	533
2004	N/A	N/A	616	117	3547	465
2005	N/A	N/A	89	20	2943	376
2006	N/A	N/A	4	0	3314	409
Total	1478	323	39306	4835	39306	4835

firms' financial information from initial registration statements. I download the S-1 forms from the SEC's EDGAR service, which is available from 1996. For IPO firms, I rely on standard financial databases such as COMPUSTAT and CapitalIQ to collect firm financial information. This allows me to compare withdrawn and IPO firms' characteristics at the time of filing.

I collect additional information on firm characteristics from various sources. I obtain data on venture capital (VC) funding from SDC, VentureXpert, and registration statements. I supplement the data with information on firms' age at the time of the IPO filing and its underwriters' ranking obtained from registration forms, VentureXpert, Jay Ritter's webpage, and the SDC database. Finally, I collect information on firms' exits, i.e., events in which firms were acquired, went public in a second attempt (for withdrawn firms), or filed for bankruptcy. I use COMPUSTAT and CapitalIQ to search for acquisitions and bankruptcies, and the SDC database to identify second IPOs of withdrawn firms. I perform extensive checks to verify the nature of private firms' exits using the Deal Pipeline database, Lexis-Nexis and web searches.

Panel E compares the characteristics of IPO firms and withdrawn firms at the time of filing. I find no significant differences in firm size (measured by log firm assets) and R&D spending (normalized by firm assets). However, withdrawn firms have a higher Cash-to-Assets ratio and are lower net income to assets.

The literature often uses the reputation of the lead underwriter as a proxy for firm quality, based on the rationale that higher-quality firms are more likely to be matched with a higher quality underwriter.[27] I find no significant differences

---

[27] The underwriter ranking is based on a scale of 0 to 9, where 9 implies highest underwriter prestige. The ranking is compiled by Carter and Manaster (1990), Carter, Dark, and Singh (1998), and Loughran and Ritter (2004). I use the rating that covers the particular time period when the firm

between the two groups using this firm quality proxy. Moreover, there is no significant difference in firm age at the time of filing.[28] However, I find that withdrawn firms are slightly more likely to be backed by VC funds (51 percent relative to 46 percent of IPO firms). This difference is significant at the 10 percent threshold. Finally, there are no significant differences in the location within the IPO wave.[29]

There are stark differences, however, in the NASDAQ fluctuations that firms experience after the IPO filing. Specifically, firms that went public experienced on average a 3 percent increase in the two-month NASDAQ returns following the IPO filing, while firms that selected to withdraw experienced, on average, a sharp drop of 6 percent over a similar period. However, the differences in NASDAQ returns in the three months prior to the IPO filing are fairly small (5 percent increase for firms that ultimately remained private versus 7 percent for those that went public). Given the importance of NASDAQ fluctuations around the time of the IPO filing in this analysis, I discuss these differences separately in the next section.

1.3.4 IPO Filings and NASDAQ Fluctuations

Issuers are highly sensitive to stock market fluctuations during the book-building phase (Busaba, Benveniste, and Guo, 2001; Benveniste et al., 2003; Dunbar, 1998; Dunbar and Foerster, 2008; Edelen and Kadlec, 2005). Stock market fluctuations

---

went public. If the rating for that period is not available, I employ the rating in the most proximate period.

[28] Firm age is calculated from founding date. The firm age of issuers that went public is kindly available at Jay Ritter's webpage. I collected firms' age of issuers that remained private from IPO prospectuses.

[29] Beneveniste et al. (2003) demonstrate that differences in the location within the IPO wave may be associated with the probability of IPO completion. I follow their methodology and define a firm as a "pioneer" if its filing is not preceded by filings in the same Fama-French industry in the previous 180 days (using all IPO filings, irrespective of patenting activity). "Early followers" are those that file within 180 days of a pioneer's filing date.

shift both investors' willingness to pay and issuers' reservation value. If a firm only partially adjusts its reservation value, stock market declines may lead to an issuer's withdrawal.[30]

If NASDAQ fluctuations change investors' willingness to pay, why wouldn't firms simply wait for more favorable market conditions rather than withdraw the filing? There are several reasons. First, a filing registration automatically expires 270 days after the last amendment of the IPO filing, which limits the time to complete the IPO filing (Lerner, 1994). Additionally, waiting is costly: as long as the application is pending, firms cannot issue private placements, and are forbidden to disclose new information to specific investors or banks. Any new information disclosed must be incorporated into the public registration statement. In fact, firms are required to update the registration statement periodically to reflect the current affairs of the company irrespective of raising alternative means of capital. These considerations lead firms to withdraw at an even earlier date prior to the automatic expiration of the IPO filing.

Figure 1 illustrates the sensitivity of issuers to market movements over the time period of the sample. I plot the fraction of monthly filings that ultimately withdrew against the two months of NASDAQ returns calculated from the middle of each month, which approximates the stock market fluctuations during the initial part of the book-building phase. The figure demonstrates a strong and negative correlation between NASDAQ movements and IPO withdrawals, even when focusing on the pre-2000 period.[31]

---

[30] The summer of 2011 illustrates the sensitivity of issuers to market fluctuations. For example, in the week of August 8th, U.S. stocks plummeted following the downgrade of U.S. treasuries and the debate over the U.S. debt ceiling. During the same week, 12 IPOs were planned but only one completed the process.

[31] The correlation of the two plots equals -0.44, or -0.34 if considering only the pre-2000 period.

Figure 1 - NASDAQ Fluctuations and IPO Withdrawals

The figure illustrates the sensitivity of IPO filings to NASDAQ fluctuations. The sample includes all IPO filings from 1985 through 2003 in the United States, after excluding unit investment trusts, Closed-end funds, REITs, Limited partnerships, and financial companies are excluded from the sample. Overall there are 8563 IPO filings, with 6958 complete registrations and 1605 withdrawn registrations. The dashed line is the fraction of monthly filings that ultimately withdrew their registration. The solid line is the two-Month NASDAQ returns calculated from the middle of each month. The correlation of the two plots is -0.44, and -0.34 before 2000. Both correlations are significantly different from zero at 0.01%.

26

In light of the costs associated with preparing for an IPO filing, this sensitivity might be surprising. However, Ritter and Welch (2002) argue that "market conditions are the most important factor in the decision to go public"; therefore it is not surprising that firms are likely to withdraw following a deterioration in market conditions. A survey by Brau and Fawcett (2006) indeed finds that CFOs that withdrew an IPO registration recognized that market conditions "played a decisive role in their decision." Additionally, Welch (1992) argues that "information cascades" can induce later investors to rely on earlier investors' choices, which may lead to rapid failure of the issue offerings in cases of market declines during the initial period of the book-building phase.

Panel F describes firm exit events in the five years following the IPO filing. These include acquisitions, bankruptcies, or IPOs of withdrawn firms. As discussed earlier, I am able to capture the patenting activity of firms in the five years following their IPO filing, even after either acquisitions or second IPOs of withdrawn firms. I find that 18 percent of the withdrawn firms ultimately go public in a second attempt in the five years following the IPO filing. Additionally, 29 percent of the withdrawn firms and 24 percent of the IPO firms are acquired over this period.

The resulting low rate of return to public equity markets was highlighted in the literature (Dunbar and Foerster, 2008; Busaba, Beneviste, and Guo, 2001). However, when incorporating alternative exit options of withdrawn firms, the fraction of firms that exit in the five years following the event rises to 50 percent. There are several explanations for the low rate of return to public equity markets. Returning to the IPO markets in a second attempt may be difficult as the window of

Both correlations are significantly different from zero at the 0.01% level.

opportunities may be closed due to the boom and bust nature of the IPO markets (Ibbotson and Ritter 1995). Brau and Fawcett (2006) interview CFOs and found that those that withdrew an IPO expressed greater concern about the uncertainty and costs associated with the IPO process. These perceptions may deter firms from a second attempt at going public. Brau and Fawcett (2006) find additionally that the most important signal when going public is a firm's past historical earnings. If going public requires several years of fast growth to attract investors' attention, such growth may be difficult to re-generate in a second attempt. Finally, Dunbar and Foerster (2008) suggest that there are reputational costs associated with the decision to withdraw which prevent firms from returning to equity markets.

1.3.5   Instrumental Variable Related Tests

Having introduced the data, this section presents the results briefly discussed in section 2.D to explore the validity of the instrumental variables approach. The first set of results is presented in Table 2. I explore whether firms experiencing NASDAQ drops are significantly different from other filing firms in the same year. A firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns after the IPO filing are within the bottom 10 percent of filers in a given year. I repeat the same exercise with the bottom 25 percent, and median as alternative cutoff thresholds. I explore whether firms that experienced NASDAQ drops are significantly different across various observables such as firm financial information at the time of filing, age, VC backing, IPO filing characteristics, and pre-filing innovation measures. I find no differences between the two groups when thresholds reflect a substantial drop in NASDAQ conditions, i.e., at the bottom 10 percent or 25 percent threshold. When using medians as a cutoff, I find a weak difference

28

in the VC backing variable.

In the second set of results, I conduct a placebo test. The exclusion restriction requires that the two-month NASDAQ returns affect innovation only through the ownership choice channel. If this is the case, we should expect that two-month NASDAQ returns *after* the IPO completion choice would have no effect on long-run innovation. In Table 1 of the Appendix, I explore whether the two-month NASDAQ returns can predict future innovation once the ownership structure is fixed, i.e., immediately after the decision to either issue equity or withdraw filing. I find that once ownership is determined, NASDAQ fluctuations do not significantly predict long-run innovative performance, in contrast to the two-month returns immediately after the IPO filing.

Finally, I investigate directly whether the instrument can explain changes in innovative trends in the core technologies of firms at the time of IPO filing.[32] I use all patents granted by the USPTO to calculate changes in innovative trends in these technologies.[33] In Table 2 of the Appendix, I find that the instrument does not predict changes in any of these innovative trends. Clearly, firms may switch to different technologies following the IPO. However, this test suggests that whether or not such a switch occurred, it is not likely to be driven by the two-month change in the NASDAQ.

---

[32] I define a technology class as a core technology if the share of patents in that class, in the three years before the IPO filing, is above the median share of patents across all the technology classes of the firm.

[33] Specifically, the change in average patent quality of each core technology is the average scaled citations of all patents in the specific technology class in the five years after the IPO filing, divided by the average scaled citations in the three years prior to the IPO filing in the corresponding technology class. Similarly, I construct the change in the total number of patents in the core technology, and also the change in the weighted number of patents, when patents are weighted by the number of citations. Since firms may have multiple core technologies, I weight the measures outlined above by the number of patents a firm produced in each core technology class.

Table 2 – NASDAQ Drops and Firm Characteristics

The table presents differences in firm characteristics between IPO filers that experienced a NASDAQ drop and other filers in the same year that did not experience a NASDAQ drop. In column (1), a firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns after the IPO filing are within the bottom 10 percent of all filers in the same year. In column (2) the threshold used is bottom 25 percent, and 50 percent threshold in column (3). Innovative measures are based on the three years up to (and through) the IPO filing. Variables are defined in Section A of the Appendix. *, **, and *** indicate that differences in means are statistically significant at the 10%, 5%, and 1% levels.

NASDAQ Drop Threshold (annual):	Bottom 10%	Bottom 25%	Bottom 50%
Financials Information at IPO filing			
Log Total Assets	0.104	0.077	0.001
R&D / Assets	0.007	−0.018	−0.004
Net Income / Assets	0.008	−0.014	−0.007
Cash / Assets	0.036	0.013	0.015
IPO Characteristics			
Lead Underwriter Ranking	0.124	0.110	0.067
Firm age	0.068	−0.946	−0.016
VC-backed	0.061	0.053	0.053*
Pre-filing Innovation measures:			
Average Citations	0.905	0.064	0.262
Average Scaled Citations	−0.071	0.072	0.101
Number of Patents	0.603	−1.354	−2.204
Scaled Number of Patents	0.330	−0.326	−0.530
Average Scaled Generality	−0.021	0.023	0.024
Average Scaled Originality	−0.039	−0.017	0.001
Scaled Best patent	−0.197	0.277	0.158

1.4 Results

1.4.1 Within-Firm relationship between IPOs and Innovation

Before turning to the instrumental variables analysis, in this section I explore the within-firm changes in innovation of firms that successfully completed the IPO filing. The specification presented in Table 3 uses the various innovation measures as dependent variables and has the following form:

$$Y_{it} = \beta_0 + \sum_{\substack{k=-3 \\ k \neq 0}}^{k=5} \gamma_k EventYear_{i,k} + \tau_i + \mu_t + \varepsilon_{i,t}$$

$EventYear_{i,k}$ is a dummy variable indicating the relative year around the IPO in which a patent was applied for approval (year zero is the year of the IPO and the omitted category). All specifications are estimated using OLS and include firm fixed effects (τ_i) and year fixed effects (μ_t). Standard errors are clustered at the firm level.[34]

The unit of observation in columns (1) to (6) of Table 3 is at the patent level. The dependent variable in column (1) is the raw count of patent citations. I find a monotonic decline in patents' novelty that starts two years before the IPO event, and continues in the five years thereafter. Since citations vary over time and between technology classes, in column (2) I use the scaled citations measure. The coefficients represent the change in relative innovation quality, and demonstrate a similar pattern to the one found in column (1). The post-IPO decline in scaled citations is displayed in Figure 2. The magnitude of the effect is substantial. For

---

[34] In an unreported analysis I verify that these results remain unchanged when the estimated model is quasi maximum likelihood Poisson, the standard model used in count data analysis. The model is discussed in the Appendix.

example, the coefficient of the year dummy three years after the IPO equals -0.597, implying a decline of 31.64 percent in innovation quality relative to the pre-IPO filing period (average scaled citations is 1.89).

In column (3) I repeat the same specification, but use patent originality as a dependent variable. Patent originality deteriorates significantly, starting two years after the IPO event. In column (4) the effect becomes even more significant when I estimate it using scaled originality. In columns (5) and (6), similar patterns arise when I estimate the effects on generality and scaled generality. Lastly, in columns (7) and (8) I consider changes in innovation measured by number of patents per year in the years around the IPO event. I find no change in the number of patents produced after the IPO, measured by either simple patent counts or scaled number of patents.

Taken together, the results indicate a change in the composition of patents around the IPO. The quality of innovation declines, as do the generality and originality measures, indicating that research becomes less fundamental. Additionally, I find no evidence for an increase in innovative scale following the IPO. However, these results could be driven by reversion to the mean and life cycle effects, irrespective of the IPO filing. To better understand whether this decline is driven by the IPO, next sections present the results using the instrumental variables approach.

1.4.2   Internal Innovation

In this section I use the instrumental variables approach, described in Section 2 to study the effects of going public on internally generated firm innovation.

Table 3 – Within-firm relationship between IPOs and Innovation

The table presents within-firm changes in innovative activity around the IPO of firms that completed the IPO filing. The dependent variables are stated at the top of each column. In columns (1) to (6), a patent is the unit of observation, while in columns (7) and (8) firm-year is the unit of observation. *Event Year* are dummy variables indicating the relative year around the IPO event (the omitted category is the year of the IPO). Variables are defined in Section A of the Appendix. The estimated model is Ordinary Least Squares (OLS), and standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Citations	Scaled Citations	Originality	Scaled Originality	Generality	Scaled Generality	Patents	Scaled Patents
Event Year -3	3.086***	0.209	0.014	0.048	0.033**	0.053	−0.330	−0.215*
	(1.035)	(0.185)	(0.021)	(0.039)	(0.014)	(0.047)	(0.438)	(0.113)
Event Year -2	3.752***	0.406***	0.022**	0.065***	0.019*	0.041	−0.192	−0.141
	(0.843)	(0.135)	(0.011)	(0.025)	(0.010)	(0.029)	(0.345)	(0.092)
Event Year -1	1.873***	0.214**	0.002	0.006	0.008	0.009	0.022	−0.039
	(0.475)	(0.089)	(0.012)	(0.027)	(0.008)	(0.026)	(0.282)	(0.065)
Event Year 1	−2.422***	−0.342***	−0.009	−0.018	−0.007	−0.001	0.069	0.060
	(0.450)	(0.077)	(0.006)	(0.016)	(0.007)	(0.023)	(0.209)	(0.062)
Event Year 2	−3.677***	−0.384***	−0.017**	−0.046***	−0.015*	−0.024	−0.265	−0.049
	(0.558)	(0.086)	(0.007)	(0.018)	(0.007)	(0.024)	(0.428)	(0.113)
Event Year 3	−4.748***	−0.597***	−0.017**	−0.054***	−0.026***	−0.063**	−0.197	−0.049
	(0.635)	(0.094)	(0.008)	(0.020)	(0.009)	(0.029)	(0.468)	(0.132)
Event Year 4	−5.739***	−0.662***	−0.022**	−0.072***	−0.032***	−0.063*	0.091	−0.002
	(0.789)	(0.110)	(0.009)	(0.022)	(0.011)	(0.036)	(0.486)	(0.150)
Event Year 5	−6.991***	−0.719***	−0.024**	−0.075***	−0.029**	−0.046	−0.216	−0.100
	(0.870)	(0.121)	(0.010)	(0.024)	(0.013)	(0.045)	(0.433)	(0.152)
Observations	39,306	39,306	38,093	38,093	35,232	35,232	13,302	13,302
R-squared	0.039	0.014	0.010	0.002	0.017	0.002	0.037	0.045
Firm FE	yes	yes	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes	yes	yes

Figure 2 - Quality of Innovation around the IPO Event

The figure presents the changes in patent quality, measured by scaled citations in the years around the IPO (year zero is the year of the IPO event). The chart estimates and confidence intervals are taken from the year dummy variables in the second column of Table 3.

The first-stage results, presented in Table 4, demonstrate the effect of the two-month NASDAQ returns on IPO completion. The dependent variable is equal to one if a firm completed the IPO filing, and zero otherwise. All specifications include filing year and industry fixed effects using OLS.[35] In column (1), I find that the coefficient of the two-month NASDAQ returns equals 0.704 and is significant at 1 percent. A change of one standard deviation in NASDAQ returns translates into a decline of 8.72 percent in the likelihood of completing the IPO. Moreover, the F-statistic equals 47.79 and exceeds the threshold of $F = 10$ which suggests that the instrument is strong and unlikely to be biased toward the OLS estimates (Bound, Jaueger, and Baker, 1995; Staiger and Stock 1997).

A concern with the post-IPO filing returns is that its variation may be either capturing the pre-IPO filing fluctuations that motivate firms to submit the initial registration statement, or reflecting the state of the IPO market. Therefore, I add additional control variables such as the three-month NASDAQ returns prior to the IPO filing and the location of the filer within the IPO wave. I also control for the number of pre-filing patents, and a dummy variable indicating whether the firm is backed by a VC fund and re-estimate the model in column (2). The coefficient of the post-IPO filing NASDAQ returns is still significant at 1 percent with a higher F-statistic of 52.03 reflecting the greater accuracy of the first stage. The sensitivity to market fluctuations slightly increases, and equals 0.763. This result suggests that the two-month NASDAQ returns play an important role in determining IPO completion, and is almost orthogonal to the added control variables, confirming the findings in Table 2.

---

[35] Probit model generates similar estimates.

Table 4 – First Stage

The table reports the first-stage estimation of the instrumental variables analysis. The dependent variable is a dummy variable which is equal to one if a firm completed the IPO filing, and zero otherwise. The *NASDAQ returns* variable is constructed differently across specifications. In the *Two Months* specification (columns (1) to (4)), NASDAQ returns are the two-month returns after the IPO filing date. In columns (5) and (6), *All* specification indicates that NASDAQ returns are calculated over the entire book-building period, i.e., from the date of the initial registration statement to the completion or withdrawal dates. Finally, *Binary* in columns (7) and (8) uses a dummy variable and is equal to one if a firm has experienced a NASDAQ drop. A firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns from the date of the IPO filing are within the bottom 25 percent of all filers in the same year. In columns (3) and (4) the sample is restricted to IPO filings before the year 2000. When control variables are included, the following variables are added to the specification: three-month NASDAQ returns prior to the IPO filing, number of patents in the three years before the IPO filing, VC-backed dummy, Pioneer and Early Follower variables. The variables are defined in Section A of the Appendix. The estimated model is Ordinary Least Squares (OLS), and robust standard errors are calculated in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample	Full	Full	Pre-2000	Pre-2000	Full	Full	Full	Full
Instrument	Two Months	Two Months	Two Months	Two Months	All	All	Binary	Binary
NADSAQ returns	0.704***	0.763***	0.690***	0.723***	0.381***	0.400***	−0.106***	−0.111***
	(0.102)	(0.106)	(0.128)	(0.132)	(0.080)	(0.081)	(0.022)	(0.022)
Observations	1,801	1,801	1,458	1,458	1,801	1,801	1,801	1,801
R-squared	0.138	0.149	0.082	0.089	0.127	0.136	0.124	0.134
Filing year FE	yes	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes	yes
Control variables	no	yes	no	yes	no	yes	no	yes
F-stat	47.79	52.03	28.9	29.9	22.63	24.13	24.16	25.99

In columns (3) and (4) I verify that the variation of the instrument is not driven only by the year 2000. I repeat the specification above, but limit the sample to pre-2000 years. The sensitivity of IPO completion to market fluctuations remains strongly significant at 1 percent, with only a slight change in magnitude (0.690 relative to 0.704 estimated in column (1)), illustrating that sensitivity of filers to market conditions is not a unique phenomena to the burst of the internet bubble, as additionally demonstrated in Figure 1.

In the remainder of the table I explore alternative specifications of the instrument. In columns (5) and (6) I calculate the NASDAQ returns over the entire book-building period, from the first day of the IPO filing until the IPO completion or withdrawal dates.[36] Although the coefficient is still significant at 1 percent, and the *F*-statistic is sufficiently high, the magnitude of the coefficient declines, and one standard deviation change reflects a 6.17 percent change in the likelihood that the firm will complete the IPO filing. The weaker effect reflects the importance of the first months in the book-building period, where most of the marketing efforts are concentrated. This is consistent with Welch's (1992) argument of "information cascades": later investors are more likely to rely on earlier investors' choices, leading to the rapid success or failure of the equity offering.

In columns (7) and (8), I construct a dummy variable that equals one if the two-month NASDAQ returns experienced by a filer are among the lowest 25 percent of all filers within the same year. The dummy variable is highly significant, reflecting a 10.6 percent decline in the likelihood that a firm will complete the IPO filing. Finally, Figure 3 illustrates the non-parametric relation between the two-month NASDAQ fluctuations and the likelihood to complete the IPO filing. The

---

[36] When the IPO withdrwal date is not available, I calculate it as the 270 days after the last IPO filing amendment (Lerner 1994)

figure shows that as long as the NASDAQ fluctuations are negative, there is a positive and monotonic association between NASDAQ returns and the likelihood to complete the IPO filing. When NASDAQ returns are positive, filers become less sensitive to market conditions and the likelihood to complete an IPO filing becomes more or less stable around 85%.

Overall, the first-stage results indicate that NASDAQ fluctuations have a strong effect on IPO completion, particularly when NASDAQ declines. Moreover, the two-month NASDAQ effect seem to be orthogonal to the added control variables.

Simple Illustration of Reduced Form Results

Before proceeding to the multivariate analysis, I illustrate the results by a simple comparison of the post-IPO innovative performance of firms that experienced a NASDAQ drop relative to other filers within the same year. This comparison is equivalent to the reduced-form estimation illustrated in the example in Section 2.B when the instrument is binary and equals one if a firm experienced a NASDAQ drop. This approach is attractive because of its simplicity and the absence of any distributional or functional form assumptions. If experiencing a NASDAQ decline affects the decision to complete the IPO but does not affect the long-run scaled measures of innovation, differences in averages illustrate the effects of going public on innovative activity.

For this analysis, a firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns after the IPO filing are within the bottom 25 percent of filers in a given year. Column (2) of Table 2 illustrates that there are no significant differences between the two groups in any of the firm characteristics and innovation measures at the time of the IPO filing. However, a comparison of post-IPO

38

Figure 3 – Two-month NASDAQ fluctuations and IPO completion likelihood

The figure presents the non-parametric association of the two-month post-IPO filing NASDAQ returns and the likelihood to complete the IPO filing of firms in the sample.

Table 5 - Reduced Form

The table reports differences in the five-year innovative performance following the IPO filing between filers that experienced a NASDAQ drop and other filers in the same year that did not experience a NASDAQ drop. A firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns after the IPO filing are within the bottom 25 percent of all filers in the same year. This comparison is equivalent to a reduced form estimation when the instrument is binary and equals one if a firm experienced a NASDAQ drop. *IPO* is a dummy variable that is equal to one if a firm completed its IPO filing, and zero otherwise. Variables are described in section A of the Appendix. *, **, and *** indicate that the difference in means is statistically signicant at the 10%, 5%, and 1% levels.

| | NASDAQ Drop | | | No NASDAQ Drop | | | |
	Mean	Median	S.D.	Mean	Median	S.D.	Difference
IPO	0.74	1.00	0.44	0.85	1.00	0.36	-0.111***
Scaled Citations	1.59	1.19	2.05	1.34	1.09	1.15	0.247***
Scaled Number of Patents	5.56	1.91	12.42	5.91	1.49	16.64	-0.351
Scaled Generality	1.10	1.10	0.67	1.10	1.09	0.67	-0.005
Scaled Originality	1.09	1.09	0.39	1.04	1.06	0.43	0.047*
Scaled Best Patent	5.36	3.14	7.92	4.14	2.69	4.99	1.215***

filing performances reveals significant differences.

Table 5 illustrates a strong correlation between two-month NASDAQ declines and subsequent five-year innovative performance. The likelihood that the IPO will be completed declines by 11.1 percent for firms experiencing low NASDAQ returns. These firms produce patents with higher average scaled citations in the subsequent five years (the difference is significant at a 1 percent level) and generate patents with higher average scaled originality. The difference in patent quality is also apparent when one considers the most-cited patent produced after the IPO filing (rather than the average citation rates). I find no differences in the number of patents produced following the IPO filing.

These reduced-form results demonstrate that going public affects firms' innovative activity, and it leads to more incremental type of innovation. The rest of the section makes use of the continuous value instrument, using the entire variation in the two-month NASDAQ returns, and studies separately each of the innovative performance measures.

Innovation Novelty

The first set of results explores the effect of IPO on innovation novelty. The dependent variable is the average scaled citations of patents in the five years following the IPO filing. I control for the equivalent measure in the three years prior to the IPO filing. All specifications follow the model described in Section 2.B, controlling for filing year and industry fixed effects. Additionally, I control for the three-month pre-IPO filing NASDAQ returns, a dummy variable indicating whether the issuer is backed by a VC, and Pioneer and Early Follower indicators that capture the location within the IPO wave. Robust standard errors are reported in parentheses.[37]

In column (1) of Table 6, I report the endogenous OLS model and find no differences between IPO firms and withdrawn firms as the IPO coefficient is insignificant and close to zero. Column (2) presents the reduced-form estimation, obtained by substituting the endogenous IPO variable with the instrument. I find a strong and negative correlation between two-month NASDAQ returns and average scaled citations in the subsequent five years. This strong correlation is plausibly generated

---

[37] It may be natural to cluster standard errors at the level of the quarter since the selection to complete the IPO filing may be correlated across issuers filing in proximity to one another. In an unreported analysis I run this specification and find that in fact clustered standard errors decline relative to the robust estimates. This may indicate that there is no need to cluster firms at that level. As illustrated by Kezdi (2004), clustering may generate a bias toward over-rejection and overestimated t-statistics when there is no need for clustering. Using a robust standard errors in my setting may be a more conservative approach with lower t-statistics.

Table 6 - Innovation Novelty

The table reports the effect of an IPO on innovation novelty. The dependent variable is the average scaled citations per patent in the five years after the IPO filing. *IPO* is a dummy variable equals to one if a firm completed the IPO filing, and zero otherwise. *NASDAQ returns* variable is the two-month NASDAQ returns calculated from the IPO filing date. Control variables included in the regressions are: pre-filing average scaled citations per patent, pre-filing average scaled number of patents per year, Pioneer, Early follower, VC-backed dummy, and the three-month NASDAQ returns before the IPO filing. Variables are described in section A of the Appendix. In columns (1) and (2) the estimated model is Ordinary Least Squares (OLS), and Two-stage Least Squares (2SLS) in column (3). Column (4) estimates the instrumental variables approach using a quasi maximum likelihood Poisson model, which is discussed in Section B of the Appendix. In all specifications, marginal effects are reported. The standard errors in column (4) are corrected using the delta method. *Magnitude* is the ratio of the *IPO* coefficient to the pre-filing average of scaled citations per patent. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable	Scaled Citations	Scaled Citations	Scaled Citations	Scaled Citations
Model	OLS	OLS	2SLS-IV	Poisson-IV
IPO	−0.019		−0.831**	−0.980**
	(0.069)		(0.409)	(0.427)
NASDAQ returns		−0.498**		
		(0.239)		
Magnitude	-1.02%	-	-43.51%	-52.41%
Observations	1,079	1,079	1,079	1,079
R-squared	0.239	0.242	0.128	0.148
Filing year FE	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes
Control variables	yes	yes	yes	yes

through the effect of the two-month NASDAQ fluctuations on the decision of firms whether to complete the IPO filing or not. This result corresponds to the findings in Table 5. In column (3), I report the estimates of the two-stage least squares. The coefficient of the IPO variable is significant and equals -0.831, implying that average scaled citations per patent of IPO firms drops after the event by 43.51 percent (=0.83/1.91, when 1.91 is the average number of scaled citations in pre-event years). In column (4) I use the quasi maximum likelihood (QML) Poisson model to estimate the IV specification. The estimates are similar to column (3): the coefficient of interest is significant, negative, and of a similar magnitude.

It is interesting to note that the OLS coefficient overestimates the effect of going public on the quality of innovation, compared to the IV estimate. As illustrated in the example in Section 2.B, this suggests that the selection bias associated with the decision to complete the IPO filing is positive, and on average, more innovative firms are more likely to complete the IPO filing.

The Fundamental Nature of Research

In this part I explore whether the decline in patent citations is associated with a change in the nature of projects. Specifically, firms that pursue less basic or fundamental research may produce less influential innovations. In Table 7, I use the originality and generality measures to capture the fundamental nature of patents. The estimation follows the same specification used in the previous section, substituting average scaled citations with average scaled originality or generality.

