Bubbles and the Value of Innovation[†]

Valentin Haddad
UCLA and NBER

Paul Ho
Federal Reserve Bank of Richmond

Erik Loualiche
University of Minnesota

February 28, 2022

Abstract

Booming innovation often coincides with intense speculation in financial markets. Using over a million patents, we document two ways the market valuation of innovation and its economic impact become disconnected during bubbles. Specifically, an innovation raises the stock price of its creator by 40% more than is justified by future outcomes. In contrast, competitors' stock prices move little despite their profits suffering. We develop a theory of investor disagreement about which firms will succeed that reconciles both the facts, unlike existing models of bubbles. Optimal innovation policy during bubbles must account for the disconnect.

[†]Haddad: valentin.haddad@anderson.ucla.edu; UCLA Anderson School of Management, 110 Westwood Plaza, Los Angeles CA 90095; phone: 310-825-1814. Ho: paul.ho@rich.frb.org. Loualiche: eloualic@umn.edu. We thank Andrew Atkeson, Markus Brunnermeier, Edouard Challe, Alexandre Corhay, Emmanuel Farhi, Nicolae Garleanu, Harrison Hong, Pablo Kurlat, Deborah Lucas, Alberto Martin, Stavros Panageas, Gill Segal, Alp Simsek, Avanidhar Subrahmanyam, Jaume Ventura, and seminar participants at Columbia University, Federal Reserve Bank of Minneapolis, Federal Reserve Bank of Richmond, HKUST, London Business School, Minnesota, MIT Sloan, Princeton, UBC Sauder, UCLA, University of Amsterdam, University of Luxembourg, Wharton, Banque de France, Bank of England, AFA, Finance in the Cloud, FIRS, Macro Finance Society meetings, NBER Asset Pricing, SED, and WFA for useful comments. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

1 Introduction

Measuring the impact of innovation is important for understanding and quantifying economic growth, firm dynamics, and numerous other questions in economics and finance. A growing literature advocates for asset prices as a real-time measure comparable across industries and time. Stock price reactions to the announcement of an innovation reveal the present value of the impact to the innovator and spillovers to its competitors (Kogan, Papanikolaou, Seru, and Stoffman (2017); Bloom, Schankerman, and Van Reenen (2013)). This paper challenges this approach empirically and theoretically: intense speculation on financial markets, a common occurrence in innovation booms (e.g., Scheinkman, 2014), disconnects the responses of stock prices and real outcomes to an innovation.

We establish the disconnect empirically and propose a theory that explains its structure. Two main facts emerge in the data during bubbles. First, the impact of an innovation on the stock price of its creator increases by 40% relative to the real outcomes it will generate. Next, even though innovation in a firm damages profits of its competitors, these negative spillovers have no impact on the stock prices of the competing firms. Existing theories do not jointly match these facts. We thus propose a theory of disagreement among investors *about which firms will succeed* that resolves this tension. In particular, to capture the behavior of competitive spillovers, it is key that the investors who are optimistic about a firm differ from those who are optimistic about its competitors. Consequently, innovation policy during bubbles must not only account for the fact that inflated prices stimulate more innovation, but also that the evaluation of externalities is distorted.

The existing literature argues that asset prices *always* reflect the economic impact of innovation, but only establishes this result *on average*. Our empirical analysis pushes back against this view: the relation is systematically altered during episodes of bubbles. This rejection is meaningful because speculative episodes often coincide with large volumes of innovation. Operationally, we define bubbles as industry-years that have experienced a sharp run-up in stock prices, following Greenwood, Shleifer, and You (2018), and study how market-based measures of the value of innovation change during bubbles. Although this measure is likely imperfect—the mere existence of bubbles is debated—the imperfection biases our estimates against finding a disconnect.

The first disconnect that we document arises in the direct value of innovation to the innovator herself. The seminal work of Kogan et al. (2017) shows that the change in stock price of the innovator in the few days after a patent is approved reflects the eventual economic impact of the innovation, measured, for example, using sales or profits. This relation breaks down during bubbles: the market value of innovation increases by 30% at the patent level and between 40% and 50% at the firm level despite the absence of corresponding cash-flow improvements. These market value increases are greater for patents that receive more citations even though the relationship between cash flow and citations remains unchanged during bubbles.

We document an additional disconnect when measuring the impact of innovation on competing firms. Whereas speculation inflates asset-price measures of the direct value of innovation, asset-price measures of spillovers are muted during bubbles. Bloom, Schankerman, and Van Reenen (2013) establish that both market value and real outcomes of a firm are hurt by competitors' innovation on average. In contrast, we find that there is no negative effect on the stock price during bubbles even though the negative impact on real outcomes remains unchanged.

We need a theory that rationalizes these facts in order to understand the economics of the disconnect and its consequences. For example, how should innovation policy respond to the presence of speculation? While it is well understood that speculation can create abnormal patterns in asset prices, existing theories of bubbles struggle to account for the empirical results. If speculation simply adds a random component to prices—as with rational bubbles or unanchored beliefs—we would see neither the overreaction of the innovators' stock prices nor the underreaction of competitors' stock prices. Alternatively, if investors are overly optimistic about the innovative sector, then valuations react more strongly to announcement of innovation, lining up with the direct effect but not the competitive spillovers we find. Disagreement about the general prospects of the innovative sector lead to the same issue because the optimistic views dominate for prices.

In our theory, investors speculate due to disagreement about which firms will succeed. For example, investors disagree on whether Twitter or Facebook is going to be the most successful social network, even though they agree on the total value of the social network industry. Similar to aggregate disagreement, the investors in each firm are optimistic about its prospects. This leads to the overreaction of a firm's stock price to its innovation. However, a new feature of this form of disagreement is that different investors are optimistic about different firms. This leads to specialization in their portfolios—an effect we document in the data—resulting in the price of different firms being pinned down by different investors. This creates an asymmetry in how interactions among firms are perceived. While the owner of an innovation believes it will displace its competitors, the owners of competitors do not. They believe their own firms are better and do not perceive innovation by others as a threat. Due to this asymmetry, the model generates underreaction of stock prices to competitive spillovers (priced by owners of competitors) while maintaining overreaction in the direct effects (priced by owners of the innovator). We obtain simple characterizations of these results thanks to a new parametrization of beliefs that overcomes the complexity of entertaining disagreement about a large number of firms.

Finally, our model also shows how to interpret asset prices despite the disconnect we find during bubbles. For example, a planner trying to decide on optimal innovation must not only lean against inflated prices to prevent excessive innovation, but also realize that the market assessment of externalities is distorted. The distortion of externalities is an expression of our new mechanism: the beliefs of those receiving the externality differ from those originating it.

Related Literature. The idea that asset prices are useful to evaluate the economic value of innovation goes back to Pakes (1985) and Griliches, Pakes, and Hall (1986). Notable implementations of this idea include Harhoff et al. (1999), Hall, Jaffe, and Trajtenberg (2005), and Nicholas (2008). In more recent work, Kogan et al. (2017) use the stock price reaction to a patent approval to measure the value of an innovation to its creator. Bloom, Schankerman, and Van Reenen (2013) use market values to measure the spillovers of innovations to other firms. We challenge this view and show that bubbles disconnect asset prices from the value of innovation. The impact of bubbles on the process of innovation can also be found in other dimensions. Hombert and Matray (2020) document distortions in labor market outcomes; Dong, Hirshleifer, and Teoh (2020) in the quantity and ambitiousness of innovation at the firm level.[1]

Our theoretical explanation of the disconnect builds on the work studying the effect of belief disagreement on asset prices as in Miller (1977), Harrison and Kreps (1978), and Scheinkman and Xiong (2003).[2] More recent contributions include Barberis, Greenwood, Jin, and Shleifer (2018), who emphasize the role of extrapolation, Chinco (2020), who incorporates social interactions, and Hirshleifer and Plotkin (2020), who present an evolutionary theory of cultural traits leading to investment booms. Unlike the existing literature, disagreement in our model occurs at the firm level, which leads to new predictions about how markets value firm interactions.

There are other approaches to studying the relation between financial markets and innovation, building on classic models of technological growth (e.g., Acemoglu, 2008). One strand of the literature centers on uncertainty and risk compensation rather than speculation. Pastor and Veronesi (2005, 2009) focus on learning during innovation booms, Gârleanu, Kogan, and Panageas (2012) and Kogan, Papanikolaou, and Stoffman (2020) highlight displacement risk, Kung and Schmid (2015) study the long-run risk associated with innovation, and Corhay, Kung, and Schmid (2015) and Loualiche (2020) measure risks from competition across firms. Another stream of work focuses on rational bubbles (Samuelson (1958); Tirole (1985)) and whether they crowd in or crowd out investment. Recent examples include Farhi and Tirole (2011) and Martin and Ventura (2012). We argue that neither of these approaches can reconcile the two facets of the disconnect.

Finally, we contribute to the literature discussing optimal policy in the presence of disagreement. We follow the welfare criterion of Brunnermeier, Simsek, and Xiong (2014), who circumvent the need to choose a specific belief by studying allocations that are efficient across all convex combinations of agents' beliefs. Davila (2020) also follows this approach in the context of financial-transaction taxes. Kondor and Köszegi (2017) consider the regulation of sophisticated financial products. Dávila and

[1] van Binsbergen and Opp (2019) quantify the effects of mispricing more generally.

[2] Brunnermeier and Oehmke (2013) and Scheinkman (2014) survey the literature on bubbles more broadly. Janeway (2012) offers a first-person account of the interaction of bubbles and innovation.

