Risk Management Failures*

Matthieu Bouvard† Samuel Lee‡

May 2019

Abstract

We model risk management as information acquisition that delays trading decisions. In markets with preemptive competition, this can lead to a race to the bottom where prioritizing trade execution over risk management is optimal for each firm, but collectively inefficient. As time competition intensifies, mean trading profit supplants risk concerns as the main driver of risk management quality, causing risk misallocation to rise with trading speed and volume. This pathology of risk management failure—the trio of time-consuming risk assessment, preemptive competition, and boom markets—has distinct regulatory implications.

JEL Classification: G20; G32; G01.

*We thank Bruno Biais, Philip Bond, Adolfo de Motta, Nadya Malenko, Albert Menkveld, Konstantin Milbradt, Martin Oehmke, Adriano Rampini, Hersh Shefrin, Boris Vallée, Basil Williams and Liyan Yang for helpful discussions. We also thank participants in seminars at the Tinbergen Institute, Université Laval, McGill University, the Toulouse School of Economics, Santa Clara University, the University of Calgary, the 24th Mitsui Finance Symposium, the 2015 China International Conference in Finance, the 2015 OxFIT conference, the 2015 Wharton Conference on Liquidity and Financial Crises, the 2015 Washington University Corporate Finance Conference, and the 2016 AFA Meetings. Matthieu Bouvard acknowledges financial support from FRQSC (Young Researcher Grant #175728). This project was partly completed at the Stanford Institute for Economic Policy Research, which Samuel Lee thanks for its hospitality.

†McGill University, Desautels Faculty of Management. E-mail: matthieu.bouvard@mcgill.ca.

‡Santa Clara University, Leavey School of Business, ECGI, and Swedish House of Finance. E-mail: sslee@scu.edu.

There are three ways to make a living in this business: Be first, be smarter, or cheat. Well I don't cheat, and even though I like to think we have got some pretty smart people in this building, of the two remaining options it sure is a hell of a lot easier to just be first.

(*Margin Call*, J.C. Chandor, 2011)

1 Introduction

The role of risk management is to collect, aggregate, and analyze information about risk exposures and then to manage those risks. Most risk management theories focus on the second stage, wherein a firm grapples with *known* risk exposures. The relevant literature shows that financial constraints and agency problems can lead to a constrained efficient lack of hedging (e.g., Smith and Stulz, 1985; Froot, Scharfstein and Stein, 1993; Rampini and Viswanathan, 2010).

A failure to manage risk can also originate in the first stage. Stulz (2008) describes real-world cases that are symptomatic of communication or information deficiencies inside firms. In five of six categories of risk management failure that Stulz identifies such failure involves *learning* of critical risk exposures.[1] In this paper, we argue that learning about risk should be viewed not only as a single-firm problem but also as a set of interrelated decisions that firms competing in the same market must make. Focusing on financial markets, the idea is essentially that the value of risk management to a firm depends on the intensity of the competition it faces, which in turn depends on the overall level of risk management across all firms. This loop creates externalities in learning about risk so that, absent coordination, risk management failures can result from a *market* failure.

The key premise underlying this approach is that risk management introduces frictions in a financial institution's operations that affect its ability to compete in capital markets. This is based on two observations. First, according to finance practitioners, risk management

[1] Those five categories are (1) mismeasurement of known risks, (2) failure to take risks into account, (3) failure to communicate risks to top management, (4) failure to monitor risks, and (5) failure to use appropriate metrics. On this issue, some argue that risk-shifting incentives can cause underinvestment in information (Maćkowiak and Wiederholt, 2012) and that risks of a certain nature are likely to be neglected due to behavioral biases (Shefrin 2015).

> reduces to the basics of getting the right information, at the right time, to the right people, such that those people can make the most informed judgments possible.[2]

That this organization-wide effort to gather, transmit, aggregate, and analyze information about risk exposures must be not only accurate but also *timely* implies that, if the process is too slow, it imposes costly delays or is less useful.

Second, the time demand on risk management may stem from a fast-paced competitive market environment. In a survey on risk management in financial firms, the most frequently cited concern after "higher-quality and more timely reports" is the "balance between a sales-driven front-office culture and a risk-focused culture" (Ernst & Young, 2013, 14). Relatedly, Stulz (2014) argues that more granular monitoring processes prevent the accumulation of large unmonitored pockets of risk, but also impede traders' ability to react quickly to market opportunities. This trade-off is the entry point of our analysis: we assume that risk assessment "costs" time, which delays trades. To firms that face short-lived trading opportunities, this entails costs that never fade in importance. On the contrary, opportunity costs scale with the size or frequency of investments.

To explore the implications of this premise, we construct a simple model of trading under time pressure in which competing firms search for scarce trading opportunities, but before beginning their search, each decides whether to activate a risk management system. An active system investigates any discovered opportunity on its "fit" with the firm's risk profile, which improves trading decisions but takes time. This creates tension between trade execution and risk management that depends on the time pressure a firm faces. While time pressure makes any trading strategy less profitable, it undermines strategies that include risk management more severely for two reasons: (1) the information value of risk management depends on the ability to execute trades, and (2) the execution of slower strategies is more time-sensitive. Consequently, a firm eschews risk management if time pressure surpasses a certain threshold, in which case prioritizing trade execution over risk management maximizes the firm's value despite potential

[2] We take this definition from a report entitled "Containing Systemic Risk: The Road to Reform" (CRMPG III, 2008, 70), which was submitted to the US Treasury and the Financial Stability Board by a policy group composed of top executives and chief risk officers of leading global banks. The report also describes risk management as a process by which *"critical information flowing into and out of risk monitoring processes can be distilled and compiled in a coherent and timely manner and made available, not only to the risk managers, but to key business leaders across the institution and to top management."*

"risk management failures"—losses from risks that the firm would have mitigated had it properly monitored them.

While individually optimal, this can lead to a collectively inefficient race to the bottom: Firms that reduce risk management to trade more quickly increase time pressure in the market, which encourages other firms to do the same. Thus, there exist equilibria in which competitive forces undermine risk management incentives to the detriment of all firms. The root of this market failure is that every firm, while optimizing its own time use, fails to consider its contribution to time pressure in the market. To examine the conditions under which such failures are likely, we use a global game treatment to eliminate equilibrium multiplicity. The comparative statics reveal a fundamental shift in the determinants of risk management: Absent time pressure, whether risk management is activated depends on its information value, that is, how critical it is for a firm to monitor its risk exposures, but it is independent of market conditions, that is, how appealing trades are ex ante without such information. When time pressure increases as a result of structural factors (e.g., technology), market conditions gradually replace information value as the main driver of risk management choices. At the limit, they become the exclusive consideration: firms activate risk management only during "busts" and abandon it during "booms," irrespective of the information value.

Our theory identifies a set of three conditions as the catalyst for such failures: (1) firms in which collecting and aggregating risk data is time-intensive, for example, due to decentralized risk-taking decisions, (2) markets in which investment opportunities are short-lived, for example, due to preemption risk, and (3) "hot markets" where investments have high ex ante values (even without risk management) but left-tail risks. This confluence of organizational features, market structure, and market conditions is distinct from the determinants of risk management posited in theories centered on risk shifting or financial constraints, such as leverage, implicit bailouts, counterparty risk, or collateral. This should help to distinguish our theory empirically.

Another distinct aspect of our theory lies in its focus on the connection between risk management and market quality. Strategic complementarities amplify firm-level tensions between trade execution and risk management into market-wide tension between trading activity and risk allocation. In equilibrium, trading activity is inversely related to allocative efficiency. Causality runs both ways. Not only can heavier trading activity reduce risk management, but the avoidance of risk management boosts trading activity: Firms that reduce risk management trade more quickly and less selectively, both of which increase trading per time unit. Liquidity—in the sense of trading speed and volume—can thus be the cause and consequence of real inefficiency.

Since coordination failures are constrained inefficient, there is scope for regulation. Unlike in much of the existing literature on bank regulation, capital or liquidity requirements do not address the core problem. While capital structures affect the value that firms place on information about risk exposures, a *failure* to process such information arises, in our model, from externalities in the markets in which firms compete. Insofar as the externalities operate through fast and excessive trading, Pigovian taxes to moderate or decelerate trading can help but, in general, are a double-edged sword. Rather, our analysis strongly supports mandatory standards for the quality of risk management processes.[3] Such standards already play a role in the stress tests the Federal Reserve requires banks to pass, and in regulatory investigations of (the role played by risk management in) trading scandals.[4]

The approach adopted in this paper differs from existing theories in two ways. First, as mentioned, risk management comprises two broad functions: measuring risks and choosing which to hedge. Previous theories focus on the latter function, analyzing whether and how firms should hedge *known* risk exposures (Smith and Stulz, 1985; Froot, Scharfstein and Stein, 1993; Rampini and Viswanathan, 2010).[5] To use risk management jargon, they study a firm's optimal "risk appetite." Our theory focuses on the information process, or "risk assessment," that uncovers

[3]This is akin to the rationale for corporate governance regulation (Dicks, 2012; Raff, 2011; Cheng, 2011; Acharya, Pagano and Volpin, 2013), which recognizes that agency problems alone—leading to second-best outcomes—do not justify regulatory intervention. Relevant studies focus on pecuniary and information externalities in labor markets as the sources of constrained inefficiency in explaining the rationale for regulation.

[4]For example, after its "London Whale" trading loss of more than $6 billion, JP Morgan Chase was fined nearly $1 billion by U.S. and U.K. regulators for "unsafe and unsound practices" in overseeing its trading operations.

[5]In reduced form, the benefit of risk management is a private value (a firm-specific benefit of hedging idiosyncratic risk) of entering a financial contract that is traded at a common value (the market price of a hedging contract). (In earlier work, Stulz (1984) analyzes optimal hedging policies from the perspective of a risk-averse manager.) Rampini and Viswanathan (2010) argue that hedging should be subject to the same frictions as financing: collateral committed to hedging contracts reduces a firm's capacity to finance current investments. Our paper shares the focus on the costs of risk management, but the resource that firms commit to risk management is time and opportunity costs arise from preemption risk in financial markets.

unknown misalignments between a firm's actual risk exposure and its risk appetite, and prompts corrective action. Large losses due to overlooked misalignments is what we term risk management failures. Second, previous theories focus on financing frictions and agency problems as sources of inefficiency, and forgoing (some) risk management in equilibrium is second-best efficient. In our model, externalities are the root friction, and the equilibrium outcome can be constrained inefficient. For these reasons, our predictions about risk management failures pertain to the confluence of specific market and firm structures, rather than contracting frictions, and discussing the possible scope for regulation is more pertinent.[6]

Technically, our framework resembles bank-run models where the first-come-first-served rule embodies preemption (Diamond and Dybvig, 1983). This similarity allows us to use global games methods developed for bank-run equilibria (Goldstein and Pauzner, 2005) to obtain equilibrium uniqueness and hence sharp comparative statics predictions. An, Song and Zhang (2018) show that OTC markets with search frictions generate preemption motives for dealers, inducing them to build excessive exposure to risky assets. The analysis of preemptive competition is also central to the literature on innovation and patent races.[7]

2 Risk management in equilibrium

2.1 Baseline model

A mass M of risk-neutral firms (traders), indexed by k, competes for trading opportunities. Time is continuous, and a generic trading opportunity takes the form of a mispricing $\pi > 0$ that appears in the market at $t = 0$. For example, demand shocks to agents in segmented markets can create price discrepancies between assets with correlated cash flows, as in Shleifer and Vishny

[6]Links between competition and risk taking have been studied in the banking literature with a focus on the effects of competition on bank franchise values (Keeley, 1990; Hellmann, Murdock and Stiglitz, 2000; Boyd and De Nicoló, 2005; Martinez-Miera and Repullo, 2010) and returns to screening (Ruckes, 2004; Dell'Ariccia and Marquez, 2004). An alternative perspective is adopted in Parlour and Rajan (2001), where lenders to the same borrower exert negative externalities on each other by raising the borrower's overall default incentives.