Columns (1)-(3) provide the results with respect to average scaled originality of patents in the five years following the IPO filing. In column (1), I estimate the endogenous variable specification. I find no significant difference between with-

Table 7 - Fundamental Nature of Research

The table reports the effect of an IPO on the fundamental nature of research. The dependent variable is the average Scaled Originality per patent in the five years after the IPO filing in columns (1) to (3) and average Scaled Generality in columns (4) to (6). *IPO* is a dummy variable equals to one if a firm completed the IPO filing, and zero otherwise. *NASDAQ returns* variable is the two-month NASDAQ returns calculated from the IPO filing date. In columns (1) to (3) I control for the pre-filing average scaled originality, and in columns (4) to (6) I control for the corresponding generality measure. Additional control variables are: pre-filing average scaled citations, pre-filing average scaled patents per year, Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ returns before the IPO filing. Variables are described in section A of the Appendix. The estimated model is OLS, and two-stage least squares in columns (3) and (6). *Magnitude* is the ratio of *IPO* coefficient to the pre-filing average of scaled originality or scaled generality per patent. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	(1) Scaled Originality	(2) Scaled Originality	(3) Scaled Originality	(4) Scaled Generality	(5) Scaled Generality	(6) Scaled Generality
Model	OLS	OLS	2SLS - IV	OLS	OLS	2SLS - IV
IPO	−0.006 (0.010)		−0.137** (0.068)	−0.001 (0.016)		−0.087 (0.092)
NASDAQ returns		−0.081** (0.036)			−0.050 (0.051)	
Magnitude	-0.10%	-	-13%	0%	-	-8%
Observations	1,079	1,079	1,079	1,079	1,079	1,079
R-squared	0.231	0.234	0.102	0.226	0.226	0.206
Filing year FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Control variables	yes	yes	yes	yes	yes	yes

44

drawn firms and IPO firms. The reduced-form estimation in column (2), which substitutes the IPO variable for the instrument, shows that the instrument is significant at -0.081. The two-stage least squares estimates in column (3) demonstrate that the post-filing average originality of firms that completed the IPO significantly declines as the IPO coefficient equals -0.137 reflecting a decline of 13 percent ($=-\frac{0.13}{1.06}$, the average scaled originality in pre-event years is 1.06). These findings suggest that issuers who remained private produce patents that rely on a broader set of technologies. In columns (5)-(8) I repeat the analysis this time with respect to average scaled generality measure, and results demonstrate no significant effects.

Scale of Innovation

The decline in innovation novelty may be driven by an increase in the scale of innovation, measured by number of patents. In that case, addition of low-quality innovative projects may generate the results rather than a repositioning of research to lower impact topics. The analysis in Table 8 addresses this conjecture by exploring changes in innovative scale. The dependent variable is the average scaled number of patents per year after the IPO filing. I control for the pre-IPO filing corresponding measure. The specification is identical to the estimation in the previous sections. One complication in this analysis is coming from the attrition problem that may arise due to patent approval lags, particularly toward the end of the sample. Patents applied toward the end of the sample period may have not yet been approved and therefore are not considered in the analysis. In that regard, scaling patent counts is important not only to account for variations in patent filings but also because it alleviates the attrition problem. The attrition problem is further mitigated by the fact that patent approval lags are likely to affect similarly both IPO

45

firms and withdrawn firms.

The endogenous model in column (1) indicates that IPO firms produce significantly more patents per year following the IPO filing with a 37.75 percent increase relative to the pre-IPO average. Column (2), however, indicates that the above effect is insignificant when the reduced form specification is estimated. The 2SLS estimate in column (3) indicates that the coefficient of the IPO variable is insignificant and the magnitude declines to 28.17 percent. In fact, when using the IV Poisson specification in column (4), the coefficient of the IPO variable is close to zero and insignificant.

Given the length of research projects, the magnitude of increase in scale may appear only several years after the IPO. In column (5), I use the innovative scale measure over years two to five after the IPO filing, and control for the scaled number of patents per year in prior years (in the three years before the IPO filing and one year thereafter). Similar to the results in column (4), I find no evidence of an increase in the number of patents produced by IPO firms. Overall, the results suggest that there is no causal evidence of an increase in the scale of innovation.

Patent Portfolio

Since the change in patent quality is not driven by changes in the number of patent filings, it is natural to further investigate the nature of the change in firms' research following the IPO. In this part, I study the structure of the patent portfolio.

In the first analysis I investigate the dispersion of patents across different technology classes, using the Herfindahl index. The lower the Herfindahl measure, the higher the concentration of patents in a specific set of technologies. To allow a meaningful calculation of the Herfindahl measure, I restrict the analysis to firms

Table 8 - Innovation Scale

The table reports the effect of an IPO on innovation scale. The dependent variable is the average scaled number of patents per year in the five years after the IPO filing. *IPO* is a dummy variable equal to one if a firm completed the IPO filing, and zero otherwise. *NASDAQ returns* variable is the two-month NASDAQ returns calculated from the IPO filing date. Control variables included in regressions are: pre-filing average scaled citations, pre-filing average scaled number of patents per year, Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ returns before the IPO filing. Variables are described in section A of the Appendix. In columns (1) to (4), the pre-filing period is within the range of [-3,0] years around the IPO filing, while the post-IPO corresponds to the years [1,5]. In column (5), the pre-filing period covers the years [-3,1] while the years [2,5] used to calculate the post-IPO filing measure. The estimated model is OLS in columns (1) and (2), and two-stage least squares in column (3). Columns (4) and (5) estimate the specification using a quasi maximum likelihood Poisson model discussed in Section B of the Appendix. In all specifications, marginal effects are reported. In columns (5)-(6) standard errors are corrected using the delta method. *Magnitude* is equal to the ratio of the *IPO* coefficient, divided by the pre-filing scaled number of patents per year. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Sample	post	post	post	post	post plus
Dependent Variable	Scaled Patents	Scaled Patents	Scaled Patents	Scaled Patents	Scaled Patents
Model	OLS	OLS	2SLS - IV	Poisson IV	Poisson IV
IPO	0.268***		0.200	0.002	−0.003
	(0.066)		(0.474)	(0.662)	(1.067)
NASDAQ returns		0.127			
		(0.305)			
Magnitude	37.75%		28.17%	0.28%	-0.12%
Observations	1,801	1,801	1,801	1,801	1,458
R-squared	0.184	0.178	0.184	0.168	0.174
Filing year FE	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes
Control Variables	yes	yes	yes	yes	yes

that have at least two patents before and two patents after the IPO filing.[38] The dependent variable is the Herfindahl measure of all patents applied in the five years subsequent the IPO filing. I control for the pre-IPO filing corresponding measure, and the other standard control variables described in previous sections. In column (1) of Table 9 I estimate the 2SLS-IV specification. The coefficient of the IPO variable is significant and equals to -0.287, which is equivalent to a 58 percent decline in the dispersion of patents across technology classes relative to the pre-IPO filing period. This finding suggests that following the IPO, firms' patent portfolio becomes more focused on a narrower set of technologies.

I obtain further insights into firm patenting activity by exploring changes in the quality of patents in core technologies and in expanded technology classes. I divide patents in two ways. First, I divide patents into those in core technologies versus those in non-core technologies. I define a technology as a *(non-) core technology* if the share of patents in a certain technology before the IPO filing is above (below) the median share of patents across classes in the firm. Second, I divide patents that belong to expanded technology classes versus non-expanded classes. I consider a technology class as an *expanded class* if the share of patents in a class increases following the IPO relative to its share before the IPO.

The results of this analysis are presented in the remaining columns of Table 9. The dependent variable in column (2) is the average scaled citations of patents within core technologies in the five years following the IPO filing. I control for the pre-IPO filing patent quality within the same technologies. Estimating the 2SLS model, I find that the IPO coefficient equals -0.910 and is significant at a 5 percent

---

[38] Similar results are obtained even when the sample is restricted to firms with at least four patents before and four after the IPO filing in order to get a more precise Herfindahl measure, although the results are noisier due to the smaller sample size.

Table 9 - Patent Portfolio

The table reports the effect of an IPO on patent portfolio composition in the five years following the IPO filing. In column (1), the dependent variable is the Herfindahl measure to capture patents dispersion across firm technology classes. In columns (2) to (5) the dependent variable is the average scaled citations of patents within the (non-) core technologies or (non) expanded classes. I define a technology class as a (non-) core technology if the share of patents in a certain technology class before the IPO filing is above (below) the median share of patents across technology classes in a firm. Additionally, a technology class is considered (non-) expanded if the share of patents in a class (did not) increase following the IPO relative to the share of patents before the IPO filing. In all specifications I control for the average scaled citations before the IPO filing in the correponding partition. In column (1) I control also for the pre-IPO filing Herfindahl of patents generated in the three years before the IPO filing. *IPO* is a dummy variable equals to one if a firm completed the IPO filing, and zero otherwise. The instrumental variable is the two-month NASDAQ returns calculated from the IPO filing date. Additional control variables included in all regressions are: pre-filing average scaled number of patents, Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ returns before the IPO filing. Variables are described in section A of the Appendix. The models are estimated using two-stage least squares. *Magnitude* is equal to the *IPO* coefficient, divided by the pre-filing average scaled citations in the respective partition. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	Herfindahl	Scaled Citations	Scaled Citations	Scaled Citations	Scaled Citations
Sample	Large Portfolio	Core Tech	Non-Core Tech	Expanded Classes	Non-Expanded Classes
Model	2SLS - IV	2SLS - IV	2SLS - IV	2SLS - IV	2SLS - IV
IPO	−0.287**	−0.910**	−0.383	−0.846*	−0.095
	(0.142)	(0.458)	(0.450)	(0.464)	(0.377)
Magnitude	-58.37%	-48.66%	-22.31%	-50.66%	-6.51%
Observations	792	1,079	898	1,079	670
R-squared	0.158	0.171	0.141	0.078	0.248
Filing year FE	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes
Control Variables	yes	yes	yes	yes	yes

level. This estimation reveals large differences in the post-IPO quality of patents within the core technologies, as the quality of patents of IPO firms is lower by 48.6 percent relative to the pre-IPO filing average quality of patents at core technologies. I re-estimate this model in column (3), but focus on innovation novelty in non-core technologies. While the IPO coefficient is negative, I don't find significant differences between IPO firms and withdrawn firms. Similarly, I repeat the analysis for expanded and non-expanded classes in columns (6) and (7). I find that the decline in the quality of the patents of IPO firms is concentrated in expanded technology classes.

The results suggest that IPO firms focus on a narrower set of technologies, while those that remained private are more likely to experiment in a broader set of technologies. Moreover, IPO firms produce lower quality patents particularly in core technology classes and technology classes that were expanded following the IPO.

Robustness Checks

In the following section I summarize the results of several unreported supplemental analyses that test the robustness of the findings and explore alternative explanations. I start by considering more carefully the hypothesis that IPO firms have a lower threshold of filing patent applications, which leads to the addition of low-quality patents and hence the decline in average quality. However, the best (most-cited) patent is unlikely to be affected by such addition of low-quality patent filings. Studying changes in the best patent, I find that the quality of the best patent declines following the IPO, with comparable magnitude to the decline in the average innovation quality reported in Table 6. This evidence, which adds up to the finding of the overall number of patents, suggests that going public affects the en-

50

tire patent distribution rather than simply driving average performance down by the addition of low-quality projects.

Second, I examine when differences between IPO firms and withdrawn firms first emerge. Since research is a long-term process, the effect should not take place immediately after the IPO. I repeat the instrumental variable estimation separately for each year in the years following the IPO filing. I find that, as expected, the differences in quality between IPO firms and withdrawn firms become significant only from the second year onward after the IPO filing.

Third, I explore whether the results are mostly driven by the year 2000. As illustrated in Table 4, the instrument strongly predicts IPO completion even when all firms that filed in 2000 onward are excluded. I re-estimate the innovation novelty regressions after excluding all firms that filed to go public during the internet bubble in the years of 1999 – 2000. Naturally, standard errors increase due to the decline in sample size, but the results remain significant and qualitatively the same.

Fourth, I verify that the results are robust to different citations horizons. As noted earlier, Akcigit and Kerr (2011) find that citations are concentrated in the first few years following a patent's approval; therefore, results should not vary substantially when using different citation horizons. I repeat the analysis, using citation horizons of two and four years after the patent's approval. I find that the results are qualitatively similar.

Finally, a common caveat in interpreting instrumental variables results is that the estimates apply only to a subset of firms who respond to variations in the instrument. Since capital-dependent firms are likely to complete the IPO irrespective of NASDAQ fluctuations, the IV estimate may underestimate the average treatment effects of IPO on innovation.

To explore this caveat in detail, it is useful first to recognize that the fraction of sensitive firms varies with NASDAQ fluctuations. The larger the NASDAQ drops are, the larger the fraction of firms that are likely to withdraw (i.e., the larger the sensitive group). In fact, at the limit, all firms are likely to be sensitive. Therefore, I repeat the IV analysis, using only firms that experienced extreme fluctuations in the NASDAQ (using both tails of the NASDAQ returns distribution). While this decreases the size of the sample, it increases the external validity of the results, since the fraction of sensitive firms is larger. As expected, I find that the fraction of firms that respond to such variations increases. Importantly, when using the extreme values of NASDAQ as an instrument, the effect of IPO on innovation novelty remain similar to previous findings. This evidence suggests that the results are not driven by a unique unrepresentative set of firms, but rather relevant to a broader set of firms in the population.

1.4.3 *Inventor Mobility and Productivity Changes*

A substantial portion of the R&D investment is in the form of wages for highly educated scientists and engineers. Their efforts generate intangible assets, which encompass the firm's knowledge. To the extent that this knowledge is "tacit," it is embedded in the firm's human capital, and departure of inventors may lead to knowledge loss. Therefore, firms tend to smooth their R&D spending over time in an effort to reduce the risk of human capital loss (Hall, Griliches, and Hausman, 1986; Lach and Schankerman, 1989). Changes associated with the transition to public equity markets may have substantial ramifications for the firm's human capital. Retaining key employees may become difficult following the IPO as options are vested, and disparities in wealth between employees may affect their incentives.

Additionally, dilution in ownership and changes in firm governance may affect employees as well. Given the decline in innovation novelty and the importance of inventors, it is natural to explore the human capital channel. In this section, I study mobility choices and productivity changes of inventors following the IPO.

Inventor Level Data

The patent database provides an interesting opportunity to track inventors' mobility across firms, as each patent application includes both the name of the inventor and its assignee (most often the inventor's employer). The analysis of inventor-level data is, however, complicated for several reasons. First, patents are associated with inventors based on their name and geographic location. Inventors' names are unreliable, as first names can be abbreviated and different inventors may have similar or even identical names. Second, attempting to detect inventor mobility using patents is necessarily inexact. While it is possible to infer that an inventor changed firms (e.g., patented for company A in 1987 and for company B in 1989), the precise date of the relocation is unavailable. Additionally, in transitions in which inventors did not produce patents in the new location are not observable. Nevertheless, this method identifies relocations of the more creative inventors who patent frequently and presumably matter the most.

To overcome the hurdle of name matching, I use the Harvard Business School patenting database, which includes unique inventor identifiers. The unique identifiers are based on refined disambiguation algorithms that separate similar inventors based on various characteristics (Lai, D'Amour, and Fleming, 2009). I attribute a patent equally to each inventor of a patent. Overall, I have information on approximately 36,000 inventors in my sample. I restrict the analysis to inventors that

produced at least a single patent before and after the IPO filing and explore the patenting behavior of inventors in the three years before and five years after the IPO filing. I identify three inventor types:

1. Stayer – an inventor with at least a single patent before and after the IPO filing at the same sample firm.

2. Leaver – an inventor with at least a single patent at a sample firm before the IPO filing, and at least a single patent in a different company after the IPO filing.[39]

3. Newcomer – an inventor that has at least a single patent after the IPO filing at a sample firm, but no patents before, and has at least a single patent at a different firm before the IPO filing.

Out of the 36,000 inventors in my sample, I classify 13,300 inventors by the above categories. These inventors account for approximately 65 percent of the patents in the sample.

In Panel A of Table 10, I compare the patenting activity of stayers and leavers in the three years before the IPO filing.[40] I first consider only IPO firms, and find that leavers produced more novel patents, measured by either raw or scaled citations. These differences are significant at a 1 percent level. Additionally, leavers generate slightly more patents, when accounting for variations in propensity to patent across technologies and over time. Interestingly, these patterns are reversed for withdrawn firms. Those who remained at the firm produced higher-quality

---

[39] I verify that all inventor relocations are not mistakenly associated with acquisitions and name changes.

[40] If an inventor's status corresponds to the definitions of both a stayer and a leaver, I classify her as a leaver. The results do not change in a meaningful way if I classify her as a stayer instead.

patents measured by scaled citations, while no significant differences arise in terms of number of patents.

Next, I compare the post-IPO filing patents generated by stayers and newcomers. Newcomers in IPO firms produce more cited patents than stayers, and differences are significant at 1 percent when I compare either raw or scaled citations. Additionally, newcomers produce fewer patents than stayers, although this may result mechanically from the shorter time period they stayed at the firm following the IPO. Again, I find opposite results when considering withdrawn firms. The quality of patents produced by newcomers is lower than those who remained at the firm, when considering either raw or scaled citations. These differences are strong and significant at a 1 percent level. Interestingly, the quality of patents generated by leavers is significantly higher than the quality of patents generated by newcomers for both IPO firms and withdrawn firms.

Inventor Level Analysis

I explore the changes in inventor level activity using the instrumental variable approach introduced in Section 2.B. I start by investigating changes in innovation quality of stayers. Then, I examine inventor mobility by studying inventors' likelihood to leave or join the firm following the IPO filing.

The results are reported Table 11, when the unit of observation is at the level of the inventor. In columns (1) and (2) I focus on the set of inventors that remained at the firm, and the dependent variable is the average scaled citations per patent produced by inventors in the five years after the IPO filing. I control for the inventor's pre-IPO filing citations per patent, as well as filing year and industry fixed effects, VC-backed dummy, pre-IPO filing NASDAQ returns, and location within

55

Table 10 - Inventor Summary Statistics

The table reports summary statistics of innovative activity of 16,108 inventors in the sample with at least a single patent application before and after the IPO filing date. Panel A compares the pre-IPO filing patents of inventors who either remained at the firm or left after the IPO filing. Panel B compares the post IPO-filing innovative activity of inventors that remained at the firm relative to newcomers,i.e., inventors that joined the firm following the IPO filing. A *stayer* is an inventor with at least a single patent before and a single patent after the IPO filing at the same sample firm. A *leaver* is an inventor with at least a single patent at a sample firm before the IPO filing, and at least a single patent in a different company after the IPO filing. Finally, a *newcomer* is an inventor who has at least a single patent after the IPO filing at a sample firm, but no patents before, and has at least a single patent at a different firm before the IPO filing. Variables are described in section A of the Appendix. * , **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

Panel A - Pre IPO Filing

| | IPO Firms | | | Withdrawn Firms | | |
| | Leavers | Stayers | | Leavers | Stayers | |
	count mean	count mean	mean difference	count mean	count mean	mean difference
Citations	3743 14.44	3806 11.71	2.731***	708 11.49	558 11.85	−0.354
Scaled Citations	3743 2.37	3806 2.12	0.253***	708 2.36	558 2.74	−0.374**
Number of patents	3743 2.96	3806 2.86	0.107	708 3.35	558 3.36	−0.009
Scaled Number of patents	3743 1.1	3806 1.01	0.088***	708 1.21	558 1.29	−0.085

Panel B - Post IPO Filing

| | IPO Firms | | | Withdrawn Firms | | |
| | Newcomers | Stayers | | Newcomers | Stayers | |
	count mean	count mean	difference	count mean	count mean	mean difference
Citations	6787 7.58	3806 5.61	1.968***	506 4.61	558 7.08	-2.466***
Scaled Citations	6787 1.62	3806 1.41	0.210***	506 1.4	558 3.11	-1.709***
Number of patents	6787 2.49	3806 3.52	−1.033***	506 2.37	558 3.17	-0.803***
Scaled Number of patents	6787 0.86	3806 1.28	−0.423**	506 0.86	558 1.14	-0.274***

Table 11 – Inventor Mobility and Changes in Innovative Productivity

The table reports the effects of an IPO on inventors' mobility and innovative activity. In columns (1) and (2) the sample is restricted to stayers and the dependent variable is the average scaled citations after the IPO filing of stayers. In columns (3) and (4), the sample includes inventors who are either stayers or leavers, and the dependent variable equals to one if inventor left the firm. In columns (5) and (6) the sample includes inventors who are either stayers or newcomers, and the dependent variable equals to one if the inventor joined the firm. *Late Leavers* includes in the sample only leavers who patented in a different firm for the first time three years after the IPO filing. *Late Newcomers* includes in the sample only newcomers that produced their first patent in a sample firm at least three years after the IPO filing. Inventor classifications of stayer, leaver and newcomer are defined in Table 10 and in the text. *IPO* is a dummy variable equals to one if a firm completed the IPO filing, and zero otherwise. The instrument is the two-month NASDAQ returns calculated from the IPO filing date. In all specifications I control for the average scaled citations and scaled number of patents before the IPO filing of the inventor. Additional control variables are: Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ return before the IPO filing. Variables are described in Section A of the Appendix. All models, except column (2), are estimated using two-stage least squares. Column (2) estimates the instrumental variable approach using a quasi maximum likelihood Poisson model which is discussed in Section B of the Appendix. *Magnitude* is equal to the *IPO* coefficient, divided by the pre-filing average scaled citations. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

Dependent variable	(1) Citations of Stayers	(2) Citations of Stayers	(3) Likelihood to leave	(4) Likelihood to leave	(5) Likelihood to hire	(6) Likelihood to hire
Description	Full Sample	Full Sample	Full Sample	Late Leavers	Full Sample	Late Newcomers
Model	2SLS - IV	Poisson-IV	2SLS - IV	2SLS - IV	2SLS - IV	2SLS - IV
IPO	−1.094**	−1.169***	0.183***	0.275***	0.388***	0.351***
	(0.457)	(0.397)	(0.062)	(0.070)	(0.078)	(0.069)
Magnitude	-47.94%	-51.23%	-	-	-	-
Observations	6,657	6,657	8,773	5,678	11,678	9,334
R-squared	0.203	0.245	0.017	0.043	0.058	0.084
Filing year FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Control Variables	yes	yes	yes	yes	yes	yes

the IPO wave. Standard errors are clustered at the level of the firm, to allow for correlations between inventors in the same firm. I estimate the 2SLS-IV in column (1), and find that the IPO coefficient equals -1.094 and is significant at a 1 percent level. The magnitude of this coefficient is large, corresponding to a 48 percent decline in inventor's innovation novelty in IPO firms relative to the pre-IPO filing period. I repeat the analysis in column (2) using the Poisson specification, and find a similar result. These findings suggest that the decline in IPO firms' innovative activity could be at least partially attributed to the change in quality of innovation produced by inventors who remained at the firm.

To estimate whether inventors are more or less likely to leave the firm after the IPO filing I focus only on the subset of stayers and leavers in column (3). The dependent variable equals one if the inventor is classified as a leaver. I control for the average quality of patents produced by an inventor in the pre-filing period, the number of patents produced, as well as the other control variables used in previous specifications. Standard errors are clustered at the level of the firm. The 2SLS-IV estimates of column (3) illustrate that inventors in IPO firms are 18 percent more likely to leave the firm after the IPO, and coefficient is significant at 1 percent.

A natural concern regarding the validity of the instrument in this setup is that NASDAQ returns may affect labor market conditions and thus correlate with the likelihood that an inventor will leave the firm. However, since the empirical exercise compares firms that filed in the same year and given the lengthy process of the job search, it may be reasonable to assume that employees of firms that filed to go public at the same year will face similar labor market conditions in the five years following the IPO filing. To verify the robustness of the results, I restrict the sample further by focusing only on late leavers, i.e., inventors who produced patents

in a different firm for the first time at least three years after the IPO filing. This lag between the IPO filing event and relocations may reduce the likelihood that the two-month NASDAQ change is correlated with future labor market conditions. I estimate this specification in column (4) and find that, in fact, the magnitude of the coefficient becomes larger, and employees at firms that went public are 27.5 percent more likely to leave the firm relative to withdrawn firms. These results demonstrate that the decline in the quality of innovation of IPO firms is potentially driven also by the departure of inventors.

Finally, I explore whether IPO firms are more likely to attract new inventors. In order to address this question, I restrict the analysis to stayers and newcomers. The dependent variable in column (5) is a dummy variable indicating that an inventor is a newcomer. Using the 2SLS-IV specification I find that IPO firms are substantially more likely to hire new inventors. The magnitude of the coefficient is large, corresponding to a 38.8 percent increase. In column (6), I repeat the same exercise as in column (4) and restrict attention to late newcomers who produce their first patent at least three years after the IPO filing. I find that the coefficient slightly decreases, but is still highly significant, corresponding to a 35 percent increase in the likelihood to hire newcomers.

The results reveal that the transition to public equity markets has important implications for the human capital accumulation process as it shapes firms' ability to retain and attract inventors. Following the IPO, there is an exodus of inventors leaving the firm, and importantly, these inventors are those who are responsible for the more novel innovations before the IPO. Additionally, going public affects the productivity of the inventors who remained at the firm. The average quality of patents produced by stayers decline substantially at IPO firms. These two effects

59

can explain the decline in the innovative quality of IPO firms. However, the effect is partially mitigated by the ability of IPO firms to attract new inventors who produce patents of higher quality than the inventors who remained at the firm.

1.4.4 Acquisition of External Technology

The transition to public equity markets allows firms to acquire companies more easily by exploiting the improved access to capital and the potentially overvalued stock (Shleifer and Vishny 2003). Acquisition of ready-made technologies is also attractive since it is easier to communicate to shareholders, quicker to implement, and less prone to failures relative to a long process of internal innovation. This section shows that following the IPO, firms are more likely to rely on external technologies.

Figure 4 illustrates the annual acquisition likelihood of at least a single target in the years around the IPO filing. IPO firms exhibit a sharp increase in likelihood following the IPO, while there is no meaningful effect for withdrawn firms. In Panel A of Table 12, I find that the acquisition likelihood of IPO firms increases from 9 percent in the three years prior to the IPO, to 66 percent in the five years following the event. The comparable change for withdrawn firms is from 10 percent to 24 percent following the IPO filing, and this change is not significant. These findings confirm the results of Celikyurt, Sevilir, and Shivdasani (2010) who find that IPO firms are more prolific acquirers even than mature public firms within their industry, and their average expenditure on acquisitions is substantially greater than either capital expenditures or R&D.

Acquisitions, however, are used for a variety of reasons. The question remains whether acquisitions are used to buy external technologies. I collect information

Figure 4 - Acquisition Likelihood

The figure presents the annual probability to acquire at least a single firm in the three years before and five years after the IPO filing. The solid line describes filers that completed the IPO filing, and the dashed line corresponds to withdrawn filers.

61

Table 12 - Acquisition of External Technologies

The table reports summary statistics of firm acquisitions in the three years before and five years after the IPO filing. Panel A compares IPO firms and withdrawn firms and their respective likelihood to engage in at least a single acquisition. Panel B details the ownership status of target firms. Panel C describes the summary statistics of acquisitions of targets with patents. Panel D is a simplified reduced form table, illustrating differences in likelihood to acquire external patents between filers that experienced a NASDAQ drop and other filers in the same year. A firm is said to have experienced a NASDAQ drop if the two-month NASDAQ returns after the IPO filing is within the bottom 25 percent of all filers in a given year. Panel E compares internal patents generated by IPO firms after they went public with the external patents they acquired through mergers and acquisitions. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

Panel A - Acquisitions before and after IPO filing

	Complete	Withdrawn	Difference
Three years pre-IPO filing			
Total number of acquisitions	178	46	-
Avg. number of acquisitions per firm	0.12	0.14	-0.022
Acquisition likelihood	0.09	0.10	-0.009
Amount spent on acquisitions	3.94	7.05	-3.113
Five years post-IPO filing			
Total number of acquisitions	4043	428	-
Avg. number of acquisitions per firm	2.27	0.59	1.688***
Acquisition likelihood	0.66	0.24	0.419***
Amount spent on acquisitions	173.47	41.64	131.8***

Panel B - Target ownership status

Ownership Status		
Public	324	7.98%
Public Sub.	604	14.88%
Private Sub.	585	14.41%
Private	2,547	62.73%
Total Public	928	22.86%
Total Private	3,132	77.14%

on patents generated by target firms in the years prior to the acquisition. A complication arises since, as demonstrated in Panel B, approximately 30 percent of the acquisition targets are firm subsidiaries. In these cases, it is difficult to identify whether assigned patents are generated by the parent firm or by the subsidiary. Therefore, I collect patent information on independent firms only (approximately 90 percent of these are privately owned). Given that almost all of the subsidiaries are acquired by IPO firms, the results underestimate the true contribution of acquisitions to the IPO firms' innovation and provide only a lower bound.

The number of external patents acquired by public firms in the five years following the IPO is substantial. As illustrated in panel C, approximately 7,500 patents were acquired through mergers and acquisitions, relative to approximately 30,000 patents produced. Before the IPO filing, both withdrawn and IPO firms rarely acquire external patents through M&A (the fraction of external patents are 3 percent and 1 percent for withdrawn and IPO firms respectively). However, in the five years following the IPO filing there is a drastic change. The fraction of external patents in IPO firms' portfolio increases to 31 percent while it remains small for withdrawn firms (8 percent). This pattern is illustrated in Figure 5, demonstrating the annual likelihood to acquire at least a single external patent per year.

The patterns described so far demonstrate a sharp increase in dependence on external technologies following the IPO. Similar patterns arise when using the instrumental variable approach. For example, panel D shows that firms that experienced two-month NASDAQ returns within the bottom 25 percent of all filers in the same year acquire significantly fewer external patents relative to the rest of filers in the same year (1.27 versus 4.70 patents in the subsequent five years). Similar results arise when using the multivariate IV analysis, even when controlling for industry

Figure 5 - Acquisition Likelihood of External Patents

The figure presents the annual probability to acquire at least a single external patent through M&A in the three years before and five years after the IPO filing. The solid line describes firms that completed the IPO filing, and the dashed line corresponds to withdrawn filers.

acquisition propensities.

Given the substantial reliance on external patents, it is interesting to compare the external and internal patents of IPO firms. Panel E details these differences. On average, external patents exhibit higher quality than patents generated internally and are more likely to be in new technologies for which the firm has no patents before the IPO filing (and less likely to be in core technology classes) relative to the patents generated within the firm.

1.5   Discussion

The empirical findings illustrate that going public has substantial effects on the manner in which firms pursue innovation. The financing view suggests that the improved access to capital may allow firms to enhance their innovative activities. While I find that the transition to public equity markets enables firms to acquire external technologies, the financing view by itself cannot explain the decline in the quality of internal innovation following the IPO, nor the departure of key inventors from the firm.

The incentives view, however, is consistent with the main empirical findings. This view suggests that in addition to the improved access to capital, the transition to public equity market affects managers' and inventors' incentives. This translates into a selection of less novel projects and departure of key inventors. In this section, I explore two incentives-related explanations that are consistent with the empirical findings.

1.5.1 *Managerial Incentives*

Going public may affect managers' incentives which may consequently lead to a change in the type of innovative projects selected and to a greater reliance on acquisitions of external technologies. Evidence shows that stock markets seem to misvalue innovation, even when outcomes are persistent and predictable (Cohen, Diether, and Malloy, 2011). As argued by Aghion, Van Reenen, and Zingales (2009), this may be driven by the weaker incentives of dispersed shareholders to fully understand complex projects pursued by the firm relative to more concentrated ownership structures. Career concerns and takeover threats may pressure managers to select more conventional projects which can be more easily communicated to the stock market (Stein, 1989; Ferreira, Manso, and Silva, 2010). Concerns regarding such adverse effects of market pressures are often raised by CEOs and entrepreneurs. For example, when explaining the delay in Facebook's IPO, Mark Zuckerberg, CEO and founder, claimed that "being private is better for us right now because of some of the big risks we want to take in developing new products. ... Managing the company through launching controversial services is tricky, but I can only imagine it would be even more difficult if we had a public stock price bouncing around."[41]

The difficulty in conveying complex projects to the stock market may lead managers to exploit the improved access to capital and potentially overvalued stock in order to acquire technologies externally, rather than developing them within the firm. The former strategy is attractive since acquisitions are easily observed, potentially less prone to failures, and quicker to implement. The shift in the focus toward more incremental projects internally and the greater reliance on external

---

[41] Facebook Blog, September 2010.

technologies may explain the departure of skilled entrepreneurial inventors.