Walther (2021) study prudential policy with distorted beliefs. Our model sheds light on optimal innovation policy with speculation.

The paper proceeds as follows. In Section 2, we present the main stylized facts about the value of innovation during a bubble. We introduce our model of speculation with business-stealing in Section 3 and show in Section 4 that the model predictions match the data. Section 5 studies the planner problem and Section 6 presents further predictions and supporting evidence for the model. Finally, Section 7 concludes.

2 Disconnect Between Market and Outcome Based Value of Innovation During Bubbles

While the nature of patents and the style of innovations change over time, the market value of an innovation should consistently reflect its economic impact. A growing literature exploits this property and shows how to use asset prices to measure the value of innovation (e.g. Kogan et al. (2017)).

In this section, we document that this idea breaks down during bubbles. Specifically, we show two ways in which the market value of innovation gets disconnected from its economic impact. First, the impact of an innovation on the stock price of its creator increases sharply relative to the real outcomes it will generate. Second, even though innovation in a firm damages the profits of its competitors, these negative spillovers have no impact on the stock prices of the competing firms. We will see that these two facts combined are particularly informative to separate theories of bubbles.

2.1 Methodology

We present our data and empirical methodology. Table IA.1 provides summary statistics.

Bubbles. We identify episodes of bubbles using the empirical definition of Greenwood, Shleifer, and You (2018). We split firms into 49 industries following the classification of Fama and French (1997). An industry-month is defined as being in a bubble if it satisfies three conditions simultaneously. First, the value-weighted portfolio of the corresponding industry experienced a return of 100% or more over the previous two years. Second, this industry's value-weighted return also exceeds the return of the market by at least 100% over the past two years. Third, the industry's value-weighted return over the past five years is larger than 50%. We aggregate to the industry-year level—the coarseness of the remainder of our data—by considering an industry-year in a bubble if the industry is in a bubble for at least a month in the year. In practice, multiple months are always in a bubble year due to the persistence of the bubble criteria. This approach identifies 74 industry-years in a bubble between 1962 and 2017.

As pointed out by Greenwood, Shleifer, and You (2018), this classification is not perfect: it is likely to miss some bubbles but also classify some actual productive booms as bubbles. Our interpretation of the evidence relies on the selected episodes being more likely to be bubbles than a typical industry-year. In line with this assumption, Greenwood, Shleifer, and You (2018) show that these episodes exhibit a constellation of characteristics consistent with speculative bubbles. In particular, in line with our focus on innovation, they find that "price run-ups ... involving younger firms [and] having higher relative returns among the younger firms ... [are] more likely to crash." Foreshadowing our theory, these industries with young and innovative firms are likely to have scarce information, making them more susceptible to the disagreement that we model.[3]

We also confirm the relevance of the bubble classification for innovation empirically. In Internet Appendix Table IA.2, we show that there are 1.4 more patents issued within a USPTO technology class during bubbles, which translates to a 15% increase in the number of patents created in a patent class during a bubble.

The value of an innovation. We use the measures of private value of innovation from Kogan et al. (2017).[4] Their dataset combines stock market and patent data for U.S. firms for the period from 1926 through 2010. They measure the stock market response in the three-day window after a firm has issued a new patent, controlling for the return on the market portfolio during that period. Kogan et al. (2017) also aggregate the stock market value from all the patents of a given firm every year to measure the value of innovation at the firm level. To account for the ultimate quality of each patent, Kogan et al. (2017) use the forward-looking number of citations generated by a patent for their patent-level analysis and the number of citations generated by all the patents produced by a firm in a given year for their firm-level analysis. Finally, we follow their measure of the impact of a patent on firm outcomes by studying the response of sales, profits, and productivity at horizons up to five years.

Spillovers to competitors. We also measure how a firm suffers when its competitors innovate. To that end, we follow the methodology of Bloom, Schankerman, and Van Reenen (2013), who regress a firm-level outcome—log market value or log future sales—on the quantity of innovation by groups of "neighboring" firms. To identify competitive spillovers separately from knowledge spillovers, in which the firm learns from an innovator, they distinguish between firms in neighboring industries that compete against each other and firms issuing patents in the same USPTO technology class that learn from each other.[5] Specifically, they construct distances between firms in

[3]Other characteristics could favor creativity, such as the culture of the CEO as demonstrated in He and Hirshleifer (2020).

[4]We thank Dimitris Papanikolaou for graciously sharing his data with us.

[5]We discuss the implications of our theory for knowledge spillovers in Internet Appendix D.3.

each of these two spaces and, for each firm-year, compute distance-weighted stocks of innovative capital from all other firms. The resulting firm-level exposures for competitors and related innovators are *spillsic* and *spilltech*, respectively, and regressions on these quantities measure competitive and knowledge spillovers. They use three different measures of distance: Jaffe, Mahalanobis and an instrument for the Jaffe distance. We describe these measures in Internet Appendix A. When the dependent variable is log market value, the regression identifies the market-based spillover.[6] When focusing on log sales, the coefficients identify the outcome-based spillover.

2.2 The Value of Innovation During Bubbles

The market value of innovation increases in bubbles. To study how the value of innovation changes during bubbles, we consider the following regression specification:

$$\log \xi_{j,t} = \beta B_{j,t} + \gamma Z_{j,t} + \varepsilon_{j,t}, \tag{1}$$

where $\xi_{j,t}$ is the market value of patent j issued during year t using the market-based measure from Kogan et al. (2017). The variable $B_{j,t}$ is an indicator for whether the firm issuing patent j was in an industry experiencing a bubble during year t. As in Kogan et al. (2017), the controls $Z_{j,t}$ include the logarithm of the number of citations for the patent as well as firm and year fixed effects. The firm fixed effects ensure that our results are not driven by the type of firms or industries that go through bubbles. We run a similar regression at the firm level, replacing the value of innovation and citation variables with their firm-level analogues.

Table 1 shows that the private value of innovation increases sharply during bubbles. In a bubble, the private value of innovation is 30% higher at the patent level (columns 1 and 2) and 40% to 50% higher at the firm level (columns 4 and 5). These effects are both economically and statistically significant. Broadly inflated asset prices across the industry during a bubble need not translate into high valuations of patents specifically. In particular, our results reject the common view that prices are independent of fundamentals during bubbles. Under this view, valuations could be high overall but would not be responsive to news of patent issuance.

In columns 3 and 6, we provide further evidence that prices do react to fundamental news during bubbles, albeit to a degree that is unjustified by ex-post real outcomes, by considering the effect of patent quality on its valuation. In particular, we show that the response of market value to citations is even stronger during bubbles. Rather than removing any link between asset prices and innovation, bubbles exacerbate how markets value innovations. Finally, because volatility is likely to be higher during bubbles, we consider specifications controlling for firm-level and

[6] Taking the log ensures that we are considering a semi-elasticity, the counterpart to the spillover in our model.

Table 1
Market Value of Innovation in Bubbles

	Patent Level			Firm Level		
	(1)	(2)	(3)	(4)	(5)	(6)
Bubble	0.317***	0.315***	0.186	0.514***	0.427***	0.265**
	(0.094)	(0.094)	(0.132)	(0.114)	(0.123)	(0.134)
Log Citations		0.016***	0.014***		0.823***	0.821***
(forward looking)		(0.004)	(0.005)		(0.011)	(0.011)
Log Citations x Bubble			0.054**			0.061***
			(0.025)			(0.012)
Fixed Effects	Y, F	Y, F	Y, F	Y, F	Y, F	Y, F
Observations	1,171,806	1,171,806	1,171,806	47,886	47,886	47,886
R^2	0.68	0.68	0.68	0.89	0.94	0.94

Note: Table 1 presents panel regressions of the value of innovation, as measured in Kogan et al. (2017) at the patent and firm levels, on a dummy from Greenwood, Shleifer, and You (2018) that captures whether the firm is in an industry that is in a bubble or not. We include fixed effects for firm F and patent grant year Y. Standard errors clustered at the year level are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

industry-level volatility in Internet Appendix Table IA.3. This analysis confirms that our results are not driven by changes in volatility.

The effect of innovation on real outcomes during bubbles. The increased valuation of innovation during bubbles is not justified by the effect on real outcomes. We consider the response of real firm outcomes to innovation using a counterpart to regression (1) where we replace the market value on the left-hand side with profit, sales, and productivity. Table 2 reports the results for profits; we obtain similar results for sales and productivity in Internet Appendix Table IA.4. We focus on firm-level specifications since real outcomes are observed at low frequency. Under the view that higher prices reflect improved fundamentals, the effect of innovation on prices corresponds to a higher present value of future cash flows. Therefore we measure not only the instantaneous response of quantities but also the response at longer horizons, up to five years.

Innovation does affect real outcomes. Consistent with the literature, we find that more cited patents go along with larger profits, sales, and productivity at horizons up to five years. However, the relation between citations and cash flows does not change during bubbles. The coefficient on the interaction of citations and the bubble indicator is tightly estimated around zero at horizon up to three years and remains economically small and statistically insignificant in all specifications. In addition, we

Table 2
Firm Profits at Different Horizons Following Bubbles

	Horizon (years)				
	1	2	3	4	5
Citations	0.014***	0.021***	0.027***	0.033***	0.042***
	(0.003)	(0.004)	(0.005)	(0.006)	(0.007)
Bubble	0.032	−0.002	−0.023	−0.034	−0.037
	(0.025)	(0.035)	(0.035)	(0.042)	(0.042)
Citations x Bubble	0.001	−0.002	0.009	0.008	−0.022
	(0.012)	(0.011)	(0.014)	(0.018)	(0.020)
Fixed Effects	Y, I	Y, I	Y, I	Y, I	Y, I
Observations	121,017	109,523	99,248	90,123	81,895
R^2	0.09	0.10	0.11	0.12	0.13

Note: Table 2 presents panel regressions of future profits for horizons of one to five years on a bubble dummy and a citation measure of patent value as in Kogan et al. (2017). We use industry (I) and year (Y) fixed effects. Standard errors clustered at the firm-year level are presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1%, respectively.

do not observe a level effect during bubbles: none of the five coefficients is significantly different from zero.