[7]Most studies on this subject use sequential games or real options models in which strategic choices coincide with the acts of preemption (see, e.g., Leahy (1993), Caballero and Pindyck (1996), Grenadier (2002), and Aguerrevere (2009)). In contrast, strategic choices are made ex ante in our model. The work that is closest to ours is Askenazy, Thesmar and Thoenig (2006) where competing innovators choose ex ante between organizational designs that differ in production efficiency and time-to-market.

(1997) or Gromb and Vayanos (2002).

A friction interferes with traders' ability to instantly take advantage of such mispricing. The friction can take either of two forms. First, discovery can take time: a delay between the appearance of mispricing and the time at which a trader becomes aware of it can occur. Second, execution can take time: a delay between a trader's decision to trade and its physical execution can occur. Because these two frictions are equivalent in our framework, we model only the first one.[8] Specifically, traders discover the trading opportunity at random times that are identically and independently distributed according to an exponential distribution:

$$\tilde{t}_k \sim \text{Exp}(\lambda^{-1}).$$

The traders' discount factors are normalized to 1.

Upon locating an opportunity, a trader can request a trade (of one unit). A trade pays the sum of the mispricing π, which is a common value across traders, and a private value α_k. We interpret α_k as the "fit" between the trade and the risk profile of that particular trader, desk, or firm. For example, the firm may prefer trades that hedge rather than increase existing exposures, due to frictions that amplify the impact of negative cash-flow shocks. Importantly, we assume that such frictions reduce welfare and, accordingly, that α_ks represent social gains or losses (c.f. Section 4.3). Such losses may come from deadweight bankruptcy costs or capital constraints, as in Froot, Scharfstein and Stein (1993).[9] Similarly, α_k can reflect k's shadow cost of mobilizing collateral to guarantee positions, as in Rampini and Viswanathan (2010).

There is uncertainty about the private values. At the time of discovery, k merely knows that

$$\alpha_k = \begin{cases} \alpha_+ & \text{with probability } \rho \\ \alpha_- & \text{with probability } 1-\rho. \end{cases}$$

[8]More precisely, a specification of our framework in which all traders locate the trading opportunity immediately and execution introduces latency is isomorphic to one in which locating the opportunity introduces latency and execution is instantaneous.

[9]Among the most important operational risk management metrics are so-called concentration limits, the role of which is to ensure that the bank is not exposed too heavily to one particular idiosyncratic risk (Ernst & Young, 2013, 20).

The private values $\{\alpha_k\}_{k\in[0,M]}$ have a mean of zero and are independent across traders. We also assume $-\alpha_- > \pi > -\alpha_+$, that is, a trade is profitable if and only if $\alpha_k = \alpha_+$. We call risk management the process of producing information about α_k. Specifically, before $t = 0$, each firm simultaneously decides whether to activate a risk management technology. The technology investigates any requested trade and executes it only if $\alpha_k = \alpha_+$. Investigation takes a deterministic time ι, however, and hence delays execution.

The way we model risk management infuses the private values with two more interpretations. First, the randomness of α_k implies fundamental risk that risk management cares about. Mean-preserving spreads of the α_k-distribution could represent variation in this risk resulting from market or price volatilities. Second, the learning element implies that ex ante uncertainty about α_k also gauges the need for collecting information, or conversely, how much the risk is a priori unknown. (Indeed, the value of this information process in our model is a function of ρ and α_-.) In practice, this need is great(er) when risk-taking decisions are decentralized. Thus, the α_k-distribution also measures, in part, the degree to which a firm's organizational structure calls for internal risk coordination. We will return to these interpretations in Section 3.

Finally, the mispricing is sensitive to trading pressure: it disappears once the mass of trades exploiting this opportunity reaches I. For example, I could be the net order flow that eliminates differences in local demands across segmented markets, as in Kondor (2009).[10] The finite size of the trading opportunity creates preemptive competition between the traders, the intensity of which is captured by the ratio

$$i \equiv \frac{I}{M}.$$

The smaller this ratio is, the more intense is the competition. To focus on cases where the finite size of the trading opportunity generates concern about preemption, we restrict I to be strictly smaller than the mass of traders ρM for whom $\alpha_k = \alpha_+$, that is, $i < \rho$.[11] This restriction simplifies the exposition for now but is not essential: we later endogenize M by adding an entry

[10] In Section 4, we discuss model variants in which the mispricing decreases smoothly in the mass of trades, or depends directly on the time t since it appeared in the market.

[11] For $i < \rho$, the trading opportunity is exhausted in finite time in any equilibrium. The results are qualitatively the same if $\rho \leq i < 1$. For $i > 1$, time pressure disappears.

stage.

We conclude the description of this basic framework with a few remarks on our assumptions.

First, risk management blocks trades that create unwanted risks. In practice, some risks can be hedged. Allowing for hedges does not change the gist of our results as long as identifying risks and implementing hedges takes time. As a pre-trade process it would delay trade execution, while as a post-trade process it would leave firms exposed to the risks in the interim. The latter point also applies when those who ex post identify that they hold positions with α_--values can retrade with those for whom the positions (may) have α_+-values (c.f. Section 4.4).

Second, one can interpret risk management in our model in terms of position limits. Suppose a firm limits a trader's position to n_s units, but allows it to be expanded by n_r units, subject to risk management approval. The firm requires approval to assess whether the increased exposure is desirable. Our model is a reduced-form version of such a setting, with n_s and n_r normalized to 0 and 1, respectively.[12] According to this interpretation, firms in our model reduce risk management by relaxing position limits.

Third, activating risk management is an ex ante firm-level decision. The idea is that it requires organizational choices to collect and aggregate information from decentralized trading decisions.[13] In Section 4.4, we discuss a model variant in which firms can make risk management choices "on the fly" as trading opportunities are discovered. In a companion paper (Bouvard and Lee, 2019), we allow traders that operate within a firm's chosen risk management framework to manipulate the protocols.[14]

Fourth, the notion that material time pressure for traders accompanies competition in financial markets is widespread. Employee survey data from the U.S. Department of Labor's Occupational Information Network (O*Net) ranks professions on a scale of 1-100 along multiple dimensions. At 89, the average score of traders along dimensions reflecting exposure to competition, the impact of decisions on an organization, and time pressure, is the highest of any

[12] See Online Appendix for a formalization of this argument.

[13] The Basel Committee uses the term risk management *framework* to describe systematic risk controls in firms.

[14] There, we also generalize the insights of this paper to a setting in which new mispricings emerge in the market continuously over time.

profession and a statistical outlier.[15] Mehta et al. (2012) report that most major U.S. banks use historical simulations to assess risk because the Monte Carlo method, although known to provide a more comprehensive and precise picture of (tail) risks, is too time-consuming.

Finally, assuming a continuum of firms and i.i.d. random variables renders the model highly tractable: How many firms in total locate an opportunity over time and how many of those firms will have positive private values is deterministic and commonly known. In fact, the only aggregate uncertainty is strategic: To infer how many trades are executed over time, a firm must form beliefs about everyone else's risk management activity. In Section 3, we add exogenous aggregate uncertainty to eliminate equilibrium multiplicity and sharpen the predictions of the model.

2.2 Privately optimal risk management

Consider a trader who believes that trading opportunities stay alive for a period of length T, which we take as exogenous for now. In the absence of risk management, his expected profit is π conditional on locating the opportunity before T, which happens with probability

$$p_h(T) \equiv 1 - e^{-T/\lambda}.$$

We refer to this strategy as "hasty." This strategy is obviously irrelevant when $\pi < 0$. It is only when hasty trading is, on average, profitable that risk management incurs opportunity costs.

When engaging in risk management, the trader's expected profit is $\rho(\pi + \alpha_+)$ conditional on locating the opportunity and identifying α_k before T, which happens with probability

$$p_d(T) \equiv \begin{cases} 0 & \text{if } T < \iota, \\ 1 - e^{-(T-\iota)/\lambda} & \text{otherwise.} \end{cases}$$

We refer to this strategy as "deliberate." We also refer to either strategy as "implemented" once

[15] Source: O*Net Online, `http://www.onetonline.org/`. Traders are subsumed under the category *Sales Agents, Securities and Commodities.*

it is possible for the trader to execute the trade. For the hasty strategy, implementation amounts to locating the opportunity. For the deliberate strategy, it further requires identifying α_k, and does not entail execution if $\alpha_k = \alpha_-$.

The difference between the unconditional expected profits of the two strategies is

$$\begin{aligned} \Delta(T) &\equiv p_d(T)\rho(\pi + \alpha_+) - p_h(T)\pi \\ &= p_h(T)(1-\rho)|\pi + \alpha_-| - [p_h(T) - p_d(T)]\rho|\pi + \alpha_+|. \end{aligned} \tag{1}$$

$\Delta(.)$ represents a firm's private net value of risk management. On the bottom line, the first term reflects the benefit of risk management: avoiding bad trades that would occur with probability $p_h(T)(1-\rho)$ under the hasty strategy. This benefit depends on the implementation probability $p_h(.)$, which means that the value of risk management is contingent on the option of executing the trade. The second term reflects the cost of risk management: failing to capture good trades with probability $[p_h(T) - p_d(T)]\rho$ that would be executed under the hasty strategy. This opportunity cost depends on the difference between the implementation probabilities, $p_h(T) - p_d(T)$, and hence the relative speeds of the two strategies.