Overall, a change in managerial incentives can explain the three main findings in the paper: decline in internal innovation novelty, departure of inventors and the increased reliance on external technologies.

1.5.2   Inventor Incentives

Going public may affect inventors' incentives as well. For example, the dilution in ownership claims of future innovations may lead inventors to pursue less ambitious projects, or alternatively may lead inventors to leave the firm to implement their ideas in a private firm setting in which they can capture a larger fraction of the returns for their innovation.

Another difficulty in retaining inventors is due to the improved ability of inventors to cash out through their stock options, which may lead to their departure. Google's prospectus provides some anecdotal evidence. As claimed in the risk factors section in its IPO filing: "the initial option grants to many of our senior management and key employees are fully vested. Therefore, these employees may not have sufficient financial incentives to stay with us."[42] This naturally raises the question why couldn't Google provide even stronger financial incentives to prevent the departure of key employees. While Google provides some additional grants, these are relatively mild and not sufficiently strong to retain employees. The reason for the relatively mild additional compensation is due to Google's attempt to avoid generating even greater gaps in pay between employees. Specifically, the filing states that "this offering may create disparities in wealth among Google employees, which may adversely impact relations among employees and our corporate

---

[42] Google's prospectus, p. 13

culture in general."[43] This anecdotal evidence is consistent with broader evidence suggesting that firms' wage setting is constrained by workers' views about what constitutes a fair wage (Blinder and Choi, 1990; Agell and Lundborg, 1995; Campbell and Kamlani, 1997).

Inventors' incentives may also be affected by the improved ability of firms to acquire external technologies following the IPO. Rotemberg and Saloner (1994) discuss the incentives benefits from having a narrow business, which increases the likelihood of implementation ideas generated by employees and therefore increases their ex-ante incentives. Acquisitions may adversely affect the likelihood of implementing inventors' innovative projects and weaken their incentives to pursue ambitious and novel projects.

This discussion suggests that following the IPO it may be more difficult to provide appropriate incentives for inventors and therefore less feasible to induce them to pursue high-quality innovation. This, in turn, may force managers, regardless of the change in their incentives, to rely more heavily on the acquisition of external technologies. Hence, changes in inventors' incentives, associated with the transition to public equity markets, can be similarly consistent with the findings of a decline in novelty of innovation, departure of skilled inventors, and the greater reliance on acquisitions.

1.5.3 Suggestive Evidence

While both theories can explain the empirical findings, they have different implications. The managerial incentives explanation suggests that firms can pursue high-quality innovation, but corporate governance considerations translate into

---

[43] Google's prospectus, p. 9

managerial career concerns and prevent managers from doing so. The inventor's incentives theory suggests that providing appropriate incentives to inventors is difficult in a public firm setup and therefore, irrespective of managerial preferences, this setting is less productive for innovation. In this section I provide some suggestive evidence supporting the managerial incentives theory. However, this evidence does not rule out the inventors incentives theory.

To explore whether managerial incentives affect innovation, I start by considering the case of managerial entrenchment. A more entrenched CEO may be harder to replace, and thus less likely to be sensitive to market pressures. I capture managerial entrenchment by investigating whether the CEO is also the chairman of the board (Shleifer, and Vishny, 1989). The CEO's dual role as chief executive and chairman of the board implies that the CEO can direct board initiatives affecting the CEO's job security and compensation, as well as responding to takeover threats. Inventors' incentives, however, are plausibly not affected directly by whether the CEO is also the chairman of the board. Thus, if CEO entrenchment is correlated with a higher quality of innovation, this may provide evidence for the importance of managerial incentives and stock market pressures.

I collect information on board characteristics from S-1 filings, to determine whether the CEO is also the chairman at the time of the IPO.[44] Since S-1 filings are available through the SEC Edgar system from 1996, the number of observations in this analysis is smaller. In Table 13, I repeat the IV analysis to explore the effect of going public on innovation novelty separately for IPO firms with and without an entrenched CEO. In column (1), I find that when the CEO is the chairman of the board, the decline in innovation novelty following the IPO is not significant with

---

[44] Execucomp database collects information about executives from S&P 1000 firms only.

69

a magnitude of a 20.1 percent decline relative to the pre-IPO period. In column (2) I contrast this result with the case where the CEO is not the chairman of the board: here, going public is associated with a decline of 64 percent in the novelty of patents produced in the five years following the IPO, significant at 5 percent.[45] In columns (3) and (4) I repeat the analysis with respect to the likelihood of inventors to leave the firm. In column (3), I find that when the CEO is the chairman, the likelihood of inventors to leave the firm is in fact negative, yet insignificant, relative to firms that remained private. When the CEO is not the chairman, however, column (4) demonstrates that inventors are 10.8 percent more likely to leave, consistent with the decline in innovation quality. These results provide some evidence of the importance of managerial incentives in generating innovation, and its subsequent effect on inventors turnover.

In order to test whether dilution of ownership claims on innovation and cashing out affect inventors incentives and departure choices, it is necessary to have information on their compensation within the firm. In absence of this type of data, I consider whether acquisitions adversely affect inventors. If acquisitions reduce the likelihood of implementation of internal projects, and adversely affect inventors, I expect to find a substitution effect between acquisitions and internal innovation where greater reliance on acquisitions may be correlated with lower quality of internal innovation. Also, we may expect to find a greater departure of inventors when firms rely more heavily on acquisitions.

I distinguish between firms that acquired external technologies in the five years

---

[45] I estimate columns (1) and (2) separately, instead of using an interaction term of IPO variable and CEO entrenchment dummy. An interaction term will require using an additional instrumental variable. While it is possible to use the interaction of NASDAQ returns and entrenched CEO dummy as an instrument, this has limited power, and is particularly problematic given the small number of observations.

Table 13 - Empirical Evidence of Alternative Theories

The dependent variables are listed separately in each column. In columns (1)-(2) and (5)-(6), the unit of observation is at the firm level and the dependent variable is the average scaled citations per patent in the five years after the IPO filing. In columns (3)-(4) and (7)-(8), the unit of observation is at the individual level, inventors are included in the sample only if they are either a *stayer* or *leaver*, and the dependent variable is a dummy indicating whether an individual is a leaver. Stayer and leaver classifications are defined in Table 10 or in the text. In sub-sample *Chair*, the sample includes all withdrawn firms and only IPO firms that at the time of the IPO filing the CEO acts as the chairman of the board. The *No Chair* sub-sample includes the all withdrawn firms and only IPO firms that at the time of the IPO filing the CEO is not the chairman of the board. Information about CEO position is collected from initial registration statements which are available from 1996. The sub-sample *M&A* includes all withdrawn firms and only IPO firms that acquired at least one firm in the five years following the IPO filing. *No M&A* is constructed similarly, but includes only IPO firms that did not acquire target firms. *IPO* is a dummy variable equals to one if a firm completed the IPO filing, and zero otherwise. The instrument is the two-month NASDAQ returns calculated from the IPO filing date. All specifications add the following control variables: average scaled citations before the IPO filing, pre-filing average scaled number of patents, Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ return before the IPO filing. Variables are described in section A of the Appendix. All models are estimated using two-stage least squares. *Magnitude* equals to the *IPO* coefficient divided by the pre-filing average scaled citations of the firms in the respective sample. Robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	Scaled Citations	Scaled Citations	Likelihood to Leave	Likelihood to Leave	Scaled Citations	Scaled Citations	Likelihood to Leave	Likelihood to Leave
Sub-sample	Chair	Not Chair	Chair	Not Chair	M&A	No M&A	M&A	No M&A
IPO	-0.359	-1.193**	-0.140	0.108*	-0.555	-0.898**	0.031	0.164**
	(0.529)	(0.558)	(0.086)	(0.065)	(0.405)	(0.416)	(0.059)	(0.070)
Magnitude	-20.17%	-64.14%	-	-	-28.46%	-50.17%	-	-
Observations	325	428	2,626	4,292	759	490	6,232	3,803
R-squared	0.207	0.247	0.049	0.032	0.145	0.135	0.039	0.029
Filing year FE	yes	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes	yes
Control variables	yes	yes	yes	yes	yes	yes	yes	yes

71

following the IPO and firms that did not engage in such acquisitions, and run the IV estimation separately for each group, using innovation novelty as a dependent variable. In column (5) I find no significant decline in innovation quality for firms that acquired external technologies (although the coefficient is negative), while in column (6) I find a significant decline in the innovation quality of firms that did not engage in such acquisitions. Additionally, in columns (7) and (8) I study the likelihood of inventors leaving the firm, and find that inventors are more likely to leave firms that did not acquire external technologies. These findings are consistent with Sevilir and Tian (2011) who find that acquisitions complement innovation and improve the acquiring company's innovation. Hence, I find no direct evidence of an adverse effect of acquisitions on internal innovation and inventors' departure. However, these findings should be interpreted lightly as they merely reflect correlations and do not test for alternative channels through which inventors' incentives might be affected following the IPO.

Overall, this section provides suggestive evidence regarding the underlying mechanisms that generate the decline in firm-level innovation. More precisely, I find that managerial incentives play an important role in leading to a decline in quality of innovation and departure of skilled inventors.

1.6 Conclusion

In this paper, I investigate an important but understudied aspect of initial public offerings, namely, the effect on firm innovation. I find that the transition to public equity markets has a substantial effect on firms' innovative activities along three dimensions. First, the projects selected within the firm are less novel, and rely on a narrower set of technologies. Second, key inventors are more likely to leave

the firm, and the productivity of remaining inventors declines, while at the same time firms attract new human capital to the firm. Finally, firms rely more heavily on acquisition of external technologies.

I consider two views in which going public may matter for innovation. On the one hand, the financing view suggests that improved access to capital may enhance innovation. On the other hand, the incentives view suggests that, in addition to access to capital, going public affects managers' and inventors' incentives. This may lead to a selection of more conventional projects. I find that although the financing view is consistent with some aspects of the empirical findings (the increased reliance on external technologies), it cannot explain the decline in the novelty of internal innovation and the departure of key inventors following the IPO. In contrast, the incentives view explains the effects of going public along all three dimensions of the empirical findings.

Estimating the effects of going public on innovation is challenging due to its inherent selection bias. My empirical strategy compares firms that went public with firms that intended to go public, but ultimately withdrew their IPO filing and remained private. I use NASDAQ fluctuations during the book-building phase as an instrument for the decision to complete the IPO filing.

The findings in this paper reveal a complex trade-off between public and private ownership forms. While private firms are able to generate more novel innovation and retain skilled inventors, public firms can rely on acquisitions of external technologies and attract human capital. These results have implications for determining the optimal point at which a firm should go public in its life cycle.

The results draw attention to the effects of IPO on both the ability of firms to retain and attract human capital and on the productivity of the remaining inven-

tors. Seru's (2010) study of the impact of mergers on innovation has found that mergers affect mostly the productivity of inventors remaining at the firm, rather than affecting their likelihood to leave. The difference in results suggests that productivity changes that coincide with various corporate events such as mergers and IPOs are nuanced, heterogeneous, and require better understanding.

This paper does not address the general equilibrium effects of the IPO market on innovation and its corresponding welfare consequences. Yet, the results suggest that there may be important complementarities between public and private ownership structures. While private ownership may allow firms to pursue more ambitious innovations, improved access to capital may allow public firms to acquire technologies, mostly from private firms. This suggests that ownership structure plays an important role in shaping the market for technologies.

Finally, corporate managers, bankers, and policy makers alike have expressed concerns that the recent dearth of IPOs marks a breakdown in the engine of innovation and growth (Weild and Kim, 2009). Some blame the Sarbanes-Oxley Act (SOX) for raising the costs of compliance for publicly traded firms.[46] Regardless of the role of SOX in explaining the recent IPO cycle, policy prescriptions of this sort raise the question of whether the transition to public equity markets affects innovation and if so how. This paper contributes to the debate by demonstrating that IPOs affect innovation, but that their effects may be indirect. While innovation novelty declines following the IPO, it allows public firms to acquire entrepreneurial firms, and thus, potentially facilitates innovation through increased demand for new technologies.

---

[46] In the hope that IPO market stimulation will "jumpstart innovation and job creation," President Obama's Council on Jobs and Competitiveness has urged Congress to amend the Sarbanes-Oxley Act to allow small companies to tap public equity markets.

2. PRIVATE EQUITY AND INDUSTRY PERFORMANCE

This chapter is coauthored with Josh Lerner, Morten Sorensen, and Per Stromberg

2.1 Introduction

In response to the global financial crisis that began in 2007, governments are re-thinking their approach to regulating financial institutions, with private equity (PE) funds in particular being targeted by regulators. Most dramatically, in 2010 the European Commission adopted the Alternative Investment Fund Managers directive (European Commission 2010), which contains a sweeping set of rules regulating the PE industry.

Regulators, politicians, and labor organizers have long expressed concern about the impact of PE funds, pointing to their need to rapidly return capital to investors and the potentially deleterious effects of such practices as the extensive leveraging of firms. Critics have pointed to case studies that illustrate the negative conse-quences of the transactions. For instance, Rasmussen (2008) points to the buyout of Britain's Automobile Association, which led to large-scale layoffs and service dis-ruptions while generating substantial profits for the transaction's sponsor, Permira. The Service Employees International Union (2007, 2008) presents studies that show the deleterious effect that excessive leverage, cost-cutting, and poor managerial de-cisions by PE groups can have on firms and industries in cases such as Hawaiian

75

Telecom, Intelsat, KB Toys, and TDC.

A central hypothesis in the finance literature since Jensen (1989), however, has been that PE has the ability to improve the operations of firms. By closely monitoring managers, restricting free cash flow through the use of leverage, and incentivizing managers with equity, it is argued, PE-backed firms are able to improve operations in the firms they finance.

Several case and clinical studies illustrate Jensen's (1989) hypothesis. For instance, in the Hertz buyout, the PE investor Clayton, Dubilier & Rice (CD&R) addressed inefficiencies in pre-existing operations procedures to help increase the profitability of Hertz. Specifically, CD&R created value by lowering overhead costs, reducing inefficient labor expenses, cutting non-capital investments down to industry standard levels, and aligning managerial incentives with return on capital (Luehrman 2007). Similarly, the buyout of O.M. Scott & Sons led to substantial operating improvements in the firm's existing operations, in part due to powerful management incentives, as well as the active involvement by the PE investors (Baker and Wruck 1989).

This paper investigates the impact of PE investments on aggregate growth and cyclicality. Specifically, we examine the relationship between the presence of PE investments and the growth rates of total production, employment, and capital formation across 20 industries in 26 major nations between 1991 and 2007. The magnitude of PE investments is substantial: in a given year and country, we estimate that approximately 4% of the average industry is acquired by PE investors, measured in terms of sales, which is significant given the extended holding periods of these investments.

For our production and employment measures, we find that PE investments are

associated with faster growth. Industries where PE funds have been active in the past five years grow more rapidly than other industries, whether measured using total production, value added, total wages, or employment. One concern is that this growth may come at the expense of greater cyclicality, which could translate into greater risks for investors and stakeholders. Thus, we also examine whether economic fluctuations are exacerbated by the presence of PE investments, but we find little evidence that this is the case. Activity in industries with PE backing appears to be no more volatile in the face of industry cycles than in other industries, and sometimes less so. The reduced volatility is particularly apparent in total wages and employment. These patterns continue to hold when we focus on the impact of PE in Continental Europe, where concerns about these investments have been most often expressed.

In our baseline empirical specification, we include country-industry, industry-year, and country-year fixed effects (FEs), so the impact of PE activity is measured relative to the average performance in a given country, industry, and year. For instance, if the Swedish steel industry has more PE investment than the Finnish one, we examine whether the steel industry in these two countries performs better or worse over time relative to the average performance of the steel industry across all the countries in our sample, and whether the variations in performance over the industry cycles are more or less dramatic.

We believe it is unlikely that these results are driven by reverse causality, i.e., PE funds selecting to invest in industries that are growing faster and/or are less volatile. The results are essentially unchanged if we only consider the impact of PE investments made two to five years earlier on industry performance. Granger causality tests suggest that past PE investment precedes subsequent improvements

in industry performance, while past industry performance has no impact on future PE investment. The results continue to hold when we use an instrumental variables technique employing the size of the private pension and insurance company asset pool in the nation and year as a percentage of GDP.

This paper is related to the modest and mixed literature on the competitive effects of PE. Chevalier (1995a, 1995b) shows that buyouts of supermarket chains lead to positive outcomes for local rivals. These rivals are more likely to enter or expand in an urban region, if there are a number of firms that have undergone buyouts and charge higher prices in these markets. She suggests that these results are consistent with "softer" product market competition. Similarly, Oxman and Yildrim (2008) suggest that PE corporate governance practices spill over on competitors after a buyout. In contrast, Hsu, Reed, and Rocholl (2010) find that rivals experience a decrease in both their stock prices and their operating performance around the time of PE investments in their industry. This different result may be due to the authors focusing on a subset of isolated transactions and including so-called private investments in public equity (PIPEs) in their analysis.

Our analysis has some limitations. First, the question of economic growth and volatility is only one of many questions that regulators encounter when assessing the impact of PE investment. Among the unaddressed topics are the impact on productivity, the distribution of wealth across society, and the competitive dynamics across industries. Second, it is too early to assess the consequences of the economic downturn in 2008 and 2009, a period where the decrease of investment and absolute volume of distressed PE-backed assets was greater than in earlier cycles. Third, our results suggest that spillovers from PE-backed companies are important, but data limitations prevent us from exploring them in more detail here.

This paper proceeds as follows: In Section 2.2, we develop the hypotheses. Section 2.3 describes the construction of the dataset, and the results are presented in the Section 2.4. Section 2.5 presents concluding remarks.

2.2 *Industry Performance and Private Equity*

Several alternative perspectives have been offered as to how PE investments affect the prospects of an industry. In this section, we begin by reviewing the suggestions about changes regarding overall performance; we then discuss hypotheses regarding the interaction between economic cycles and PE investments.

2.2.1 *The impact of PE investments on industry performance*

Our initial examination compares the performance of industries where PE funds have been more or less active.

The Jensen (1989) hypothesis that PE-backed firms have improved operations has been supported by a number of empirical studies that focus on the effects on the individual PE-backed companies. Kaplan (1989) examines changes in accounting performance for 76 large management buyouts of public companies between 1980 and 1986. He shows that in the three years after the transaction, operating income, cash flow, and market value all increase. He argues that these increases reflect the impact of improved incentives rather than layoffs. Looking at more recent public-to-private transactions in the United States, however, Guo, Hotchkiss, and Song (2009) find only weak evidence that gains in operating performance of bought-out firms exceed those of their peers. Muscarella and Vetsuypens (1990) examine 72 "reverse LBOs" (RLBOs), that is, companies taken private that went public once again. These firms experienced a dramatic increase in profitability,

which they argue is a reflection of cost reductions. John, Lang, and Netter (1992) present supporting empirical evidence that the threat of takeover spurs firms to voluntarily undertake restructurings.

More recent studies have used large samples and a variety of performance measures to more directly assess whether PE makes a difference in the management of the firms in which they invest. Bloom, Sadun, and Van Reenen (2009) survey over 4,000 firms in Asia, Europe, and the U.S. to assess their management practices. They show that PE-backed firms are on average the best-managed ownership group in the sample, though they cannot rule out the possibility that these firms were better managed before the PE transaction. Davis et al. (2008) compare all U.S.-based manufacturing establishments that received PE investments between 1980 and 2005 with similar establishments that did not receive PE investments. They show that PE-backed firms experience a substantial productivity growth advantage (about two percentage points) in the two years following the transaction; about two-thirds of this differential is due to improved productivity. Cao and Lerner (2009) examine the three- and five-year stock performance of 496 RLBOs between 1980 and 2002. RLBOs appear to consistently outperform other IPOs and the stock market as a whole. Large RLBOs that are backed by PE firms with more capital under management perform better, while quick flips — when PE firms sell off an investment soon after acquisition — underperform.

These findings might suggest that we would see superior performance for PE firms, regardless of the economic conditions. Moreover, if PE firms represent a significant fraction of the activity in certain industries (as shown below), there may also be a positive effect at the industry level. Potentially, 'contagion' effects might arise if improvements in bought-out firms spur their competitors to improve. This

effect is difficult to document empirically, and our analysis provides no direct evidence on the channels through which PE transactions affect the industries.

2.2.2 *The impact of economic cycles*

Numerous practitioner accounts have suggested that the PE industry is highly cyclical, with periods of easy financing (often in response to the successes of earlier transactions) leading to an acceleration of deal volume, greater use of leverage, higher valuations, and ultimately more troubled investments (akin to the well-known 'corn-hog cycle' in agricultural economics). This pattern is corroborated in several academic studies. Axelson et al. (2010) document the cyclical use of leverage in buyouts. Using a sample of 1,157 transactions completed by major groups worldwide between 1985 and 2008, they show that the level of leverage is driven by the cost of debt, rather than the industry- and firm-specific factors that affect leverage in publicly traded firms. The use of leverage is also strongly associated with higher valuation levels and lower PE fund returns. Kaplan and Stein (1993) document that the 1980s buyout boom saw an increase in valuations, reliance on public debt, and incentive problems (e.g., parties cashing out at the time of transaction). Moreover, in the transactions done at the market peak, the outcomes were disappointing: of the 66 largest buyouts completed between 1986 and 1988, 38% experienced financial distress, which they define as default or an actual or attempted restructuring of debt obligations due to difficulties in making payments; 27% did default on debt repayments, often in conjunction with a Chapter 11 filing. Kaplan and Schoar (2005) show that fund performance is negatively correlated with inflows into these funds. Private equity funds raised during periods of high capital inflows, which are typically associated with market peaks, perform far worse than

their peers.

These findings corroborate the suggestions that the availability of financing impacts booms and busts in the PE market. If firms completing buyouts at market peaks employ excessive leverage, we may expect industries where a significant fraction of firms have undergone buyouts to experience more intense subsequent downturns. Moreover, the effects of this overinvestment would be exacerbated if PE investments drive rivals not backed by PE to aggressively invest and leverage themselves. Chevalier (1995b) shows that in regions with supermarkets receiving PE investments, the rivals responded by adding and expanding stores.

An alternative perspective, suggested by some recent events in the PE industry, is that PE-backed firms may do better during economic downturns because their investors constitute a concentrated shareholder base, which can continue to provide equity financing in a way that might be difficult to arrange for other companies during downturns, as frequently happened during the recent recession. This perspective implies that PE-backed companies may outperform their peers during downturns, as they have access to equity financing that other firms do not have. The presence of PE investors as shareholders may lead to fewer failures in difficult economic conditions.

A related argument, originally proposed by Jensen (1989), is that the high levels of debt in PE transactions force firms to respond earlier and more forcefully to negative shocks to their business. As a result, PE-backed firms may be forced to adjust their operations more rapidly at the beginning of an industry downturn, enabling them to better weather a recession. Even if some PE-backed firms eventually end up in financial distress, their underlying operations may thus be in better shape than their peers. This facilitates an efficient restructuring of their capital structure

and lowers the deadweight costs on the economy. Consistent with this argument, Andrade and Kaplan (1998) study 31 distressed leveraged buyouts from the 1980s that subsequently became financially distressed, and found that the value of the firms post-distress was slightly higher than the value before the buyout, suggesting that even the leveraged buyouts that were hit most severely by adverse shocks added some economic value.

Finally, institutional differences between PE funds and other financial institutions may make PE funds less susceptible to industry shocks. A major source of concern for financial institutions is the so-called 'run on the bank' phenomenon. Runs occur when holders of short-term liabilities, such as depositors or repo counterparties, simultaneously refuse to provide additional financing and demand their money back. Other versions of this phenomenon arise when companies simultaneously draw down lines of credit, hedge fund investors simultaneously ask for redemptions, or a freeze in the market for commercial paper prevents structured investment vehicles (SIVs) from rolling over short-term commercial paper. However, it is unlikely that PE investments create dangers through this mechanism. Private equity funds are typically prevented from borrowing themselves, and the funds' only claimants are their limited partners (LPs), which are typically bound by ten-year lock-up agreements. Hence, the funds have no short-term creditors that can run. Still, extensive loans may be provided to the individual portfolio companies. However, these loans are typically made by a concentrated set of lenders and are without recourse to other portfolio companies or the fund generally. Hence, an individual creditor's ability to be repaid is largely unaffected by the actions of other creditors, mitigating the incentive to run.

2.3 Data Sources and Sample Construction

We combine two datasets in order to analyze how PE investments affect industries. One dataset contains information about PE investments compiled by Capital IQ, and another contains industry activity and performance across the Organisation for Economic Cooperation and Development (OECD) member countries that are included in the OECD's Structural Analysis Database (STAN).

2.3.1 PE investment sample

We use the Capital IQ (CIQ) database to construct a base sample of PE transactions. This database is recognized as the most comprehensive database of worldwide PE transactions. Strömberg (2008) compares CIQ LBO data during the 1980s with older LBO studies using 1980s data and estimates that during this early period, well before Capital IQ's formation, the database's coverage was somewhere between 70% and 85%. The base sample contains all private placements and M&A transactions in CIQ where (a) the list of acquirers includes (at least) one investment firm, (b) where the transaction is classified as 'leveraged buyout,' 'management buyout,' or 'going private,' (c) that were announced between January 1986 and December 2007, and (d) where the target company is located in an OECD country included in the STAN database. Thus, we only look at later-stage buyout transactions, and do not include venture capital investments. We exclude transactions that were announced but not completed as of December 2007, as well as transactions that did not involve a financial investor (e.g., a buyout executed by the management team itself was excluded). This results in a sample with about 14,300 transactions, involving 13,100 distinct firms.

We use various measures of PE activity relative to the size of the industry. For

most of our analyses we use an indicator variable that equals one if there are any PE investments during the previous five years. This has the advantage of being well-defined even absent information about deal sizes and the total size of the industry. For some analyses we use more refined measures of PE activity. We only have information about the deal size for 50% of our transactions (though more of the larger transactions), so for those analyses we impute missing deal sizes by constructing fitted values from a regression of deal size on fixed effects for country, investment year, and target industry. Using the imputed transaction sizes, we generate aggregate country-year-industry measures of PE volume in the form of summed deal sizes. We then scale the total deal size calculated in this way by the total industry production as reported by STAN (see below) to construct a relative measure of PE investments in the industry. Since the imputations are noisy, we do not use this measure directly. Rather, to reduce the noise, we construct indicators for whether PE activity is above or below the median amount, or in the different quartiles, based on this measure.

2.3.2   Industry Data

The STAN database provides industry data across OECD countries compiled from national statistics offices. It contains economic information at the country, year, and industry level. Thus, a typical observation would be the German transport equipment industry in 1999. STAN includes measures of total production, employment, and capital formation, as described in Table 1.

Throughout this paper, we focus on the following measures of industry activity:

- Production (gross output), the value of goods and/or services produced in a year, whether sold or stocked, in current prices.

Table 14: Descriptions of OECD STAN industry variables

Industry Variable	Description
Production (gross output)	Value of goods and/or services produced in a year, whether sold or stocked, measured at current prices.
Value added	Industry contribution to national GDP. Value added comprises labor costs, consumption of fixed capital, taxes less subsidies, measured at current prices.
Labor costs (compensation of employees)	Wages and salaries of employees paid by producers as well as supplements such as contributions to social security, private pensions, health insurance, life insurance and similar schemes.
Number of employees	Persons engaged in domestic production excluding self-employed and unpaid family workers.
Gross fixed capital formation	Acquisitions, less disposals, of new tangible assets (such as machinery and equipment, transport equipment, livestock, constructions) and new intangible assets (such as mineral exploration and computer software) to be used for more than one year, measured at current prices.
Consumption of fixed capital	Reduction in the value of fixed assets used in production resulting from physical deterioration, normal obsolescence or normal accidental damage.

Source: OECD, STAN database, 2003.

- Value added represents the industry's contribution to national GDP, i.e., output net of materials purchased. While the methodology for constructing this measure differs across nations, our focus here is on differences across time, which should reduce the effect of national differences in the measure.

- Labor costs, which comprise wages and salaries of employees paid by producers, as well as supplements such as contributions to social security, private pensions, health insurance, life insurance, and similar schemes.

- Number of employees, which is the traditional measure of employment, excluding self-employed and unpaid family members working in the business.

- Gross capital formation is acquisitions, less disposals, of new tangible assets, as well as such intangible assets as mineral exploration and computer software. This variable is the closest aggregate to capital expenditures.

- Consumption of fixed capital measures: the reduction in the value of fixed assets used in production resulting from physical deterioration or normal obsolescence.

2.3.3   Mapping Capital IQ to STAN industries

We have to rely on the OECD/STAN industry classification, since the endogenous variables are only defined at this level. The STAN database and Capital IQ, however, rely on different industry classifications. Industries in the STAN database are classified by the International Standard Industrial Classification (ISIC) Code, which does not map directly to the Capital IQ classification. To overcome this issue, we first use the mapping from the CIQ industry classification into SIC Codes, and then use another existing mapping from SIC to ISIC industries.

The mapping of CIQ industry classifications to SIC Codes includes only matches for the most detailed levels of the CIQ classifications to four-digit SIC Codes. Whenever possible we use this matching to get equivalent SIC Codes. However, PE transactions are often defined by CIQ at a more aggregated industry level classification (hence includes multiple refined categories), for which no direct mapping to SIC, and ultimately to ISIC, exists. In these cases, we used all SIC Codes that belong to the sub-categories of the industry classification of CIQ and therefore had multiple four-digit SIC Codes for a single CIQ (upper level) industry classification.

In some cases, the mapping of a single aggregated level CIQ industry to multiple four-digit SIC Codes generated no conflict as all of the four-digit SICs corresponded to the same ISIC industry classification, creating a one-to-one mapping. In cases where the four-digit SIC Codes corresponded to different industries in the ISIC scheme, we considered the particular deals and selected the most suitable industry. In 390 transactions, we were not able to determine with certainty the appropriate match in ISIC, and those transactions were dropped, leaving us with 13,910 PE transactions with ISIC classifications.

Finally, we group ISIC sub-industries to balance PE activity across industries. For example, there are 520 PE transactions within the "Food products and beverages" sub-industry classification, and only two transactions in the "Tobacco" industry. The ISIC parent category of these two classifications is "Food products, beverages, and Tobacco." Therefore, we use this aggregate category rather than the two more refined ones. As a result, the industry classification we use is a refined ISIC classification, but in cases of small PE activity we are using a more aggregate industry level. In unreported analyses we verify that the results hold using the refined (non-grouped) industry classifications.