Combining the results on market values and real outcomes, we obtain our first disconnect. In times of bubbles the market value of innovation overstates its ultimate impact on sales and profits. The disconnect is substantial: the 40% increase in market value goes along with virtually no improvement in real outcomes.

2.3 The Spillovers of Innovation on Competitors During Bubbles

We now turn to the second disconnect: while innovation has negative spillovers on competitors, markets ignore these spillovers during bubbles. To establish this result empirically, we enrich the specification of Bloom, Schankerman, and Van Reenen (2013) to estimate spillovers conditional on a bubble:

$$\begin{aligned}\log X_{i,t} = \; & \beta \left(B_{i,t} \times \log spillsic_{i,t}\right) + \gamma_1 \log spillsic_{i,t} \\ & + \gamma_2 \log spilltech_{i,t} + \gamma_3 B_{i,t} + \delta Z_{i,t} + \varepsilon_{i,t},\end{aligned} \tag{2}$$

where $X_{i,t}$ is either the market value (Tobin's q) or output (sales normalized by an industry price index) of firm i in year t. As before, $B_{i,t}$ is an indicator of whether firm i is in an industry that experienced a bubble in year t. The controls $Z_{i,t}$ include firm and year fixed effects. We measure market-based spillovers by taking $X_{i,t}$ to be the market value of the firm and measure outcome-based spillovers by taking $X_{i,t}$ to be firm output. The main coefficient of interest is β, which measures how different competitive spillovers are during bubbles.

Table 3 reports the results. Columns 1 to 3 consider the effect of innovation by competitors on market value for the three different approaches to measuring spillovers. The significant and negative coefficient on *spillsic* indicates that outside of bubbles, competitive spillovers have a negative effect on valuations, consistent with the results in Bloom, Schankerman, and Van Reenen (2013). The coefficient on the interaction of *spillsic* with bubbles is positive, indicating a reduction in the effect of competitive spillovers. The estimated competitive spillover in a bubble (the sum of coefficients on *spillsic* and the interaction term) is positive for the Jaffe and Mahalanobis spillover measures but negative when we instrument for the spillover. In all cases this represents a considerable reduction in how the market assesses competitive spillovers. When a firm innovates during a bubble, the market value of its competitors does not change much.

We contrast our results on firm value to the spillovers on sales in columns 4 to 6. We find a negative coefficient on *spillsic*, statistically significant in all but one specification. Outside of bubbles, firm fundamentals react negatively to innovations by competitors. Unlike the results on market valuations, the presence of a bubble does not have any effect on the spillovers to real quantities. The coefficient on the interaction of bubbles and spillovers is statistically insignificant (columns 4 and 5) and small relative to the main coefficient (columns 4, 5, and 6).[7]

For both the direct effect and spillovers, we observe a divorce between asset prices and real quantities. This challenges the notion that the two approaches convey the same information about the value of innovation. Further, prices are not just uniformly inflated in bubbles; we uncover a tension in the behavior of asset prices: the price of the innovator overreacts while competitors underreact. On the one hand, bubbles lead to higher direct market value of innovations in a discerning way: more impactful patents receive even higher valuations during bubbles. On the other hand, market values mostly ignore negative spillovers to competitors. In what follows, we present a theory that ties these facts together and argue that the facts cannot be explained by alternative theories of bubbles in the literature.

3 A Model of Disagreement and Innovation

We argue that the critical element to understanding the disconnect is that investors disagree about which specific firms will succeed. For example, investors disagree not about the total value of social networks but rather the value of Facebook versus Twit-

[7] In Internet Appendix Table IA.5, we confirm the robustness of our results with a different approach to measuring spillovers based on the unnormalized covariance distance metrics.

Table 3
Competitive Spillovers in Times of Bubbles

	Market-based Spillovers			Outcome-based Spillovers		
	Jaffe	Mahalanobis	IV Jaffe	Jaffe	Mahalanobis	IV Jaffe
	(1)	(2)	(3)	(4)	(5)	(6)
Bubble x Spill-SIC	0.152***	0.200***	0.178***	0.004	−0.000	0.007***
	(0.027)	(0.037)	(0.038)	(0.009)	(0.013)	(0.002)
Spill-SIC	−0.088***	−0.103***	−0.314***	−0.021***	−0.021**	−0.044
	(0.016)	(0.033)	(0.104)	(0.006)	(0.010)	(0.046)
Spill-Tech	0.405***	0.844***	1.214***	0.175***	0.159***	0.188**
	(0.145)	(0.174)	(0.171)	(0.025)	(0.040)	(0.074)
Fixed Effects	Y, F	Y, F	Y, F	Y, F	Y, F	Y, F
Observations	8,896	8,946	8,896	8,775	8,825	8,775
R^2	0.74	0.74	0.73	0.99	0.99	0.99

Note: Table 3 presents panel regressions of firm value (log of sales or Tobin's q) on a measure of competition from Bloom, Schankerman, and Van Reenen (2013) interacted with a bubble dummy, measured as in Greenwood, Shleifer, and You (2018) that captures whether the firm is an industry that is in a bubble state or not. We control for the technological spillover measure that corresponds to the public firms that issue patent in similar technological space. We follow the specification from Table of 3 and 5 of Bloom, Schankerman, and Van Reenen (2013) and include the same controls. We use firm F and year Y fixed effects. Standard errors clustered at the year level are presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1%, respectively.

ter. We present a model combining investor disagreement with a standard framework of firm competition to this effect. Like the existing literature, bubbles stem from investor disagreement (Miller, 1977; Harrison and Kreps, 1978; Scheinkman and Xiong, 2003), but we show that our new type of disagreement is crucial to rationalize the empirical results of Section 2.

3.1 Setup

At date 0, households work to create blueprints or new ideas. They sell the blueprints to firm creators, who implement them as firms. On financial markets, households who disagree with each other speculate on the claims to these firms' production. At date 1, firms compete and produce; households receive the payoffs from their positions in firms and consume. Appendix Figure A.1 illustrates the overall structure of the model; the proofs are in the appendix. We now detail each of these steps.

Firms. A continuum of firms, indexed by i and with total mass M_e, is created in equilibrium at date 0. Formally, these firms constitute the whole economy, but one can interpret the model as representing one sector of an economy.[8] At date 1, firms enter the production stage, and their productivity a_i is revealed. To capture competition across these firms, we assume that only a fixed mass M of the most productive firms is able to produce. M_e is an endogenous outcome, while M is a fixed parameter.[9] Given the cumulative distribution function of productivities in the population F, only firms above a cutoff $\underline{a}$ are able to produce, with

$$\underline{a} := F^{-1}\left(1 - \frac{M}{M_e}\right). \tag{3}$$

The profits of a firm with productivity a are then given by

$$\pi(a) = a^{\eta} \cdot \mathbf{1}\left\{a \geq \underline{a}\right\}, \tag{4}$$

where η determines how differences in productivities translate into differences in generated profits, and the indicator function captures whether the firm produces or not.[10] We concentrate on situations in which $M < M_e$, so that fewer firms produce than are created; the case of $M = M_e$ is straightforward. This allocation of production slots is a simple way to generate the competitive spillovers we observe in the data. Indeed, firms do not internalize that they might take over the slot of another firm. Internet Appendix B.1 provides foundations for this competitive structure. Section 4.2.2 confirms the robustness of our conclusions to alternative modelling of competition.

Households. Households play two roles in the model. They work to create the blueprints for new firms and speculate on which of the new firms will succeed. Formally, there is a unit mass of households indexed by j. At date 0, household j is endowed with a fixed unit of consumption good c_0 and its share of firm creators, which we describe in the next section. In addition, each household decides how many blueprints to supply, b_j. Blueprints are produced at a convex cost $W(b_j) = f_e b_j^{\theta+1} M^{-\theta}/(\theta+1)$, where the parameter θ is the elasticity of supply of blueprints and f_e controls the level of production costs.[11] Finally, each household also decides on the number of shares to invest in each firm on the financial market, $\{s_i^j\}_i$. Households have heterogeneous beliefs about the distribution of productivity a_i for each firm i, which we describe in detail below. We assume that they can only take long positions in claims to firms.[12]

[8] Our results would be unchanged as long as firms do not interact across sectors because we will assume the presence of a quasilinear good.

[9] Internet Appendix D.4 considers an extension with a varying number of active firms.

[10] The fact that the marginal active firm collects positive profits improves tractability but is not crucial to our conclusions. We show in Internet Appendix D.2 that our results also hold in a variant of the model where the marginal firm earns zero profits.

[11] We introduce the total mass of firms M in the blueprint production function to simplify the algebra. None of our results rely on this assumption. This assumption would only matter if one were to do comparative statics with respect to the parameter M.