Since both implementation probabilities $p_h(.)$ and $p_d(.)$ are increasing functions, time pressure affects both hasty and deliberate traders negatively. The slope of $\Delta(.)$ is thus a priori ambiguous. For $T < \iota$, time pressure affects only the hasty strategy: the implementation probability is zero under the deliberate strategy, while it strictly increases with T under the hasty strategy. But for $T > \iota$, the deliberate strategy is more sensitive to time pressure than the hasty strategy for reasons related to both the benefit and the cost of risk management:

a. *Value of information.* Conditional on implementation, the deliberate strategy is more profitable than the hasty strategy. Hence, even if the implementation probabilities were to decrease equally, increasing time pressure would lower the unconditional expected profit on the margin to a greater extent under the deliberate strategy. The difference in conditional profits, $\rho(\pi + \alpha_+) - \pi = (1-\rho)|\pi + \alpha_-| > 0$, is precisely the conditional benefit of risk management: avoiding bad trades. The unconditional value of this benefit shrinks as the probability of trad-

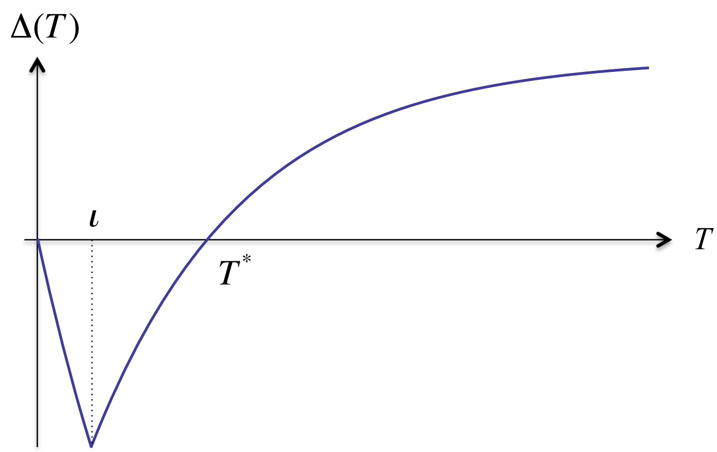

Figure 1: Time pressure and net private value of risk management.

T is the date at which the trading opportunity disappears; hence, it is an inverse measure of time pressure. ι is the risk management delay. $\Delta(T)$ is the firm's expected profit from trading while engaging in risk management relative to trading without risk management.

ing decreases with time pressure, indicating that the information value of risk management depends on execution.

b. *Value of time.* The implementation probability is more sensitive to changes in the deadline under the deliberate strategy than under the hasty strategy: $p_d'(T) > p_h'(T)$. To implement her strategy, a hasty trader must find the opportunity before T, whereas a deliberate trader must find it before $T - \iota$. A trader who can search from 0 to T gains less from a marginal increase in search time than one who can search only from 0 to $T - \iota$. Intuitively, additional time matters more to those who have less to begin with. So, the difference $p_h(.) - p_d(.)$ and hence the opportunity cost of risk management both increase with time pressure.

Thus, $\Delta(.)$ is U-shaped and reaches its minimum at $T = \iota$ (see Figure 1).[16] Furthermore, $\Delta(.)$ falls to 0 as $T \to 0$, indicating that no strategy can be implemented when there is no time. Conversely, when $T \to \infty$, $\Delta(T)$ converges to $(1-\rho)|\pi + \alpha_-| > 0$ because then implementation is certain, so the opportunity cost of risk management vanishes. In sum, these properties imply

[16] $\Delta(.)$ is similarly U-shaped when investigation time is random and exponentially distributed, as search time is. The non-monotonicity of $\Delta(.)$ will later complicate the global games refinement (Section 3).

that there is a *unique* point $T^* > 0$ at which a trader is indifferent between the two strategies: $\Delta(T^*) = 0$. When time pressure is high, i.e., $T < T^*$, it is optimal for a trader to abandon risk management and prioritize execution, i.e., $\Delta(T) < 0$. In contrast, when time pressure is low, i.e., $T > T^*$, the benefit of informed decision-making under risk management outweighs the loss in execution speed, i.e., $\Delta(T) > 0$.

Lemma 1. *The private value of risk management $\Delta(T)$ is strictly decreasing for $T < \iota$ and strictly increasing for $T > \iota$. Furthermore, there exists a unique threshold T^* such that firms activate risk management if and only if $T > T^*$, and T^* is an increasing function of π.*

The main implication of Lemma 1 contrasts with what follows from theories of risk management deficiencies based on risk-shifting arguments. Here, the choice to forgo risk management is not driven by differences in the financial claims held by various stakeholders (e.g., managers, shareholders, or debtholders), nor does it pit their interests against one another's. On the contrary, a firm forgoes risk management to stay competitive under time pressure and thereby maximizes firm value for all stakeholders.

Lemma 1 also establishes that T^* increases in π, that is, risk management is less attractive when trading is ex ante more attractive. This is already visible in (1): When π increases, losses from bad trades, $|\pi+\alpha_-|$, shrink, while forgone profits from good trades, $[p_h(T)-p_d(T)]\rho(\pi+\alpha_+)$, grow. Thus, the benefit of risk management decreases, while its opportunity cost increases.

By Lemma 1, a firm's best response to time pressure is given by

$$\Pr(h) = q(T) \equiv \begin{cases} 1 & \text{if } T < T^*(\pi) \\ [0,1] & \text{if } T = T^*(\pi) \\ 0 & \text{otherwise.} \end{cases}$$

Given a continuum of ex ante identical traders, q also denotes the fraction of hasty traders.

Notice that $q(T)$ is monotonic, which will be key to generating strategic complementarity. Also, T^* is independent of scale in the sense that multiplying the payoff $\pi + \tilde{\alpha}_k$ (by any positive scalar) does not change the sign of $\Delta(.)$—the costs (and benefits) of risk management scale up with the size of the trading opportunity. Last, any and all interaction between firms goes through T but every firm, being infinitesimal, takes T as given.

2.3 Collectively inefficient risk management

In equilibrium, the deadline T is endogenously determined. Let q denote the fraction of traders that play the hasty strategy. The time T by which the trading opportunity is exhausted satisfies

$$qp_h(T) + (1 - q)\rho p_d(T) = i. \tag{2}$$

Let $T(q)$ denote the solution to (2). $T(q)$ is decreasing: the deadline moves forward when the proportion of hasty traders grows. The reason is twofold. First, hasty traders implement more quickly because they avoid the risk management delay (execution speed). Second, conditional on locating the opportunity, rather than trading only when $\alpha = \alpha_+$, that is, with probability $\rho < 1$, hasty traders execute with probability 1 (indiscriminateness). As a result, the opportunity to trade is depleted at a faster rate.

We let

$$T_h \equiv T(1) \qquad \text{and} \qquad T_d \equiv T(0)$$

denote, respectively, the nearest possible deadline (all traders are hasty), and the farthest possible deadline (all traders are deliberate). T_h and T_d bound the range of deadlines that can arise in the market.

These bounds can be used to identify parameter regions in which strategic dominance arises. We know from Lemma 1 that the threshold T^* below which the private value of risk management is negative increases with π. For high enough π, there may be an equilibrium in which all traders are hasty (hereafter, hasty equilibrium). Indeed, let $\overline{\pi}$ be defined by $T^*(\overline{\pi}) = T_d$. If $\pi > \overline{\pi}$, then $T_d < T^*(\pi)$: Even if everyone else were to be deliberate, T_d would be lower than the threshold T^* below which the hasty strategy is the best response, making it strictly dominant for any trader to be hasty. An analogous argument applies to low π. If $\pi < \underline{\pi}$, where $T^*(\underline{\pi}) = T_h$ if $T_h > \iota$ or else

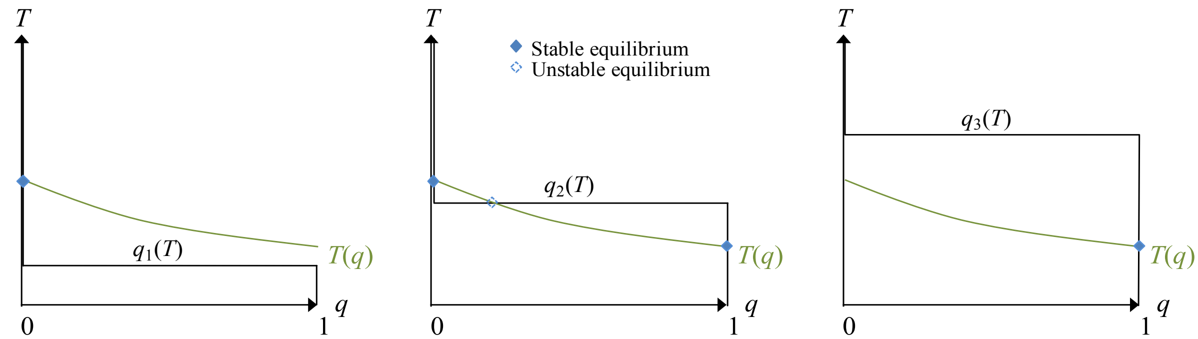

Figure 2: Strategic complementarities in risk management.

$T(q)$ depicts how time pressure depends on firms' risk management choices, while the step function $q(T)$ depicts firms' optimal risk management choice as a function of time pressure for various levels of the common value π (i.e., the ex ante value of a trade). Higher $q(T)$ correspond to higher values of π. Intersections between $q(T)$ and $T(q)$ constitute equilibria.

$\underline{\pi} = 0$, being deliberate is strictly dominant for all traders (hereafter, deliberate equilibrium).

When π is in the intermediate region $(\underline{\pi}, \overline{\pi})$, the equilibrium prediction is ambiguous: $q(T)$ and $T(q)$ intersect at three (fixed) points. Suppose a trader believes everyone else is hasty. The implied deadline T_h then occurs sooner than $T^*(\pi)$ so that his best response is also to be hasty. Thus, a hasty equilibrium exists. If the trader instead believes that everyone else is deliberate, the implied deadline T_d occurs later than $T^*(\pi)$ and supports a deliberate equilibrium. In the third equilibrium, the fraction of hasty traders q is precisely such that the implied deadline renders everyone indifferent between the two strategies, i.e., $T(q) = T^*(\pi)$ or $\Delta(T(q)) = 0$. The various equilibrium configurations are illustrated in Figure 2.

Proposition 1. *There exists a unique interval* $[\underline{\pi}, \overline{\pi}] \neq \emptyset$ *such that the hasty equilibrium exists for all* $\pi \geq \underline{\pi}$*, the deliberate equilibrium exists for all* $\pi \leq \overline{\pi}$*, and one equilibrium in which both strategies are used exists for all* $\pi \in (\underline{\pi}, \overline{\pi})$*. There are no other equilibria.*

The source of equilibrium multiplicity is that $T(q)$ and $q(T)$ are both decreasing, which creates a positive feedback loop: When more traders are deliberate, the opportunity to trade is depleted at a slower pace as risk management delays execution and blocks executions whenever $\alpha_k = \alpha_-$. This lowers the time pressure on others, making them more inclined to be deliberate as well. Conversely, greater hastiness increases time pressure, which begets more hastiness.[17]

This strategic complementarity also makes the interior equilibrium unstable: Any shock that causes $T(q)$ to deviate from T^* breaks the indifference condition and pushes all traders to adopt the same strategy.