This results in a sample of 11,135 country-industry-year observations during the years 1986 to 2007. For each country-industry-year, we measure PE activity as the volume of PE deals occurring during the previous five years in this country and industry. In particular, an observation is a PE industry if it had at least one PE investment during those five years. This definition was motivated by the holding periods reported by Strömberg (2008). With this definition, we can only compare activity from 1991 to 2007, leaving us with 8,596 country-industry-year observations.

Tables 15, 16, and 17 present the distribution of deals across industries, years, and countries. Several patterns are visible: (1) the heavy representation of buyouts as a share of economic activity in traditional industries, such as 'textiles, textile products, leather,' 'machinery and equipment,' 'pulp, paper, paper products, printing,' 'electrical and optical equipment,' and 'chemical, rubber, plastics and fuel products'; (2) the acceleration in buyout activity, first modestly during the late 1980s and then especially in the mid-2000s; and (3) the greater level of activity in a handful of traditional hubs for PE funds, including the U,S., the Netherlands, Sweden, and the U.K.

Although the concentration of PE activity across certain industries, years, and countries may have been a potential concern, our analysis includes year, industry, and country fixed effects. This, together with the fact that country-industry-year is the unit of observation, ensures that our results are not driven by a few industry, year, or country outliers.

Table 18 is a comparison of the growth the industry measures for PE and non-PE industries. PE industries grow more quickly in terms of output, value added, and employment; however, the PE industries have a slower growth rate for gross

89

Table 15: Distribution of deals by industry

The sample consists of 8,596 country-industry-year observations of OECD countries between 1991 and 2007. Observations is the number of observations in the industry. *PE* Industries contains the number of observations classified as *PE industries*. An industry is a *PE industry* if it had at least one PE investment during the previous five years. Deals is the number of deals, and Deal Volume is the combined size of the deals (normalized to 2008 US$ billions). Imputed Deal Volume imputes the size for deals with missing size information.

Industry	Observations	PE Industries	Deals	Deal Volume	Imputed Deal Volume
Agriculture, hunting, forestry and fishing	432	84	52	6.18	9.88
Basic metals and fabricated metal products	431	234	740	67.06	116.43
Chemical, rubber, plastics and fuel products	431	223	724	112.28	161.18
Community, social and personal services	430	216	1,124	311.70	378.35
Construction	430	173	318	28.20	47.07
Electrical and optical equipment	431	229	837	138.14	180.80
Electricity, gas and water supply	431	84	109	100.90	123.29
Financial intermediation	426	232	551	150.24	201.25
Food products, beverages and tobacco	431	221	548	66.33	106.60
Hotels and restaurants	426	171	436	132.45	155.52
Machinery and equipment	431	255	1,256	124.90	204.62
Manufacturing and recycling	431	166	367	29.74	54.62
Mining and quarrying	429	98	148	28.12	39.85
Other non-metallic mineral products	431	131	160	18.78	29.39
Pulp, paper, paper products, printing, publishing	431	216	520	111.22	141.51
Real estate, renting and business activities	426	284	2,662	369.24	512.65
Textiles, textile products, leather	431	213	425	30.42	63.45
Transport equipment	431	113	106	14.74	21.94
Transport, storage and communications	430	231	578	247.70	286.87
Wholesale and retail trade – repairs	426	279	1,605	328.07	441.77
Total	8,596	3,853	13,266	2,416.42	3,277.05

Table 16: Distribution of deals by year

Observations is the number of country-industry-year observations per year. PE Industries contains the number of observations classified as *PE industries*. An industry is a *PE industry* if it had at least one PE investment during the previous five years. Deals is the number of deals, and Deal Volume is the combined size of the deals (normalized to 2008 US$ billions). Imputed Deal Volume imputes the deal size for deals with missing size information.

Year	Observations	PE Industries	Deals	Deal Volume	Imputed Deal Volume
1986	n/a	n/a	95	19.56	27.15
1987	n/a	n/a	109	18.51	27.43
1988	n/a	n/a	157	42.83	60.77
1989	n/a	n/a	137	59.75	68.07
1990	n/a	n/a	120	21.41	32.47
1991	456	116	158	13.29	21.88
1992	469	139	178	15.73	26.80
1993	509	177	197	16.44	29.61
1994	516	191	262	15.57	25.68
1995	520	202	347	35.05	49.86
1996	520	204	431	43.53	57.30
1997	520	206	655	55.41	86.12
1998	520	202	871	94.46	144.40
1999	520	217	824	86.41	131.17
2000	520	228	780	105.44	138.76
2001	520	251	687	80.83	102.62
2002	520	269	722	93.28	122.11
2003	520	276	945	145.73	178.78
2004	520	293	1,217	203.73	278.14
2005	520	293	1,428	258.58	368.21
2006	520	316	1,788	404.54	552.20
2007	406	273	1,776	748.42	963.42
Total	8,596	3,853	13,884	2,578.48	3,492.93

Table 17: Distribution of deals by country

The sample consists of 8,596 country-industry-year observations of OECD countries between 1991 and 2007. Observations is the number of observations in each country. PE Industries contains the number of observations classified as *PE industries*. An industry is a *PE industry* if it had at least one PE investment during the previous five years. Deals is the number of deals, and Deal Volume is the combined size of the deals (normalized to 2008 US$ billions). Imputed Deal Volume imputes the size for deals with missing size information.

Country	Observations	PE Industries	Deals	Deal Volume	Imputed Deal Volume
Australia	320	125	122	14.64	18.55
Austria	340	77	53	1.78	3.93
Belgium	340	129	117	13.00	22.56
Canada	340	218	292	98.98	117.12
Czech Republic	300	158	37	5.06	5.89
Denmark	340	94	142	9.79	17.30
Finland	340	161	192	7.66	16.06
France	339	274	1,273	121.04	176.37
Germany	340	220	598	109.79	187.06
Greece	324	30	7	4.45	6.14
Hungary	320	142	18	1.15	3.39
Ireland	340	104	4	0.00	0.01
Iceland	339	6	46	19.09	20.69
Italy	340	210	335	42.21	57.69
Japan	328	70	73	20.79	26.71
Netherlands	340	204	320	84.87	125.66
Norway	340	73	71	5.00	9.53
Poland	286	171	41	2.33	2.61
Portugal	320	63	27	0.25	0.33
Slovakia	300	111	13	0.18	0.93
South Korea	340	47	20	4.81	4.81
Spain	320	171	217	38.93	42.58
Sweden	340	186	267	43.07	57.60
Switzerland	340	158	111	17.66	31.46
United Kingdom	340	318	2,194	377.13	423.60
United States	340	333	6,676	1,372.78	1,898.46
Total	8,596	3,853	13,266	2,416.42	3,277.05

fixed capital formation.

One natural question is whether the volume of buyouts during our sample period is sufficiently large to have a material impact on the industries in which the funds invest. The most direct approach is to look at the implied share of PE investments in the industries in our sample. We wish to compute the mean share of total industry value represented by PE transactions annually.

Because enterprise value is not available for privately-held firms, we must approximate this measure. In particular, we compute a "revenue multiple" from the publicly traded firms in Global Compustat for each industry and year as the ratio between the aggregate enterprise value (the sum of the market value of equity, plus the book value of debt and preferred stock) of all publicly traded firms across all sample nations and the revenues for the same set of firms. We then assume that this ratio also characterizes the privately-held firms in each industry in the same year. Thus, we estimate the ratio of the aggregate volume of PE investments in each industry and year (not using imputed deals, in order to be conservative), as well as the product of the estimated revenue multiple and the aggregate production by public and private firms, as estimated by the OECD.

These ratios vary by year, reflecting the ebb and flow of PE activity. If we examine the average annual share of PE activity across the entire sample period by industry, it varies from 0.9% (for transport equipment) to 13.5% (for machinery and equipment). The weighted average across all industries is 4.35%, with an interquartile range from 2.5% to 7.1%. This suggests that for the typical industry, the impact of PE over this period is quite substantial, especially in light of the five-to-seven year holding period, which characterizes the typical PE investment (Strömberg 2008). This measure may understate the volume of PE activity. Not only are

Table 18: Industry growth variables

The sample consists of 8,596 country-industry-year observations of OECD countries between 1991 and 2007. An industry is considered as a *PE industry* if it had at least a single PE deal in the previous five years. *P*-value provides the p-value of a test of equality of the means of PE and non-PE industries. See Table 1 for variable definitions.

	All Industries			PE Industries			Non-PE Industries			
	Observations	Average Growth	Std. Dev.	Observations	Average Growth	Std. Dev.	Observations	Average Growth	Std. Dev.	P-value
Production (gross output)	7,351	5.9	8.8	3,318	6.2	8.5	4,033	5.7	9.1	0.03
Value added	8,238	5.6	10.2	3,635	5.8	9.8	4,603	5.5	10.5	0.17
Labor costs (compensation of employees)	7,831	5.1	7.5	3,398	5.3	7.4	4,433	5.0	7.6	0.18
Number of employees	6,269	0.0	5.0	2,862	0.3	4.1	3,407	-0.3	5.6	0.00
Gross fixed capital formation	7,004	7.1	76.6	3,223	6.8	27.6	3,781	7.5	101.1	0.67
Consumption of fixed capital	7,351	5.9	8.8	3,318	6.2	8.5	4,033	5.7	9.1	0.03

transactions with missing data excluded, but as discussed above, CIQ's coverage is incomplete.

Moreover, it is likely that having a significant fraction of firms in an industry under buyout ownership has a substantial effect on competitors as well. As discussed in the introduction, earlier work suggests that the impact of PE extends beyond the bought-out firms.

2.4 Analysis

2.4.1 Industry performance

We begin by examining the relationship between various industry characteristics and the role of PE in the industry. In each case, an observation is an industry-country-year triple, and the dependent variable is the growth rate of a given economic variable (e.g., employment).

We employ several specifications. First, we include an indicator that denotes whether the industry is a PE industry or not (defined, as noted above, as an industry with at least one PE investment during the previous five years). Note that this definition does not use the imputed deal values, since it only depends on the presence of PE deals. Second, we use two indicators to capture whether an industry is a low or high PE industry. A low PE industry (PE5 Low) is a PE industry where the fraction of total imputed PE investments divided by total production (both normalized to 2008 U.S. dollars) is smaller than the median (conditional on having a non-zero level of PE investment), while a high PE industry has PE investments to production ratio above the median. We also perform the analysis dividing PE activity into quartiles to better measure the differential effects of different activity levels. For both the median and the quartile dummy specifications, industries with

no PE activity is the omitted group, i.e., coefficients should be interpreted as relative to observations with zero PE investment in a given country-industry-year. All specifications include country-industry fixed effects.

To control for common shocks across industries and countries, we include industry-year, country-industry, and country-year fixed effects in our specifications. Hence, we estimate the fixed-effect panel regression:

$$y_{ciy} = PE_{ciy}\beta + \eta_{ci} + \xi_{iy} + \mu_{cy} + \varepsilon_{ciy}$$

where y_{ciy} is the endogenous variable of interest, e.g., the growth rate of employment; PE_{ciy} is an indicator for whether the industry is a PE industry; η_{ci} is a country-industry fixed effect; ξ_{iy} is an industry-year fixed effect; μ_{cy} is a country-year fixed effect; and ε_{ciy} is the residual error term.

The results in Table 19 indicate that industries with PE deals have significantly higher growth rates of production and value added. For instance, in the first regression, the coefficient of 1.368 implies that the total production of an average PE industry grows at an annual rate that is 1.368% higher than a non-PE industry. The average growth rate is 5.9%. When we include country-year fixed effects, the coefficient is still statistically significant but declines to 0.541; i.e., the excess growth in PE industries is 0.541% per year. This drop in the magnitude of the effect may indicate that PE investors invest in countries during periods of above-average growth.

In Table 9, we report the significance of a statistical test for differences between high- and low-PE industries, as well as differences between the four quartiles of PE activity (reported as PEL = PEH). Without country-year FEs, we find some evidence that the effect is stronger for industries with more PE activity. With country-year FEs, the effect, although not statistically significant, appears slightly stronger for

96

Table 19: PE activity and growth rate of productivity

The table contains OLS panel regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the growth rate of production or value added (as defined by OECD). The exogenous variables are an indicator for positive PE activity over the previous five years at the country-industry level (PE_5), indicators for whether the measured PE activity is below or above the median activity level (PE_5 Low and PE_5 High), and indicators for quartiles. The omitted base category is no PE activity over the previous five years. Country-industry (C-I FE), industry-year (I-Y FE), and country-year (C-Y FE) fixed effects are included as indicated. Robust standard errors are in parentheses. $PE_L = PE_H$ contains the significance level of a Wald test of equality of the *PE Low* and *PE High* coefficients or all the quartile coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Production (gross output)	Production (gross output)	Production (gross output)	Production (gross output)	Production (gross output)	Value Added	Value Added	Value Added	Value Added	Value Added
PE_5	1.368***	0.541**				1.259***	0.448			
	(0.292)	(0.251)				(0.324)	(0.286)			
PE_5 Low			1.197***	0.615**				1.008***	0.450	
			(0.298)	(0.258)				(0.332)	(0.300)	
PE_5 High			1.712***	0.381				1.764***	0.445	
			(0.371)	(0.328)				(0.407)	(0.377)	
PE_5 Q1					0.522*					0.436
					(0.301)					(0.332)
PE_5 Q2					0.723**					0.463
					(0.297)					(0.370)
PE_5 Q3					0.447					0.430
					(0.335)					(0.391)
PE_5 Q4					0.351					0.492
					(0.404)					(0.472)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
I-Y FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
C-Y FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
$PE_L = PE_H$			0.091*	0.405	0.734			0.025**	0.988	0.998
Observations	6,976	6,976	6,976	6,976	6,976	7,013	7,013	7,013	7,013	7,013

industries with less PE activity. The large number of country-year FEs reduces the statistical power and statistical significance. Similarly, the data do not appear to contain sufficient information to separate the effects when the level of PE activity is broken down by quartile. All coefficients are positive, but not statistically significantly different.

In Table 19, for value added, we also find that the PE investments are associated with faster growth. Without country-year FEs, the relation is particularly striking, with industries with lower levels of PE activity growing 1.008% faster per year than industries without PE activity, and industries with more PE activity growing 1.764% faster on average. These coefficients remain positive, but muted, when including country-year FEs. Statistical significance also declines, although the loss in statistical power is a potential reason for this decline, as mentioned above.

A natural concern is the direction of causality. It is possible that PE investors pick industries that have the potential to grow, and our results may reflect this industry choice rather than the causal effect of the investments on the industry. To mitigate this concern, we change our definition of the PE industry measure to only include investments during the period from two-to-five years prior to the observation, called the twice-lagged measure (the original PE measure included all five years prior to the observation). The results are reported in Table 20. We find that the results are very similar, indicating that the effect that we find is unlikely to be driven by PE investors entering countries and industries where they expect stronger immediate growth.

Table 21 considers measures of employment. PE industries appear to grow significantly faster in terms of labor costs and the number of employees. In the specifications without country-year fixed effects, the annual growth rate of total labor cost

Table 20: Twice-lagged PE activity and growth rate of productivity

The table contains OLS panel regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the growth rate of production or value added (as defined by OECD). The exogenous variables are an indicator for positive PE activity over the previous four years -2 to -5, i.e., *not* including the year prior to the year where the growth in the endogenous variable is measured (PE_{2-5}), indicators for whether the measured PE activity is below or above the median activity level (PE_5 *Low* and PE_5 *High*), and indicators for quartiles. The omitted base category is no PE activity over the previous five years. Country-industry (*C-I FE*), industry-year (*I-Y FE*), and country-year (*C-Y FE*) fixed effects are included as indicated. Robust standard errors are in parentheses. $PE_L = PE_H$ contains the significance level of a Wald test of equality of the *PE Low* and *PE High* coefficients or all the quartile coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1) Production (gross output)	(2) Production (gross output)	(3) Production (gross output)	(4) Production (gross output)	(5) Production (gross output)	(6) Value Added	(7) Value Added	(8) Value Added	(9) Value Added	(10) Value Added
PE_5	1.208***	0.449*				1.247***	0.506*			
	(0.269)	(0.238)				(0.299)	(0.278)			
PE_5 Low			1.076***	0.454*				0.974***	0.416	
			(0.273)	(0.244)				(0.313)	(0.295)	
PE_5 High			1.434***	0.441				1.715***	0.670*	
			(0.354)	(0.318)				(0.398)	(0.388)	
PE_5 Q1					0.334					0.309
					(0.281)					(0.336)
PE_5 Q2					0.587*					0.527
					(0.306)					(0.377)
PE_5 Q3					0.514					0.686*
					(0.339)					(0.404)
PE_5 Q4					0.413					0.732
					(0.370)					(0.466)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
I-Y FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
C-Y FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
$PE_L = PE_H$			0.233	0.964	0.854			0.046*	0.502	0.819
Observations	6,976	6,976	6,976	6,976	6,976	6,976	7,013	7,013	7,013	7,013

is 0.779 percentage points greater for PE industries, while the number of employees grows at an annual rate that is 0.845 percentage points greater. With country-year fixed effects, these estimates decline to 0.16 percentage points for total labor cost and to 0.4 percentage points for the number of employees, neither of which is statistically significant.

These findings may be surprising, since a common concern is that PE investors act aggressively to reduce costs with little concern for employees. This concern is not necessarily inconsistent with our results, since we are looking at the industry rather than the firm level. Even if buyouts may lead to initial employment reductions at PE-backed firms (as found in Davis et al. (2009) for the U.S.), the greater subsequent growth in total production, observed in Table 6, may lead to subsequent employment growth in the industry overall. Considering the specifications with PE activity quartiles, the growth rate of labor costs and number of employees is fastest in industries with moderate levels of PE activity. This suggests that the increase in employment is not primarily driven by increases at the PE-backed firms themselves but driven by the spillover effects at other firms.

As above, we are concerned about the direction of causality. Table 22 repeats the analysis using the twice-lagged PE measure. The magnitudes in Tables 21 and 22 are largely similar, suggesting that the effect we identify is not mainly driven by PE investors picking industries with expectations of immediate employment growth. If anything, the results of the twice-lagged measure suggest that the growth in the number of employees is more robust than the growth in labor costs.

Finally, in Table 23 we examine measures of fixed capital formation and consumption of fixed capital. These measures appear much more volatile than the production and employment measures, with substantially larger standard errors,

Table 21: PE activity and growth rate of employment

The table contains OLS panel regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the annual growth rate of labor costs or total employment (as defined by OECD). The exogenous variables are an indicator for positive PE activity over the previous five years at the country-industry level (PE_5), indicators for whether the measured PE activity is below or above the median activity level (PE_5 Low and PE_5 $High$), and indicators for quartiles. The omitted base category is no PE activity over the previous five years. Country-industry (C-I FE), industry-year (I-Y FE), and country-year (C-Y FE) fixed effects are included as indicated. Robust standard errors are in parentheses. PE_L = PE_H contains the significance level of a Wald test of equality of the PE Low and PE $High$ coefficients or all the quartile coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Number of Employees	Number of Employees	Number of Employees	Number of Employees	Number of Employees
PE_5	0.779***	0.160				0.845***	0.400**			
	(0.269)	(0.197)				(0.210)	(0.181)			
PE_5 Low			0.575**	0.241				0.954***	0.469**	
			(0.270)	(0.211)				(0.214)	(0.193)	
PE_5 High			1.192***	-0.016				0.605**	0.241	
			(0.337)	(0.244)				(0.272)	(0.232)	
PE_5 Q1					0.080					0.415*
					(0.249)					(0.217)
PE_5 Q2					0.424*					0.545**
					(0.230)					(0.230)
PE_5 Q3					0.047					0.347
					(0.266)					(0.246)
PE_5 Q4					0.045					0.069
					(0.298)					(0.289)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
I-Y FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
C-Y FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
$PE_L = PE_H$			0.016**	0.247	0.231			0.123	0.285	0.393
Observations	6,743	6,743	6,743	6,743	6,743	5,771	5,771	5,771	5,771	5,771

Table 22: Twice-lagged PE activity and growth rate of employment

The table contains OLS panel regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the annual growth rate of labor costs or total employment (as defined by OECD). The exogenous variables are an indicator for positive PE activity over the previous four years –2 to –5, i.e., *not* including the year previous to the year where the growth in the endogenous variable is measured (PE_{2-5}), indicators for whether the measured PE activity is below or above the median activity level (PE_5 *Low* and PE_5 *High*), and indicators for quartiles. The omitted base category is no PE activity over the previous five years. Country-industry (*C-I FE*), industry-year (*I-Y FE*), and country-year (*C-Y FE*) fixed effects are included as indicated. Robust standard errors are in parentheses. $PE_L = PE_H$ contains the significance level of a Wald test of equality of the *PE Low* and *PE High* coefficients or all the quartile coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Number of Employees	Number of Employees	Number of Employees	Number of Employees	Number of Employees
PE_5	0.594**	-0.038				0.781***	0.341**			
	(0.251)	(0.193)				(0.190)	(0.170)			
PE_5 Low			0.359	-0.038				0.860***	0.344*	
			(0.258)	(0.204)				(0.201)	(0.181)	
PE_5 High			1.007***	-0.038				0.630***	0.336	
			(0.302)	(0.243)				(0.243)	(0.221)	
PE_5 Q1					-0.179					0.261
					(0.233)					(0.212)
PE_5 Q2					0.134					0.464**
					(0.226)					(0.206)
PE_5 Q3					0.128					0.525**
					(0.259)					(0.231)
PE_5 Q4					-0.217					0.028
					(0.309)					(0.275)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
I-Y FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
C-Y FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
$PE_L = PE_H$			0.006***	0.999	0.263			0.285	0.969	0.078*
Observations	6743	6743	6743	6743	6743	5771	5771	5771	5771	5771

102

making it difficult to discern any relationship between PE investments and capital formation. If anything, the results suggest that PE investments reduce the gross fixed capital formation and consumption of fixed capital, but these results are more tentative.

2.4.2   Cyclical Patterns

We next analyze how PE relates to industry cycles. For each industry and year, we average the growth rate of the production and employment measures across countries to attain the average growth rate. This rate measures the annual aggregate shock in these variables (e.g., production output in the steel industry fell by 2% on average in 2002 across the nations in our sample). We then investigate whether PE industries are more or less exposed to this shock by including the PE measure interacted with this average growth measure in the regressions. In particular, we estimate the specification:

$$y_{ciy} - \overline{y}_{iy} = PE_{ciy}\beta + (PE_{ciy} \times y_{iy})\gamma + \eta_{ci} + \varepsilon_{ciy}$$

where y_{ciy} is the endogenous variable of interest (e.g., the growth rate of employment); y_{iy} is the mean of the endogenous variable across countries (e.g., the average growth rate of employment in industry i during year y); PE_{ciy} is an indicator for whether the industry is a PE industry; η_{ci} is a country-industry fixed effect; and ε_{ciy} is the residual error term. Note that this specification does not permit us to include individual year controls, since demeaning by subtracting industry-year averages also removes any aggregate year variation. To capture any remaining serial correlation and cyclicality, we also estimate specifications that allow the error terms to follow AR(1) processes, as indicated.

Table 23: PE activity and growth rate of capital formation

The table contains OLS panel regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the annual growth rate of gross fixed capital formation or consumption of fixed capital (as defined by OECD). The exogenous variables are an indicator for positive PE activity over the previous five years at the country-industry level (PE_5), indicators for whether the measured PE activity is below or above the median activity level (PE_5 Low and PE_5 $High$), and indicators for quartiles. The omitted base category is no PE activity over the previous five years. Country-industry (C-I FE), industry-year (I-Y FE), and country-year (C-Y FE) fixed effects are included as indicated. Robust standard errors are in parentheses. $PE_L = PE_H$ contains the significance level of a Wald test of equality of the PE Low and PE $High$ coefficients or all the quartile coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Gross Fixed Capital Formation	Gross Fixed Capital Formation	Gross Fixed Capital Formation	Gross Fixed Capital Formation	Gross Fixed Capital Formation	Consumption of Fixed Capital	Consumption of Fixed Capital	Consumption of Fixed Capital	Consumption of Fixed Capital	Consumption of Fixed Capital
PE_5	-0.103	-3.499				0.165	-0.494			
	(2.463)	(3.346)				(0.415)	(0.323)			
PE_5 Low			-0.092	-3.565				-0.353	-0.435	
			(2.550)	(3.653)				(0.469)	(0.367)	
PE_5 High			-0.124	-3.362				1.096**	-0.611*	
			(2.665)	(3.044)				(0.467)	(0.357)	
PE_5 Q1					-1.190					-0.657**
					(3.256)					(0.328)
PE_5 Q2					-6.172					-0.181
					(4.718)					(0.520)
PE_5 Q3					-4.792					-0.568
					(3.258)					(0.382)
PE_5 Q4					-2.965					-0.468
					(4.070)					(0.490)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
I-Y FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
C-Y FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
$PE_L = PE_H$			0.985	0.918	0.301			0.001***	0.619	0.739
Observations	6,074	6,074	6,074	6,074	6,074	4,712	4,712	4,712	4,712	4,712

If PE and non-PE industries were equally sensitive to economic conditions, we would expect the coefficient on the interaction term, γ, to be zero. For example, if the average growth rate of employment first increases by 2% and this increase is equally large for PE and non-PE industries and then decreases by 2% and this decrease is also equally large, then γ is zero. In contrast, imagine that the growth rate of employment increases by 2% on average, but this increase is distributed such that PE industries grow by 3% and non-PE industries grow by only 1%, and this is followed by a 2% decline in growth rate, but this decline is distributed such that PE industries decline by 3% and non-PE industries decline by 1%, then the coefficient γ is positive and we interpret this as PE investments amplifying the exposure to the aggregate shocks.

In Tables 24 and 25, we examine the impact on production and employment. Across all the regressions, the interaction terms are negative, which suggests that PE industries are less exposed to industry shocks than non-PE industries.

To interpret the coefficients, using the estimates in the first regression in Table 25, if an industry on average experiences a 5% increase in total labor costs in a given year (the aggregate shock), a PE industry will experience, on average, a 5.576% increase (5% + 1.992% + 5% x -0.214 = 5.922%). Conversely, following a 5% decrease in labor costs, a PE industry will only experience, on average, a 2.394% decline (5% + 1.591% + (–5%) x –0.203 = –1.938%). Hence, an aggregate swing from +5% to –5% (a 10% difference) in aggregate growth rates translates into a swing from 5.9% to –1.9% (a 7.8% difference) in the growth rates for PE industries. For the production and employment analyses (not value added), the coefficients are significantly negative in the simple specification and many of the coefficients in the employment analysis remain statistically significant when high and low PE

Table 24: PE activity and productivity cycles

The table contains OLS panel regression coefficients. An observation is the annual growth rate of the indicated productivity measure (subtracting its average growth rate across countries) at the country-industry-year level. The exogenous variable $PE_5 \, x \, Avg \, growth$ contains the interaction between PE_5 and the average growth rate of the endogenous variable, averaged over countries. PE_5 is an indicator for positive PE activity in the country-industry during the previous five years. The variables $PE_5 \, Low \, x \, Avg \, growth$ and $PE_5 \, High \, x \, Avg \, growth$ are constructed similarly, where $PE_5 \, Low$ and $PE_5 \, High$ are indicators for below or above median PE activity. The regressions contain country-industry (C-I FE) fixed effects. Standard errors in parentheses are robust and calculated allowing for AR(1) serial correlation as indicated along with the autocorrelation coefficient. $PA_L = PA_H$ contains the significance level of a Wald test of equality of the $PE_5 \, Low \, x \, Avg \, growth$ and $PE_5 \, High \, x \, Avg \, growth$ coefficients. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Production (gross output)	Production (gross output)	Production (gross output)	Value Added	Value Added	Value Added
PE$_5$ x Avg growth	-0.059** (0.029)			-0.066 (0.043)		
PE$_5$ Low x Avg growth		-0.088 (0.057)	-0.080* (0.042)		-0.132 (0.103)	-0.133** (0.063)
PE$_5$ High x Avg growth		-0.025 (0.052)	-0.040 (0.039)		0.008 (0.075)	0.008 (0.059)
PE$_5$	1.417*** (0.315)			1.567*** (0.389)		
PE$_5$ Low		1.572*** (0.407)	1.522*** (0.397)		1.673*** (0.620)	1.761*** (0.514)
PE$_5$ High		1.562*** (0.378)	1.349*** (0.394)		1.662*** (0.425)	1.474*** (0.485)
C-I FE	Yes	Yes	Yes	Yes	Yes	Yes
AR(1)	Yes 0.206	No 0.344	Yes 0.206	Yes 0.099	No 0.234	Yes 0.097
PA$_L$=PA$_H$			0.476			0.094*
Observations	6,499	6,976	6,999	6,536	7,013	6,536

Table 25: PE activity and employment cycles. The table contains OLS regression coefficients. An observation is the annual growth rate of the indicated employment measure (subtracting its average growth rate across countries) at the country-industry-year level. The exogenous variable *PE x Avg growth* contains the interaction between *PE* and the average growth rate of the endogenous variable, averaged over countries. *PE* is an indicator for positive PE activity in the country-industry during the previous five years. The variables *PE Low x Avg growth* and *PE High x Avg growth* are constructed similarly, where *PE Low* and *PE High* are indicators for below or above median PE activity. The regressions contain industry, country and country-industry (*Co-Ind FE*) fixed effects as indicated. Standard errors are calculated with clustering at the country-year level and presented in parenthesis. $PA_L = PA_H$ contains the significance level of a Wald test of equality of the *PE Low x Avg growth* and *PE High x Avg growth* coefficients. Statistical significance at the 1%, 5% and 10% levels are indicated by ***, ** and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Labor costs (compensation of employees)	Labor costs (compensation of employees)	Labor costs (compensation of employees)	Number of employees	Number of employees	Number of employees
PE$_5$ x Avg growth	-0.203***			-0.098**		
	(0.041)			(0.045)		
PE$_5$ Low x Avg growth		-0.277***	-0.229***		-0.172***	-0.114**
		(0.049)	(0.055)		(0.050)	(0.054)
PE$_5$ High x Avg growth		-0.112**	-0.111*		-0.039	-0.036
		(0.050)	(0.059)		(0.055)	(0.063)
PE$_5$	1.591***			0.538***		
	(0.306)			(0.171)		
PE$_5$ Low		1.910***	1.657***		0.750***	0.792***
		(0.361)	(0.415)		(0.173)	(0.206)
PE$_5$ High		1.295***	1.517***		0.324	0.493*
		(0.345)	(0.431)		(0.215)	(0.282)
Industry FE	Yes	Yes	No	Yes	Yes	No
Country FE	Yes	Yes	No	Yes	Yes	No
Co-Ind FE	No	No	Yes	No	No	Yes
P$_L$xA=P$_H$xA	0.004***	0.080*		0.016**	0.213	
Observations	6,743	6,743	6,743	5,771	5,771	5,771
R-squared	0.228	0.228	0.306	0.068	0.069	0.196

107

industries are included separately. Overall, it appears that PE activity translates into smaller employment fluctuations than average, but industries with a higher amount of PE activity may follow a growth pattern that is closer to that of the industry as a whole.