[12] Such a hard constraint facilitates the analysis, but the important assumption is some limit or cost to the ability to take short positions. The importance of limits to short-selling in models

Households behave competitively and take prices as given. Hence, household j solves the problem

$$\max_{c_0, s_i^j \geq 0, b_j} \; c_0 + \mathbf{E}^j \left\{ \int s_i^j \pi_i di \right\} - W(b_j) \tag{5}$$

$$\text{s.t.} \;\; c_0 + \int s_i^j \; p_i \; di \leq 1 + p_b b_j + \Pi, \tag{6}$$

where p_i and p_b are the respective prices of firm i and blueprints, $\mathbf{E}^j$ is household j's expectation, and Π denotes firm creators' aggregate profits.[13]

Firm Creators. Finally, firm creators connect the steps of innovation and trading in financial markets. They pay households to create blueprints, and issue claims to the corresponding firms on the financial market. Formally, there is a continuum of short-lived firm creators. At date 0, each firm creator can use a unit blueprint to create a new firm, which is then sold on competitive financial markets. They participate in competitive markets for blueprints and firms, taking their respective prices p_b and p_i as given.[14] The firm creator's problem at time $t = 0$ is therefore:

$$\max_{c \in \{0,1\}} c \cdot (p_i - p_b). \tag{7}$$

Equilibrium Conditions. The competitive equilibrium of the economy is defined as follows. Firm creators maximize profits from selling their firms, taking prices as given. Households maximize their perceived expected utility by choosing their optimal blueprint discovery effort and an optimal portfolio allocation, taking the prices of blueprints and firms as given. Firms maximize profits given their production status. Finally, the markets for blueprints and claims to firms' profits (the stock market) clear:

$$\int b_j dj = M_e, \tag{8}$$

$$\forall i \in [0, M_e], \quad \int s_i^j dj = 1. \tag{9}$$

of disagreement has been well-known since the work of Miller (1977), Harrison and Kreps (1978), or more recently Scheinkman and Xiong (2003). D'avolio (2002) and Jones and Lamont (2002) document empirical evidence of such short-sale constraints, with the former showing that these constraints increase with disagreement. A theoretical justification is provided by Duffie, Garleanu, and Pedersen (2002), who microfound short-sale costs using a search model.

[13]The assumption of risk neutrality does not play any role in the analysis. Formally, all our results would be identical if the objective of households were: $c_0 + \mathcal{U}^{-1}\left(\mathbf{E}^j \left\{\mathcal{U}(\int s_i^j \pi_i di)\right\}\right) - W(b_j)$, where $\mathcal{U}(.)$ is an increasing and concave function.

[14]We assume firm creators do not have any information about the firms they create.

In a symmetric equilibrium, this leads to the conditions:

$$W'(M_e) = p_b = p_i. \tag{10}$$

The marginal effort cost to create a blueprint is equal to its price — household effort optimization. The price of a blueprint is equal to the price at which a firm trades on financial markets — firm creator optimization. Finally, we need to determine the willingness of households to buy claims to firms, which depends on their beliefs.

3.2 Beliefs

The critical element to match our empirical findings is our model of beliefs. In particular, we assume that households disagree about which firms will be successful. Such disagreement arises naturally in innovative episodes: households must rely on their priors to evaluate new firms (or ideas) and, in the absence of data on these firms, agree to disagree.[15] In general, it is challenging to keep track of heterogeneous beliefs across many firms. We overcome this issue by imposing structure on the distribution of beliefs. Beliefs are governed by two exogenous parameters: the actual population distribution F and a scalar n for the intensity of disagreement. We assume that even though agents disagree about each individual firm, they agree on the population distribution of firm productivity.[16] We also assume that the population distribution F follows a Pareto distribution: $F(a) = 1 - a^{-\gamma}$ for $a \geq 1$.

A simple narrative for our specification of individual beliefs is as follows. Each household organizes firms into a continuum of packets containing n firms each and believes it knows the exact ranking of productivity draws within each packet. We assume that the composition of packets and the order of firms within packets is drawn in an i.i.d. equiprobable fashion across agents and firms, and that each firm can only be in one packet.[17] The parameter n controls the intensity of disagreement. When $n = 1$, households consider all firms to be the same, with their productivity drawn from F. As n increases, households have views on the comparison of more firms and thus have a stronger prior that the best firm in each packet will have high productivity.

[15] Morris (1995) discusses arguments for heterogeneous priors, many of which are particularly applicable in our setting.

[16] A motivation for focusing on disagreement across firms rather than about the aggregate is the empirical observation that high firm entry in a sector often follows disruptive innovation, either through a large technological change or the introduction of new products. New firms then conduct micro innovations to take advantage of a macro innovation (Mokyr, 1992). Empirically, these episodes appear particularly relevant to economic growth, as discussed for instance in Abernathy (1978) and Freeman (1982).

[17] Formally, there is no choice of packets for investors. This narrative is equivalent to assuming that each investor has beliefs from the cumulative distribution of the first order statistic from F for a sample of size n for a fraction $1/n$ of firms, the second order statistic for a separate fraction $1/n$, the third order statistic... The composition of these groups of size M_e/n is independent across investors.

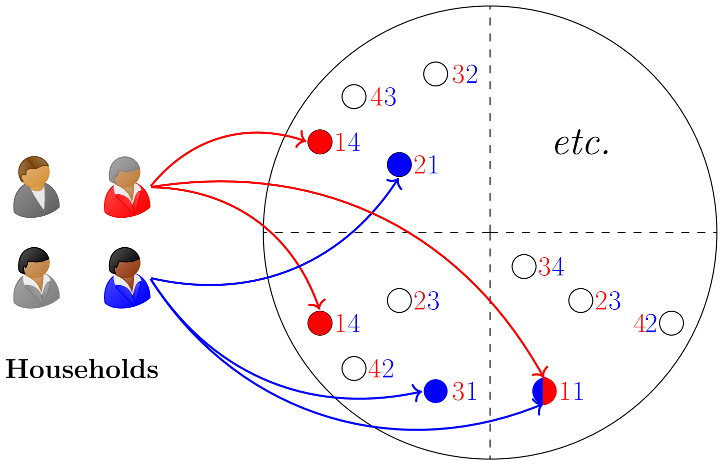

Figure 1
Households' beliefs and choice of firms to invest in.
The figure illustrates the case of $n = 4$. For this figure, the composition of each packet is identical across agents for exposition purposes. This choice does not affect the model results.

In equilibrium, household j only invests in the subset of firms that it considers to be the most productive in its packets. These firms are perceived by household j to have productivity drawn from F^n, the distribution of the maximum of n independent draws from F. Since households rank firms differently, they have different beliefs about the productivity distribution of any given firm and invest in different sets of firms, as illustrated in Figure 1. Each household believes that the firms it invests in are, in expectation, more productive than the average firm in the economy. As a result, the price of each firm is the expected value of its profits under the distribution F^n:

$$W'(M_e) = V^{(n)}(M_e) = \int_{\underline{a}}^{\infty} \pi(a) dF^n(a) = \int_{F^{-1}(1-M/M_e)}^{\infty} a^\eta dF^n(a). \tag{11}$$

Equation (11) pins down the quantity of firms created M_e and their price $p_i = V^{(n)}(M_e)$. While the clean expressions in equation (11) and Section 4 highlight the tractability arising from the symmetry in our model, our main insights only require that different investors are optimistic about about different firms. For example, investors could herd and be optimistic about similar firms (correlated rankings) as long as they are not all optimistic about the same firms. In that view, n proxies for the number of investor cliques.

4 Implications for the Value of Innovation

Our theory provides an explanation for the two forms of disconnect documented in Section 2. When there is intense speculation (large n), the value of an innovation reproduces our empirical facts. First, the market value of an innovation is inflated relative to its economic impact. Second, the market value of competitors reacts minimally to an innovation, even though they ultimately suffer. Finally, we show that other theories of bubbles cannot explain these facts.

4.1 Private Value

Our theory predicts a large number of new innovations and high prices on financial markets during periods of intense speculation, consistent with the boom phase of bubbles. The market price of an innovation is the market-based value $p_i = V^{(n)}(M_e)$. Because each innovation coincides with a firm in the model, this corresponds to what we measured empirically by the change in stock price of a firm following an innovation.[18] As n increases, the demand for firms $V^n(.)$ shifts up. When investors disagree, they tilt their respective portfolios toward their favorite firms. As disagreement increases, this specialization strengthens, and investors in each given firm become more optimistic about the firm's prospects. The higher demand from financial market participants increases the price of firms p_i and leads to more new firms M_e.

The outcome-based private value of innovation does not increase as much as market value in the model, mirroring the bust phase of bubbles. In the model, the outcome-based private value of innovation is the realized output of the firm, the counterpart to what we measured empirically by changes in sales or profits following an innovation. At date 1, not all investors can be right, and the average output of a firm is $V^{(1)}(M_e)$ — investors actually agree about this at date 0. In the absence of disagreement, $n = 1$, this coincides with the market-based private value $V^{(n)}(M_e)$. However, as n increases, the market-based and outcome-based values diverge and the ratio $V^{(n)}(M_e)/V^{(1)}(M_e)$ increases. In other words, the increase in market value due to speculation does not reflect a commensurate increase in fundamentals.

The market value of an innovation increases during a bubble, but this increase is not justified by a change in outcomes. These two predictions rationalize the empirical results in Table 1 and Table 2 for the market-based and outcome-based values of innovation, respectively. On the theoretical side, these results are the defining properties of bubbles applied to innovation. While these patterns also arise in other theories of bubbles, we show below that our specific type of disagreement leads to a unique behavior of competitive spillovers.