Corollary 1. *Only corner equilibria are stable.*

In every equilibrium, payoffs are symmetric across traders. Because the aggregate payoff decreases in the mass of executed trades with $\alpha_k = \alpha_-$, we obtain a simple Pareto ranking:

Corollary 2. *Everyone being deliberate strictly Pareto-dominates any non-deliberate equilibrium.*

That is, the social optimum in our model is achieved when every trader is deliberate, while any non-deliberate equilibrium is a coordination failure – a "race to the bottom." Traders can be trapped in the hasty equilibrium because their individual objective functions fail to internalize the social value of risk management. Consider trader k's net gain from deviating to the deliberate strategy when everyone else is hasty:

$$\Delta(T_h) = p_d(T_h)(1-\rho)|\pi + \alpha_-| - [p_h(T_h) - p_d(T_h)]\pi. \tag{3}$$

The second term captures the preemption motive: By switching to the deliberate strategy, the trader becomes less likely to capture the common value π. This private loss to k is not a social loss, as another trader will capture the common value in lieu of k. The first term corresponds to the private benefit of risk management: Under the deliberate strategy, k can avoid a loss of $|\pi + \alpha_-|$. But the social benefit of risk management in this event further depends on the private value of the (hasty) trader k' who then executes in lieu of k: The social gain is zero if $\alpha_{k'} = \alpha_-$ but otherwise $\alpha_+ - \alpha_-$, which happens with probability ρ. Since $\rho(\alpha_+ - \alpha_-) > |\pi + \alpha_-|$, trader k does not fully internalize the allocative efficiency gain of risk management.

Proposition 1 and its corollaries are noteworthy for two main reasons: Most theories of risk management problems focus on the moral hazard posed by agents who do not fully bear

[17] Compared with bank runs or currency attacks, the strategic complementarities in our model operate through learning incentives as opposed to direct payoff externalities: by choosing to be hasty, traders raise the opportunity cost of acquiring information for other traders.

potential losses. By contrast, such externalization of losses does not drive our results. Hence, as a matter of positive theory, the key predictions about risk management incentives do not pertain to executive compensation, capital structure, bailout policies, financial contagion, or other sources of risk-shifting incentives. Rather, our analysis implicates factors such as competition, market structure and conditions, and speed of information processing as important determinants of risk management quality – thereby pointing to a novel, alternative set of potentially testable empirical relationships.

From a normative perspective, it is worth noting that in our model forgoing risk management does not constitute a second-best outcome. Instead, any non-deliberate equilibrium is constrained Pareto-inefficient, as a result of coordination failure, and thus creates scope for Pareto-improving regulation. What is more, such regulation may have to target (some of) the determinants of risk management quality mentioned above, or impose penalties on firms for risk management failures—even if none of a firm's financial losses is externalized. This is because, here, the externalities of risk management choices operate through competitive market pressure.

3 Co-determinants of risk management failures

While equilibrium multiplicity highlights strategic complementarity, it generates ambiguity about the influence of structural parameters on firms' risk management choices. In this section we use global games techniques to resolve this indeterminacy (Carlsson and van Damme, 1993), and with respect to the resultant unique equilibrium, explore how parameter changes impact the level of risk management. The main purpose of these comparative statics is to outline the set of circumstances under which risk management failures are likely to be observed according to our theory, and also the predicted relationship between financial market characteristics and risk management quality.

3.1 Global Games treatment

Standard global games techniques are applicable in settings that feature *global* strategic complementarity (see, e.g., Morris and Shin, 2003). In our model, this would mean that increasing the fraction of deliberate traders always raises the private value of risk management. But this is not the case: While T is monotonically increasing in the fraction of deliberate traders, $\Delta(.)$ is not monotonic in T. Our model does, however, satisfy a weaker form of strategic complementarity: By Lemma 1, $\Delta(.)$ crosses 0 once, and is monotonic when positive. These properties define *one-sided* strategic complementarity, as in Goldstein and Pauzner (2005), whose method we use to obtain equilibrium uniqueness in our setting.

To this end, we add aggregate uncertainty in the model by assuming that the common value π, instead of being a fixed parameter, is a random variable uniformly distributed over $(\hat{\pi}-\delta, \hat{\pi}+\delta)$, about the realization of which traders have dispersed information.[18] Specifically, before deciding whether to activate risk management, each privately observes a noisy signal, $s_k \equiv \pi + \xi_k$, where $\{\xi_k\}_{k\in[0,1]}$ are uniformly and independently distributed on $[-\varepsilon, +\varepsilon]$. Therefore, this information structure keeps traders from knowing precisely what others know and thereby from perfectly coordinating on one strategy.

The equilibrium derivation is similar to Goldstein and Pauzner (2005). Equilibrium strategies take the threshold form typical of global games: Every trader is hasty if his signal s_k lies below a unique common threshold s^* and deliberate otherwise. Because the signals are heterogeneous, hasty and deliberate traders generically coexist. It is, however, characteristic of global games that equilibrium uniqueness is preserved even in information structures that are arbitrarily close to common knowledge of π. That is, even as the signal noise vanishes, i.e., $\varepsilon \to 0$, the equilibrium remains unique although everyone converges on the same strategy. In the following proposition, we denote the net value of risk management as a two-variable function $\Delta(T, \pi)$, with slight abuse of notation, to highlight the role of π.

Proposition 2. *For all $\varepsilon > 0$, a unique equilibrium exists. For $\varepsilon \to 0$, the equilibrium converges to: All traders play the hasty strategy if $\pi > \pi^*$ and the deliberate strategy if $\pi < \pi^*$, where π^**

[18]As is usual in global games, we assume that δ is large enough to guarantee the existence of lower- and higher-dominance regions, i.e., $\delta > \max\{\hat{\pi} + \alpha_+, \alpha_- - \hat{\pi}\}$. This allows for realizations of π as in Section 2 where $-\alpha_+ < \pi < -\alpha_-$ and under symmetric information, multiple equilibria may coexist, as well as realization of π where $\pi > \alpha_-$ or $\pi < -\alpha_+$ where the symmetric-information equilibrium is unique.

is strictly positive and satisfies

$$\int_0^1 \Delta[T(q), \pi^*] dq = 0. \tag{4}$$

Proposition 2 is the counterpart of Proposition 1 in the richer environment of global games. We build on this prediction to explore the determinants of risk management failure. Because the likelihood that a firm plays the deliberate strategy increases with π^*, we henceforth refer to π^* as "risk management quality."

3.2 Comparative statics

We now describe how π^*, as defined in (4), depends on exogenous model parameters. We begin with risk parameters that determine the ex ante value of a trade and the value of the information on which a trade is conditioned if risk management is active. We then turn to parameters that directly affect the time pressure a firm experiences, exploring their interaction with the risk parameters as well. Last, we endogenize traders' entry and analyze how entry costs affect both market and risk management quality.

Risk moments. A firm's incentives to run risk management depend, first and foremost, on the structure of its potential trading risk, as this determines the value of monitoring (i.e., learning about) its risk exposure. In our model, the payoff from a trade has a mean equal to the common value π and higher-order moments given by the distribution of the private value α_k. To capture the latter moments, we parametrize the private values as $\alpha_+ \equiv \frac{\sigma}{\rho}$ and $\alpha_- \equiv \frac{-\sigma}{1-\rho}$ with $\sigma > 0$. This parametrization preserves the property that the private value has a mean $\mathbb{E}(\alpha_k) = \rho\frac{\sigma}{\rho} + (1-\rho)\frac{-\sigma}{1-\rho}$ equal to 0 for any values of σ and ρ, and it yields simple expressions for variance $\rho\left(\frac{\sigma}{\rho}\right)^2 + (1-\rho)\left(\frac{-\sigma}{1-\rho}\right)^2 = \sigma^2$ and skewness $\rho\left(\frac{1}{\rho}\right)^3 + (1-\rho)\left(\frac{-1}{1-\rho}\right)^3 = \rho^2 - (1-\rho)^2 = 2\rho - 1$. A high σ means high volatility, while a high ρ means that the probability of an undesirable outcome is small but its realization is costly. Conveniently, a change in π, σ, and ρ affects, respectively, the mean, variance, and skewness of the payoff distribution while leaving the other moments unchanged. The next result is a direct consequence of Proposition 2.

Proposition 3 (Risk moments)**.** *Risk management quality, ceteris paribus,* *a. decreases for high realizations of the mean π (common value).*

b. increases with volatility σ and skewness ρ (private value).

Part b. of Proposition 3 indicates that traders partly internalize the benefit of risk management. Active risk management provides downside protection by preventing trades with negative private value. Under the above parametrization, this information value of risk management—potentially preventing bad trades—is $\mathbf{I} \equiv (1-\rho)|\pi + \alpha_-| = \sigma - (1-\rho)\pi$ (see Section 2.2, page 11), which increases in σ and ρ. That is, like a put option, the information generated by risk management is more valuable when volatility or left-tail risk increases. Along this dimension, private incentives align with welfare maximization: risk management creates a social surplus equal to $\rho(\alpha_+ - \alpha_-) = \frac{\sigma}{1-\rho}$, which also increases in σ and ρ.

Surplus is independent of π, however, because the preemption game between traders is zero-sum in the common value dimension. Hence, part a. of Proposition 3 captures the coordination failure across traders. When the common value π is higher, the opportunity cost of missing out on the trading opportunity rises. Because traders compete on speed for this common value, higher realizations of π exacerbate preemption motives, which may then dominate the risk considerations mentioned in b., leading traders to give up risk management and increase time pressure on all other traders.

This suggests two alternative reasons for lack of risk monitoring. For firms with low σ and ρ, risks are of little concern or already known, so they see no need for costly, disruptive protocols. By the same token, such firms are unlikely to suffer losses from "overlooked" risks that would in hindsight be classified as risk management failures. Alternatively, firms with high σ and ρ may neglect risk monitoring to accelerate trading. This scenario requires high π such that faster trading without risk monitoring is attractive enough to override the risk concerns.

Empirically, under which circumstances should we observe severe risk management failures? The answer is somewhat nuanced because, while higher π undermine risk management, they also mitigate the potential damage to a firm. To see this, note that the information value of risk management to an individual firm, $\mathbf{I} = \sigma - (1-\rho)\pi$, decreases in π. In fact, for π such that $\mathbf{I}$ is no longer positive, the information is worthless because the firm finds the trade attractive regardless. Socially, this is still inefficient inasmuch as, if $\alpha_k = \alpha_-$ for this firm, another is better-suited to bear that risk. Thus, there can be misallocation of risk even in the absence of significant individual losses. However, when $\rho \to 1$, π gradually ceases to affect the information value of risk management but still affects the opportunity cost of risk management, $[p_h(T) - p_d(T)]\rho|\pi + \alpha_+|$. Hence, given enough *skewness* in the risk, a fraction of firms in equilibrium is hit by large losses for high realizations of π. So, severe failures occur (only) in firms that, despite generally placing a lot of value on information about their risk exposure, wind up vying in "hot" markets for trades that involve low-probability but "catastrophic" left-tail risks.

This has the counterintuitive implication that better market conditions can harm all traders. A trader's expected payoff in the equilibrium defined by (4) is

$$\begin{aligned} & \Pr(\pi \geq \pi^*) i \mathbb{E}(\pi | \pi \geq \pi^*) + \Pr(\pi < \pi^*) i \mathbb{E}\left(\rho\pi + \sigma | \pi < \pi^*\right) \\ = \quad & i \left\{ \mathbb{E}(\pi) + \Pr(\pi < \pi^*) \mathbb{E}\left(\mathbf{I} | \pi < \pi^*\right) \right\}. \end{aligned} \quad (5)$$

On the bottom line, $\mathbb{E}(\pi)$ is the unconditional value of a trade, and $\mathbb{E}\left(\mathbf{I}|\pi < \pi^*\right)$ is the expected information value of risk management over the range of π for which risk management is active. Consider a rightward shift of the π-distribution.[19] This increases $\mathbb{E}(\pi)$ but also reduces the probability $\Pr(\pi < \pi^*)$ that risk monitoring takes place. If $\mathbf{I}$ is high as a result of large σ and ρ, the internal loss of information can outweigh the external gain in profitability, so that the overall effect is negative.[20] Intuitively, even the firms themselves may feel that a market "heats" up too much for their own good, wary that the pressure to run along with a market-wide rush might cause them to overlook critical risks.