2.4.3 Geographic Patterns

It is interesting to explore whether the impact of PE is different in Continental Europe than in the U.S. and U.K. Not only is the level of PE activity higher in the U.S. and U.K. than in most other nations, but the industry is more established. Thus, we repeat the analysis, separating Continental Europe from the U.S. and U.K.

In unreported results, we repeat the base specifications reported in Tables 19 and 21 with the sample restricted to Continental European countries. All the main effects remain largely unchanged for the Continental Europe sample, suggesting that the effects are not primarily driven by the U.S. and U.K. Moreover, we find that that the effects are not statistically different for Continental Europe and the U.S./U.K., although the U.S./U.K. subsample is naturally a smaller sample, with reduced statistical power to distinguish the effect of PE investments.

2.4.4 Addressing causality concerns

One natural concern relates to the interpretation of these results. While it appears that PE is associated with more rapid growth at an industry level in our sample, it is natural to wonder which way the causation runs. Does the presence of PE lead to higher production, or do PE investors invest where they anticipate industries will grow? We respond to this question in several ways.

First, our baseline analysis considers PE investments during the five years prior

to the observed growth in total production and employment. As discussed above, we also narrowed our measure to only include investments in the years two through five prior to the investment. If our effects are due to PE investors anticipating growth in particular sectors, they would have to be quite prescient to anticipate growth two years in advance.

Second, we address this concern using an instrumental variables technique. To identify exogenous variation, we use the size of the private pension and insurance company asset pool in the nation and year, expressed as a percentage of GDP. This kind of identification strategy has been employed in other papers in the venture capital literature, such as Kortum and Lerner (2000) and Mollica and Zingales (2007). The basic idea is that in nations with larger pension and insurance pools (institutional assets), domestic PE funds are more likely to raise capital and invest it locally. Moreover, pension policy and insurance regulation are driven by broader socio-economic considerations, rather than a desire to impact the local PE industry or current investment opportunities in this industry. The intuition is that when a country's institutional assets increase, this increase leads to an increase in PE activity across all industries in this country, and the IV estimates the marginal change in industry growth resulting from this increase.

Although the instrument only varies at the country-year level and PE investors invest at the country-year-industry level, the identification follows from standard arguments for identification using instrument variables. Specifically, identification of the local average treatment effect (LATE) follows from an exclusion restriction and a monotonicity condition (conditions 1 and 2 in Imbens and Angrist (1994)).

The exclusion restriction requires that changes in pension assets are independent of the error term in the regression. While this is difficult to establish em-

pirically, pension funds primarily change as a result of pension reforms, and we have reviewed changes in pension policies in Germany and the U.K. These reviews suggest that a wide array of considerations drive reforms in the rules governing long-term savings, including demographic pressures and the consequent dangers of running out of funding, the presence of perceived disparities (e.g., between white- and blue-collar workers, in the treatment of stay-at-home mothers), and the desire to increase the labor supply. We found no evidence that these changes are motivated by a perception that PE investments offered particularly attractive investment opportunities. One concern with respect to the exclusion restriction, however, is that the motivation to increase the labor supply could potentially generate some of the results that we see. It is unlikely, however, that such reforms should be concentrated in industries where PE firms are active.

The monotonicity condition states that an increase in institutional assets in a given country must be associated with a weakly increasing amount of PE activity in each of the industries in this country. In other words, an increase in institutional assets cannot lead to a decline in PE activity for any industry, which seems reasonable.

To estimate this model, we supplement the dataset with data on financial assets held by domestic pension funds and insurance corporations from the OECD. We only include funded pension obligations, excluding for instance, public pension plans that hold very few investable assets but are funded on a "pay as you go" basis. Table 13 presents the distribution of financial assets across countries.

The instruments for the PE variable we employ are financial assets normalized with the country's GDP, along with country and industry fixed effects. The results of this analysis are shown in Table 27, which also includes regular OLS estimates

110

Table 26: Distribution of financial assets by country

Observations is the number of country-year pairs for which financial assets data is available (since 1990). Financial Assets is the value of assets held by domestic autonomous pension funds and insurance corporations (in 2008 US$ billions). Financial Assets to GDP Ratio is the fraction of financial assets normalized by country's GDP.

Country	Observations	Financial Assets (2008 US$ billions)		Financial Assets to GDP Ratio	
		Average	Std. Dev.	Average	Std. Dev.
Australia	18	480.14	269.90	0.72	0.18
Austria	18	86.32	37.07	0.28	0.09
Belgium	18	152.77	93.40	0.40	0.19
Canada	18	809.71	296.95	0.77	0.11
Switzerland	9	725.27	136.57	1.84	0.12
Czech Republic	13	12.45	7.87	0.10	0.03
Germany	17	1,493.70	465.62	0.49	0.13
Denmark	14	260.46	76.98	1.00	0.15
Spain	18	233.03	143.10	0.23	0.09
Finland	13	45.13	15.38	0.23	0.05
France	14	1,456.34	600.87	0.66	0.18
United Kingdom	18	3,062.95	1,075.63	1.49	0.28
Greece	13	10.49	4.62	0.05	0.01
Hungary	17	8.65	9.12	0.08	0.06
Ireland	7	253.39	102.63	1.15	0.21
Iceland	7	19.25	7.95	1.14	0.20
Italy	13	473.46	246.79	0.26	0.10
Japan	18	3,327.69	374.91	0.59	0.08
South Korea	6	368.09	87.98	0.42	0.04
Netherlands	18	899.32	319.63	1.48	0.28
Norway	13	105.06	38.10	0.39	0.03
Poland	16	26.21	32.77	0.08	0.07
Portugal	13	62.29	27.26	0.33	0.09
Slovakia	12	3.02	2.37	0.06	0.02
Sweden	13	302.44	95.06	0.80	0.15
United States	18	12,900.00	3,246.92	1.08	0.14

for comparison. With the exception of number of employees, the previous results of a positive impact of PE investment on industry performance are robust. The coefficients on the PE investment variable actually increase substantially in magnitude. Interpreting this estimate as a LATE suggests that local PE investors, who are more affected by the instrument, have a particularly large effect on growth rates. In unreported analyses, we also repeat this exercise using lagged and twice lagged assets-to-GDP as the instrument, and the results remain consistent.

Third, we address the endogeneity issue using Granger (1969) causality. This empirical approach investigates the relative timing of related time series, in our case PE investments, relative to production and employment growth. PE investments will Granger-cause production growth (or employment growth) if a previous increase in PE investments is associated with a subsequent increase in production, but a previous increase in production growth is unrelated with subsequent changes in PE investments. Granger causality has been widely studied and applied in macroeconomics, and there has been substantial debate over the interpretation of the causality concept. The concerns and caveats are well understood. Since we have separate time series for each country-industry pair, we adopt a panel Granger analysis. This is a more recent extension of the traditional approach and is less established (see Hartwig, 2009). We adopt a natural parsimonious empirical specification.

We estimate a three-equation system of linear equations. The endogenous variables are the total production and employment growth rates and an indicator of PE activity (we use an indicator for PE activity in each year, not the past five years as used above). The exogenous variables are lags of the endogenous variables, in addition to country and industry fixed effects. We first estimate the system using

Table 27: Instrumental variables analysis

The table contains OLS and 2SLS regression coefficients. An observation is a country-industry-year pair. The endogenous variable is the deviation of the annual growth rate of production, value added, labor costs, and total employment (as defined by OECD) relative to the average rate in the same industry and year. The exogenous variables are an indicator for positive PE activity over the previous five years at the country-industry level (PE_5), and industry- and country-fixed effects as indicated. The 2SLS specifications use the fraction of assets held by domestic institutional investors to GDP to instrument PE, along with country and industry fixed effects in the first stage. Standard errors in parentheses are calculated with clustering at the country-year level (both at first and second stage) and are presented in parentheses. Statistical significance at the 1%, 5% and 10% levels are indicated by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Production (gross output)	Production (gross output)	Value Added	Value Added	Labor Costs (compensation of employees)	Labor Costs (compensation of employees)	Number of Employees	Number of Employees
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
PE_5	0.906***	2.414**	1.117***	2.838**	0.684***	2.309**	0.546***	0.102
	(0.241)	(1.167)	(0.270)	(1.108)	(0.253)	(1.079)	(0.177)	(0.892)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,976	5,918	7,013	5,951	6,743	5,745	5,771	5,098
R-squared	0.177	0.186	0.130	0.134	0.225	0.259	0.067	0.082

GLS (SUR), taking into account cross-equation correlations in the error terms.

Estimated coefficients using two different specifications of the lags are reported in Table 28. Note that the coefficients from the system regressions are identical to the coefficients one would obtain from estimating single-equation OLS regressions. The standard errors, however, adjust for cross-equation correlations in the error terms. In the first equation, we see evidence that PE investments Granger-cause production and employment growth. In the third equation, however, we find no evidence that increases in production or employment growth are associated with subsequent PE investments. Indeed, the individual coefficients are all insignificant and Wald tests for the joint significance of either the production or labor coefficients do not reject the hypothesis that they are zero. Combined, this evidence indicates that the direction of causality likely flows from PE investments to total production and employment growth.

It is well known that including fixed effects in dynamic panel models can lead to biased and inconsistent estimators. Given our long panel — from 1991 to 2007 — we suspect that this problem is small. Nevertheless, Table 16 reports estimates of the system-GMM procedures proposed by Arellano and Bond (1991) and Blundell and Bond (1998) to overcome this problem. These procedures involve first-differencing of the endogenous variables, which is similar to including country-industry fixed effects. The table reports joint significance tests of the lagged PE indicators, lagged labor growth, and lagged production growth. In addition, we report tests for first- and second-order serial correlation in the error terms. First differencing, by construction, generates first-order autocorrelation, as reported, which is fully consistent with the specification. We find some evidence of higher-order autocorrelation for labor growth, however, but not for production growth and the PE

114

Table 28: Granger causality (SUR)

The table contains coefficients from two specifications of a linear SUR/VAR model, each with three
equations. Columns 1-3 contain estimates of the first specification; columns 4-6 contain the second o
Endogenous variables are productivity growth (deviation from annual industry average), labor growt
(deviation from annual industry average), and a PE indicator (equals one for each country-industry-y
with any PE activity). Exogenous variables are lagged endogenous variables with the lag in parenthe
and industry- and country-fixed effects. Note that the coefficients, but not the standard errors, are
identical to those obtained from six single-equation OLS regressions. Reported standard errors allow
cross-equation correlations. Significance levels of Wald tests of the joint significance of the PE,
productivity, and labor coefficients are reported. Statistical significance at the 1%, 5%, and 10% leve
are indicated by ***, **, and *, respectively.

	(1) Prod.(%)	(2) Labr.(%)	(3) PE(0/1)	(4) Prod.(%)	(5) Labr.(%)	(6) PE(0/1)
PE_1 (1)	0.918***	0.562***	0.096***	0.894***	0.607***	0.098***
	(0.297)	(0.218)	(0.015)	(0.312)	(0.217)	(0.015)
PE_1 (2)	0.000	-0.051	0.115***	0.126	0.012	0.115***
	(0.307)	(0.225)	(0.015)	(0.321)	(0.224)	(0.015)
PE_1 (3)	-0.433	-0.007	0.144***	-0.342	0.015	0.151***
	(0.312)	(0.228)	(0.016)	(0.327)	(0.228)	(0.015)
PE_1 (4)	0.213	-0.287	0.136***	0.210	-0.281	0.133***
	(0.323)	(0.237)	(0.016)	(0.339)	(0.236)	(0.016)
PE_1 (5)	-0.221	0.039	0.101***	-0.151	-0.000	0.102***
	(0.333)	(0.244)	(0.017)	(0.349)	(0.244)	(0.016)
Prod. Growth (1)	0.137***	0.161***	0.000	0.154***	0.166***	0.000
	(0.015)	(0.011)	(0.001)	(0.015)	(0.011)	(0.001)
Prod. Growth (2)	-0.110***	0.041***	-0.000	-0.051***	0.055***	-0.000
	(0.016)	(0.011)	(0.001)	(0.016)	(0.011)	(0.001)
Prod. Growth (3)	0.093***	0.076***	-0.000	0.086***	0.065***	-0.001
	(0.015)	(0.011)	(0.001)	(0.015)	(0.011)	(0.001)
Prod. Growth (4)	0.005	0.006	-0.000			
	(0.015)	(0.011)	(0.001)			
Prod. Growth (5)	0.145***	0.065***	0.001			
	(0.015)	(0.011)	(0.001)			
Labor Growth (1)	0.071***	0.089***	-0.000	0.068***	0.088***	-0.001
	(0.022)	(0.016)	(0.001)	(0.022)	(0.016)	(0.001)
Labor Growth (2)	-0.008	-0.032**	-0.000	-0.014	-0.026*	-0.000
	(0.022)	(0.016)	(0.001)	(0.022)	(0.015)	(0.001)
Labor Growth (3)	-0.007	-0.054***	-0.001	0.009	-0.032**	-0.001
	(0.021)	(0.015)	(0.001)	(0.020)	(0.014)	(0.001)
Labor Growth (4)	-0.065***	-0.025*	-0.002*			
	(0.020)	(0.014)	(0.001)			
Labor Growth (5)	-0.036**	-0.026**	0.000			
	(0.016)	(0.012)	(0.001)			
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Prob[PE = 0]	0.040*	0.181	0.000***	0.028**	0.043*	0.000***
Prob[Prod. = 0]	0.000***	0.000***	0.813	0.000***	0.000***	0.892
Prob[Labor = 0]	0.000***	0.000***	0.413	0.024**	0.000***	0.537
Observations	4,730	4,730	4,730	4,939	4,939	4,939
R-squared	0.228	0.366	0.500	0.227	0.387	0.497

115

Table 29: System GMM

The table contains coefficients from six single-equation GMM specifications, using first differences to eliminate country-industry fixed effects. Specifications 1-3 use moments and instruments suggested by Arellano and Bond (1991); specifications 4-6 follow Blundell and Bond (1998). Endogenous variables are productivity growth (deviation from annual industry average), labor growth (deviation from annual industry average), and a PE indicator (equals one for each country-industry-year with any PE activity). Exogenous variables are these variables lagged with the lag in parentheses. Robust standard errors are in parentheses. Significance levels of Wald tests of the joint significance of the PE, productivity, and labor coefficients are reported. Significance levels for Arellano-Bond tests for first- and second-order autocorrelation in the error terms are reported. The null hypothesis is no autocorrelation (by construction, first-differencing creates first-order autocorrelation as reported and fully consistent with the specification). Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *, respectively.

	(1) Prod.(%)	(2) Labr.(%)	(3) PE(0/1)	(4) Prod.(%)	(5) Labr.(%)	(6) PE(0/1)
PE_1 (1)	5.680***	1.954*	0.104***	4.369***	1.748*	0.174***
	(1.441)	(1.007)	(0.038)	(1.441)	(0.920)	(0.021)
PE_1 (2)	-1.210	-0.536	0.127***	-3.481***	-1.557*	0.186***
	(1.363)	(1.042)	(0.031)	(1.279)	(0.920)	(0.021)
PE_1 (3)	1.277	2.000**	0.157***	-0.317	0.789	0.201***
	(1.368)	(1.014)	(0.027)	(1.309)	(0.911)	(0.022)
PE_1 (4)	-1.497	-3.217***	0.153***	-2.731**	-3.664***	0.185***
	(1.456)	(1.105)	(0.026)	(1.333)	(1.040)	(0.023)
PE_1 (5)	0.350	-0.310	0.113***	-0.112	-0.082	0.140***
	(0.498)	(0.298)	(0.021)	(0.396)	(0.256)	(0.020)
Prod. Growth (1)	0.120**	0.267***	0.001	0.154***	0.239***	-0.001
	(0.049)	(0.046)	(0.002)	(0.043)	(0.043)	(0.002)
Prod. Growth (2)	-0.086	-0.002	0.002	-0.067	-0.020	0.001
	(0.060)	(0.034)	(0.002)	(0.054)	(0.031)	(0.002)
Prod. Growth (3)	0.144***	0.141***	0.006***	0.147***	0.128***	0.003*
	(0.036)	(0.039)	(0.002)	(0.034)	(0.034)	(0.002)
Prod. Growth (4)	0.075**	0.076**	0.002	0.054*	0.047	-0.000
	(0.035)	(0.034)	(0.002)	(0.031)	(0.030)	(0.002)
Prod. Growth (5)	0.171***	0.075***	0.003***	0.175***	0.076***	0.002
	(0.029)	(0.019)	(0.001)	(0.028)	(0.017)	(0.001)
Labor Growth (1)	-0.019	-0.085	-0.005	-0.010	0.025	-0.002
	(0.093)	(0.052)	(0.003)	(0.087)	(0.042)	(0.003)
Labor Growth (2)	-0.120	-0.089**	-0.006**	-0.115	-0.044	-0.003
	(0.088)	(0.036)	(0.003)	(0.084)	(0.033)	(0.003)
Labor Growth (3)	-0.158*	-0.131***	-0.010***	-0.127*	-0.098***	-0.006**
	(0.082)	(0.033)	(0.003)	(0.077)	(0.029)	(0.003)
Labor Growth (4)	-0.189**	-0.097***	-0.005**	-0.114*	-0.061**	-0.002
	(0.075)	(0.031)	(0.002)	(0.066)	(0.030)	(0.002)
Labor Growth (5)	-0.018	-0.032	0.002*	-0.022	-0.026	0.001
	(0.026)	(0.025)	(0.001)	(0.023)	(0.023)	(0.001)
Prob[PE = 0]	0.001***	0.014**	0.000***	0.000***	0.000***	0.000***
Prob[Prod. = 0]	0.000***	0.000***	0.039**	0.000***	0.000***	0.316
Prob[Labor = 0]	0.001***	0.001***	0.007***	0.010**	0.001***	0.234
Prob[AR(1)]	0.000***	0.000***	0.000***	0.000***	0.000***	0.000
Prob[AR(2)]	0.302	0.005***	0.774	0.317	0.001***	0.933
Observations	4,309	4,301	4,432	4,766	4,758	4,889

indicator.

The Blundell and Bond specification (Table 29, columns (4) to (6)) provides particularly consistent evidence of Granger causality. In this specification, the lagged PE indicator is significant for predicting future labor and production growth. However, the joint tests suggest that lagged labor- and production-growth are not statistically significant predictors of future PE investments. One remaining concern is that the labor growth process may include higher-order autocorrelation, which may exaggerate the statistical significance and introduce biases in the coefficients. Finally, our PE indicator is discrete, and estimating dynamic models with lagged endogenous discrete variables may introduce additional complications for shorter panels (see Honoré and Kyriazidou 2000). Addressing these technical econometric issues is beyond the scope of this analysis.

2.5 Conclusion

The growth of the PE industry world-wide has spurred concerns about its potential impact on the global economy. This study looks across nations and industries to assess the impact of PE on industry performance.

The key results are, first, that industries where PE funds have invested in the past five years have grown more quickly. There are few significant differences between industries with low and high PE activity, suggesting that the results are at least partly driven by spillover effects from PE-backed firms to other firms in the industry. Second, it is hard to find support for claims that economic activity in industries with PE backing is more exposed to aggregate shocks. Various approaches suggest that the results are not driven by reverse causality. Finally, these patterns are not driven solely by common law nations such as the U.K. and U.S., but also

117

hold in Continental Europe.

Our findings suggest a number of avenues for future research. First, it would be interesting to look at finer data on certain critical aspects of industry performance, such as the rates of layoffs, plant closings and openings, and product and process innovations. Second, it is important to understand the mechanisms by which the presence of PE-backed firms affects their peers. While Chevalier's (1995b) study of the supermarket industry during the 1980s was an important first step, much more remains to be explored. Finally, we are limited by the available data. The buyout boom of the mid 2000s was so massive, and the subsequent crash so dramatic, that the consequences may have been substantially different from other economic cycles (see Kosman 2009). The impact of the recent financial crisis is an important issue to explore in the future.

3. CONTRACTING WITH HETEROGENEOUS EXTERNALITIES

This chapter is coauthored with Eyal Winter

3.1 Introduction

What is the optimal structure of contracts to induce a group of agents to participate in a joint activity? How should these contracts take into account the complex externalities that prevail among the agents? These questions arise in various settings. Governments around the world seek to foster growth and innovation by emulating the success of Silicon Valley and creating planned science parks.[1] To attract companies, policy makers devote substantial resources.[2] Can governments lower the costs of establishing science parks by exploiting the heterogeneous externalities that arise between companies? Mall owners use such a strategy when leasing stores. Gould et al. (2005) show that national brand stores (which attract the most consumer traffic to malls) are being used to attract leases of smaller stores. These smaller stores generate most of the mall owners' leasing revenue.

In many situations a group member decision to participate depends on the choices of others. These relations are hardly symmetrical; in particular, partici-

---

[1] The International Association of Science Parks (IASP) currently has members in 49 countries outside the United States. According to the IASP, the number of science parks in the U.S. alone has increased from 16 in 1980 to 170 in 2010.

[2] For example, Hong Kong spent more than $2 billion to develop a planned research and development park (Cheng 1999).

pation choices may depend not only on how many members decide to participate, but also on the identity of the other participating agents. In a mall, a small store substantially gains from the presence of national brand stores, which attract a high volume of buyer traffic. The opposite externality, induced by the small store, has hardly any effect. The recruitment of a senior star to an academic department can easily attract a junior researcher to apply to that department. Invited party guests base their participation decisions on the participation of their close friends. In all of these examples the relations between the agents should be taken into account when structuring incentives.

In this paper we analyze a principal's problem of coordinating participation given heterogeneous externalities between group members. We explore a project initiated by a principal, when its success depends on the participation of a group. The principal structures a set of incentive contracts to coordinate the group members' participation. Such incentives can be tax credits, discounts, gifts, celebrities' participation, or any other benefits that are conditional on an agent's participation. We characterize the optimal, i.e., the least expensive, contracts that induce the participation of the group members.

In our model, the heterogeneous externalities are additive and described by a matrix whose entry $w_i(j)$ represents the extent to which agent i benefits from joint participation with agent j.[3] Following Segal (2003) we focus on situations in which the principal cannot coordinate agents to his preferred equilibrium in a multiple-equilibria setting. That is, we mainly focus on contracts that sustain agents' participation in a unique Nash equilibrium.[4] This set of contracts is of the form of

---

[3] We consider non-additive externalities in Section 3 of the paper

[4] Recent experimental papers (see Brandt and Cooper 2005) indicate that in an environment of positive externalities agents typically are trapped in the bad equilibrium of no-participation.

divide and conquer.[5] For any given ranking of the agents, divide-and-conquer contracts are structured in the following way: offer each agent a reward that would convince him to participate in the belief that the agents who precede him in the ranking participate, and all subsequent agents abstain. Thus, the optimal contract is achieved by the ranking (henceforth, optimal ranking) that produces the least expensive divide-and-conquer incentive scheme.

Given the complex relations between the agents due to heterogeneity we ask: (1) Who should be getting a higher-powered incentive for participation? In other words, how should we determine the optimal ranking of the agents? (2) How do changes in the structure of externalities affect the principal's cost of sustaining the group's participation?

We show that the optimal ranking can be constructed using a virtual popularity tournament between the agents. In this tournament, we say that agent i beats agent j if agent j's benefit from i's participation is greater than i's benefit from j's participation. This binary relation is described by a directed graph. We use basic graph theory arguments to characterize the optimal ranking which depends on the number of winnings in the virtual tournament.[6] Hence, the agents' payoffs are determined with respect to their success in the tournament. This idea that agents who induce higher externalities receive higher-powered incentive rewards is supported by an empirical paper by Gould et al. (2005) who demonstrate that while national brand stores occupy over 58% of the total leasable space in shopping malls they

---

[5] Segal (2003) uses a similar structure to characterize a setting of homogeneous externalities. Che and Yoo (2001) show that a similar structure arises as an optimal mechanism in a moral hazard in team setups.

[6] The ranking is directly determined by the number of winnings if the directed graph is acyclic. If the graph is cyclic, ranking depends on the number of winnings as well as on the differences between agents' externalities.

pay only 10% of the total rent collected by the mall owners.

A key characteristic of group externalities is the level of asymmetry[7] between the pairs of agents, which we show to reduce the principal's cost. Greater asymmetry offers the principal more leverage in exploiting the externalities to lower costs. This result has a significant implication on the principal's choice of group for the initiative in the selection stage.

Our problem surprisingly connects to two quite distinct topics: (1) ranking sport teams based on tournament results, which has been discussed in the Operations Research literature, and (2) ranking candidates in a voting problem based on the outcomes of pairwise elections, which was suggested by Condorcet (1785). Condorcet's solution uses a similar approach to ours, where candidates are the nodes in the graph, and the arcs' directions are the election results of pair-wise voting.

This work is part of an extensive body of literature on multi-agent contracting in which externalities arise between the agents.[8] Our general approach is closely related to the seminal papers by Segal (1999, 2003) on contracting with externalities. Segal (2003) introduced a general model of trade contracts that admit externalities among agents. He shows that increasing externalities implies that the principal gains from using a divide-and-conquer mechanism, when he cannot coordinate

_____

[7] By asymmetry we refer to the sum of differences in bilateral externalities.

[8] To give a few examples, these applications include vertical contracting models (Katz and Shapiro 1986a; Kamien, Oren, and Tauman 1992) in which the principal supplies an intermediate good to N identical downstream firms (agents), which then produce substitute consumer goods; employment models (Levin 2002) in which a principal provides wages to induce effort in a joint production of a group of workers; exclusive dealing models (Rasmusen, Ramseyer, and Wiley 1991; Segal and Whinston 2000) in which the principal is an incumbent monopolist who offers exclusive dealing contracts to N identical buyers (agents) in order to deter the entry of a rival; acquisition for monopoly models (Lewis 1983; Kamien and Zang 1990; Krishna 1993) in which the principal makes acquisition offers to N capacity owners (agents); and network externalities models (Katz and Shapiro 1986b).

122

players to play his most-preferred equilibrium. Segal's model is sufficiently general to fit nicely into a variety of IO applications (like takeovers, vertical contracting, exclusive dealing, and network externalities). While Segal (2003) defines the divide-and-conquer mechanism in a general contracting setup that allows for heterogeneity, he doesn't solve for the optimal mechanism except for special cases such as the symmetric case (although he is able to obtain some comparative static results without deriving the optimal mechanism explicitly). Our objective here is to solve for the optimal mechanism for any matrix of externalities. While our environment is more restrictive than Segal's in the sense that agents' choices are binary (participate or not), we develop a sharper characterization by allowing externalities to be heterogeneous and thus capture the contracting implications of complex relations between the agents.

In fact, most of the literature assumes that externalities are homogeneous (in our setting, such an assumption implies that the benefit from joint participation depends on the number of participants, and not on the identity of the agents). Exceptions to this assumption are Jehiel and Moldovanu (1996) and Jehiel, Moldovanu, and Stachetti (1996), who consider an auction in which a single indivisible object is sold to multiple heterogeneous agents. Jehiel and Moldovanu (1999) introduce resale markets and consider the implications of the identity of the initial owner of the good to the initial consumer. Our paper is also related to Milgrom and Roberts (1990) who pointed out that a principal can gain from collusion or coordination among his agents in an interaction that gives rise to strategic complementarity.

We consider several extensions to verify the robustness of our assumptions. First, we study situations in which agents' choices are sequential and we show that our solution applies when the principal is interested in implementing effort via a

123

stronger solution concept that admits a dominant strategy for each player at his relevant subgame. We show that the analysis remains valid when we allow the externalities to affect agents' outside options, as well as for more complicated contingent contracts. We consider more general externalities structures. In particular, we allow externalities to be both negative and positive, and provide the conditions under which the solution for the mixed externalities participation problem can be derived by decomposing the mixed problem to two problems one of which is positive and the other negative. Finally, we consider the case of a non-additive externalities structure.

The rest of the paper is organized as follows. In Section 3.2 we introduce the general model. Section 3.3 provides the solution to a participation problem with positive externalities between the agents. In Section 3.4 we consider several extensions of the model, by which we demonstrate that our results apply in more general settings. Section 3.5 demonstrates how the model can be used to solve selection problems and Section 3.6 concludes. Proofs are presented in the Appendix.

3.2   *The Model*

A participation problem is given by a triple $(N, \mathbf{w}, \mathbf{c})$ where N is a set of n agents. The agents' decision is binary: participate in the initiative or not. The structure of externalities $\mathbf{w}$ is an $n \times n$ matrix specifying the bilateral externalities between the agents. An entry $w_i(j)$ represents agent i's added value from participation in the initiative jointly with agent j. Agents gain no additional benefit from their own participation, i.e., $w_i(i) = 0$. Agents' preferences are additively separable; i.e., agent i's utility from participating jointly with a group of agents M is $\sum_{j \in M} w_i(j)$ for every $M \subseteq N$. In the extensions section we consider a model in which agents'

124

preferences are non-additive; i.e., externalities are defined over all subsets of agents in group N.

We assume that the externalities structure $\mathbf{w}$ is fixed and exogenous. Also, $\mathbf{c}$ is the vector of the outside options of the agents. For simplicity, and with a slight abuse of notation, we assume that every outside option is constant and equals c for all agents. In the extensions section we demonstrate that our results hold also when the outside options are affected by the participation choices of the agents.

We assume that contracts offered by the principal are simple and descriptive in the sense that the principal cannot provide payoffs that are contingent on the participation behavior of other agents. Many of the examples discussed above seem to share this feature. Based on the data used by Gould et al. (2005) which includes contractual provisions of over 2,500 stores in 35 large shopping malls in the U.S., there is no evidence that contracts make use of such contingencies. The theoretical foundation for the absence of such contracts is beyond the scope of this paper. One possible explanation is the complexity of such contracts. In Section 3.4 we demonstrate that our analysis remains valid even if we allow contingencies to be added to the contracts.

The set of contracts offered by the principal can be described as an incentives vector $\mathbf{v} = (v_1, v_2, ..., v_n)$ in which agent i receives a payoff of v_i if he decides to participate and zero otherwise. v_i is not constrained in sign and the principal can either pay or charge the agents but he cannot punish them for not participating (limited liability). Given a contracting scheme $\mathbf{v}$, agents face a normal-form game

$G(\mathbf{v})$.[9,10] Each agent has two strategies in the game: participation or abstention. For a given set M of participating agents, each agent $i \in M$ earns $\sum_{j \in M} w_i(j) + v_i$ and each agent $j \notin M$ earns c, his outside option. We define *full implementation* contracts to be contracts that induce group participation as a unique Nash equilibrium. Alternatively, *partial implementation* contracts induce the group to participate in a Nash equilibrium, which is not necessarily unique.

3.3   Contracting with Positive Externalities

Positive externalities are likely to arise in many contracting situations. Network goods, opening stores in a mall and attracting customers, and contributing to public goods are a few such examples. In this section we consider situations in which agents benefit from the participation of the other agents in the group. Suppose that $w_i(j) > 0$ for all $i, j \in N$, such that $i \neq j$. In this case, agents are more attracted to the initiative as the set of participants grows.