4.2 Competitive Spillovers

In Section 2.3, we showed that when a firm innovates during bubbles, its competitors' stock prices do not decrease, even though their cash-flows ultimately suffer. We construct the counterpart of this experiment in the model and ask what happens to the value and profit to all other firms in the economy when a new firm innovates and

[18] In Section 6, we come back to the distinction between innovation and firm.

enters the market. Spillovers correspond to the change in the value of all competitors relative to the private value of a new firm. Formally, the market-based spillover is

$$spill_{mkt}(n) = \frac{M_e V^{(n)\prime}(M_e)}{V^{(n)}(M_e)} \tag{12}$$

and the outcome-based spillover is

$$spill_{out} = \frac{M_e V^{(1)\prime}(M_e)}{V^{(1)}(M_e)} \tag{13}$$

In each case, the numerator is the effect on the competitors and the denominator is the private value. Internet Appendix B.2 explains the calculation of spillovers in more detail. In line with the evidence, when there is more speculation, market-based spillovers disappear while outcome-based spillovers are unchanged.

4.2.1 Outcome-Based Spillover

Consider first the spillovers on outcomes, which do not depend on individual beliefs. Using the definition of $V^{(n)}$ in equation (11), we can show that the outcome-based spillover from equation (13) is

$$spill_{out} = -\frac{\overbrace{\int_{\underline{a}}^{\infty} \pi(\underline{a}) dF(a)}^{\text{profits displaced}}}{\underbrace{\int_{\underline{a}}^{\infty} \pi(a) dF(a)}_{\text{profits earned}}} = -\frac{\gamma - \eta}{\gamma} \tag{14}$$

The ratio of the two integrals has an intuitive explanation. The numerator represents the expected amount of profits displaced by the introduction of a new firm. It is the product of the output $\pi(\underline{a})$ of the marginal producing firm (the firm that will get displaced) with the probability that that firm gets displaced, $\int_{\underline{a}}^{\infty} dF(a)$. The denominator is the expected output of a new firm.

Displacement lowers aggregate output, thus $spill_{out}$ is negative. However, displacement only occurs when the new firm is more productive than an existing firm, thus $|spill_{out}| < 1$. There is a smaller spillover if the distribution of firm productivity is more dispersed (lower values of γ) or when productivity differences translate into larger output differences (larger values of η). Finally, the outcome-based spillover $spill_{out}$ does not depend on the amount of speculation n, either directly or through the equilibrium number of firms M_e, in line with our empirical results.[19]

[19]Technically, this result is the consequence of specifying a power production function and a Pareto productivity distribution. Under a Pareto distribution, the ratio between marginal and average productivity is independent of the lower cutoff.

4.2.2 Market-Based Spillover

For market prices, investor beliefs matter. When investors agree ($n = 1$), market value and average output coincide; market-based and outcome-based measures of spillovers are identical. However, with disagreement, the two measures diverge. Using equation (12), we obtain

$$spill_{mkt}(n) = \frac{M_e V^{(n)\prime}(M_e)}{V^{(n)}(M_e)} = -\frac{\overbrace{\int_{\underline{a}}^{\infty} \pi(\underline{a}) \frac{F^{n\prime}}{F^{\prime}}(\underline{a}) dF(a)}^{\text{perceived profits displaced by others}}}{\underbrace{\int_{\underline{a}}^{\infty} \pi(a) \frac{F^{n\prime}}{F^{\prime}}(a) dF(a)}_{\text{perceived profits earned}}}. \tag{15}$$

Similar to the spillover on outcomes, the spillover on market values is a ratio of displaced profits to profits earned. However, it is important to note who values each of these quantities in the market, which is captured by the different distributions on the denominator and numerator in equation (15): $F^{n\prime}/F^{\prime}(\underline{a})$ versus $F^{n\prime}/F^{\prime}(a)$. The denominator, which captures the baseline value of a firm, is driven by investors who are the most optimistic about it, as we discuss in Section 3.2. However, the spillovers of an innovation—the numerator—are not valued by those most optimistic about the innovation, but rather by those receiving the spillovers. For those other investors, the new innovation is one of many, when their firm is better than a typical firm. Therefore, the spillovers are perceived by those receiving it as much weaker than they actually are. In particular, they believe that their own firms' productivities are drawn from the distribution F^n whereas the competing innovation is drawn from the population distribution F.

Increasing disagreement, larger n, corresponds to shifting the perceived distribution of productivities to the right; the change in probability weights is increasing as we move to higher productivities. Such an increase affects expected profits more strongly than the value of displaced firms, decreasing the wedge. The following proposition summarizes this dampening.

Proposition 1. *Speculation lowers the intensity of market-based spillovers:*

$$|spill_{mkt}(n)| < |spill_{mkt}(1)|, \quad \text{for } n > 1. \tag{16}$$

This proposition gives rise to a clear prediction: the market-based spillover is lower in speculative periods. This corresponds to the result we find in Section 2.3 where we measure these spillovers by gauging the reaction of firms' valuations to variation in the amount of innovation by their competitors. In Internet Appendix C, we show that the spillover to asset prices can even disappear altogether in the limit of large disagreement with $n \to \infty$. We can also compare market-based and outcome-based spillovers. An immediate corollary of Proposition 1 is that the presence of speculation reduces market-based spillovers relative to outcome-based spillovers. In Section 4.3 below, we show how these results are specific to our model of bubbles and differ from other theories.

Robustness. This result is robust to changes in how we model competitive spillovers. For simplicity, our baseline takes a stark view of competition: either the firm is in the top mass M and produces at full capacity, or it is excluded altogether. We consider a number of generalizations of these assumptions; the disappearance of spillovers to market value occurs in all of them.

In Internet Appendix D.1, we consider a smoother form of competition: profits decrease in an arbitrary way with the productivity rank of a firm. When a new firm enters, all the firms it overtakes suffer as their productivity ranks get knocked down. In this setting, the spillovers arise throughout the firm distribution. However, investors are still more optimistic about the firms they invest in than those inflicting the spillover onto them.

Another consideration is that the marginal firm earns positive profits in our baseline; in practice, this would generate incentives for inframarginal firms to invest resources to acquire this production slot. Internet Appendix D.2 entertains this possibility by allowing firms to compete for the production slots by spending on advertisement. In equilibrium, the most productive firms end up being active like in our baseline. However, they must spend the potential profits of the first non-producing firm on advertisement to secure their slot. In this setting, the spillovers of a new firm no longer occur through the exclusion of the marginal firms — which earns zero net profit — but through increasing the necessary advertisement spending to enter. Still, we show that the total spillovers are equal to those in our baseline and therefore that they disappear in market value with more disagreement.

We also consider three sets of explicit foundations for the profit functions in Internet Appendix D.3. One approach is to assume that firms must hire workers, assumed to be in fixed aggregate supply, and have decreasing return to scale production functions. The second approach is to assume that firms produce goods that are imperfect substitutes, as in Dixit and Stiglitz (1977). The third approach is to allow for a firm's productivity to depend on the overall productivity of the economy, capturing the knowledge spillovers emphasized by Bloom, Schankerman, and Van Reenen (2013). Entertaining these foundations leads to additional sources of spillovers, on workers, aggregate demand, and productivity, respectively. However, beyond these other equilibrium effects, they lead to the firm profit function of equation (4).[20] Hence, the competitive spillovers behave like in our baseline.

Finally, we study the impact of variations in the theory of how firms are created and become active. Allowing some response in the mass of production slots M to the amount of blueprint creation (Internet Appendix D.4) mitigates the actual compet-

[20] In these settings, the elasticity of profits to productivity η is a function of the degree of decreasing returns to scale and the elasticity of substitution across goods.

itive spillovers. However, it does not change how speculation affects these spillovers in market value. Relatedly, requiring an input in fixed supply to create firms (e.g., investments in infrastructure) creates congestion effects in entry (Internet Appendix D.5) but does not affect the effect firms have on each other's profits.

4.3 Comparison to Other Theories of Bubbles

Our specific model of disagreement is uniquely suited to understand the behavior of the value of innovation. While our theory shares a number of properties with existing models of bubbles, it distinguishes itself in some key dimensions, notably the valuation of spillovers.[21]

The first theory we consider is that the supposed bubbles are not actually bubbles. Instead, they are periods of high valuation that arise because innovation waves have a large positive impact on the economy. In our model, this would correspond to no disagreement ($n = 1$) and a uniform increase in profits. That is, we replace equation (4) by

$$\pi(a) = A \cdot a^{\eta} \cdot \mathbf{1}\left\{a \geq \underline{a}\right\}, \tag{17}$$

where $A > 1$ represents an aggregate productivity increase. In this case, all valuations and real outcomes are shifted up by the same factor A. This would imply no differential change in market-based and outcome-based measures during bubbles, at odds with all of our empirical results.

Second, the high valuations during bubbles could reflect irrationally high expectations about the impact of innovation on the economy. Alternatively, households could have lower discount rates when evaluating the firms (Kogan, Papanikolaou, and Stoffman, 2020). These two views can be represented in our model in the same way. Output at date 1 is given by the baseline specification of equation (4). However decisions at date 0 are taken by multiplying this output by a common factor as in equation (17). In the overoptimism interpretation, A represents how much people overestimate the productivity. In the discount rate interpretation, $A = 1/(1 + r)$ represents the discounting of cash flows. Both low discount rates and high optimism lead to inflated market values relative to output by the factor A, generating our results on the private value of innovation. However, this theory fails to match the disconnect between market-based and output-based spillovers. Because all investors overestimate the value of all firms, their perceptions of interactions between firms is unaffected. Formally, because the market-based spillover (15) is a ratio of valuations, it is unaffected by the common factor A and remains equal to the output-based spillover.