Time pressure. The remaining model parameters—the size I of the trading opportunity, the mass M of traders, the average search time λ, and risk management delay ι—can be collapsed into two ratios: an (inverse) measure $i \equiv \frac{I}{M}$ of competition, and an external-internal speed ratio $\nu \equiv \frac{1/\lambda}{1/\iota} = \frac{\iota}{\lambda}$, which reflects how fast search is relative to risk management. The ratios highlight

[19] I.e., let $\pi' = \pi + \delta$ where δ is a constant such that the support of π' still satisfies the conditions specified in Section 3.1. Note that this change does not affect the threshold π^*.

[20] Since π has support on $[\hat{\pi} - \delta, \hat{\pi} + \delta]$, a rightward shift in the distribution is an increase in $\hat{\pi}$ that keeps the constraint $\delta > \max\{\hat{\pi} + \alpha_+, \alpha_- - \hat{\pi}\}$ satisfied (see footnote 18). This always is possible if δ is high enough.

that all underlying parameters, essentially, affect the time pressure faced by a firm.

Proposition 4 (Time pressure)**.** *Risk management quality, ceteris paribus,*

a. decreases when external speed increases relative to internal speed, i.e., when ν increases.

b. decreases when competition intensifies, i.e., when i decreases.

Part a. of Proposition 4 emphasizes that two types of speed matter. An increase in search speed λ raises the probability that a trader finds an opportunity before any deadline $T(q)$, but does so also for everyone else, which brings all deadlines $\{T(q)\}_{q\in[0,1]}$ forward. These effects offset each other so that, under hasty strategies, no one gains or loses. The increase in λ does, however, raise the opportunity cost of risk management: The opportunity is likelier to close between $\tilde{t}_k$, when a trader discovers it, and $\tilde{t}_k + \iota$, when he can capture it under the deliberate strategy. This decreases the implementation probability of the deliberate strategy relative to that of the hasty strategy. As a result, π^* must decrease for indifference condition (4) to hold. Hence, what matters is not speed in general, but the delay risk management imposes *relative* to the time it takes traders to locate opportunities.

Part b. of Proposition 4 says that more intense competition raises the chances of coordination failure. As explained earlier, the marginal trader, who receives signal $s^* = \pi^*$, forms posterior beliefs about the distribution of deadlines $T(q)$, which spans $[T_h(i), T_d(i)]$ and includes T^*. By (2), when i decreases, $T(q)$ shifts down for all q, that is, the deadline moves forward for every realization of q. This in turn shifts more probability mass into the region below T^* where the marginal trader prefers being hasty (i.e., $\Delta[T(q), \pi^*] < 0$), so that the integral in (4) turns negative and π^* must decrease to maintain indifference.[21]

This set of comparative statics highlights aspects that most strongly distinguish our theory from others in the existing literature. Rather than focusing on a firm's liabilities and investments, the main focus of most theories of risk management, they are concerned with technological or institutional aspects. For example, advances in information technology affect ι and λ. According

[21] The fact that $\Delta(\pi^*, .)$ is negative below T^* and positive otherwise drives the monotonicity of $\pi^*(i)$. The proof cannot rely solely on this observation, however, because $\Delta(., \pi^*)$ is non-monotonic below T^* (see Appendix).

to our theory, such advances are a double-edged sword: even in a world where information systems improve at identifying risks quickly, risk management need not improve in the firms to which this matters—those competing on speed to seize trading opportunities—as trading also becomes faster. Technological progress can also affect informational barriers or costs of entry, thereby affecting competition.[22] Institutional features influence speed, too. For example, ι likely increases in the size and scope of the firm to which a trader belongs, because reconciling risks is harder across a more complex organization, while λ, the trader's ability to discover mispricings, and i, the intensity of search competition, depend on market features such as transparency, market segmentation, and regulation.

Interaction. Empirically, one may want to run a "horse race" between risk-moment variables (liabilities, assets, and investment opportunities) and time-pressure variables (competition, technology, and institutions) to see which has greater explanatory power for risk management quality. But importantly, the effects of these forces are not independent of each other, as the following cross-derivatives show:

Proposition 5 (Interaction)**.** *An increase in the speed ratio ν or in competition intensity i^{-1} weakens the impact of volatility σ and skewness ρ on risk management quality. That is,*

$$\frac{\partial \pi^*}{\partial \sigma \partial \nu} < 0,\ \frac{\partial \pi^*}{\partial \rho \partial \nu} < 0,\ \frac{\partial \pi^*}{\partial \sigma \partial i} > 0 \text{ and } \frac{\partial \pi^*}{\partial \rho \partial i} > 0.$$

In other words, an increase in external (search) speed relative to internal (risk-management) speed, or an increase in competition intensity, not only has the direct effect on risk management described in Proposition 4 but also makes traders less responsive to changes in the riskiness of their trades. That is, in an environment where time pressure is high, a rise in volatility or tail risk is less likely to cause a transition from a hasty equilibrium to a deliberate one. To see why,

[22] We explicitly analyze entry costs further below.

write out the net value of risk management for a given π using the σ-ρ-parametrization:

$$\begin{aligned}\Delta(T) &= p_d(T)(\rho\pi + \sigma) - p_h(T)\pi \\ &= p_d(T)\mathbf{I} - [p_h(T) - p_d(T)]\pi. \end{aligned} \tag{6}$$

To make our point cleanly, we give the last line (6) a slightly different structure than we used in equation (1) of our analysis in Section 2.2. Recall that Δ is the net benefit of risk management, that is, of switching from the hasty strategy to the deliberate strategy. The second term isolates the effect of a pure *loss in execution speed*: the decrease $-[p_h(T)-p_d(T)]$ in execution probability times the foregone trade value π. The first term adds to this the *gain in execution quality*, the information value $\mathbf{I}$, at the lower execution probability $p_d(T)$. Now note how the weights of these effects shift as $T \to \iota$: On the one hand, $p_d(T) \to 0$, so the influence of $\mathbf{I}$ vanishes. On the other hand, $p_h(T) - p_d(T)$ grows, and with it the influence of π.[23] Thus, time pressure regulates the importance of information value relative to that of opportunity costs in risk management decisions.

Proposition 5 thus describes a fundamental shift: If technology or institutional factors cause speed to matter more, common-value *market conditions* π (part b. of Proposition 3) will supplant private-value *risk concerns* σ and ρ (part a. of Proposition 3) as firms' main consideration in setting risk management quality. This point is clearest at the extremes. For $T \to \infty$, $\Delta(T) = \mathbf{I}$. Here, only the information value of risk management counts, and that depends on σ and ρ. In contrast, for $T \leq \iota$, $\Delta(T) = -[p_h(T) - p_d(T)]\pi$. Here, the information value is irrelevant. Instead, firms' decisions to activate risk management depend exclusively on whether $\pi > 0$ or $\pi \leq 0$. For empirical purposes, this offers a much more specific, nuanced prediction regarding the interaction of time pressure and risk moments than the simple claim that both can matter.

Endogenous entry. For our last comparative statics result, we endogenize the mass of traders to describe how entry barriers affect the quality both of the market and of firms' risk management. Suppose traders must pay a fixed cost χ to enter the market. In equilibrium, competition intensity

[23] In essence, these two effects represent points a. and b. in Section 2.2.

i^{-1} must then satisfy the zero-profit condition

$$i\{\mathbb{E}(\pi) + \Pr[\pi < \pi^*(i)]\alpha_+\} = \chi. \tag{7}$$

Equilibrium is now determined by linking zero-profit condition (7) to indifference condition (4), jointly pinning down competition intensity i^{-1} and equilibrium threshold $\pi^*(i)$.

Corollary 3. *A decrease in the entry cost χ increases the mass of traders, moves the expected deadline forward, and decreases risk management quality.*

Lower entry costs mechanically raise market speed: With a larger mass of active traders, more trades occur in any time interval. In addition, there is an equilibrium effect: A faster market lowers the threshold π^* above which traders forgo risk management, which further accelerates the market as fewer trades are delayed or blocked. Thus, market speed is *inversely* related to allocative efficiency, and owing to the feedback loop, both *cause* and *consequence* of inefficiency. Corollary 3 also shows that the failure to coordinate on the Pareto-dominant equilibrium exists even in an environment where traders make zero-profit ex ante.

According to our theory, the pathology of risk management failure involves multiple factors—market conditions, market structure, competition, organizational structure, and technology—inducing in concert (1) "hot" markets, (2) time-based competition, and (3) risk management demands that stall trading. While the specific combination of factors may vary from case to case, the same basic intuition applies: Because severe risk management failures transpire only when risks are of significant concern but this is also precisely when (the information generated by) risk management is valuable, the pursuit of trading opportunities must be sufficiently attractive and risk management sufficiently disruptive to override those concerns.

3.3 Discussion

Some precursors of risk management failures identified in the previous section can be related to structural shifts in financial markets and financial institutions induced by deregulation and technical changes (c.f., Kroszner and Strahan, 1999; Rajan, 2006).

In the U.S., the deregulation of the banking industry since the 1970s has progressively allowed banks once confined within state borders to compete across states (Black and Strahan, 2002). Progress in communication and information technology has, in addition, lowered organizational and informational barriers to entry: it has made it easier for banks to operate farther from their bases (Petersen and Rajan, 2002; Alessandrini, Presbitero and Zazzaro, 2008) and reduced local information advantages, enabling them increasingly to rely on "hard" information (Rajan, Seru and Vig, 2015). These changes have effectively lowered the cost of entry into local markets, thus increasing competition.

This in turn has spurred the growth of the most successful banks into enormous institutions that hold a disproportionate share of total banking assets in complex organizational structures (Black and Strahan, 2002; Cetorelli, McAndrews and Traina, 2014). As heightened competition shrunk margins in the traditional deposit business, these banks also shifted toward new business models and lines (e.g., securitization, structured finance), assuming risks that are harder to assess and hedge (Rajan, 2006; Acharya, Pagano and Volpin, 2013).[24] The burden on risk management grew accordingly. Indeed, a 2008 report issued by the Banking Supervision and Regulation arm of the Federal Reserve Board opened with, *"In recent years, banking organizations have greatly expanded the scope, complexity, and global nature of their business activities. . . As a result, organizations have confronted significant risk management and corporate governance challenges, particularly with respect to compliance risks that transcend business lines, legal entities, and jurisdictions of operation."*[25] These challenges have a significant time aspect. A 2013 survey of major U.S. banks finds that comprehensive stress tests take several months to complete and that the delay is a barrier to using the tests as an effective management tool (Ernst & Young, 2013).[26]

At the same time, technological progress has enabled banks to identify and exploit trading opportunities more quickly (Frame and White, 2004). In a 2009 speech on financial innovation,

[24] Consistent with this, Demsetz and Strahan (1997) show that larger banks, whose diversified activities might be expected to lower their total risk, have in fact more volatile earnings than smaller, less diversified banks.