As a first step toward characterizing optimal full implementation contracts, we show in Proposition 1 that an optimal contracting scheme is part of a general set of contracts characterized by the *divide-and-conquer*[11] property. This set of contracts is constructed by ranking agents in an arbitrary fashion, and by offering each agent

---

[9] We view the participation problem as a reduced form of the global optimization problem faced by the principal, which involves both the selection of the optimal group for the initiative and the design of incentives. Specifically, let U be a (finite) universe of potential participants. For each $N \subseteq U$ let $v^*(N)$ be the total payment made in an optimal mechanism that sustains the participation of the set of agents N. The principal will maximize the level of net benefit he can guarantee himself, which is given by the following optimization problem: $\max_{N \subseteq U}[u(N) - v^*(N)]$, where $u(N)$ is the principal's gross benefit from the participation of the set N of agents and is assumed to be strictly monotonic with respect to inclusion; i.e., if $T \subsetneq S$, then $u(T) < u(S)$.

[10] In the extensions section we also consider the case of a sequential offers game.

[11] See Segal (2003) and Winter (2004) for a similarly structured optimal incentive mechanism in a setting of homogeneous externalities.

a reward that would induce him to participate in the belief that all the agents who precede him in the ranking participate and all subsequent agents abstain. Due to positive externalities, "later" agents are induced to participate (implicitly) by the participation of others and thus can be offered smaller (explicit) incentives. More formally, the *divide-and-conquer* (*DAC*) contracts have the following structure:

$$\mathbf{v} = (c, c - w_{i_2}(i_1), c - w_{i_3}(i_1) - w_{i_3}(i_2), ..., c - \sum_k w_{i_n}(i_k))$$

where $\boldsymbol{\varphi} = (i_1, i_2, ..., i_n)$ is an arbitrary order of agents. We say that $\mathbf{v}$ is a DAC contracting scheme with respect to the ranking $\boldsymbol{\varphi}$. The following proposition, which is similar to the analysis in Segal (2003, subsection 4.1.1) provides a necessary condition for optimal contracts.

Proposition 1 *If $\mathbf{v}$ is an optimal full implementation contracting scheme then it is a divide-and-conquer contracting scheme.*

Note that given contracting scheme v, agent i_1 has a dominant strategy in the game $G(\mathbf{v})$ to participate.[12] Given the strategy of agent i_1, agent i_2 has a dominant strategy to participate as well. Agent i_k has a dominant strategy to participate provided that agents i_1 to i_{k-1} participate as well. Therefore, contracting scheme $\mathbf{v}$ sustains full participation through an iterative elimination of dominated strategies.

---

[12] Since rewards take continuous values we assume that if an agent is indifferent then he chooses to participate. Alternatively, we can define an optimal contract to be a vector of payments v such that any arbitrary increase of the payoffs of all players will result in a unique equilibrium in which all agents participate. Note also that while the indifferences may generate multiple Nash equilibria, full participation is a unique Nash equilibrium with iterative elimination of weakly dominated strategies (without assuming that agents participate whenever they are indifferent).

3.3.1 Optimal Ranking

The optimal contracting scheme satisfies the divide-and-conquer property with the ranking that minimizes the principal's payment. The optimal ranking is determined by a virtual popularity tournament among the agents, in which each agent is "challenged" by all the other agents. The results of the matches between all pairs of agents are described by a *simple* and *complete*[13] directed graph $G(N, A)$, where N is the set of nodes and A is the set of arcs. N represents the agents, and $A \subset N \times N$ represents the results of the matches, which is a binary relation on N. We refer to such graphs as *tournaments*.[14] More precisely, the set of arcs in tournament $G(N, A)$ is as follows:

(1) $w_i(j) < w_j(i) \iff (i, j) \in A$

(2) $w_i(j) = w_j(i) \iff (i, j) \in A$ and $(j, i) \in A$

The interpretation of a directed arc (i, j) in the tournament G is that agent j values mutual participation with agent i more than agent i values mutual participation with agent j. We simply say that agent i *beats* agent j whenever $w_i(j) < w_j(i)$. In the case of a two-sided arc, i.e., $w_i(j) = w_j(i)$, we say that agent i is *even* with agent j and the match ends in a tie.

In characterizing the optimal contracts we distinguish between cyclic and acyclic tournaments. We say that a tournament is *cyclic* if there exists at least one node v for which there is a directed path starting and ending at v, and *acyclic* if no such path exists for all nodes.[15] The solution for cyclic tournaments relies on the acyclic solution, and therefore the acyclic tournament is a natural first step.

---

[13] A directed graph $G(N, A)$ is *simple* if $(i, i) \notin A$ for every $i \in N$ and *complete* if for every $i, j \in N$ at least $(i, j) \in A$ or $(j, i) \in A$.

[14] We allow that (i, j) and (j, i) are both in A.

[15] By definition, if $(i, j) \in A$ and $(j, i) \in A$, then the tournament is cyclic.

3.3.2 Optimal Ranking for Acyclic Tournaments

A ranking φ is said to be *consistent* with tournament $G(N, A)$ if for every pair $i, j \in N$, if i is ranked before j in φ, then i beats j. In other words, if agent i is ranked higher than agent j in a consistent ranking, then agent j values agent i more than agent i values j. We start with the following graph theory lemma:

Lemma 1 *If tournament $G(N, A)$ is acyclic, then there exists a unique ranking that is consistent with $G(N, A)$.*

We refer to the unique consistent ranking proposed in Lemma 1 as the *tournament ranking*.[16] In the tournament ranking, each agent's location in the tournament ranking is determined by the number of his wins. Hence, the agent ranked first is the agent who won all matches and the agent ranked last lost all matches. As we demonstrate later, there may be multiple solutions when tournament $G(N, A)$ is cyclic. Proposition 2 provides the solution for participation problems with acyclic tournaments, and shows that the solution is unique.

Proposition 2 *Let $(N, \mathbf{w}, \mathbf{c})$ be a participation problem for which the corresponding tournament $G(N, A)$ is acyclic. Let φ be the tournament ranking of $G(N, A)$. The optimal full implementation contracting scheme is given by the DAC with respect to φ.*

The intuition behind Proposition 2 is based on the notion that if agents $i, j \in N$ satisfy $w_i(j) < w_j(i)$ then the principal is able to reduce the cost of incentives by $w_j(i)$, rather than by only $w_i(j)$, by giving preferential treatment to i and placing him higher in the ranking. Applying this notion to all pairs of agents minimizes the

---

[16] The tournament ranking is actually the ordering of the nodes in the unique Hamiltonian path of tournament $G(N, A)$.

129

principal's total payment to the agents, since it maximizes the inherent value of the participants from the participation of the other agents.

The optimal contracting scheme can be viewed as follows. First the principal pays the outside option c for each one of his agents. The winner of each match in the virtual tournament is the agent who imposes a higher externality on his competitor. The loser of each match pays the principal an amount equal to the benefit that he gets from mutually participating with his competitor. The total amount paid depends on the size of bilateral externalities and not merely on the number of winning matches. However, the higher agent i is located in the tournament, the lower is the total amount paid to the principal.

An intuitive solution for the participation problem is to reward agents according to their level of popularity in the group, such that the most popular agents would be the most rewarded. A possible interpretation of popularity in our context would be the sum of externalities imposed on others by participation, i.e., $\sum_{j=1}^{n} w_j(i)$. However, as we have seen, agents' ranking in the optimal contracting scheme is determined by something more refined than this standard definition of popularity. Agent i's position in the ranking depends on the set of peers that value agent i's participation more than i values theirs. This two-way comparison may result in a different ranking than the one imposed by a standard definition of popularity. This can be illustrated in the following example in which agent 3 is ranked first in the optimal contracting scheme despite being less "popular" than agent 1.

Example Consider a group of four agents with an identical outside option $c = 20$. The externalities structure of the agents is given by matrix $\mathbf{w}$, as shown in Figure 6. The tournament G is acyclic and the tournament ranking is $\boldsymbol{\varphi} = $

$(3, 1, 2, 4)$. Consequently, the set of optimal contracts is $\mathbf{v} = (20, 17, 14, 10)$, which is the *divide-and-conquer* contracting scheme with respect to the tournament ranking. Note that agent 3 who is ranked first is not the agent who has the maximal $\sum_{j=1}^{n} w_j(i)$.

Figure 6

The derivation of the optimal contracting scheme requires the rather elaborate step of constructing the virtual tournament. However, it turns out that a substantially simpler formula can derive the cost of the optimal contracts. Two terms play a role in this formula: the first measures the aggregate level of externalities, i.e., $K_{agg} = \sum_{i \, j} w_i(j)$; the second measures the bilateral asymmetry between the agents, i.e., $K_{asym} = \sum_{i<j} |w_i(j) - w_j(i)|$. Hence, K_{asym} stands for the extent to which agents induce mutual externalities on each other. The smaller the value of K_{asym} the higher the degree of mutuality of the agents. Proposition 3 shows that the cost of the optimal contracting scheme is additive and declining in these two measures.

Proposition 3 *Let $(N, \mathbf{w}, \mathbf{c})$ be a participation problem and V_{full} be the principal's cost of the optimal full implementation contracts. If the corresponding tournament $G(N, A)$ is acyclic then $V_{full} = n \cdot c - \frac{1}{2} \left(K_{agg} + K_{asym} \right)$.*

An interesting consequence of Proposition 3 is that for a given level of aggregate

131

externalities, the principal's payment is decreasing with a greater level of asymmetry among the agents, as stated in Corollary 3.1.

Corollary 3.1 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem with an acyclic tournament. Let* V_{full} *be the principal's cost of the optimal full implementation contracts. For a given level of aggregate externalities,* V_{full} *is strictly decreasing with the asymmetry level of the externalities within the group of agents.*

The intuition behind this result is related to the virtual tournament discussed above. In each match the principal extracts "fines" from the losing agents. It is clear that these fines are increasing with the level of asymmetry (assuming $w_i(j) + w_j(i)$ is kept constant). Hence, a higher level of asymmetry allows the principal more leverage in exploiting the externalities. This observation has important implications for the principal's selection stage.

Consider the comparison between the optimal full and partial implementation contracts, where in the latter the principal suffices with the existence of a full participation equilibrium, not necessarily unique. With partial implementation, the cost for the principal in the optimal contracting scheme is substantially lower. More specifically, in the least costly contracting scheme that induces full participation, each agent i receives $v_i = c - \sum_j w_i(j)$. However, these contracts entail a no-participation equilibrium as well; hence coordination is required. The total cost of the partial implementation contracts is $V_{partial} = n \cdot c - \sum_{i\,j} w_i(j)$ and the principal can extract the full revenue generated by the externalities.[17]

---

[17] Our emphasis on full implementation is motivated by the fact that under most circumstances the principal cannot coordinate the agent to play his most-preferred equilibrium. Brandts and Cooper (2005) report experimental results that speak to this effect. Agents' skepticism about the prospects of the participation of others trap the group in the worst possible equilibrium even when the group is small. Nevertheless, one might be interested in evaluating the cost of moving from partial to full implementation.

It is worth mentioning that for a fixed level of aggregate externalities, the difference between full and partial implementation contracts, $V_{full} - V_{partial}$, is stricty decreasing with the level of asymmetry of the externalities within the group. In the extreme case where $K_{asym} = 0$ ($w_i(j) = w_j(i)$ for all pairs), the cost of moving from partial to full implementation is the most expensive. On the other hand, when externalities are always one-sided, i.e., for each pair of agents $i, j \in N$ satisfies that either[18] $w_i(j) = 0$ or $w_j(i) = 0$,then the additional cost is zero and full implementation is as expensive as partial implementation.

Note that increasing the aggregate level of externalities will not necessarily increase the principal extraction of revenue in the optimal contracting scheme. For example, in an asymmetric two-person problem raising slightly the externality that the less attractive agent induces on the other one will not change the principal's revenue.[19] From the perspective of the agents, their reward is *not* a continuous increasing function of the externalities they impose on the others. However, it is possible that a slight change in these externalities may increase rewards significantly, since a minor change in externalities may change the optimal ranking and thus affect agents' payoffs.

The extreme asymmetric case nicely contrasts with the completely symmetric case, where the principal's surplus increases with any slight increase of the externalities. With partial implementation, which allows the principal full extraction of surplus, the principal's revenue is sensitive to the values of externalities whether the problem is symmetric or asymmetric.

---

[18] Since this section deals with positive externalities, assume that $w_i(j) = \varepsilon$ or $w_j(i) = \varepsilon$ when ε is very small.

[19] It can be shown that in an *n*-person asymmetric problem one can raise the externalities in half of the matrix's entries (excluding the diagonal) without affecting the principal's surplus extraction.

3.3.3 *Optimal Ranking of Cyclic Tournaments*

In the previous section we demonstrated that optimal full implementation contracts are derived from a virtual tournament among the agents in which agent i beats agent j if $w_i(j) < w_j(i)$. However, the discussion was based on the tournament being acyclic. If the tournament is cyclic, the choice of the optimal DAC contracting scheme (i.e., the optimal ranking) is more delicate since Lemma 1 does not hold. Any ranking is prone to inconsistencies in the sense that there must be a pair i, j such that i is ranked above j although j beats i in the tournament. To illustrate this point, consider a three-agent example where agent i beats j, agent j beats k, and agent k beats i. The tournament is cyclic and any ranking of these agents necessarily yields inconsistencies. For example, take the ranking $\{i, j, k\}$, which yields an inconsistency involving the pair (k, i) since k beats i and i is ranked above agent k. This applies to all possible rankings of the three agents.

The inconsistent ranking problem is similar to problems in sports tournaments, which involve bilateral matches that may turn out to yield cyclic outcomes. Various sports organizations (such as the National Collegiate Athletic Association - NCAA) nevertheless provide rankings of teams/players based on the cyclic tournament outcome. Extensive literature in Operations Research suggests solution procedures for determining the "minimum violation ranking" (e.g., Kendall 1955, Ali et al. 1986, Cook and Kress 1990, and Coleman 2005) that selects the ranking for which the number of inconsistencies is minimized. It can be shown that this ranking is obtained as follows. Take the cyclic (directed) graph obtained by the tournament and find the smallest set of arcs such that reversing the direction of these arcs results in an acyclic graph. The desired ranking is taken to be the consis-

tent ranking (per Lemma 1) with respect to the resulting acyclic graph.[20]

One may argue that this procedure can be improved by assigning weights to arcs in the tournament depending on the score by which team i beats team j and then look for the acyclic graph that minimizes the total weighted inconsistencies. In fact this approach goes back to Condorcet's (1785) classical voting paper in which he proposed a method for ranking multiple candidates. In the voting game, the set of nodes is the group of candidates, the arcs' directions are the results of pairwise voting, and the weights are the plurality in the voting. The solution to our problem follows the same path. In our framework arcs are not homogeneous and so they will be assigned weights determined by the difference in the bilateral externalities. As in Condorcet's voting paper, we will look for the set of arcs such that their reversal turns the graph into an acyclic one. While Young (1988) characterized Condorcet's method axiomatically, our solution results from a completely different approach, i.e., the design of optimal incentives to maximize revenues.

Formally, we define the weight of each arc $(i, j) \in A$ by $t(i, j) = w_j(i) - w_i(j)$. Note that weights are always non-negative as an arc (i, j) refers to a situation in which j favors i more than i favors j. Hence $t(i, j)$ refers to the extent of the one-sidedness of the externalities between the pairs of agents. If an inconsistency in the ranking arises due to an arc (i, j), then this implies that agent j precedes agent i despite the fact that i beats j. Relative to consistent rankings, inconsistencies generate additional costs for the principal. More precisely, the principal has to pay an additional $t(i, j)$ when inconsistency is due to arc $(i, j) \in A$.

For each subset of arcs $S = \{(i_1, j_1), (i_2, j_2), ..., (i_k, j_k)\}$ we define $t(S) = \sum_{(i,j) \in S} t(i, j)$, which is the total weight of the arcs in S. For each graph G and subset of arcs S we

---

[20] Multiple rankings may result from this method.

135

denote by G_{-S} the graph obtained from G by reversing the arcs in the subset S. Consider a cyclic graph G and let S^* be a subset of arcs that satisfies the following:

(1) G_{-S^*} is acyclic.

(2) $t(S^*) \leq t(S)$ for all S such that G_{-S^*} is acyclic.

Then, G_{-S^*} is the acyclic graph obtained from G by reversing the set of arcs with the minimal total weight, and S^* is the set of pairs of agents that satisfies inconsistencies in the tournament ranking of G_{-S^*}. Proposition 4 shows that the optimal ranking of G is the tournament ranking of G_{-S^*} since the additional cost from inconsistencies, $t(S^*)$, is the lowest.

Proposition 4 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem with a cyclic tournament G. Let $\boldsymbol{\varphi}$ be the tournament ranking of G_{-S^*}. Then, the optimal full implementation contracts are the DAC with respect to $\boldsymbol{\varphi}$.*

In the symmetric case, the principal cannot exploit the externalities among the agents, as $K_{asym} = 0$, and the total payment made by the principal is identical for all rankings. This can be seen to follow from Proposition 4 as well by noting that the tournament has two-way arcs connecting all pairs of agents, and $t(i, j) = 0$ for all i, j and $t(S)$ is uniformly zero. An intriguing feature of the symmetric case is that all optimal contracting schemes are discriminative in spite of the fact that all agents are identical.

Corollary 4.1 *When the externalities structure $\mathbf{w}$ is symmetric then all DAC contracts are optimal.*

We can now provide the analogue version of Proposition 3 for the cyclic case. In this case, the optimal ranking has an additional term $K_{cyclic} = t(S^*)$ representing

the cost of making the tournament acyclic, i.e., the cost borne by the principal due to inconsistencies.

Proposition 5 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem. Let* V_{full} *be the principal's optimal cost of a full implementation contract. Then* $V_{full} = n \cdot c - \frac{1}{2}(K_{agg} + K_{asym}) + K_{cyclic}$.

Corollary 3.1 still holds for pairs of agents that are not in S^*. More specifically, if we increase the level of asymmetry between pairs of agents that are outside S^*, we reduce the total expenses that the principal incurs in the optimal contracting scheme.

3.4 Extensions

In this section we discuss the implications of the assumptions we made so far. We demonstrate that the optimal contracts remain optimal if we assume sequential participation choices when the principal desires to implement participation in a subgame perfect equilibrium with the property that each player has a dominant strategy on the subgame that he plays. In addition, we show that even when the outside option is affected by the agents' participation choices, the construction of the optimal contracts remains unchanged. We demonstrate that when contracts can be contingent on the participation of a subset of the agents, then the optimal contracts are closely related to the analysis above. Our analysis is valid in more general setups in which externalities can be either negative or positive. Moreover, the solution is also relevant to non-additive externalities structures.

3.4.1 Sequential Participation Decisions

We first point out that our analysis applies to any sequential game except for one of perfect information, i.e., when each player is fully informed about all the participation decisions of his predecessors. Indeed, this extreme case of perfect information is a strong assumption as agents rarely possess the participation decisions of *all* their predecessors. Any partial information environment implies that some actions are taken simultaneously, and therefore the divide-and-conquer contracting scheme and the virtual tournament apply.

Nevertheless, it is interesting to point out that our analysis is also relevant to the extreme case of perfect information. Consider a game in which players have to decide sequentially about their participation based on a given order. Suppose that the principal wishes to implement the full participation in a subgame perfect equilibrium with the additional requirement that each player has a dominant strategy on the subgame in which he has to play.[21] It is easily verified that the optimal contracting scheme in this framework is the DAC applied to the order of moves; i.e., the first moving player is paid c and the last player is paid $c - \sum_{j \in N} w_i(j)$. Under this contracting scheme each player has a dominant strategy on each subgame. Assume now that the principal can control the order of moves (which he can do by making the offers sequentially and setting a deadline on agents' decisions). Then the optimal sequential contracting scheme is exactly identical to the one discussed in previous sections for the simultaneous case. If the principal suffices with a standard subgame perfect equilibrium (without the strategy dominance condition), then the optimal contracting scheme will allow him to extract more and he

---

[21] Such a requirement may reflect the principal's concern that a player will fail to apply complex backward induction reasoning.

will pay $c - \sum_{j \in N} w_i(j)$ to all agents.

3.4.2 *Participation-dependent Outside Options*

In many situations non-participating agents are affected by the participation choices of other agents. Consider the case of a corporate raider who needs to acquire the shares of N identical shareholders to gain control (similar to Grossman and Hart 1980). If the raider is enhancing the value of the firm when he holds a larger stake in the firm, then selling shareholders impose positive externalities on non-participating agents. If the raider gains private benefits from the firm which will decrease its value, then selling shareholders induce negative externalities on the non-participating agents.

In this section we consider the case in which the agents' outside option is partly determined by the agents who choose to participate. For a given group of agents $P \subseteq N$ who participate, we define the outside option of non-participants as $c + \eta \sum_{j \in P} w_i(j)$. In the former analysis we assumed $\eta = 0$.[22] Segal (2003) defines externalities as increasing (decreasing) when an agent is more (less) eager to participate when more agents participate. In our setup, eagerness to participate is identity-dependent. When $\eta \leq 1$, we say that agents are more eager to participate when highly valued agents choose to participate. If $\eta > 1$, the benefits of non-participation outweigh the benefits of participation when highly valued agents choose to participate; hence agents are less eager to participate. In Segal's terminology, the former case is equivalent to *increasing* externalities, while the latter is

---

[22] The following analysis can be generalized by specifying an externalities matrix q that distinguishes between the benefits of participating and non-participating agents. It can be shown that in such a case our analysis remains unchanged. However, we choose to use the simpler and more intuitive outside option form of $c + \eta \sum_{j \in C} w_i(j)$.

equivalent to *decreasing* externalities.

Following the analysis of Proposition 1, if $\mathbf{v}$ is an optimal full implementation contracting scheme then it is easy to verify that under the current setup, $\mathbf{v}$ is a DAC of the form:

$$\mathbf{v} = (c, c - (1 - \eta)w_{i_2}(i_1), ..., c - (1 - \eta)\sum_k w_{i_n}(i_k))$$

where $\varphi = (i_1, i_2, ..., i_n)$ is an arbitrary ranking. Instead if $\eta > 1$, the participation problem is identical to a standard participation problem (with fixed outside option) where externalities are $(1 - \eta)w_i(j) < 0$. In these negative externalities problems the DAC mechanism does not apply and the optimal scheme requires that the principal reimburse the agents for their total burden, i.e., $c - (1 - \eta)\sum_j w_i(j))$, which is a positive number whenever the outside option and $w_i(j)$ are positive. Finally, the case of $\eta = 1$ corresponds to an environment of no externalities at all and the optimal scheme requires simply to reimburse agents for their outside option. We can summarize with the following proposition:

Proposition 6 *Let* $(N, \mathbf{w}, \mathbf{c}^*)$ *be a participation problem where* $c_i^* = c + \eta \sum_{j \in P} w_i(j)$ *and* $P \subseteq N$ *is a group of participating agents. Let* $G(N, A)$ *be the equivalent tournament. The optimal full implementation contracts are given as follows:*

(1) for $\eta < 1$, *DAC contracts with respect to the optimal ranking;*[23]

(2) for $\eta \geq 1$, *the optimal mechanism pays agent i the payoff* $c - (1 - \eta)\sum_j w_i(j))$, *which is exactly c, whenever* $\eta = 1$.

Note that one could consider a different case in which the outside option of agents is a linear function of the externalities agents induce. Also in this case,

---

[23] As described in Section 3.5.

asymmetry improves the principal's rent extraction.

3.4.3 Contingent Contracts

Our model assumes that the principal cannot write contracts that make a payoff to an agent contingent on the participation of other agents. Under such contracts the principal could extract the total surplus from positive externalities among the agents.[24] We find such contracts not very descriptive. Based on the data used by Gould et al. (2005) which consists of contractual provisions of over 2,500 stores in 35 large shopping malls in the U.S., there is no evidence that contracts make use of such contingencies. Shopping malls are a natural environment for contingent contracting; the fact that these contracts are still not used makes it likely that in other, more complicated settings such contracts are exceptional as well. The theoretical foundation for the absence of such contracts is beyond the scope of this paper. However, one possible reason for their absence is the complexity of such contracts, especially in environments where participation involves long-term engagement and may be carried out by different agents at different points in time. We point out that if partial contingencies are used, i.e., participation is contingent on a subset of the group, our model and its analysis remain valid. Specifically, for each player i, let $T_i \subseteq N$ be the contingency set, i.e., the set of agents whose participation choice can appear in the contract with agent i. Let $T = (T_1, T_2, ..., T_n)$ summarize the contingency sets in the contracts. The optimal contracts under the contingency sets are closely related to the original optimal contract (when contingencies are not

---

[24] One possible contracting scheme is to offer agent i a participation reward of $v_i = c - \sum_{j \in N} w_i(j)$ if each of the other agents participates, and a reward of $v_i = c$ if any of the contingencies is violated. Such contracts will sustain full participation as a unique Nash equilibrium, and the principal extracts the entire surplus.

allowed). More precisely, let **w** be the original matrix of externalities. Denote by w^T the matrix of externalities obtained from **w** by replacing $w_i(j)$ with zero whenever $j \in T_i$. Lemma 6.1 in the Appendix shows that the optimal full implementation contracting scheme is as follows: agent i gets c if one of the agents $j \in T_i$ does not participate, i.e., the contingency requirement is violated.[25] If all agents in T_i participate, then agent i gets the payoff $v_i(N, \mathbf{w}^T, \mathbf{c}) - \sum_{j \in T_i} w_i(j)$, where $v_i(N, \mathbf{w}^T, \mathbf{c})$ is the payoff for agent i for the participation problem $(N, \mathbf{w}^T, \mathbf{c})$ under no-contingencies.

3.4.4 Mixed Externalities Structure

So far we have limited our discussion to environments in which agents' participation positively affects the willingness of other agents to participate. However, in many situations this is not the case, such as in environments of congestion. Traffic, market entry, and competition among applicants all share the property that the larger the number of agents who participate, the lower the utility of each participant. The heterogeneous property in our framework seems quite descriptive in some of these examples. In the context of competition it is clear that a more qualified candidate/firm induces a larger negative externality. It is also reasonable to assume, at least for some of these environments, that the principal desires a large number of participants in spite of the negative externalities that they induce on each other.

In Proposition 7 we demonstrate that in order to sustain full participation as a unique Nash equilibrium under negative externalities the principal has to fully compensate all agents for the participation of the others.

---

[25] In fact, the principal can offer lower payments to the agents in case of contingencies' violations, by exploiting the participation of other agents. However, these off-equilibrium payments do not affect the principal's payment in the full participation equilibrium.

Proposition 7 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem with negative externalities. Then optimal full implementation contracts* $\mathbf{v}$ *are given by* $v_i = c + \sum_{i \neq j} |w_i(j)|$ *, and* $\mathbf{v}$ *is unique.*

Naturally, real-world multi-agent contracting problems may capture both positive and negative types of externalities. In social events, individuals may greatly benefit from some of the invited guests, while preferring to avoid others. In a mall, the entry of a new store may benefit some stores by attracting more customers, but impose negative externalities on its competitors.

Our analysis of the mixed externalities case is based on the following binary relation. We say that an agent i is *non-averse* to agent j if $w_i(j) \geq 0$, and we write it as $i \succeq j$. We will assume that $\succeq$ is symmetric and transitive, i.e., $i \succeq j \implies j \succeq i$ and if $i \succeq j$ and $j \succeq k$ then $i \succeq k$. Note that this assumption does not imply any constraint on the magnitude of the externalities, but just on their sign. While the symmetry and transitivity of the *non-averse* relation seem rather intuitive assumptions, not all strategic environments satisfy them. These assumptions are particularly relevant to environments where the selected population is partitioned into social, ethnic, or political groups with animosity potentially occurring only between groups but not within groups. We analyze a specific example of this sort of environment in Section 3.5.

It turns out that the optimal solution of participation problems with symmetry and transitivity of the *non-averse* relation is derived by a decomposition of the participation problem into two separate participation problems: one that involves only positive externalities, and the other that involves only negative externalities. This is done by simply decomposing the externalities matrix into a negative and a positive matrix. In the following proposition we show that the *decomposition con-*

tracting scheme, a contract set that is the sum of the two optimal contracts of the two decomposed participation problems, is the optimal contracting scheme for the mixed externalities participation problem.

Proposition 8 *Consider a participation problem* $(N, \mathbf{w}, \mathbf{c})$. *Let* $(N, \mathbf{w}^+, \mathbf{c})$ *be a participation problem such that* $w_i^+(j) = w_i(j)$ *if* $w_i(j) \geq 0$ *and* $w_i^+(j) = 0$ *if* $w_i(j) < 0$, *and let* u^+ *be the optimal full implementation contracts of* (N, w^+, c). *Let* $(N, \mathbf{w}^-, \mathbf{0})$ *be a participation problem such that* $w_i^-(j) = w_i(j)$ *if* $w_i(j) < 0$ *and* $w_i^-(j) = 0$ *if* $w_i(j) \geq 0$, *and let* $\mathbf{u}^-$ *be the optimal full implementation contracts of* $(N, \mathbf{w}^-, \mathbf{0})$. *Then, the decomposition contracting scheme* $\mathbf{v} = \mathbf{u}^+ + \mathbf{u}^-$ *induces a unique full participation equilibrium. Moreover, if agents satisfy symmetry and transitivity with respect to the non-averse relation,* $\mathbf{v}$ *is the optimal contracting scheme.*

Proposition 8 shows that the virtual popularity tournament discussed in earlier sections plays a central role also in the mixed externalities case as it determines payoffs for the positive component of the problem. When symmetry and transitivity hold, the principal can exploit the positive externalities to reduce payments. In this tournament i beats j whenever (1) $w_j(i) \geq 0$ and $w_j(i) \geq 0$, and (2) $w_j(i) > w_i(j)$. Note that under the *non-averse* assumptions, the principal provides complete compensation for the agents who suffer from negative externalities, as in the negative externalities case. Finally, it is easy to show that equivalently to Proposition 5, the principal's cost of achieving full implementation in a mixed externalities setting is equivalent to the positive externalities setup, except that now the principal has to add the compensation for the negative externalities.

3.4.5 Non-additive Preferences

We propose here an extension of the model in which we impose no restrictions on agents' preferences; i.e., preferences are no longer assumed to be separably additive. Using an iterative procedure that makes use of the solution for the additive case allows us to narrow down the set of potential optimal incentive contracts, even when no structure is assumed.

A participation problem is described by a group of agents N and their outside option is equal to c, as noted previously. We assume a general externalities structure, which is composed of the non-additive preferences of the agents over all subsets of agents in the group N. More specifically, for each i, $v_i : 2^{N\setminus\{i\}} \longrightarrow R$. The function $v_i(S)$ stands for the benefit of agent i from the participation with the subset $S \subseteq N$. We normalize $v(\varnothing) = 0$. The condition of positive externalities now reads: for each i and subsets S, T such that $T \subset S$ we have $v_i(S) \geq v_i(T)$.

Arguments similar to those used in Proposition 1 show that the optimal contracting scheme that sustains full participation as a unique equilibrium also satisfies the divide-and-conquer property. Hence, the optimal contracts rely on the optimal ranking of the agents.

We leave the detailed description of the procedure to the proof of Proposition 9. Instead, we provide an example to illustrate the basic ideas.