Third, bubbles could arise because of aggregate disagreement. Rather than disagreeing on the profits of Facebook versus Twitter, investors disagree on the total

[21] See Simsek (2021) for a recent survey of theories of bubbles.

value of social networks. Optimists push the price of all firms up; pessimists would like to short but cannot. Formally, we assume that half of the population thinks that aggregate productivity in equation (17) is $A = A_h$, while the other half thinks it is $A = A_l$, where $A_h > A_l$ and $(A_h + A_l)/2 = 1$ such that beliefs are correct on average. In equilibrium the price of firms is determined by the optimists and is $p = A_h \mathcal{I}_1$. This setting is exactly equivalent to the previous one and therefore suffers from the same issue. It matches the rise in valuation relative to output, but fails to account for the behavior of competitive spillovers. While only some investors are optimistic, they are optimistic about all the firms. Hence, they understand spillovers; market prices reflect them. This is a key distinction from our framework, where different investors are optimistic about different firms and the dispersion of optimism distorts the market evaluation of spillovers.

Last, bubbles could be episodes in which prices have an additional component which has nothing to do with fundamentals. Such bubbles can arise as the consequence of irrational exhuberance (Shiller, 2015) or as a form of rational bubble (Blanchard, 1979; Tirole, 1985). This view is not helpful to understand the evidence. For example, while valuations are inflated during the bubble episodes we measure, the valuation also overreacts to the announcement of fundamental news, patent approvals. In addition, the quality of the patent, as measured by citations, has a larger effect on the response of prices to this news. Both of these results suggest a tight connection of the bubble component with fundamentals, at odds with these theories.[22]

Overall, while we cannot rule out a role for other models of bubbles in innovation booms, only our theory of disagreement across firms offers a unifying explanation for the link between innovation and bubbles.

5 Using the Value of Innovation for Optimal Policy

How should innovation policy respond to the disconnect of price-based and outcome-based measures of innovation? In this section, we use our theory of the disconnect to shed light on this issue. We focus on a simple policy, a subsidy or tax on firm entry.

We consider the decision of a planner maximizing social welfare by choosing a tax τ for each additional firm created from blueprints, rebated as a lump sum to households. A common challenge in presence of disagreement is the choice of belief under which to evaluate welfare. When agents disagree, which beliefs should the planner use to evaluate welfare? We follow the recommendation of Brunnermeier, Simsek, and Xiong (2014) and choose their belief-neutral social welfare criterion. Specifically the planner

[22]In all generality, rational bubbles can also be correlated with fundamentals. Obtaining our empirical results in that framework would rely on an improbable coincidence: there is a different rational bubble for each firm, correlated positively both with the fundamentals of this firm and of its competitors in a way that exactly cancels out competitive spillovers.

favors an allocation over another one if the utilitarian social welfare function — the sum of expected utilities — is larger for all convex combinations of investors' beliefs. In our model, because agents agree on the population distribution of firm productivity F, they also agree on the total expected utility. Thus, the planner problem simplifies to maximizing aggregate output net of entry costs $M_e V^{(1)}(M_e) - W(M_e)$.

This yields the optimal tax:

$$\tau = 1 - (spill_{out} + 1)\frac{\text{private value}_{out}}{\text{private value}_{mkt}}. \tag{18}$$

Two aspects drive the choice of policy. First, the planner leans against the bubble. The private value of firms on markets (the stock price $V^{(n)}$), which shapes incentives to innovate, is inflated relative to their actual output $V^{(1)}$, the private value according to the planner's beliefs. This force will push the planner towards taxing entry in order to counteract the effect of excessive market valuation, and becomes stronger with more disagreement. As such our first disconnect is not only a measurement issue, but also contributes to inefficient entry.

Second, the planner wants to address the externalities in a Pigouvian fashion. In our theory, competitive spillovers are a source of externality: new firms do not internalize their effect on others. The planner measures spillovers based on real outcomes. Our empirical results and model show that directly using estimates of spillovers measured using market values instead would lead to incorrect inference. In the planner's view, the market assessment of externalities is erroneous because the beliefs of those receiving the externality differ from those originating it. It is not only that all valuations are too large, but also that the consequences of innovations are misattributed.

Interestingly, agents in the economy do not support such a tax. In Appendix D, we show that they would favor a simple tax $\tau = -spill_{mkt}$. Despite their disagreements, they (wrongly) believe that markets accurately reflect the externalities of innovation.

6 Additional Empirical Evidence

In this section, we present additional direct evidence for the mechanism based on portfolio choices and the interaction of speculation with the boundaries of the firm.

6.1 Portfolio Holdings

Disagreement across investors about which firms will succeed has implications for portfolio decisions. When investors disagree, they concentrate their investments in the firms they are more positive about. This implies that portfolios are less diverse during bubbles. In particular, our model predicts that for a level n of disagreement, each investor only invests in one out of n firms within the industry.

We test this prediction empirically using 13F filings of portfolio positions. Each institution with over $100 million invested in stocks must report their holdings quarterly to the SEC. We obtain this information from 1980 to 2020 from the Thomson s34 dataset.[23] To focus on investors with meaningful investment in an industry, we restrict our attention to institution-industry pairs where the institution holds more than 10 stocks in this industry. The empirical counterpart to n is the ratio of the number of stocks available in an industry to the number of stocks held in the institution's portfolio for that industry.

To assess whether portfolios are more concentrated during bubbles, we regress this measure on the bubble indicator and report the results in Panel A of Table 4. We control for the number of stocks in the industry, the size of the institution, and industry volatility. Column 1 is the baseline effect, column 2 includes industry fixed effects, column 3 includes date fixed effects, and column 4 combines the two types of fixed effects. The results show consistently that portfolio concentration significantly increases during bubbles — investors hold a smaller fraction of the number of stocks available in an industry.

In Panel B of Table 4, we consider whether the increase in portfolio concentration differs across types of institutions. We repeat the specification of Panel A, allowing for type-specific intercepts and coefficients on the bubble indicator.[24] We find that the increase in portfolio concentration is pervasive across all types of institutions. The only group of institutions with a more pronounced increase in concentration is investment advisers. Last, in Internet Appendix Tables IA.6 and IA.7, we consider alternative measures of portfolio concentration using the four-firm concentration ratio (CR4) and the Herfindahl index (HHI). We find that our results are robust to using these measures: portfolio concentration increases significantly during bubbles.

6.2 Disagreement and the Boundaries of the Firm

Another prediction of our theory is that the boundaries of the firm matter for valuations. In our model, each firm is valued under the beliefs of the investors who are the most optimistic about it, but the most optimistic investors for two different firms will in general be different. As a result, the valuation of the combined firm will be lower than or equal to the sum of valuations of the standalone firms. In contrast, these boundaries are irrelevant for asset prices under the alternative theories discussed in Section 4.3: valuations are the sum of cash flows under the common belief.

Thus, our model predicts a "conglomerate discount" on the value of innovation in times of bubbles: in the presence of disagreement or bubbles, multiple-product firms experience a smaller increase in the value of their innovation than narrow firms,

[23] We apply the same filters as Koijen and Yogo (2019) and Haddad, Huebner, and Loualiche (2021).

[24] We use the classification of Koijen and Yogo (2019) which adresses misattribution in the Thomson files.

Table 4
Ownership and Bubbles

	Panel A: Fraction of stock holdings in industry			
	(1)	(2)	(3)	(4)
Bubble	2.664***	1.763***	2.374***	1.443***
	(0.263)	(0.178)	(0.344)	(0.204)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	No	Industry	Date	Industry, Date
Observations	1,449,168	1,449,168	1,449,168	1,449,168
R^2	0.48	0.56	0.49	0.56
	Panel B: Fraction of stock holdings by institution type			
	(1)	(2)	(3)	(4)
Bubble x Banks	1.872***	1.182***	1.988***	0.973***
	(0.382)	(0.286)	(0.472)	(0.327)
Bubble x Insurance	2.475***	1.868***	2.437***	1.621***
	(0.483)	(0.393)	(0.611)	(0.460)
Bubble x Inv. Advisers	5.108***	3.948***	4.366***	3.428***
	(0.449)	(0.480)	(0.346)	(0.418)
Bubble x Mutual Funds	1.071**	0.289	0.803	−0.048
	(0.464)	(0.319)	(0.568)	(0.378)
Bubble x Pension Funds	1.008	0.268	0.712	−0.052
	(0.856)	(0.763)	(0.972)	(0.834)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	No	Industry	Date	Industry, Date
Observations	1,449,168	1,449,168	1,449,168	1,449,168
R^2	0.49	0.56	0.49	0.56

Note: Table 4 presents a regression of portfolio concentration within an industry on a bubble dummy in that industry-year. Industry portfolio concentration is measured using portfolio holdings the SEC 13F filings, where we construct the ratio of the total number of firms in an industry to the number of firms held in a portfolio for this industry. In Panel B, institution types follow Koijen and Yogo (2019) and are derived from the SEC 13F filings. All specifications include controls for institution size (log of AUM), the logarithm of number of stocks in a given industry, and industry volatility. We use industry and date fixed effects depending on the specification. Panel B also includes fixed effects for institution types. Standard errors clustered at the date (quarterly) level are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

which cater to specific investors who tend to have strong belief in their success.[25] We confirm this prediction in Internet Appendix Table IA.8 and IA.9. Specifically,

[25] Reed, Saffi, and Van Wesep (2020) provide empirical evidence of a similar conglomerate discount in stock valuations. Huang et al. (2020) document a similar discount in portfolios.

we find that the overvaluation of patents during bubbles is less pronounced for firms operating across multiple industries.