[25] Fed Board Report on "Compliance Risk Management Programs and Oversight at Large Banking Organizations with Complex Compliance Profiles," October 16, 2008. This is echoed in the closing words of a Bloomberg article on Februrary 23, 2016: *"the real lesson of the London Whale [risk management scandal] is that megabanks such as JPMorgan are not only too big to fail – they may also be too big to manage and too big to regulate."*

[26] Relatedly, a 2012 McKinsey survey of major U.S. banks reports that running a value at risk on a *single* typical trading portfolio takes two to fifteen hours and potentially much longer in stressed markets (Mehta et al., 2012).

Ben Bernanke notes that credit-scoring models brought *"ever-faster evaluation of creditworthiness, identification of prospective borrowers, and management of existing accounts."*[27] Zhu (2018) provides direct evidence of a link between new technologies and the speed at which information about future earnings is incorporated into asset prices. Technological progress notwithstanding, the speed of risk management may thus have decreased *relative* to the speed of markets.

Overall, this transformation of the U.S. banking industry features two of the factors that favor equilibria with low risk management in our model: increased competition (i) due to lower entry costs (χ), and a faster progression of search speed relative to risk management speed (ν). These factors not only favor hasty strategies, but also, as implied by Proposition 5, reduce responsiveness to risk factors, such as: leverage, which amplifies volatility (σ), or maturity mismatch, which can exacerbate downside shocks (ρ) via liquidity shortages and runs. The aforementioned structural changes could therefore explain why broker-dealers which, before 2008, faced lighter constraints on leverage than commercial banks and relied on overnight short-term funding, which is much less stable than deposits, had not set up risk management protocols that were commensurate with the fragility their capital structure implied and that could have detected critical risk accumulations early enough.

Finally, in our theory, market conditions play a key role. This accords with a survey of chief risk officers, more than half of whom cited market conditions as an important factor, noting that the tendency to take on more risk increases as markets pick up (Ernst & Young, 2013, 18). The market for mortgage-based securities (MBS), in which the 2008 crisis originated, was booming in the years prior: origination by the top 15 underwriters in the U.S. doubled from 2004 to 2006.[28] Coval, Jurek and Stafford (2009) argue that the boom was fueled by trading in MBS at yields that did not fully reflect their risk. According to our theory, this mispricing created opportunity costs of risk management for those racing to capture resultant market opportunities. Two of the most aggressive firms, the top MBS underwriters from 2004 to 2007 that were heavily exposed to MBS in the end, were Lehman Brothers and Bear Stearns, which failed amid cries of risk management failure.

[27] Speech a the Federal Reserve System's Sixth Biennial Community Affairs Research Conference, April 17, 2009.

[28] See "The Tip of the Iceberg: JP Morgan Chase and Bear Stearns (A)," HBS case study 9-309-001.

Our model addresses neither the source of such booms (e.g., capital flows, saving glut) nor the type of asset involved (e.g., MBS), but it offers a rationale that explains why financial institutions may systematically fail to assess and contain the risks associated with these opportunities. This defines a notion of systemic risk that differs from the concern that the financial system may be overly exposed to a handful of large institutions deemed "too-big-to-fail" or "too-interconnected-to-fail." While size can play a role in making risk management more complex and time-consuming, as mentioned earlier, the systemic nature of risk management failures in our model stems from the fact that all participants in the same market compete under time pressure and reinforce it by weakening risk controls. That is, common market forces can distort (from a social perspective) the strategies of all traders, which in turn leads to a suboptimal allocation of risks across the entire market.

Beyond the considerations discussed above, the proposed explanation of risk management failures—firms with complex and decentralized risk-taking, time-based competition, and boom markets—could be taken to the data. It may be useful to clinical "postmortem" studies of institutional risk management failures and, more broadly, it might shed light on the relative frequency of risk management failures over time or across industries.

4 Extensions

We examine here the robustness of our model and conclusions, in terms of both predictions and welfare implications, to alternative interpretations and specifications.

4.1 Social value of speed

Our main analysis has focused only on coordination failure between traders, for whom a hasty equilibrium produces a collective welfare loss $\rho(\alpha_+ - \alpha_-)$ owing to inefficient risk management. Taking a broader view, it is possible that the speed at which the trading opportunity is exploited has social value. Whether this is the case depends on the nature of the trading opportunity.

Our primary interpretation of π regards it as a mispricing in the informational sense, that is, a discrepancy between the market prices of assets and their true values. In this case, our model captures informed traders who profit at the expense of other (unmodelled) traders. Restricting attention to these sets of agents, the trading profit π is a zero-sum transfer, and the only welfare loss stems from inefficient risk management among informed traders. Informativeness of asset prices can, however, entail broader effects on real decisions (Bond, Edmans and Goldstein, 2012). Insofar as these effects depend on the *speed* at which prices incorporate private information, our results imply a tension between allocative efficiency across traders and informational efficiency.

Alternatively, π can be the product of gains from trade. For example, traders could function as intermediaries who provide liquidity to agents and capture part of the surplus created in a transaction. In this case, speed is valuable insofar as liquidity demanders value immediacy (e.g., Grossman and Miller, 1988), as they may have to exit the market or default if their demand is not met quickly enough. Here, our model implies greater misallocation of risk across intermediaries in equilibria where they provide greater immediacy.

One parsimonious model extension that adds such a tension is to let the lifetime of a trading opportunity depend not only on trading but also *directly* on time. Formally, this can be done by adding a probability θ that the opportunity vanishes within any small interval of time dt before it is exhausted. In this setup, which we analyze in the Online Appendix, risk management need not be Pareto-optimal even for traders, because they partly internalize the social benefit of fast execution (by seizing π). Generally, however, concerns about coordination failure persist. As in the baseline model, there exist parameter regions with multiple equilibria—and in those regions, the equilibria without risk management are Pareto-dominated.

This extension also delivers volume implications. In expectation, more trades will be executed in a hasty equilibrium than in a deliberate one. This implies an ambiguous relationship between trading volume and welfare, as higher trading volume reflects the social benefits of faster execution and social costs of reduced risk management. This too exemplifies the point that richer versions of our model where speed has social value capture tensions between distinct functions of financial markets, such as between risk allocation on the one hand and information efficiency or liquidity provision on the other.

4.2 Sources of strategic complementarity

In our model, coordination failures are driven by strategic complementarities in risk management, which arise through the following loop: Less aggregate risk management increases time pressure, which in turn increases the opportunity cost of delaying execution to manage risk. Now, suppose that risk management does not create delays ($\iota = 0$), but firms must instead pay a fixed cost k to set up or maintain a risk management system. Then a firm chooses to practice risk management if and only if

$$p_h(T)\rho(\pi + \alpha_+) - k > p_h(T)\pi.$$

With risk management now assumed to be instantaneous, the implementation probability under either strategy is $p_h(T)$. The above inequality can be rewritten as

$$p_h(T)(1 - \rho)|\pi + \alpha_-| > k, \tag{8}$$

which, as $p_h(.)$ is strictly increasing and $\pi + \alpha_- < 0$, shows that whether (8) holds depends on T. This is a scale effect: If T and thus $p_h(T)$ shrink, the benefit of risk management—which materializes only if the firm is able to trade—becomes too small relative to cost k. Interestingly, T continues to depend on risk management choices even for $\iota = 0$, that is, even if there is no delay imposed by risk management. This is because hasty firms trade less selectively and thus more often, so that a trading opportunity disappears more quickly. This closes the loop and can sustain the same type of equilibrium multiplicity and coordination failure as in the original model.

This model variation highlights that the key to generating coordination failure is that a trader who reduces risk management increases time pressure for every other trader in the market. The mechanism through which time pressure may in turn make risk management unattractive can operate through any of several channels: time cost (as in the original model), fixed cost (as above), or a combination thereof.[29]

[29] These costs could be related. Investments in technology (fixed costs) could reduce latency (opportunity costs). Moreover, firms could invest in search speed (reducing λ) or in risk management speed (reducing ι). In the Online Appendix, we explain that this can lead to an arms race in speed investment that can exacerbate the race to the

Yet, there are notable differences between these types of costs. The scale effect—investments to investigate trades do not pay off if trades are unlikely to be executed—loses bite when the fixed costs can be spread over many trades across a firm or across time, or for trades of substantial size. Moreover, it applies to ex ante costs but not to interim costs incurred only after an opportunity has been found. Whether bearing interim costs is worthwhile depends purely on the information value of risk management for a discovered opportunity, not on the time pressure a trader faced in discovering it. Interim costs hence do not generate the feedback loop needed for strategic complementarity to arise. Neither caveat applies to time costs: delays entail losses (foregone profits) that inherently scale up with the size or frequency of trading activity, and they always occur ex interim.[30]

Another difference is that changes in structural factors that affect market speed (e.g., market design, regulation, or technology) do not impact equilibrium risk management in the model variant with fixed costs, although the externalities operate through time pressure. This is because, as explained after Proposition 4, a general increase in search speed λ makes everyone faster so that each individual's implementation probability $P_h(T_h)$ remains the same even as T_h decreases, i.e., as time pressure increases. Condition (8) is hence unaffected by a change in λ. Now, compare (8) with the analogous condition in the model with time costs (cf. equation (1)), which is

$$p_h(T)(1-\rho)|\pi+\alpha_-| > [p_h(T)-p_d(T)]\rho|\pi+\alpha_+|. \tag{9}$$

While the left-hand side is exactly the same, here the difference $[p_h(T)-p_d(T)]$ on the right-hand side increases in λ because $p_d(T)$ shrinks. This effect underlies the comparative statics regarding the speed ratio ν in Propositions 4 and 5. Thus, the implication that structural determinants of external (search) speed or internal (risk management) speed—such as technology, market design, or organizational structure—affect the equilibrium level of risk management is germane only to the model variant with time costs.

bottom in risk management.

[30]It is also immaterial whether a firm's decision to bear interim time costs is, as assumed in the baseline model, made ex ante (see Section 4.4).

4.3 Strategic substitutability

We have interpreted the misallocation of risk to a trader with low private value (i.e., the difference $\alpha_+ - \alpha_-$) as a welfare loss. This is consistent with canonical theory that grounds risk management in financial frictions (Froot, Scharfstein and Stein, 1993): In imperfect capital markets, a negative cash flow shock may prevent firms from funding positive-NPV projects or trigger a socially costly bankruptcy.

An alternative is that a firm's losses from forgoing risk management benefit other agents. In this case risk management affects the *allocation* of surplus, not necessarily its aggregate size. For example, Acharya and Yorulmazer (2008) argue that banks have incentives to hold liquidity to buy assets at fire-sale prices, in case competitors become distressed. We will argue below that the strategic interactions crucially depend on the nature of risk management losses: Risk management choices tend to be strategic complements when lack of risk management destroys surplus, but tend to be strategic substitutes when risk management failures only generate transfers across traders.