A Simple Example

Consider a four-agent example. Given that the optimal solution is DAC for any given ranking of agents $\varphi = \{i_1, i_2, i_3, i_4\}$, the DAC contracts with respect to ranking φ are $(c, c - v_{i_2}(i_1), c - v_{i_3}(i_1, i_2), c - v_{i_4}(i_1, i_2, i_3))$. Instead of identifying the optimal ranking, we apply an iterative procedure of $N - 1$ steps to eliminate

rankings that we infer cannot be optimal. Our starting point is the set of all possible rankings of the agents; in this example there are 24 such rankings.

STEP 1. Let's assume that the bilateral externalities $v_i(j)$ between the agents result in the corresponding acyclic graph described below. Therefore the tournament yields the unique consistent ranking for step one when $\phi_1 = (3, 1, 2, 4)$.

Figure 7

We argue that any ranking that orders the first two agents in a way that contradicts their relative ranking in ϕ_1 cannot be the optimal ranking. To see this, consider the ranking $(4, 2, 1, 3)$, which is inconsistent with ϕ_1 with respect to the relative ranking of agents 4 and 2. We can immediately construct a cheaper ranking by reversing the position of the first two agents, and keeping the position of the remaining agents ranked lower in the same order. Hence, we can eliminate $(4, 2, 1, 3)$ from the set of potential optimal rankings. Applying this logic to the entire set of potential rankings we are left with 12 potential rankings; i.e., the optimal ranking of the original problem must start with any of the following pairs: $(3, 1)$, $(3, 2), (3, 4), (1, 2), (1, 4), (2, 4)$.

STEP 2. We now proceed to the second iteration in which for each agent located in the first position we construct a graph that is based on the bilateral relations

conditional on the participation of the first agent. In particular, we consider the case in which agent 1 is ranked first and build the graph based on agents' preferences conditional on the participation of agent 1; i.e., the externalities matrix is given by $(w_i(j) = v_i(j,1)|j \in \{2,3,4\})$.

Let's assume that preferences take the following form:

$$v_2(3,1) > v_3(2,1)$$
$$v_2(4,1) > v_4(2,1)$$
$$v_3(4,1) > v_4(3,1)$$

Since the graph is acyclic the unique consistent ranking of the second iteration, conditional on agent 1 being first, is $\boldsymbol{\phi}_{2|1} = (4,3,2)$. Again, we require rankings to be consistent with $\boldsymbol{\phi}_{2|1}$. For example, ranking $(1,2,4,3)$ cannot be optimal since $(2,4,3)$ is not consistent with $\boldsymbol{\phi}_{2|1}$ and transposing the order of 2 and 4 we get ranking $(1,4,2,3)$, which is cheaper. While there are six rankings in which agent 1 is ranked first, we can immediately eliminate three that do not agree with $\boldsymbol{\phi}_{2|1}$ and we are left with $\{(1,4,2,3),(1,4,3,2),(1,3,2,4)\}$. However, these rankings must agree with the constraints from the previous step. This is not the case for ranking $(1,3,2,4)$, as we can transpose the order of 1 and 3 and get a cheaper mechanism; thus we can eliminate it as well.[26] Hence, if the optimal ranking starts with agent 1 it must be followed by agent 4 ranked second. Rather than discussing the construction of cases where agents 2 and 3 are ranked first, we continue to explore the case where agent 1 is ranked first and proceed to step 3.

STEP 3. In this iteration we repeat and construct the graph based on agents 2

---

[26] We refer to this check as the *interface condition* and discuss it more fully in the proof of Proposition 9.

and 3's preferences, conditional on the participation of agents 1 and 4. Let's assume that $v_2(3, 1, 4) < v_3(2, 1, 4)$; hence $\boldsymbol{\phi}_{3|1,4} = \{2, 3\}$. Thus, the only ranking that can be optimal in the original problem conditional on agent 1 being first is $(1, 4, 2, 3)$.

General Result

The example above illustrates our procedure for generating the optimal incentive contracts can also be used iteratively to eliminate non-optimal rankings, when we impose no structure on agents' preferences.

The starting point is the set of all agents' rankings. We proceed with an iterative procedure of $N - 1$ steps; at each step we rule out possible rankings by constructing a graph that is based on the bilateral preferences of the agents conditional on the participation of the agents ranked above them. We assume that in each step the resulting graph is acyclic and thus generates a unique consistent ranking. We eliminate rankings that are inconsistent with the step's consistent ranking or with the constraints imposed in the previous step. The formal description of this iterative procedure is provided in the proof of Proposition 9.

Proposition 9 *Let $(N, \mathbf{c})$ be a participation problem with non-additive preferences, for which all tournaments in the iterative procedure are acyclic. Then, the set of surviving rankings is non-empty and includes the optimal ranking.*

Proposition 9 demonstrates that the fundamental logic underlying our analysis of additive externalities also underlies our construction of optimal contracts, while taking into account the complex structure of externalities among agents.

3.5 Group Identity and Selection

In this section we consider special externalities structures to demonstrate how the selection stage can be incorporated once we have solved the participation problem. Assume that the externalities take values of 0 or 1. In this environment an agent either benefits from the participation of his peer or gains no benefit. We provide three examples of group identities in which the society is partitioned into two groups and agents have hedonic preferences for members in these groups. We demonstrate how the optimal contracting scheme proposed in previous sections may affect the selection of the agents in the planning of the initiative.

(1) **Segregation** - agents benefit from participating with their own group's members and enjoy no benefit from participating with members from the other group. More specifically, consider the two groups B_1 and B_2 such that for each $i, j \in B_k$, $k = 1, 2$, we have $w_i(j) = 1$. Otherwise, $w_i(j) = 0$.

(2) **Desegregation**[27] - agents benefit from participating with the other group's members and enjoy no benefit from participating with members of their own group. More specifically, consider the two groups B_1 and B_2 such that for each $i, j \in B_k$, $k = 1, 2$, we have $w_i(j) = 0$. Otherwise, $w_i(j) = 1$.

(3) **Status** - the society is partitioned into two status groups, high and low. Each member of the society benefits from participating with each member of the high-status group and enjoys no benefit from participating with members of the low-status group. Formally, let B_1 be the high status group and set $w_i(j) = 1$ if and only if $j \in B_1$. Otherwise $w_i(j) = 0$.

---

[27] An example could be a singles party.

Proposition 10 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem. Let* n_1 *and* n_2 *be the number of agents selected from groups* B_1 *and* B_2, *respectively, such that* $n_1 + n_2 = n$. *Denote by* $v(n_1, n_2)$ *the principal's cost of incentivizing agents under the optimal contracts given that the group composition is* n_1 *and* n_2. *The following holds:*

1) under Segregation $v(n_1, n_2)$ *is decreasing with* $\mid n_1 - n_2 \mid$;

2) under Desegregation $v(n_1, n_2)$ *is increasing with* $\mid n_1 - n_2 \mid$;

3) under Status $v(n_1, n_2)$ *is decreasing with* n_1.

In the case of Segregation, the principal's cost of incentives is increasing with the mixture of groups; hence in the selection stage the principal would prefer to give precedence to one group over the other. In the Desegregation case the principal's cost is declining with mixture; hence in the selection stage the principal would like to balance between members of the groups. In the Status case the cost is declining with the number of agents recruited from B_1, who will be strongly preferred to the members of B_2.

3.6 Conclusion

In this paper we analyzed a contracting framework with heterogeneous externalities. Introducing a complicated structure of externalities allowed us to explore several aspects of the multi-agent contracting environments that are not apparent in the homogeneous case. These include the impact of externalities asymmetry on payments, the implications of externalities structure on the hierarchy of incentives, and the effect of variations in the externalities structures on both the principal's payments and the agents' rewards.

Exploring the role of heterogeneous externalities reveals the importance of externalities asymmetry within the group. More specifically, greater asymmetry between the agents' benefits reduced the principal's payment in the full implementation problem. In addition, externalities asymmetry turns out to play a role also in the selection between partial and full implementation, as it affects the premium required to sustain full participation as a unique equilibrium. Greater asymmetry decreases this premium, and thus makes full implementation more profitable.

The hierarchy of incentives is determined by a ranking that results from a virtual popularity tournament. In the simplest case, an agent i is ranked above agent j if agent i benefits less from joint participation than agent j does from it. We demonstrated that this ranking of incentives is different from the standard ranking that is based on agents' popularity.

The implications of externalities on the rankings of agents may suggest a preliminary game in which agents invest effort to increase the positive externalities that they induce on others which can ultimately increase the rewards from the principal. For example, agents can invest in their social efforts to make themselves more attractive guests at social events. A firm may invest to increase its market share in order to improve its ranking in an acquisition game. Under certain circumstances such an investment may turn out to be quite attractive as we have seen that a slight change in externalities may result in a substantial gain, due to a change in the ranking. Our analysis contributes to the understanding of how to make such a strategic investment profitable. The new game will now have two stages. In the first stage agents' investment efforts change the exterrnalities they induce on other members of the group. In the second stage, the externalities determine the network formation and the consequent incentives provided by the principal. The analysis

151

of such a game is beyond the scope of this paper but seems to be a natural next step.

APPENDIX

A. APPENDIX TO CHAPTER 1

A. Variable Definitions

Innovation Measures

1. *Citations* - Number of citations a patent receives from its grant year and the following three calendar years.

2. *Generality* - A patent that is being cited by a broader array of technology classes is viewed as having greater generality. Generality is calculated as the Herfindahl index of *citing* patents, used to capture the dispersion across technology classes of patents using the patent. To account for cases with a small number of patents within technology classes, I use the bias correction described in Jaffe and Trajtenberg (2002).

3. *Originality* - A patent that cites a broader array of technology classes is viewed as having greater originality. Originality is calculated as the Herfindahl index of *cited* patents, used to capture dispersion of the patent citations across technology classes. To account for cases with a small number of patents within technology classes, I use the bias correction described in Jaffe and Trajtenberg (2002).

4. *Scaled Citations* - Number of citations a patent receives divided by the average number of citations received by all patents granted in the same year and

technology class.

5. *Scaled Generality* - Generality measure of a patent divided by the average generality of all patents granted in the same year and technology class.

6. *Scaled Originality* - Originality measure of a patent divided by the average originality of all patents granted in the same year and technology class.

7. *Scaled Number of Patents* - Each patent is adjusted for variations in patent filings likelihood and for truncation bias. The truncation bias in patent grants stems from the lag in patent approval (of about two years). Thus, towards the end of the sample, patents under report the actual patenting since many patents, although applied for, might not have been granted. Following Hall, Jaffe, and Trajtenberg (2001), the bias is corrected by dividing each patent by the average number of patents of all firms in the same year and technology class.

8. *Technology Class* - A technology class is a detailed classification of the U.S. Patenting and Trademark Office (USPTO) which clusters patents based on similarity in the essence of their technological innovation. Technological classes are often more detailed than industry classifications, consisting of about 400 main (3-digit) patent classes, and over 120,000 patent subclasses. For example, within the communications category, there are various technology classes such as: wave transmission lines and networks, electrical communications, directive radio wave systems and devices, radio wave antennas, multiplex communications, optical wave guides, etc.

IPO Characteristics

9. *Firm Age* - Firm age at the year of the IPO filing, calculated from the founding date. Firm age of firms that went public is kindly available at Jay Ritter's webpage. I collected the firm age of withdrawn firms from registration statements.

10. *Early Follower* - An indicator variable that captures the location of a filer within the IPO wave. Following Beneveniste et al. (2003), a filer is considered an early follower if filed within 180 days of a pioneer in the same Fama-French 48 industry.

11. *Pioneer* - An indicator variable that captures the location of a filer within the IPO wave. Following Beneveniste et al. (2003), a filer is considered a pioneer if its filing is not preceded by an IPO filing in the same Fama-French 48 industry in the previous 180 days.

12. *Lead Underwriter Ranking* - A ranking of the lead underwriter on a scale of 0 to 9, where 9 is the highest underwriter prestige. The ranking is compiled by Carter and Manaster (1990), Carter, Dark, and Singh (1998), and Loughran and Ritter (2004).

13. *VC-Backed* - An indicator is equal to one if the firm was funded by a venture capital firm at the time of the IPO filing.

14. *Post-filing NASDAQ returns* - The two-month NASDAQ returns calculated from the day of the IPO filing.

15. *Pre-filing NASDAQ returns* - The three-month NASDAQ returns leading to the IPO filing date.

16. *Log Total Assets* - the natural logarithm of the total book value of assets.

17. *R&D / Assets* - the ratio of R&D expenditure to book value of assets.

18. *Net Income / Assets* - the ratio of net income to book value of assets.

19. *Cash / Assets* - the ratio of cash holdings to book value of assets.

B. *Quasi-Maximum Likelihood Poisson Model*

A standard approach in the technological innovation literature is to estimate count-data, such as patents and citations, with quasi-maximum likelihood (QML) method. This implies assuming that the conditional mean has the following structure:

$$E\left(Y_i^{post}|X_i\right) = \exp(\alpha + \beta IPO_i + \gamma X_i) \tag{A.1}$$

An important property of the QML is that the standard errors are consistent under fairly general conditions, even if the underlying data-generating process is not poisson. In addition, the estimator can be used for any non-negative dependent variables, whether integer or continuous (Santos Silva and Tenreyro, 2006) and QML standard errors are robust to arbitrary patterns of serial correlation (Wooldridge, 1997). The QML Poisson reported coefficients reflect marginal effects (derivative of dependent variable with respect to covariates) to allow easy comparison to the OLS estimates.[1]

---

[1] Standard errors are adjusted appropriately using the delta-method.

The instrumental variable Poisson estimate uses a control-function approach (Blundell and Powell, 2004). To illustrate this approach, consider first the exogenous case, in which IPO_i is not correlated with the error term that satisfies:

$$E(v_i | IPO_i, X_i) = 1$$

when v_i denotes the error term in equation (4). This will not hold if IPO_i is correlated with the error term. Assume that $NSDQ_i$ (two-month NASDAQ fluctuations) satisfy:

$$IPO_i = \alpha^o + \delta NSDQ_i + \beta^o X_i^o + v_i^o$$

with

$$E(v_i^o | X_i^o) = 1$$

then controlling for v_i^o in the conditional mean equation is sufficient to remove the endogeneity bias. In the estimation, I use the extended moment condition

$$E\left(Y_i^{post} | X_i\right) = \exp(\alpha + \beta IPO_i + \gamma X_i + v_i^o)$$

Intuitively, the residual v_i^o captures the endogenous variation within the variable IPO_i; adding it as a control variable in the second-stage estimation allows identifying β consistently.

Table A.1 - Placebo Test of Reduced Form Results

The dependent variable is the average scaled citations per patent over the five years after the IPO filing. *Post-IPO Filing NASDAQ returns* are the two-month NASDAQ returns calculated after the IPO filing date. *Post-Ownership choice NASDAQ returns* are the two-month NASDAQ returns calculated after either the date of the equity issuance or the date of the IPO filing withdrawal. When the date of IPO filing withdrawal is not available, I use the date of 270 days subsequent to the last amendment of the IPO filing (Lerner 1994). The variables included in the regressions are pre-filing average scaled citations, pre-filing number of patents, Pioneer, Early follower, VC-backed variable, and the three-month NASDAQ returns before the IPO filing. Variables are defined in Section A of the Appendix. The estimated model is Ordinary Least Squares (OLS), and robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Dependent Variable	Scaled Citations	Scaled Citations	Scaled Citations
Post-IPO filing NASDAQ returns	−0.498**		−0.482**
	(0.239)		(0.237)
Post-Ownership choice NASDAQ returns		0.150	0.162
		(0.254)	(0.248)
Observations	1,079	1,079	1,079
R-squared	0.242	0.240	0.242
Filing year FE	yes	yes	yes
Industry FE	yes	yes	yes
Control variables	yes	yes	yes

Table A.2 - NASDAQ Drops are Not Correlated with Long-run Innovation Trends

The table reports the association of the two-month NASDAQ returns after the IPO filing date with changes in innovation trends within the core technologies of filing firms. Firm's technology class is defined as a core technology if the share of patents in that class, in the three years before the IPO filing, is above the median share of patents across all the technology classes of the firm. Innovation trends of core technologies are calculated using all patents granted by the USPTO in the respective technology classes. The dependent variable in column (1) is the change in average patent quality calculated by the average scaled citations per patent of all patents approved in each filer's core technology in the five years after the IPO filing, divided by the average scaled citations in the three years prior to the IPO filing in the respective technology. In column (2), the dependent variable is the change in the total number of patents in the core technologies. In column (3), the dependent variable is the weighted change in the number of patents, when patents are weighted by number of citations. The estimated model is Ordinary Least Squares (OLS) and robust Standard errors are reported in parentheses. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	(1) Patent Novelty	(2) Patent Counts	(3) Weighted Patent Counts
Post-IPO filing NASDAQ returns	−0.007	−0.055	0.001
	(0.053)	(0.142)	(0.171)
Observations	1,372	1,372	1,372
R-squared	0.789	0.275	0.429
Industry FE	yes	yes	yes
Filing Year FE	yes	yes	yes
Control Variables	yes	yes	yes

B. APPENDIX TO CHAPTER 3

Proof of Proposition 1 Let $\mathbf{v} = (v_{i_1}, v_{i_2}, ..., v_{i_n})$ be an optimal full implementation contracting scheme of the participation problem $(N, \mathbf{w}, \mathbf{c})$. Hence, $\mathbf{v}$ generates full participation as a unique Nash equilibrium. Since no-participation is not an equilibrium, at least a single agent, say i_1, receives reward is at least as high as his outside option c. Otherwise, a no-participation equilibrium exists. Due to the optimality of $\mathbf{v}$ his payoff would be exactly c. Agent i_1 chooses to participate under any profile of other agents' decisions. Given that agent i_1 participates and an equilibrium of a single participation is not feasible, at least one other agent, say i_2, receives a reward at least as high as $c - w_{i_2}(i_1)$. Since $\mathbf{v}$ is the optimal contracting scheme, i_2's reward equals $c - w_{i_2}(i_1)$, and under any profile of decisions i_2 will participate. Applying this argument iteratively on the first $k - 1$ agents, at least one other agent, henceforth i_k, must get a payoff at least as high as $c - \sum_{j=1}^{k-1} w_{i_k}(j)$, but again, since v is optimal, the payoff for agent k must be equal to $c - \sum_{j=1}^{k-1} w_{i_k}(j)$. Hence, the optimal contracting scheme $\mathbf{v}$ must satisfy the *divide-and-conquer* property with respect to a ranking φ. ∎

Proof of Lemma 1 We will demonstrate that there is a single node with $n - 1$ outgoing arcs. Since the tournament is a complete, directed, and acyclic graph there cannot be two such nodes. If such a node does not exist, then all nodes in G have both incoming and outgoing arcs. Since the number of nodes is finite, we get a contradiction to G being acyclic. We denote this node as i_1 and place its

corresponding agent first in the ranking (hence this agent beats all other agents). Now let us consider a subgraph $G(N^1, A^1)$ that results from the removal of node i_1 and its corresponding arcs. Graph $G(N^1, A^1)$ is directed, acyclic, and complete and, therefore, following the previous argument, has a single node that has exactly $n - 2$ outgoing arcs. We denote this node as i_2, and place its corresponding agent at the second place in the ranking. Note that agent i_1 beats agent i_2 and therefore the ranking is consistent so far. After the removal of node i_2 and its arcs we get subgraph $G(N^2, A^2)$ and consequently node i_3 is the single node that has $n - 3$ outgoing arcs in subgraph $G(N^2, A^2)$. Following this construction, we can easily observe that the ranking $\varphi = (i_1, i_2, ..., i_n)$ is consistent among all pairs of agents and due to its construction is also unique. ∎

Proof of Proposition 2 According to Proposition 1 the optimal contracting scheme satisfies the DAC property. Hence the optimal contracting scheme is derived from constructing the optimal ranking and is equivalent to minimizing the sum of incentives, V_{full} :

$$
\begin{aligned}
V_{full} &= \min_{(j_1, j_2, ..., j_n)} \left[n \cdot c - \left\{ \sum_{k=1}^{1} w_{j_1}(j_k) + \sum_{k=1}^{2} w_{j_2}(j_k) + ... + \sum_{k=1}^{n} w_{j_n}(j_k) \right\} \right] \\
&= \max_{(j_1, j_2, ..., j_n)} \left[\sum_{k=1}^{1} w_{j_1}(j_k) + \sum_{k=1}^{2} w_{j_2}(j_k) + ... + \sum_{k=1}^{n} w_{j_n}(j_k) \right]
\end{aligned}
$$

Since no externalities are imposed on nonparticipants, the outside options of the agents have no role in the determination of the optimal contracting scheme. We will show that the ranking that solves the maximization problem of the principal is the tournament ranking. Let us assume, without loss of generality, that the tournament ranking φ is the identity permutation: hence $\varphi(i) = i$, and $W_\varphi = \sum_{k=1}^{2} w_2(k) + ... + \sum_{k=1}^{n} w_n(k)$, where W_φ is the principal's revenue extraction. By

162

way of contradiction, assume that there exists $\varphi \neq \sigma$ such that $W_\varphi \leq W_\sigma$. First, assume that σ is obtained from having two *adjacent* agents i and j in φ trade places such that i precedes j in φ and j precedes i in σ. By Lemma 1, agent i beats agent j; thus $W_\sigma = W_\varphi - w_j(i) + w_i(j)$ and $W_\sigma < W_\varphi$.

Note that since φ is the tournament ranking, agent 1 beats all agents, agent 2 beats all agents but agent 1, and so on. Now consider unconstrained $\sigma = \{i_1, ..., i_n\}$ such that $\varphi \neq \sigma$. If agent 1 is not located first, by a sequence of adjacent swaps $(1, i_j)$, we move agent 1 to the top of the ranking. In each of the substitutions agent 1 beats i_j. Next, if agent 2 is not located at the second place, by a sequence of adjacent substitutions $(2, i_j)$, we move agent 2 to the second place. Again, agent 2 beats all agents i_j. The process ends in at most n stages and produces the desired order φ. As demonstrated, any adjacent substitution results in a higher extraction, and so $W_\sigma < W_\varphi$. Therefore, the DAC contracting scheme with respect to the tournament ranking is unique and optimal. ∎

Proof of Proposition 3 Without loss of generality, assume that the tournament ranking φ is the identity permutation. Hence, under the optimal contracting scheme, the principal's payment is $V_{full} = n \cdot c - \left[\sum_{j=1}^{1} w_1(j) + ... + \sum_{j=1}^{n} w_n(j) \right]$. Denote $s_i(j) = [w_i(j) + w_j(i)]$ and $a_i(j) = [w_i(j) - w_j(i)]$. We can represent K_{agg} and K_{asym} in the following manner: $K_{agg} = \sum_{ij} w_i(j) = \sum_{i<j} (w_i(j) + w_j(i)) = \sum_{i<j} s_i(j)$ and $K_{asym} = \sum_{i<j} |a_i(j)|$. Since $w_i(j) = \frac{1}{2}(s_i(j) + a_i(j))$ we can rewrite the principal's payment as

$$
\begin{aligned}
V_{full} &= n \cdot c - \frac{1}{2} \left[\sum_{j=1}^{1} \{s_1(j) + a_1(j)\} + ... + \sum_{j=1}^{n} \{s_n(j) + a_n(j)\} \right] \\
&= n \cdot c - \frac{1}{2} \left(\sum_{i>j} s_i(j) + \sum_{i>j} a_i(j) \right)
\end{aligned}
$$

Note that $s_i(j) = s_j(i)$ and $a_i(j) = -a_j(i)$. In addition $a_i(j) > 0$ when $i > j$ as the tournament is acyclic and ranking is consistent. Therefore,

$$V_{full} = n \cdot c - \frac{1}{2} \left(\sum_{i<j} s_i(j) - \sum_{i<j} |a_i(j)| \right) = n \cdot c - \frac{1}{2} \left(K_{agg} + K_{asym} \right)$$

∎

Proof of Corollary 3.2 The result follows immediately from Proposition 3, where we show that $V_{full} = n \cdot c - \frac{1}{2} \sum_{ij} w_i(j) - \frac{1}{2} \sum_{i<j} |w_i(j) - w_j(i)|$, and from $V_{partial} = n \cdot c - \sum_{ij} w_i(j)$. Taken together, the two yield

$$V_{full} - V_{partial} = \frac{1}{2} \sum_{ij} w_i(j) - \frac{1}{2} \sum_{i<j} |w_i(j) - w_j(i)| = \frac{1}{2} \left(K_{agg} - K_{asym} \right)$$

∎

Proof of Proposition 4 Let $G(N, A)$ be a cyclic graph. Consider a subset of arcs S such that G_{-S} is acyclic, and the tournament ranking of G_{-S} is $\varphi = (j_1, j_2, ..., j_n)$. The payment of the principal V_{full} under the DAC contracting scheme with respect to φ is

$$V_{full} = n \cdot c - \left\{ \sum_{k=1}^{1} w_{j_1}(j_k) + \sum_{k=1}^{2} w_{j_2}(j_k) + ... + \sum_{k=1}^{n} w_{j_n}(j_k) \right\}$$

Note that each $(i, j) \in S$ satisfies an inconsistency in tournament ranking φ. More specifically, if $(i, j) \in S$, then i beats j, and agent j is positioned above agent i. In addition, $w_i(j) = w_j(i) - t(i, j)$, where $w_i(j) < w_j(i)$ and $t(i, j) > 0$. Consider the following substitution: if $(i, j) \in S$ then $w_i(j) = \hat{w}_j(i) - t(i, j)$; otherwise $w_i(j) = \hat{w}_i(j)$. This allows us to rewrite the principal's payment as $V_{full} = n \cdot c - \left\{ \sum_{k=1}^{1} \hat{w}_{j_1}(j_k) + ... + \sum_{k=1}^{n} \hat{w}_{j_n}(j_k) \right\} + t(S)$. Note that $\hat{w}_i(j) = \max(w_i(j), w_j(i))$. Therefore, different rankings affect only the level of $t(S)$, as the first two terms in

164

V_{full} remain indifferent to variations in the ranking. This implies that the subset S with the lowest $t(S)$ brings V_{full} to a minimum. Hence, the optimal contracting scheme is the DAC with respect to the tournament ranking of G_{-S^*}. ∎

Proof of Proposition 5 As demonstrated in Proposition 4, the optimal payment of the principal is the DAC contracting scheme with respect to the tournament ranking of G_{-S^*}. According to Proposition 4, this can be written as

$$V_{full} = n \cdot c - \left\{ \sum_{k=1}^{1} \widehat{w}_{j_1}(j_k) + \ldots + \sum_{k=1}^{n} \widehat{w}_{j_n}(j_k) \right\} + t(S)$$

when $\widehat{w}_i(j) = \max(w_i(j), w_j(i))$. Following the argument of Proposition 3, denote $s_i(j) = [\widehat{w}_i(j) + \widehat{w}_j(i)]$ and $a_i(j) = [\widehat{w}_i(j) - \widehat{w}_j(i)]$ and the principal's payment is $V_{full} = n \cdot c - \frac{1}{2} \left(\sum_{i<j} s_i(j) + \sum_{i<j} |a_i(j)| \right) + t(S) = n \cdot c - \frac{1}{2} \left(K_{agg} + K_{asym} \right) + K_{cyclic}$. ∎

Proof of Proposition 6 The cost of a full implementation contracting scheme is simply $V_{full} = nc - (1 - \eta) \sum_i \sum_{j<i} w_i(j)$. If $\eta = 1$, then the cost does not depend on the externalities. If $\eta < 1$, the minimal cost is obtained by selecting a ranking that maximizes $\sum_i \sum_{j<i} w_i(j)$. This is equivalent to the tournament ranking outlined in Proposition 4. If $\eta > 1$, the participation problem is identical to a standard participation problem (with fixed outside option) where externalities are $(1 - \eta)w_i(j) < 0$. In these negative externalities problems the DAC mechanism does not apply and the optimal scheme requires that the principal reimburse the agents for their total burden, i.e., $c - (1 - \eta) \sum_j w_i(j)$), which is a positive number whenever the outside option and $w_i(j)$ are positive. ∎

Lemma 6.1 *Let* $(N, \mathbf{w}, \mathbf{c})$ *be a participation problem and* $T = (T_1, \ldots, T_n)$ *define the contingency sets. Define* $\mathbf{w}^T$ *to be such that* $w_i^T(j) = w_i(j)$ *if* $j \notin T_i$ *and* $w_i^T(j) = 0$ *otherwise. Let* $\boldsymbol{\varphi}$ *be the optimal ranking of the participation problem* $(N, \mathbf{w}^T, \mathbf{c})$, *and let*

$v(N, \mathbf{w}^T, \mathbf{c})$ be the corresponding DAC payment vector. The optimal full implementation contracts set of $(N, \mathbf{w}, \mathbf{c})$ is such that it provides $\mathbf{c}$ for agent i if contingencies T_i are violated, and $v_i = v_i(N, \mathbf{w}^T, \mathbf{c}) - \sum_{j \in T_i} w_i(j)$ otherwise.

Proof of Lemma 6.1 Since externalities are positive, contingencies allow the principal to reduce payments. In particular, when exploiting all contingencies allowed in T, the contracting scheme that sustains a unique full participation Nash equilibrium offers each agent i a reward $v_i = c - \sum_{j \in T_i} w_i(j)$ if contingencies are met, and c if they are violated. If for all agents $T_i = N / \{i\}$, then full extraction of surplus is possible as a unique equilibrium. However, if only partial contingencies are allowed, i.e., for some agents $T_i \subset N / \{i\}$, then the principal can perform even better than in the contracts outlined above.

Let's define $\hat{w}_i(j) = w_i(j)$ if $j \notin T_i$ and $\hat{w}_i(j) = 0$ otherwise. Consider an arbitrary ranking of agents $\varphi = \{1, 2, ..., n\}$ in which the first agent is offered $v_1 = c - \sum_{j \in T_1} w_1(j)$ if contingencies are met, and c otherwise. Agent 1 will choose to participate. Given the participation of agent 1, we can offer agent 2 the following payment: $v_2 = c - \hat{w}_2(1) - \sum_{j \in T_2} w_2(j)$ if contingencies are met, and c otherwise. Hence, agent 2 will agree to participate given the participation of agent 1. Following the same argument, we could offer the last agent in the ranking $v_n = c - \sum_{i=1}^{n-1} \hat{w}_n(i) - \sum_{j \in T_n} w_n(j)$. This set of contracts will sustain full participation as a unique Nash equilibrium.

The optimal full implementation contracting scheme is thus achieved by obtaining the ranking of agents that will maximize $\sum_i \sum_{j>i} \hat{w}_i(j)$. Given our definition of $\hat{w}_i(j)$, this is equivalent to finding the optimal ranking of agents in the problem (N, w^T, c) when $w_i^T(j) = w_i(j)$ if $\notin T_i$ and $w_i^T(j) = 0$ otherwise. In other words, in the optimal full implementation contracting scheme, the payment for participation

for each agent will be $v_i = v_i(N, w^T, c) - \sum_{j \in T_i} w_i(j)$ if contingencies are met, and c otherwise. $\blacksquare$

Proof of Proposition 7 Given a contracting scheme **v**, participation is a dominant strategy for all agents under the worst-case scenario in which all other agents participate, since $u_i = \sum_{i=1}^{n} w_i(j) + v_i = c$ for every $i \in N$. To show that **v** is optimal, consider a contracting scheme **m** for which $m_i < v_i$ for some agents and $m_i = v_i$ for the rest. By way of contradiction, assume full participation equilibrium holds under contracting scheme **m**. Consider an agent i for which $m_i < v_i$. If all other agents are participating, then agent i's best response is to abstain, since $u_i = \sum_{i=1}^{n} w_i(j) + m_i < c$. Hence, **v** is a unique and optimal contracting scheme. $\blacksquare$

Proof of Proposition 8 See Complementary Note.