This result is also related to a literature on how structures of financial assets facilitating the expression of disagreement lead to higher valuations.[26] While that work focuses on the role of financial innovation, our results emphasize how changing disagreement interacts with the natural structure of firms.

7 Conclusion

In this paper, we establish empirically a disconnect between the economic impact of innovation and its market value during bubbles. This disconnect takes two contrasting aspects. The first is an overreaction: the impact of an innovation on the stock price of its creator increases sharply relative to the real outcomes it will generate. The second is an underreaction: even though innovation in a firm damages profits of its competitors, these negative spillovers have no impact on their stock prices.

Standard theories of bubbles do not account for these two facts simultaneously. We propose a new approach that resolves this tension: investors disagree about which firms will succeed, rather than about the aggregate value of an industry. Each investor is optimistic about the firms they invest in, so overestimate their prospects. However, investors receiving the competitive spillovers are optimistic about other firms, and therefore underestimate them. The theory highlights that the disconnect is not without real consequences. For example, innovation policy during a bubble must not only lean against the overvaluation in the direct value of innovation but also account for the externalities across firms even though they are hidden by disagreement.

[26] Simsek (2013) shows that with investor disagreement, profitable financial innovation facilitates betting on disagreement and Iachan, Nenov, and Simsek (2021) consider the implications for savings and equilibrium prices. Fostel and Geanakoplos (2016) and Broer (2018) show that tranching (e.g. securitization) increases asset value by separating optimists and pessimists and affects investment.

References

Abernathy, W.J. 1978. *The Productivity Dilemma: Roadblock to Innovation in the Automobile Industry*. Johns Hopkins University Press, Baltimore, Md.

Acemoglu, Daron. 2008. *Introduction to Modern Economic Growth*. Princeton University Press.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2018. "Extrapolation and Bubbles." *Journal of Financial Economics* 129 (2):203–227.

Blanchard, Olivier Jean. 1979. "Speculative Bubbles, Crashes and Rational Expectations." *Economics Letters* 3 (4):387–389.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4):1347–1393.

Broer, Tobias. 2018. "Securitization bubbles: Structured finance with disagreement about default risk." *Journal of Financial Economics* 127 (3):505–518.

Brunnermeier, Markus K and Martin Oehmke. 2013. "Bubbles, Financial Crises, and Systemic Risk." In *Handbook of the Economics of Finance*, vol. 2. Elsevier, 1221–1288.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong. 2014. "A Welfare Criterion for Models with Distorted Beliefs." *Quarterly Journal of Economics* 129 (4):1753–1797.

Chinco, Alex. 2020. "The Madness of Crowds and the Likelihood of Bubbles." Working paper, University of Illinois.

Corhay, Alexandre, Howard Kung, and Lukas Schmid. 2015. "Markups, Competition and Asset Prices." Working paper, Duke University.

Davila, Eduardo. 2020. "Optimal Financial Transaction Taxes." Working paper, Yale University.

Dávila, Eduardo and Ansgar Walther. 2021. "Prudential Policy with Distorted Beliefs." Working Paper 28879, National Bureau of Economic Research.

D'avolio, Gene. 2002. "The Market for Borrowing Stock." *Journal of Financial Economics* 66 (2-3):271–306.

Dixit, Avinash K and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3):297–308.

Dong, Ming, David Hirshleifer, and Siew Teoh. 2020. "Stock Market Overvaluation, Moon Shots, and Corporate Innovation." Working paper, UC Irvine.

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen. 2002. "Securities Lending, Shorting, and Pricing." *Journal of Financial Economics* 66 (2-3):307–339.

Fama, Eugene F and Kenneth R French. 1997. "Industry Costs of Equity." *Journal of Financial Economics* 43 (2):153–193.

Farhi, Emmanuel and Jean Tirole. 2011. "Bubbly Liquidity." *The Review of Economic Studies* 79 (2):678–706.

Fostel, Ana and John Geanakoplos. 2016. "Financial Innovation, Collateral, and Investment." *American Economic Journal: Macroeconomics* 8 (1):242–84.

Freeman, Christopher. 1982. *The Economics of Industrial Innovation.* MIT Press.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas. 2012. "Displacement Risk and Asset Returns." *Journal of Financial Economics* 105 (3):491–510.

Greenwood, Robin, Andrei Shleifer, and Yang You. 2018. "Bubbles for Fama." *Journal of Financial Economics* .

Griliches, Zvi, Ariel Pakes, and Bronwyn H. Hall. 1986. "The Value of Patents as Indicators of Inventive Activity." Working Paper 2083, National Bureau of Economic Research.

Haddad, Valentin, Paul Huebner, and Erik Loualiche. 2021. "How Competitive is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing." Working paper, UCLA Anderson.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg. 2005. "Market Value and Patent Citations." *The RAND Journal of Economics* 36 (1):16–38.

Harhoff, Dietmar, Francis Narin, F. M. Scherer, and Katrin Vopel. 1999. "Citation Frequency and the Value of Patented Inventions." *The Review of Economics and Statistics* 81 (3):511–515.

Harrison, J. Michael and David M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Quarterly Journal of Economics* 92 (2):323–336.

He, Zhaozhao and David Hirshleifer. 2020. "The Exploratory Mindset and Corporate Innovation." *Journal of Financial and Quantitative Analysis* :1–43.

Hirshleifer, David and Joshua Plotkin. 2020. "Moonshots, Investment Booms, and Selection Bias in the Transmission of Cultural Traits." Working Paper 27735, National Bureau of Economic Research.

Hombert, Johan and Adrien Matray. 2020. "Technology Boom, Labor Reallocation, and Human Capital Depreciation." Working paper, HEC.

Huang, Shiyang, Byoung-Hyoun Hwang, Dong Lou, and Chengxi Yin. 2020. "Offsetting Disagreement and Security Prices." *Management Science*.

Iachan, Felipe S., Plamen T. Nenov, and Alp Simsek. 2021. "The Choice Channel of Financial Innovation." *American Economic Journal: Macroeconomics* 13 (2):333–72.

Janeway, W.H. 2012. *Doing Capitalism in the Innovation Economy: Markets, Speculation and the State*. Cambridge University Press.

Jones, Charles M and Owen A Lamont. 2002. "Short-Sale Constraints and Stock Returns." *Journal of Financial Economics* 66 (2-3):207–239.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics* 132 (2):665–712.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman. 2020. "Left Behind: Creative Destruction, Inequality, and the Stock Market." *Journal of Political Economy* 128 (3):855–906.

Koijen, Ralph S. J. and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy* 127 (4):1475–1515.

Kondor, Péter and Botond Köszegi. 2017. "Cursed Financial Innovation." Working paper, CEU.

Kung, Howard and Lukas Schmid. 2015. "Innovation, Growth, and Asset Prices." *Journal of Finance* 70 (3):1001–1037.

Loualiche, Erik. 2020. "Asset Pricing with Entry and Imperfect Competition." Working paper, University of Minnesota.

Martin, Alberto and Jaume Ventura. 2012. "Economic Growth with Bubbles." *American Economic Review* 102 (6):3033–3058.

Miller, Edward M. 1977. "Risk, Uncertainty, and Divergence of Opinion." *Journal of Finance* 32 (4):1151–1168.

Mokyr, Joel. 1992. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.

Morris, Stephen. 1995. "The Common Prior Assumption in Economic Theory." *Economics and Philosophy* 11 (2):227–253.

Nicholas, Tom. 2008. "Does Innovation Cause Stock Market Runups? Evidence from the Great Crash." *American Economic Review* 98 (4):1370–96.

Pakes, Ariel. 1985. "On Patents, R & D, and the Stock Market Rate of Return." *Journal of Political Economy* 93 (2):390–409.

Pastor, Lubos and Pietro Veronesi. 2005. "Rational IPO Waves." *Journal of Finance* 60 (4):1713–1757.

———. 2009. "Technological Revolutions and Stock Prices." *American Economic Review* 99 (2):1451–1483.

Reed, Adam V., Pedro A. C. Saffi, and Edward D. Van Wesep. 2020. "Short-Sales Constraints and the Diversification Puzzle." *Management Science*.

Samuelson, Paul A. 1958. "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money." *Journal of Political Economy* 66 (6):467–482.

Scheinkman, J.A. 2014. *Speculation, Trading, and Bubbles*. Kenneth J. Arrow Lecture Series. Columbia University Press.

Scheinkman, Jose A. and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." *Journal of Political Economy* 111 (6):1183–1220.

Shiller, R.J. 2015. *Irrational Exuberance: Revised and Expanded Third Edition*. Princeton University Press.

Simsek, Alp. 2013. "Speculation and Risk Sharing with New Financial Assets." *Quarterly Journal of Economics* 128 (3):1365–1396.

———. 2021. "The Macroeconomics of Financial Speculation." *Annual Review of Economics* 13.

Tirole, Jean. 1985. "Asset Bubbles and Overlapping Generations." *Econometrica* 53 (6):1499–1528.

van Binsbergen, Jules H. and Christian C. Opp. 2019. "Real Anomalies." *Journal of Finance* 74 (4):1659–1706.