Consider a twice-repeated version of our original game with traders making risk management decisions at the start of every period $t \in \{1, 2\}$. While, as before, π_t denotes the common-value trading profit in period t, there is now no (exogenously specified) private value α. Still, every firm faces a probability of $1-\rho$ of being ill-suited to absorb the risk from a trade. A firm that executes an ill-suited trade drops out of the market at the end of that period. This is immaterial in the last period, but risk management has value in period 1 due to the (potential loss of) the period-2 continuation value. Finally, in every period, a firm derives a profit B_t from non-trading activities that cannot be taken over by other firms.[31] This assumption makes bankruptcy costly from a social perspective: If a firm goes bankrupt in period 1, its B_2 will not be realized. The game begins with a mass M_1 of traders who commonly know $\{\pi_t\}_{t\in\{1,2\}}$. The competition intensity in period t is $i_t = \frac{I}{M_t}$.

In the absence of an exogenous private value α, it is privately and socially optimal for traders

[31] A weaker assumption is that these activities can be taken over by other firms either at some cost, or only with some probability. The difficulty of reallocating the assets may stem from intangible resources (e.g., relationships, soft information) that are hard to transfer, or from transaction and integration costs.

to be hasty in period 2. The expected trading profit in that period is $i_2\pi_2$, and the total profit is

$$\Pi_2 \equiv i_2\pi_2 + B_2. \tag{10}$$

Consider now the *net* benefit of trading (rather than not trading) in period 1. If a firm can sustain the trade risk (probability ρ), then the net payoff is simply π_1. If, on the other hand, a trade leads to bankruptcy (probability $1-\rho$), the net payoff from trading is $\pi_1 - \Pi_2$.

To see how this extension relates to our original setup, suppose first that $\pi_2 = 0$ and therefore the continuation value $\Pi_2 = B_2$ only captures surplus that is destroyed in a bankruptcy. In that case, the net benefit of risk management in period 1 is

$$\Delta(T) = p_d(T)\rho\pi_1 - p_h(T)[\pi_1 - (1-\rho)B_2],$$

and the analysis is identical to the one in the original model, which corresponds to the particular case where the private-value loss to a firm is a social loss.

This suggests that, for strategic substitutability to arise, some of the value lost by firms that do not activate risk management must be recouped by other firms. To make this salient, consider the polar opposite: $\pi_2 > 0$ but $B_2 = 0$. The net benefit of risk management in period 1 is then

$$\Delta(T, q) = p_d(T)\rho\pi_1 - p_h(T)[\pi_1 - (1-\rho)i_2(q)\pi_2],$$

where $i_2 = \frac{i_1}{1-q(1-\rho)}$ increases in q, which denotes the fraction of hasty traders in period 1. That is, through competition intensity i_2, period-2 payoffs depend on period-1 decisions. This creates a source of strategic substitutability in risk management choices that coexists with the original strategic complementarities. Specifically, the marginal effect of a change in the fraction of hasty traders q on the net benefit of risk management now writes

$$\frac{\partial}{\partial q}\Delta[T(q), q] = \frac{\partial \Delta}{\partial T}[T(q), q]\frac{\partial T}{\partial q} + \frac{\partial \Delta}{\partial q}[T(q), q]. \tag{11}$$

As in our original model, the first term in (11) is negative for $T > \iota$: When more traders are hasty, time pressure and thus the opportunity cost of risk management increase. The new second term, however, is equal to $p_h(T)(1-\rho)i_2'(q)\pi_2$ and positive: When more traders are hasty, more go bankrupt at the end of period 1 so that there is less competition in period 2. This increases the profit of surviving traders in period 2 and hence the marginal benefit of risk management in period 1.

To see how these strategic substitutabilities can affect the equilibrium outcome, suppose

$$\pi_1 - i_1\pi_2 > 0 > \pi_1 - \frac{i_1}{\rho}\pi_2. \tag{12}$$

The left-hand inequality implies that, if a trader expects everyone else to be deliberate, his best response is to be hasty. If, on the contrary, the trader expects everyone else to be hasty, the right-hand inequality implies that his best response is to be deliberate provided that the delay ι is not too large. Under such parameters, any equilibrium features a mix of traders with risk management and some without, which reflects the strategic substitutability.

Overall, this extension highlights that strategic complementarities in risk management choices are more likely to arise when risk management failures entail deadweight losses. If instead risk management losses represent transfers to other traders, a decrease in aggregate risk management increases the size of these transfers, and may in turn increase the payoff from risk management. Note, though, that even in the extreme case where risk management losses are a mere reallocation of future business, i.e., $B_2 = 0$, strategic complementarities coexist with strategic substitutabilities, as is apparent from (11), and could still sustain multiple equilibria.

It seems reasonable to assume that risk management failures decrease welfare ($B_2 > 0$), since any friction that makes a reallocation of assets from distressed firms to healthy ones costly or imperfect generates social losses. Moreover, it is only when risk management losses are socially costly that risk management decisions matter in the first place: when the losses are pure transfers ($B_2 = 0$), traders' ex ante payoffs are the same in any equilibrium and equal to what they would earn if all could commit to practicing risk management—or not. In other words, coordination remains a strategic issue, but there is no constrained inefficient coordination *failure* in the absence of deadweight losses from risk management failures.

4.4 Alternative modeling of trading

Our framework features a stylized trading process. For example, firms cannot choose how many units to trade, nor can they reverse previous trades. These assumptions may seem restrictive but are not crucial for our results. Allowing all firms to trade more units tends to reinforce preemption risk and hence the fragility of the deliberate equilibrium under time pressure. In the Online Appendix, we also discuss the alternative assumption of allowing firms to trade more units only with risk management approval and make two points there: First, this can (but need not) cause time pressure to be increasing in the fraction of *deliberate* (rather than hasty) traders. When this is the case, the equilibrium is unique (even without the global games treatment), although the scope for coordination failure persists. Second, this setting involves an exogenous trade restriction on firms without risk management that is not consistent with optimal firm behavior. Settings where firms endogenously choose whether to permit large(r) trades only with risk management approval are captured in reduced form by our model (see the "position limit" interpretation of our model at the end of Section 2.1). Our results are also robust to retrading as long as trades are *partially* irreversible, (e.g., due to duplication of transaction costs). In fact, partial reversibility can make initial trades hastier by lowering the value of ex ante screening. Moreover, even if they are fully reversible, firms would be vulnerable while they hold (and investigate) acquired trade positions.

Another assumption is that risk management is chosen ex ante and cannot be adjusted during trading. This assumption is also inessential. In the Online Appendix, we consider a model version in which traders, once they have identified a trading opportunity, can decide at any point in time whether to continue investigating or to execute. We show that the results are qualitatively similar to those derived from the current model.

How trading pressure impacts the magnitude of a trading opportunity could also be modeled differently. Specifically, the common value π could decrease continuously as more traders execute trades, reflecting a price-sensitive demand for liquidity, as in Kondor (2009). This introduces two countervailing effects: On the one hand, when more traders are hasty, the common value at which a trader expects to trade is lower, which makes risk management more attractive. On the other hand, the faster pace at which the common value shrinks heightens preemption motives, which makes risk management less attractive. The latter effect reinforces strategic complementarities and dominates when, for example, π decreases linearly with the mass of executed trades.

5 Risk management regulation

In their own words, risk managers find balancing "business needs and risk appetite" challenging (Ernst & Young, 2013, 12). This is true for each firm in our model. Unregulated, however, the firms' privately optimal choices can generate a constrained inefficient market outcome. Because the source of inefficiency differs from frictions usually considered in the risk management literature, we briefly discuss regulatory alternatives.

Capital and liquidity requirements. Leverage or maturity play no explicit role in our model. Rather, any impact they may have is implicit in the private value $\tilde{\alpha}_k$. There are two possible scenarios: On the one hand, if management acts in the interest of all investors, $\tilde{\alpha}_k$ reflects deadweight losses from undesirable risks depending on a firm's capital structure. In this case, $|\tilde{\alpha}_k|$ may increase with leverage (e.g., owing to bankruptcy costs). If, on the other hand, management is biased toward shareholders, whose vulnerability alone is reflected in $\tilde{\alpha}_k$, $|\tilde{\alpha}_k|$ may decrease with leverage (e.g., as a result of risk shifting).[32] Constrained inefficiencies in our model are *conditional* on $\tilde{\alpha}_k$, regardless of whose "need" for risk management it reflects. Thus, capital structure regulation would modify the need for risk management, but not correct inefficiencies conditional on that need. For example, higher capital requirements may improve managerial incentives, pushing a firm toward the first scenario, but coordination failures can persist before or after such a shift.

Pigovian approaches. In the constrained inefficient outcome of our model, firms trade "too much" in that trades are not selective enough. One countermeasure is hence to levy a tax τ on

[32] Maćkowiak and Wiederholt (2012) show that managers with risk-shifting incentives, being less concerned about uncertainty, may underinvest in information. Correcting management incentives would in this case resolve the failure to acquire information. The converse is not true: Better information per se would not curb risk-shifting. Quite the contrary, it would enable management to shift risks more effectively.

every trade, which lowers the value of a trading opportunity to $\pi^{tax} = \pi - \tau$.[33] This reduces the opportunity cost of risk management, but it also makes trades with $\pi^{tax} + \alpha_{+} < 0$ unprofitable. The tax thus deters valuable as well as excessive trade.

Given the role of the external-internal speed ratio $\frac{\iota}{\lambda}$, one could also "tax" speed investments.[34] Discouraging external search speed (raising λ) is also a double-edged sword: The decrease in time pressure promotes risk management, but may also reduce valuable liquidity or price efficiency (as discussed in Section 4.1). In contrast, greater internal speed (reducing ι) helps risk management, and conditional thereon, also raises speed. Still, firms may lack incentives to invest in internal speed, or fail to coordinate on them (as discussed in the Online Appendix). This therefore provides a possible justification for subsidizing investments in risk management technology.

Alternatively, regulators could directly enforce risk management standards. Microprudential regulation could include speed and accuracy requirements that banks' risk management systems must meet, and exchange regulations could require safeguards to be integrated in (the design of) the market platforms traders compete on. We turn to these two points next.

Stress tests. The quantitative component of the Fed's Comprehensive Capital Analysis and Review (CCAR) assesses banks' risk exposures. The *qualitative* component assesses their protocols, methods, and technology for measuring and monitoring risk. In 2014, the Fed rejected the capital plans of Citibank, HSBC, RBS Citizens, and Santander, citing deficiencies in their risk management processes. The concerns about Citibank, for example, included "its ability to develop scenarios for its internal stress testing that adequately reflect and stress its full range of business activities and exposures."[35]

Thus, regulators can improve risk management incentives through mandates. Indeed, as risk management regulation was tightened in the wake of the 2008 crisis, technology investments for compliance purposes became a priority in major U.S. banks (Ernst & Young, 2013, 68f):

> Systems and data vied for the top spot on the challenges to internal transparency...

[33] Tobin (1978) proposed transaction taxes based on the argument that excess currency speculation undermines the allocative role of exchange rates for hedging purposes, i.e., risk allocation.