Proof of Proposition 9 We proceed in three steps. First we describe the iterative procedure formally. Then we show that the iterative procedure has a finite number of steps with a non-empty set of outcomes. Finally, we demonstrate that the optimal ranking of the original problem is among those orders that survive the procedure.

Formal Description of the Iterative Procedure. The starting point is the set of all possible rankings of the N agents. We start with the first two positions in the ranking. We construct a tournament ranking in which $w_{ij} = v_i(j)$ and eliminate all the rankings in which the first two agents are ordered in a manner that is inconsistent with this ranking. All rankings that survived the elimination provide a possible assignment for the first position in the order. Suppose we have implemented the procedure k-1 steps and obtained a subset of the assignment to the first $k-1$ positions of the order. Let W_{k-1} be the set of sub-orders for the first $k-1$ agents obtained in step $k-1$. For each $w \in W_{k-1}$ we denote by S_w the set of agents assigned to one

167

of the first $k-1$ slots in the assignment w. We now define a tournament T_w on the set of agents $N\backslash S_w$ in such a manner that $w_{ij} = v_i(S_w \cup j)$ defines the externalities matrix. Assuming the graph is acyclic we denote by π_w the uniquely consistent order of the agents in N/S_w.

Next, we construct a subset of assignments to the first k slots based on the set W_{k-1} and the orders π_w for each $w \in W_{k-1}$. We first look at the set of all orders of $N\backslash S_w$ and eliminate all orders in which the first two agents are ordered in a manner that is inconsistent with π_w. We denote by P_w the set of all ordered pairs that survive this elimination.

We now do the following check which we refer to as the *interface condition*: Take a pair $p_w \in Pw$. Suppose this pair is i_w, j_w and let k_w be the last player in **w**. If k_w beats i_w then the suborder w, i_w is a permissible suborder in step k and is added to W_k; otherwise it is excluded. We now proceed in a similar way for every $w \in W_{k-1}$ and the set of permissible suborders of length k which defines W_k.

Claim 1: The process ends in $N-1$ steps and results in a non-empty set of permissible orders.

Proof: We define inductively an order that survives all the steps of this procedure. The first agent in the order is the one ranked first under the tournament $w_{ij} = v_i(j)$ of the first step; call him i_1. The second agent is the one ranked first under the tournament T_w where $w = i_1$.

Suppose that the first $k-1$ slots of the order have been defined. The k-th agent in the order is the one ranked first under the tournament T_w , where $w = i_1, ... i_{k-1}$

Clearly, the interface condition that we defined earlier will never be violated since at each step the agent who is added beats all the agents who are not yet ordered. This implies that the process yields a non-empty set of orders and the

number of steps is finite.

Claim 2: The optimal order is among those orders that survive the procedure.

Proof: Suppose w.l.o.g. that the optimal order is the identity, i.e., $1, 2, 3, ..., n$. We denote this order by π. Suppose by way of contradiction that the optimal order is eliminated at some step k. This means that the interface condition between $k - 1$ and k is violated or that the order was eliminated because the next two agents (say $i, i + 1$) in the order π are not consistent with their order in the tournament ranking of $N \backslash S_w$ of the current stage. This means that $i + 1$ beats i in this tournament on $N \backslash S_w$. This implies that reversing their order will increase the principal's revenue in the divide-and-conquer scheme. Consider the tournament in stage k which is on the set of agents $k, k + 1, ..., n$ and $w_{ij} = v_i(\{1, ...k - 1\} \cup j)$. Let S be the subset of agents in $k, k + 1, ..., n$ such that for each $j \in S$ the agent j is not placed last under the consistent order of the tournament $w_{ij} = v_i(\{1, ...k - 1\} \cup j)$. Any such agent can be the next to be ordered and appear immediately after player $k - 1$. By our assumption any such agent will violate the interface condition. This means that all players in S win against $k - 1$ in the tournament defined in step $k - 1$. We now distinguish between two cases. Case 1: $k \in S$. In this case consider the order obtained by switching the positions of $k - 1$ and k in the original identity order. We denote this order by π'. We note that payments in the divide-and-conquer mechanism for the orders π and π' differ only in terms of players $k - 1$ and k. Furthermore, since k is in S, k beats $k - 1$ in the tournament defined in step $k - 1$. Hence the total payment under π' is less than that under π, which contradicts the optimality of π. We now move to Case 2: k is not in S. In this case k must be ranked last under the tournament $w_{ij} = v_i(\{1, ...k - 1\} \cup j)$. In particular $k + 1$ beats k in this tournament. Consider now the order π' which is identical to π except that $k + 1$

appears before k. As in the argument made earlier, payments to all players but k and $k+1$ are identical in π and π' and because $k+1$ beats k under the tournament $w_{ij} = v_i(\{1,...k-1\} \cup j)$, the order π' corresponds to lower total payments to k and $k+1$, again in contradiction to the optimality of π. $\blacksquare$

Proof of Proposition 10 In both segregated and desegregated environments the externalities structure is symmetric and, following Corollary 5.1, all rankings are optimal. Consider first the segregated environment. Since all rankings are optimal, a possible optimal contracting scheme is $\mathbf{v} = (c,...,c-(n_1-1),c,...,c-(n_2-1))$. Hence, the optimal payment for the principal is $v(n_1,n_2) = n \cdot c - \sum_{l=1}^{n_1-1} l - \sum_{k=1}^{n_2-1} k = n \cdot c - \frac{n_1(n_1-1)}{2} - \frac{(n-n_1)(n-n_1-1)}{2}$. Assuming that $v(n_1,n_2)$ is continuous with n_1, it follows that $\frac{\partial v(n_1,n_2)}{\partial n_1} = n - 2n_1$, the maximal payment is achieved at $n_1^* = n_2^* = \frac{n}{2}$, and the cost of incentivizing is declining with $\mid n_1 - n_2 \mid$. In the desegregated example, a possible optimal contracting scheme is $\mathbf{v} = (c,...,c,c-n_1,...,c-n_1)$. Therefore, the principal's payment is $v(n_1,n_2) = n \cdot c - (n-n_1) \cdot n_1$. Again, let us assume that $v(n_1,n_2)$ is continuous with n_1, in which case solving $\frac{\partial v(n_1,n_2)}{\partial n_1} = 2n_1 - n = 0$ results in the minimal payment for the principal in the desegregated environment being received at $n_1^* = n_2^* = \frac{n}{2}$, and the cost of incentivizing is increasing with $|n_1 - n_2|$. In a status environment, since group B_1 is the more esteemed group, all agents from B_1 beat all agents from B_2; therefore agents from B_1 should precede the agents from B_2 in the optimal ranking. A possible optimal ranking is $\varphi = \{i_1,.,,,i_{n_1},j_1,...,j_{n_2}\}$ when $i_k \in B_1$, $j_m \in B_2$, and $1 \le k \le n_1$, $1 \le m \le n_2$. Therefore, a possible optimal contracting scheme is $\mathbf{v} = (c,c-1...,c-(n_1-1),c-n_1,...,c-n_1)$. The principal's payment is $v(n_1,n_2) = n \cdot c - \sum_{l=1}^{n_1-1} l - n_2 \cdot n_1 = n \cdot c - \frac{n_1(n_1-1)}{2} - (n-n_1)n_1 = \frac{1}{2}n_1 - nn_1 + \frac{1}{2}n_1^2 + cn$. Again, assuming that $v(n_1,n_2)$ is continuous with n_1, $\frac{\partial v(n_1,n_2)}{\partial n_1} = n_1 + \frac{1}{2} - n = 0$ and the

170

minimal payment is achieved at $n_1^* = n - \frac{1}{2}$. Note that $V(n_1 = n) = V(n_1 = n-1)$. Therefore, the best scenario for the principal is when $n_1 = n$. Alternatively, the cost of incentivizing is decreasing with n_1. $\blacksquare$

BIBLIOGRAPHY

Agell, Jonas, and Per Lundborg. 1995. "Theories of pay and unemployment: survey evidence from Swedish manufacturing firms." The Scandinavian Journal of Economics: 295-307.

Aghion, Philippe, John Van Reenen, and Luigi Zingales. 2010. "Innovation and institutional ownership." Unpublished working Paper, Harvard University.

Akcigit, Ufuk, and William R. Kerr. 2010. "Growth Through Heterogeneous Innovations." NBER Working Paper.

Andrade, G. and S. Kaplan. 1998. "How Costly is Financial (Not Economic) Distress? Evidence from Highly Leveraged Transactions that Became Distressed." In Journal of Finance 53(5), 1443-1493.

Angrist, Joshua D., and Guido W. Imbens. 1995. "Two-stage least squares estimation of average causal effects in models with variable treatment intensity."Journal of the American Statistical Association: 431-442.

Arellano, M., and Bond, S. (1991) "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." In Review of Economic Studies 58, 277-297.

Arrow, Kenneth J. 1962. "Economic welfare and the allocation of resources for invention." The Rate and Direction of Inventive Activity, Princeton University Press and NBER.

Asker, John, Joan Farre-Mensa, and Alexander Ljungqvist. 2011. "Comparing the Investment Behavior of Public and Private Firms." Unpublished working paper, New York University.

Atanassov, Julian, Vikram K. Nanda, and Amit Seru. 2007. "Finance and innovation: the case of publicly traded firms." Ross School of Business Paper No. 970.

Axelson, U., Stromberg, P., Jenkinson, T., and M. Weisbach. 2010. "Borrow Cheap, Buy High? The Determinants of Leverage and Pricing in Buyouts." Available at http://ssrn.com/abstract=1596019.

Baker, G.P. and K.H. Wruck. 1989. "Organizational Changes and Value Creation in Leveraged Buyouts: The Case of The O.M. Scott & Sons Company." In Journal of Financial Economics 25, 163-190.

Belenzon, Sharon, Tomer Berkovitz, and Partick Bolton. 2011. "Governance and Innovation." Unpublished working paper, Columbia University Graduate School of Business.

Benveniste, Lawrence M., Alexander Ljungqvist, William J. Wilhelm Jr, and Xiaoyun Yu. 2003. "Evidence of information spillovers in the production of investment banking services." The Journal of Finance 58 (2): 577–608.

Bhattacharya, Sudipto, and Sergei Guriev. 2006. "Patents vs. trade secrets: Knowledge licensing and spillover." Journal of the European Economic Association 4 (6): 1112-1147.

Blinder, Alan S., and Don H. Choi. 1990. "A shred of evidence on theories of wage stickiness." The Quarterly Journal of Economics 105 (4): 1003.

Bloom, N., Sadun, R. and J. Van Reenen. 2009. "Do Private Equity-Owned Firms Have Better Management Practices?" In Gurung, A. and Lerner, J. (eds.) Globalization of Alternative Investments Working Papers Volume 2: Global Economic Impact of Private Equity 2009, New York: World Economic Forum USA, 1-23. Available at http://www.weforum.org/pdf/cgi/pe/Fullreport2.pdf.

Blundell, R., and S. Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." In Journal of Econometrics 87, 115-143.

Blundell, Richard W., and James L. Powell. 2004. "Endogeneity in semiparametric binary response models." Review of Economic Studies 71 (3): 655-679.

Bond, Stephen, and John Van Reenen. 2007. "Microeconometric models of investment and employment." Handbook of econometrics 6: 4417-4498.

Boucly, Q., Sraer, D., and D. Thesmar. forthcoming. "Growth LBOs." In Journal of Financial Economics.

Bound, John, David A. Jaeger, and Regina M. Baker. 1995. "Problems with instrumental variables estimation when the correlation between the instruments and the endogeneous explanatory variable is weak." Journal of the American statistical association: 443-450.

Brandts, Jordi and David Cooper. 2006. A change would do you good. An experimental study on how to overcome coordination failure in organizations. American Economic Review 96 (3): 669-693.

Brau, James C., and Stanley E. Fawcett. 2006. "Initial public offerings: An analysis of theory and practice." The Journal of Finance 61 (1): 399-436.

Brav, Omer. 2009. "Access to capital, capital structure, and the funding of the firm." The Journal of Finance 64 (1): 263-308.

Brown, James R., and Bruce C. Petersen. 2009. "Why has the investment-cash flow sensitivity declined so sharply? Rising R&D and equity market developments." Journal of Banking & Finance 33 (5): 971-984.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen. 2009. "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom." The Journal of Finance 64 (1) (February): 151-185.

Busaba, Walid Y. 2006. "Bookbuilding, the option to withdraw, and the timing of IPOs." Journal of Corporate Finance 12 (2): 159-186.

Busaba, Walid Y., Lawrence M. Benveniste, and Re-Jin Guo. 2001. "The option to withdraw IPOs during the premarket: empirical analysis." Journal of Financial Economics 60 (1): 73–102.

Calvo, Antoni and Matthew Jackson. 2004. The effect of social networks on employment and inequality. American Economic Review 94 (3): 426-454.

Calvo, E., J. Lasaga, and A. van den Nouweland. 1999. "Values of games with probabilistic graphs." Mathematical Social Sciences 37 (1): 79-95.

Campbell, Carl M., and Kunal S. Kamlani. 1997. "The Reasons For Wage Rigidity: Evidence From A Survey Of Firms." Quarterly Journal of Economics 112 (3): 759-789.

Cao, J. and J. Lerner. 2009. "The Performance of Reverse-Leveraged Buyouts." In Journal of Financial Economics 91 (February), 139-157.

Carter, Richard B., and Steven Manaster. 1990. "Initial public offerings and underwriter reputation." Journal of Finance: 1045-1067.

Carter, Richard B., Frederick H. Dark, and Ajai K. Singh. 1998. "Underwriter reputation, initial returns, and the long-run performance of IPO stocks." The Journal of Finance 53 (1): 285-311.

Celikyurt, Ugur, Merih Sevilir, and Anil Shivdasani. 2010. "Going public to acquire? The acquisition motive in IPOs." Journal of Financial Economics 96 (3): 345-363.

Che, Yeon-Koo and Seung-Weon Yoo. 2001. "Optimal incentives for teams." American Economic Review 91 (3): 525-541.

Chemmanur, Thomas J., and Xuan Tian. 2011. "Do Anti-Takeover Provisions Spur Corporate Innovation?" AFA 2011 Denver Meetings Paper.

Chemmanur, Thomas J., Shan He, and Debarshi K. Nandy. 2010."The going-public decision and the product market." Review of Financial Studies 23 (5): 1855.

Cheng, Maria. 1999. "Hong Kong's silicon dreams." San Francisco Chronicle. October 12.

Chevalier, J. 1995a. "Do LBO Supermarkets Charge More? An Empirical Analysis of the Effects of LBOs on Supermarket Pricing." In Journal of Finance 50 (September), 1095-1112.

Chevalier, J. (1995b) Capital Structure and Product-Market Competition: Empirical Evidence from the Supermarket Industry. In American Economic Review 85 (June), 415-435.

Cohen, Lauren, Karl Diether, and Christopher Malloy. 2011. "Misvaluing Innovation."Unpublished working Paper, Harvard University.

Condorcet, Marquis de. 1785. "Essai surl'application de l'analyse a la probabilite des decisions rendues a la probabilite des voix." Paris: De l'imprimerie royale.

Davis, S., Haltiwanger, J., Jarmin, R., Lerner, J., and J. Miranda. 2009. "Private Equity, Jobs and Productivity." In Gurung, A. and Lerner, J. (eds.) Globalization of Alternative Investments Working Papers Volume 2: Global Economic Impact of Private Equity 2009, New York: World Economic Forum USA, 43-64. Available at http://www.weforum.org/pdf/cgi/pe/Fullreport.pdf.

Degeorge, Francois, and Richard Zeckhauser. 1993. "The reverse LBO decision and firm performance: Theory and evidence."\ Journal of Finance: 1323-1348.

Dunbar, Craig G. 1998. "The Choice between Firm-Commitment and Best-Efforts Offering Methods in IPOs: The Effect of Unsuccessful Offers." Journal of Financial Intermediation 7 (1): 60-90.

Dunbar, Craig G., and Stephen R. Foerster. 2008. "Second time lucky? Withdrawn IPOs that return to the market." Journal of Financial Economics 87 (3): 610–635.

Edelen, Roger M., and Gregory B. Kadlec. 2005. "Issuer surplus and the partial adjustment of IPO prices to public information." Journal of Financial Economics.

Edmans, Alex, Itay Goldstein, and Wei Jiang. 2011. "The real effects of financial markets: The impact of prices on takeovers." Journal of Finance, forthcoming.

European Commission. 2009. Commission Staff Working Document Accompanying the Proposal for a Directive of the European Parliament and of the Council on Alternative Investment Fund Managers and amending Directives 2004/39/EC and 2009/.../EC: Impact Assessment, COM(2009) 207/SEC (2009) 577. Brussels; European Commission.

Fang, Vivian W., Xuan Tian, and Sheri Tice. 2010. "Does stock liquidity enhance or impede firm innovation." Unpublished working paper, Tulane University.

Ferreira, Daniel, Gustavo Manso, and Andre C. Silva. 2010. "Incentives to innovate and the decision to go public or private." CEPR Discussion Paper No. DP7750.

Fulghieri, Paolo, and Merih Sevilir. 2009. "Organization and financing of innovation, and the choice between corporate and independent venture capital." Journal of Financial and Quantitative Analysis 44 (06): 1291-1321.

Gao, Huasheng, Michael L. Lemmon, and Kai Li. 2010. "A comparison of CEO pay in public and private US firms." Unpublished working paper, University of British Columbia.

Gould, Eric, Peter Pashigian, and Canice Prendergast. 2005. "Contracts, externalities, and incentives in shopping malls." Review of Economics and Statistics 87 (3): 411-422.

Goyal, Sanjeev and Jose Luis Moraga. 2001. "R & D networks." RAND Journal of Economics 32 (4): 686-707.

Granger, C. W. J. 1969. "Investigating causal relations by econometric models and cross-spectral methods." Econometrica 424-438.

Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey." Journal of Economic Literature 28: 1661-1707.

Grossman, Sanford J. and Oliver D. Hart. 1980. "Takeover bids, the free-rider problem, and the theory of the corporation." The Bell Journal of Economics 11 (1): 42-64.

Guo, S., Hotchkiss, E. and W. Song. 2009. "Do Buyouts (Still) Create Value?" In Journal of Finance, forthcoming. Available at http://ssrn.com/abstract=1108808.

176

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg. 2001. "The NBER patent citation data file: Lessons, insights and methodological tools." NBER Working Paper 8498.

———. 2005. "Market value and patent citations." RAND Journal of Economics 36 (1): 16–38.

Hall, Bronwyn H., and Josh Lerner. 2010. "The financing of R&D and innovation." In Handbook of the Economics of Innovation, ed. Bronwyn H. Hall and Nathan Rosenberg, 609-639. Elsevier-North Holland.

Hall, Bronwyn H., Zvi Griliches, and Jerry A. Hausman. 1986. "Patents and R&D: Is there a lag?" International Economic Review 27 (2): 265-283.

Hanley, Kathleen W. 1993. "The underpricing of initial public offerings and the partial adjustment phenomenon." Journal of Financial Economics 34 (2): 231-250.

Hartwig, J. 2009. "A panel Granger-causality test of endogenous vs. exogenous growth." KOF Working papers 09-231, KOF Swiss Economic Institute, ETH Zurich.

Himmelberg, C. , and B. Petersen. 1994. "R&D and internal finance: A panel study of small firms in high-tech industries." The Review of Economics and Statistics: 38-51.

Hochberg, Y., and J. Rauh. 2011. "Local Overweighting and Underperformance: Evidence from Limited Partner Private Equity Investments." Unpublished Working Paper. Available at SSRN: http://ssrn.com/abstract=1798747.

Honore, B., and E. Kyriazidou. 2000. "Panel Data Choice Models with Lagged Dependent Variables." Econometrica 68, 839-874.

Hsu, S., Reed, A. and J. Rocholl. 2010. "Competitive Effects of Private Equity Investments." Unpublished Working Paper. Available at SSRN: http://ssrn.com/abstract=1571207.

Ibbotson, Roger, and Jay R. Ritter. 1995. "Initial public offerings." in Robert A. Jarrow,Vojislav Maksimovic, and William T. Ziemba, eds.: Handbooks of Operations Research and Management Science (Amsterdam: North-Holland)

Imbens, G., and J. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." In Econometrica 62, 467-475.

177

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and estimation of local average treatment effects." Econometrica: Journal of the Econometric Society: 467-475.

Jackson, Matthew. 2004. "A survey of models of network formation: Stability and efficiency." In Group formation in economics: Networks, clubs and coalitions, ed. Gabrielle Demange and Myrna Wooders, 11-57. Cambridge: Cambridge Univ. Press.

Jain, Bharat A., and Omesh Kini. 1994. "The Post-Issue Operating Performance of IPO Firms." The Journal of Finance 49 (5): 1699 - 1726.

Jehiel, Philippe and Benny Moldovanu. 1996. "Strategic nonparticipation." RAND Journal of Economics 27 (1): 84-98.

Jehiel, P. and B. Moldovanu. 1999. "Resale markets and the assignment of property rights." The Review of Economic Studies 66 (4): 971-991.

Jehiel, P., B. Moldovanu, and E. Stachetti. 1996. "How (not) to sell nuclear weapons." American Economic Review 86 (4): 814-829.

Jeng, L., and P. Wells. 2000. "The Determinants of Venture Capital Funding: Evidence Across Countries." In Journal of Corporate Finance 6 (September), 241-289.

Jensen, M. 1986. "Agency Costs of Free Cash Flow, Corporate Finance and Takeovers." In American Economic Review Papers and Proceedings 76 (May), 323-329.

Jensen, M. 1989. "The Eclipse of the Public Corporation." In Harvard Business Review 67 (September/October), 61-74.

John, K., Lang, L. and J. Netter. 1992. "The Voluntary Restructuring of Large Firms in Response to Performance Decline." In Journal of Finance 47 (July), 891-917.

Kamien, M. and I. Zang. 1990. "The limits of monopolization through acquisition." Quarterly Journal of Economics 105 (2): 465-499.

Kamien, Morton, Shmuel Oren, and Yair Tauman. 1992. "Optimal licensing of cost-reducing innovation." Journal of Mathematical Economics 21 (5): 483-508.

Kaplan, S. 1989. "The Effects of Management Buyouts on Operating Performance and Value." In Journal of Financial Economics 24 (October), 217-254.

Kaplan, S. and A., Schoar. 2005. "Private Equity Performance: Returns, Persistence and Capital Flows." In Journal of Finance 60 (August), 1791-1823.

Kaplan, S. and Stein, J. (1993) "The Evolution of Buyout Pricing and Financial Structure in the 1980s." In Quarterly Journal of Economics 108 (May), 313-357.

Kaplan, Steven K., and Bernadette A. Minton. 2008. "How has CEO turnover changed." Unpublished working paper. University of Chicago.

Katz, Michael and Carl Shapiro. 1986a. "How to license intangible property." Quarterly Journal of Economics 101 (3): 567-589.

Katz, Michael and Carl Shapiro. 1986b. "Technology adoption in the presence of network externalities." Journal of Political Economy 94 (4): 822-841.

Kendall, Maurice. 1962. "Rank correlation methods." 3rd ed. New York: Hafner.

Kortum, S. and J. Lerner. 2000. "Assessing the Contribution of Venture Capital to Innovation." In RAND Journal of Economics 31 (Winter), 674-692.

Kosman, J. 2009. "The Buyout of America: How Private Equity Will Cause the Next Great Credit Crisis." New York: Penguin.

Krishna, Kala. 1993. "Auctions with endogenous valuations: The persistence of monopoly revisited." American Economic Review 83 (1): 147-160.

Lach, Saul, and Mark Schankerman. 1989. "Dynamics of R & D and Investment in the Scientific Sector." The Journal of Political Economy: 880-904.

Lai, Ronald, Alexander D'Amour, and Lee Fleming. 2009. "The careers and co-authorship networks of US patent-holders, since 1975." Unpublished Working Paper, Harvard University.

Lanjouw, Jean O., Ariel Pakes, and Jonathan Putnam. 1998. "How to count patents and value intellectual property: The uses of patent renewal and application data." The Journal of Industrial Economics 46 (4): 405–432.

Lerner, Josh, Morten Sorensen, and Per Stromberg. 2011. "Private Equity and Long-Run Investment: The Case of Innovation." The Journal of Finance 66 (2): 445-477.

Loughran, Tim, and Jay R. Ritter. 2004. "Why has IPO underpricing increased over time." Financial Management 33 (3): 5-37.

Lerner, Josh. 1994. "Venture capitalists and the decision to go public." Journal of Financial Economics 35 (3): 293–316.

Levin, Jonathan. 2002. "Multilateral contracting and the employment relationship." Quarterly Journal of Economics 117 (3): 1075-1103.

179

Lewis, Tracy. 1983. "Preemption, divestiture, and forward contracting in a market dominated by a single firm." American Economic Review 73 (5): 1092-1101.

Luehrman, T.A. 2007. "The Hertz Corporation (A)." Harvard Business School Case 208-030.

Michaely, Roni, and Michael R. Roberts. 2007. "Corporate dividend policies: Lessons from private firms." Unpublished working paper, Cornell University.

Mikkelson, Wayne H., Megan M. Partch, and Kshitij Shaw. 1997. "Ownership and Operating Decisions of Companies that Go Public." Journal of Financial Economics 44 (3): 281-307.

Milgrom, Paul and John Roberts. 1990. "Rationalizability, learning, and equilibrium in games with strategic complementarities." Econometrica 58 (6): 1255-1277.

Mollica, M. and Zingales, L. 2007. "The Impact of Venture Capital on Innovation and the Creation of New Business." Unpublished Working Paper, University of Chicago.

Mulkay, Benoit, Bronwyn H. Hall, and Jacques Mairesse. 2001. "Firm level investment and R&D in France and the United States: A comparison." Investing Today for the World of Tomorrow: Studies on the Investment Process in Europe, New York: 229-273.

Muscarella, C. and Vetsuypens, M. 1990. "Efficiency and Organizational Structure: A Study of Reverse LBOs." In Journal of Finance 45 (December), 1389-1413.

Mutuswami, Suresh and Eyal Winter. 2002. "Subscription mechanisms for network formation." Journal of Economic Theory 106 (2): 242-264.

Oxman, J. and Yildrim Y. 2008. "Governance Effects of LBO Events. Unpublished Working Paper, Syracuse University.

Pastor, Lubos, Lucian Taylor, and Pietro Veronesi. 2009. "Entrepreneurial learning, the IPO decision, and the post-IPO drop in firm profitability." Review of Financial Studies 22 (8): 3005.

Pagano, Marc, Fabio Panetta, and Luigi Zingales. 1998. "Why do firms go public? An empirical analysis." Journal of Finance 53: 27-64.

Page, Frank, Myrna Wooders, and Samir Kamat. 2005. "Networks and farsighted stability." Journal of Economic Theory 120 (2): 257-269.

Rasmusen, Eric, J. Mark Ramseyer, and John Wiley. 1991. "Naked exclusion." American Economic Review 81 (5): 1137-1145.

Rasmussen, P. 2008. Taming the Private Equity Fund "Locusts." In Europe Today 8 (Spring), 130-133.

Ritter, Jay, and Ivo Welch. 2002. "A review of IPO activity, pricing, and allocations." The Journal of Finance 57 (4): 1795-1828.

Rotemberg, Julio J., and Garth Saloner. 1994. "Benefits of narrow business strategies." The American Economic Review: 1330-1349.

Santos Silva, J.M.C., and Silvana Tenreyro. 2006. "The log of gravity." The Review of Economics and Statistics 88 (4): 641-658.

Saunders, Anthony, and Sascha Steffen. 2011. "The costs of being private: Evidence from the loan market." Review of Financial Studies, Forthcoming.

Segal, Ilya and Michael Whinston. 2000. "Naked exclusion: Comment." American Economic Review 90 (1): 296-309.

Segal, Ilya. 1999. "Contracting with externalities." Quarterly Journal of Economics 114 (2): 337-388.

Segal, Ilya. 2003. "Coordination and discrimination in contracting with externalities: Divide and conquer?" Journal of Economic Theory 113 (2): 147-327.

Seru, Amit. 2011. "Firm boundaries matter: Evidence from conglomerates and R&D activity." Journal of Financial Economics, forthcoming.

Service Employees International Union. 2007. "Behind the Buyouts: Inside the World of Private Equity." Washington, DC: SEIU.

Service Employees International Union. 2008. "Private Equity's Appetite for Infrastructure Could Put State and Local Taxpayers and Services at Risk." Draft policy discussion paper SEIU.

Sheen, Albert. 2009. "Do public and private firms behave differently? An examination of investment in the chemical industry." Unpublished working paper, UCLA.

Shleifer, Andrei, and Robert W. Vishny. 1989. "Management entrenchment: The case of manager-specific investments." Journal of Financial Economics 25 (1): 123-139.

———. 2003. "Stock market driven acquisitions." Journal of Financial Economics 70 (3): 295-311.

Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." Econometrica 65 (3): 557-586.

Stein, Jeremy C. 1989. "Efficient capital markets, inefficient firms: A model of myopic corporate behavior." The Quarterly Journal of Economics 104 (4): 655.

Stromberg, P. 2008. "The New Demography of Private Equity." In Gurung, A. and Lerner, J. (eds.) Globalization of Alternative Investments Working Papers Volume 1: Global Economic Impact of Private Equity 2008, New York: World Economic Forum USA, 3-26. Available at http://www.weforum.org/pdf/cgi/pe/FullReport.pdf.

Tian, Xuan, and Tracy Y. Wang. "Tolerance for Failure and Corporate Innovation." Review of Financial Studies, forthcoming.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe. 1997. "University versus corporate patents: A window on the basicness of invention." Economics of Innovation and New Technology 5 (1): 19–50.

Trajtenberg, Manuel. 1990. "A penny for your quotes: patent citations and the value of innovations." The Rand Journal of Economics: 172–187.

Weild, David, and Edward Kim. 2009. "A wake-up call for America." Grant Thornton: Capital Markets Series.

Welch, Ivo. 1992. "Sequential sales, learning, and cascades." Journal of Finance: 695-732.

Winter, Eyal. 2004. "Incentives and discrimination." American Economic Review 94 (3): 764-773.

Winter, Eyal. 2006. "Optimal incentives for sequential production." RAND Journal of Economics 37 (2): 376-390.

Wooldridge, Jeffery M. 1997. "Quasi-likelihood methods for count data." Handbook of applied econometrics 2: 352-406.

Young, H. Peyton. 1988. "Condorcet's theory of voting." The American Political Science Review 82 (4): 1231-1244.