Appendix

A Equilibrium Conditions

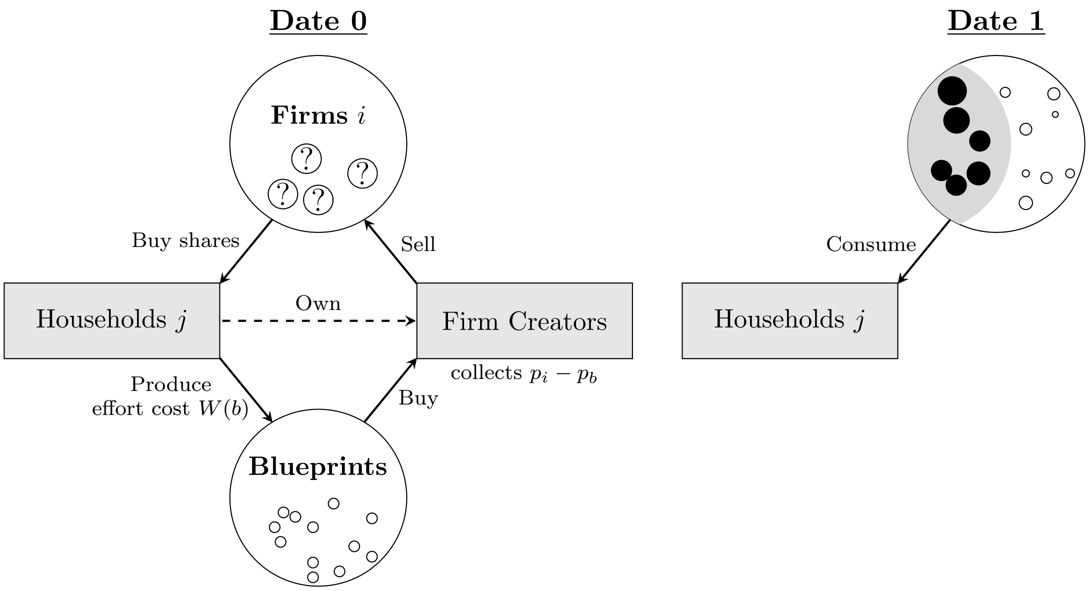

Figure A.1
Summary of model structure.
The left panel represents the stage of firm creation at date $t = 0$, when blueprints are created and households buy shares in firms on financial markets based on their beliefs. The right panel represents the production stage at $t = 1$.

Figure A.1 summarizes the model. To derive the equilibrium condition (11), we first take first-order conditions for the household problem (5)-(6). From the first-order condition for c_0, we have that the Lagrange multiplier on the budget constraint (6) is 1. Therefore, the first-order condition for b_j is

$$W'(b_j) = p_b \tag{A.1}$$

and the first-order condition for s_i^j for a firm i that household j chooses to invest in is

$$\int_{\underline{a}}^{\infty} \pi_i(a) dF^n a = p_i. \tag{A.2}$$

The integral is over the distribution F^n because households only invest in their favorite firms. The equilibrium condition (11) follows by equating (A.1) and (A.2) using the relation $p_b = p_i$ from the firm's problem (7).

B Private Value

We argued in the text that market-based private value $V^{(n)}$ and entry M_e increase with disagreement n. To see this, write the equilibrium condition (11) as:

$$f_e\left(\frac{M_e}{M}\right)^{\theta}=\int_{F^{-1}(1-M/M_e)}^{\infty}a^{\eta}dF^{n}(a) \tag{A.3}$$

The left-hand side is increasing in M_e and independent of n. The right-hand side is decreasing in M_e and increasing in n. Therefore, as n increases, the M_e satisfying (A.3) increases. Correspondingly, the right-hand side of (A.3), which is $V^{(n)}$, increases since the left-hand side increases with M_e. Since M_e is increasing with n, the outcome-based spillover $V^{(1)}$ decreases and therefore the ratio $V^{(n)}/V^{(1)}$ increases.

C Spillovers

Market-based spillover. To show that the market-based spillover defined in equation (12) can be expressed as in equation (15), write the denominator in the integral form:

$$V^{(n)}=\int_{\underline{a}}^{\infty}\pi(a)dF^{n}(a)=\int_{\underline{a}}^{\infty}\pi\left(a\right)\frac{F^{n\prime}\left(a\right)}{F^{\prime}\left(a\right)}dF(a)$$

and express the numerator as:

$$M_e\frac{dV^{(n)}}{dM_e}=\frac{M}{M_e}\cdot\pi\left(\underline{a}\right)\cdot\frac{F^{n\prime}}{F^{\prime}}\left(\underline{a}\right)=-\int_{\underline{a}}^{\infty}\pi\left(\underline{a}\right)\frac{F^{n\prime}\left(\underline{a}\right)}{F^{\prime}\left(\underline{a}\right)}dF(a). \tag{A.4}$$

Outcome-based spillover. To show that the outcome-based spillover defined in equation (13) can be expressed as in (14), we use the assumption that $F(a)=1-a^{-\gamma}$ and $\pi(a)=a^{\eta}\cdot\mathbf{1}\left\{a\geq\underline{a}\right\}$. Using the result that

$$\underline{a}=F^{-1}\left(1-\frac{M}{M_e}\right)=\left(\frac{M_e}{M}\right)^{1/\gamma},$$

we can show that the value of a firm is:

$$V^{(1)}(M_e)=\int_{\left(\frac{M_e}{M}\right)^{\frac{1}{\gamma}}}^{\infty}a^{\eta}\gamma a^{-\gamma-1}da=\frac{\gamma}{\gamma-\eta}\left(\frac{M_e}{M}\right)^{\frac{\eta-\gamma}{\gamma}}. \tag{A.5}$$

Thus, the numerator of the spillover is:

$$-M_e\frac{dV^{(1)}}{dM_e}=\left(\frac{M_e}{M}\right)^{\frac{\eta-\gamma}{\gamma}}, \tag{A.6}$$

which leads directly to the desired formula (14) for the outcome-based spillover.

Proof of Proposition 1. Proposition 1 follows directly from the following series of relations:

$$|spill_{mkt}(n;M_e)|\leq|spill_{mkt}(1;M_e)|=|spill_{out}(M_e)|=|spill_{out}(M_e^{\prime})|. \tag{A.7}$$

These hold for any M_e' and, in particular, the M_e' corresponding to the equilibrium entry under agreement. The first equality is from the definition of $spill_{mkt}$ and $spill_{out}$ and the last equality is from the fact that the outcome-based spillover (14) does not depend on M_e.

To show the first inequality, we first note that

$$V^{(n)\prime} = -\frac{1}{M_e}\frac{M}{M_e}\pi\left(\underline{a}\right) \cdot n\left(1 - \frac{M}{M_e}\right)^{n-1}.$$

Next, we bound $V^{(n)}$:

$$\begin{aligned} V^{(n)} &= \int_{\underline{a}}^{\infty} \pi(a) n F^{n-1}(a) dF(a) \geq \int_{\underline{a}}^{\infty} \pi(a) n F^{n-1}\left(\underline{a}\right) dF(a) \\ &\geq n\left(1 - \frac{M}{M_e}\right)^{n-1} \int_{\underline{a}}^{\infty} \pi(a) dF(a). \end{aligned}$$

Therefore, we are able to bound the market-based spillover for a given M_e and n:

$$|spill_{mkt}(n; M_e)| \leq -\frac{\int_{\underline{a}}^{\infty} \pi\left(\underline{a}\right) dF(a)}{\int_{\underline{a}}^{\infty} \pi(a) dF(a)} \leq |spill_{mkt}(1; M_e)|, \tag{A.8}$$

where the second inequality comes from the definition of $spill_{mkt}(1; M_e)$.

D Planner Problem

The planner's problem is:

$$\max_{\tau} 1 + M_e \mathbf{E}\{\pi_i\} - W(M_e), \tag{A.9}$$

where the expectation $\mathbf{E}\{\cdot\}$ is taken over the population distribution. The planner chooses a tax on new firms, hence firm creators now pay $(1+\tau)p_i$ for new firms. The equilibrium condition (11) now becomes:

$$W'(M_e) = (1+\tau) V^{(n)}(M_e). \tag{A.10}$$

Choosing the optimal tax is equivalent to solving for the optimal entry M_e to maximize the objective function in (A.9) subject to the equilibrium condition (A.10), which yields the first-order condition:

$$(1+\tau) V^{(n)}(M_e) = V^{(1)}(M_e) + M_e V^{(1)\prime}(M_e).$$

Rearranging, we have:

$$\begin{aligned} \tau &= 1 - \left(\frac{V^{(1)}(M_e)}{V^{(n)}(M_e)} + \frac{M_e V^{(1)\prime}(M_e)}{V^{(n)}(M_e)}\right) = 1 - \left(\frac{M_e V^{(1)\prime}(M_e)}{V^{(1)}(M_e)} + 1\right) \frac{V^{(1)}(M_e)}{V^{(n)}(M_e)} \\ &= 1 - (spill_{out} + 1) \frac{\text{private value}_{out}}{\text{private value}_{mkt}}. \end{aligned} \tag{A.11}$$

As noted in the main text, the tax in equation (A.11) is not the tax that is favored by households. Instead, households would prefer the tax implied replacing $\mathbf{E}\{\cdot\}$ in equation (A.9) with the expectation under the households' subjective distribution. This would instead yield

$$\tau = 1 - \left(\frac{V^{(n)}(M_e)}{V^{(n)}(M_e)} + \frac{M_e V^{(n)\prime}(M_e)}{V^{(n)}(M_e)} \right) = -spill_{mkt}.$$