[34] Budish, Cramton and Shim (2015) provide examples of such arms races between trading firms.

[35] Federal Reserve Board, Comprehensive Capital Analysis and Review 2014: Assessment Framework and Results, 7.

> and, indeed, have been raised as among the top challenges throughout this report. "There is a huge effort underway to redo all the plumbing, data aggregation, accuracy, quality of information," one executive said. "That's the framework in which a lot of our future-state risk systems will be addressed... a huge, multiyear, gazillion-dollar effort."

Market design. The coordination failure in our model is driven by preemptive competition, which is a consequence of time priority in market interactions or rules. This aspect of the market may become increasingly important with advances in technology and algorithmic trading. Indeed, Kirilenko and Lo (2013) warn that the gap between "machine speed" and "human speed" created by algorithmic trading may further discourage firms from practicing risk management; they suggest that a growing focus on speed and machines may necessitate *systemwide* risk management regulation "translated into computer code and executed by automated systems" with "safeguards at multiple levels of the system." In other words, interventions into market processes may become integral. Among measures debated are proposals to discretize trading time to thwart preemptive competition at very short intervals, and to build safeguards, such as trade limits and pre-trade protocols, into the system at the level of intermediaries or central counterparties, such as dealer-brokers, clearing houses, or exchanges (see, e.g, Clark, 2012; Kirilenko and Lo, 2013; Budish, Cramton and Shim, 2015).[36]

To place the case for market design interventions in context, consider Grossman and Miller's (1988) framework in which equilibrium market structure is the outcome of the tradeoff between the costs to intermediaries of maintaining a continuous presence in a market and the benefits to traders of being able to trade as quickly as possible. In that setting, speed cannot be excessive.[37] In our setting, regulating market design to moderate speed can improve welfare because

[36]In a survey on the risks of high-frequency trading, proprietary trading firms – when asked what they would change for "the betterment of the markets" – mention inter alia that (i) "requiring trading venues to uniformly apply pre-trade risk checks for all market participants would consistently apply latency to and level the playing field for all trading firms" and that (ii) "every trading venue should have limits on maximum positions, quantity per order, and credit... [and] on number of messages that can be sent to the trading venue within a specified period of time... per product/customer" (Clark and Ranjan, 2012, 13f).

[37]In Pagnotta and Philippon (2018), immediacy is determined by technological investments made by the exchanges on which investors trade, and the exchanges' choices, which shape market structure in their model, are driven by differentiation incentives. In aggregate, these investments can be too high relative to the welfare optimum

it safeguards traders' incentives to engage in risk management.

6 Conclusion

Managing risk requires monitoring it through processes that collect, aggregate, and analyze the relevant information inside a firm. These processes take time and can delay investment decisions, which creates opportunity costs that scale up with the size and frequency of a firm's investment opportunities when those opportunities are short-lived. By the same token, the degree to which investments are subjected to risk controls affects the speed at which opportunities in the market are exploited. Out of this feedback loop arises the possibility that risk management quality is not just an organizational choice but a market equilibrium outcome between competing firms. We have shown that this can cause risk management failures to result from a race to the bottom, with scope for regulation. According to this theory, such failures are catalyzed by organizational and technological aspects of firms and of the markets in which they compete, rather than by the firms' financial structures. These predictions distinguish our theory from others and may help explain differences in the rate of occurrence of risk management failures across industries and over time. Moreover, if supported by evidence, they may be useful in guiding regulation.

We think our model could be extended in two directions. First, the model could be extended to analyze how time externalities between firms interact with internal agency problems within firms. For example, micro-founding the private value of risk management in multi-divisional firms could link questions about conglomerate structures and firm boundaries to the intensity of preemptive competition. Relatedly, understanding how time pressure affects agency problems between management and traders within firms would link the optimal implementation of internal governance mechanisms (e.g., Landier, Sraer and Thesmar, 2009) to time competition, which in turn may have implications for corporate governance regulation. We study the latter questions in separate work (Bouvard and Lee, 2019).

Second, the speed-information trade-off in our model could matter for other market aspects

since they not only accelerate trading but also (are meant to) relax competition between the exchanges.

or in other types of markets. For example, we model risk management as learning about private values, but strategic complementarities could also emerge in learning about the common value of traded assets, with implications for how well markets aggregate information (see e.g., Kendall, 2018). More broadly, our formalization of time pressure extends a popular theoretical apparatus developed for bank runs and financial panics to another type of economic problem, and may be useful for modelling long-run organizational or institutional choices in a variety of settings with time-based competition such as patent races, academic publications, and news (see e.g., Askenazy, Thesmar and Thoenig, 2006, Bobtcheff, Bolte and Mariotti, 2017).

Proof of Lemma 1

Lemma 1 follows from the properties of $\Delta(T)$, which we formally state below:

1. If $\pi > 0$, there exists $T^*(\pi)$ such that $\Delta(T) < 0$ if $T < T^*(\pi)$ and $\Delta(T) > 0$ if $T > T^*(\pi)$.
2. If $\pi > 0$, $\Delta(.)$ is strictly decreasing on $(0, \iota)$ and strictly increasing on $(\iota, +\infty)$, and $T^*(\pi) > \iota$.
3. If $\pi \leq 0$, $\Delta(.)$ is strictly positive and strictly increasing on $(0, +\infty)$.
4. For any $T > 0$, $\Delta(T)$ is strictly decreasing in π.

Proof. As a reminder,

$$\begin{aligned} \Delta(T) &= \rho(\pi + \alpha_+)p_d(T) - \pi p_h(T) \\ &= \rho(\pi + \alpha_+)\max\left\{0, 1 - e^{-(T-\iota)/\lambda}\right\} - \pi\left(1 - e^{-T/\lambda}\right) \end{aligned} \tag{13}$$

Suppose $\pi > 0$. Check that (a) $\lim_{T\to 0}\Delta(T) = 0$; (b) $\lim_{T\to+\infty}\Delta(T) = \rho(\pi + \alpha_+) - \pi = -(1-\rho)(\pi + \alpha_-) > 0$; (c) $\Delta'(T) < 0$ if $T \in (0, \iota)$ and $\Delta'(T) > 0$ if $T \in (\iota, +\infty)$. Together, these facts prove points 1 and 2.

Suppose $\pi \leq 0$. We have (a) $\lim_{T\to 0}\Delta(T) = 0$; (b) $\pi + \alpha_+ > 0$ (by assumption); (c) $p_d(T)$ is weakly increasing and $p_h(T)$ is strictly increasing. Altogether, these prove point 3.

Finally, for any T, $p_h(T) > \rho p_d(T)$, which, from (13) implies $\frac{\partial\Delta}{\partial\pi} < 0$.

It remains to show that T^* is an increasing function of π when $\pi > 0$. Applying the implicit function theorem to the equation $\Delta[T^*(\pi)] = 0$, we obtain

$$\Delta'(T^*)\frac{\partial T^*}{\partial\pi} + \frac{\partial\Delta}{\partial\pi}(T^*) = 0.$$

From point 2, $\Delta'(T^*) > 0$. Furthermore, from point 4, $\frac{\partial\Delta}{\partial\pi} < 0$. Therefore, $\frac{\partial T^*}{\partial\pi} > 0$. □

Proof of Proposition 2

We derive the equilibrium of the general case in which ε can be bounded away from 0. The proof is in several steps, and we only show here the existence of a unique equilibrium *in threshold strategies*. The proof that any equilibrium is in threshold strategies follows Goldstein and Pauzner (2005) and is in the Online Appendix.

For a given realization of π, the proportion of hasty traders under a threshold strategy $\hat{s}$ is

$$q(\pi, \hat{s}) \equiv \begin{cases} 0 & \text{if } \pi \leq \hat{s} - \varepsilon, \\ \dfrac{\pi + \varepsilon - \hat{s}}{2\varepsilon} & \text{if } \hat{s} - \varepsilon < \pi < \hat{s} + \varepsilon, \\ 1 & \text{if } \pi \geq \hat{s} + \varepsilon \end{cases} \tag{14}$$

For a proportion q of hasty traders, the mass of trade executed by time T is

$$m(q, T) \equiv q p_h(T) + (1 - q)\rho p_d(T),$$

Hence the time at which the trading opportunity is exhausted, $\tau(\pi, \hat{s})$, is a solution to $m[q(\pi, \hat{s}), \tau] = i$. Finally, the net expected benefit of a deliberate strategy given a signal s_k and a threshold $\hat{s}$ is

$$u(s_k, \hat{s}) \equiv \mathbb{E}_\pi\{\Delta[\tau(\pi, \hat{s}), \pi] | s_k\} = \frac{1}{2\varepsilon} \int_{s_k - \varepsilon}^{s_k + \varepsilon} \Delta[\tau(\pi, \hat{s}), \pi] d\pi, \tag{15}$$

Step 1: Existence of a unique threshold equilibrium.

Claim 1. *$\tau(\pi, \hat{s})$ is decreasing in π and increasing in $\hat{s}$. Furthermore, $\tau(\pi + a, \hat{s} + a) = \tau(\pi, \hat{s})$.*

Proof. $q(\pi, \hat{s})$ is increasing in π and decreasing in $\hat{s}$. Furthermore, $m(q, T)$ is increasing in T, and since $p_h(T) > \rho p_d(T)$, increasing in q. Therefore, $\tau(\pi, \hat{s})$ is decreasing in π and increasing in $\hat{s}$. Finally, from (14), $q(\pi + a, \hat{s} + a) = q(\pi, \hat{s})$, which in turn implies $\tau(\pi + a, \hat{s} + a) = \tau(\pi, \hat{s})$. □

Claim 2. *There exists a unique s^* such that $u(s^*, s^*) = 0$.*

Proof. We first show the existence of s^* using upper- and lower-dominance regions. Suppose that $s < \underline{\pi} - \varepsilon$, then of any $\pi \in [s - \varepsilon, s + \varepsilon]$, $T^*(\pi) < T_h \leq \tau(\pi, s)$, therefore $\Delta[\tau(\pi, s), \pi] > 0$ and hence $u(s, s) > 0$. Similarly, if $s > \overline{\pi} + \varepsilon$, then $u(s, s) < 0$. The continuity of $u(., .)$ then implies the existence of s^*, which proves existence.

Furthermore,

$$\begin{aligned} u(s, s) &= \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s+\varepsilon} \Delta[\tau(\pi, s), \pi] d\pi \\ &= \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s+\varepsilon} \Delta[\tau(\pi + a, s + a), \pi] d\pi \\ &= \frac{1}{2\varepsilon} \int_{s+a-\varepsilon}^{s+a+\varepsilon} \Delta[\tau(\pi, s + a), \pi - a] d\pi \\ &< \frac{1}{2\varepsilon} \int_{s+a-\varepsilon}^{s+a+\varepsilon} \Delta[\tau(\pi, s + a), \pi] d\pi = u(s + a, s + a), \end{aligned}$$

where the second equality follows from Claim 1 and the last inequality follows from $\frac{\partial \Delta}{\partial \pi} < 0$. Hence, $u(s,s)$ is strictly decreasing in s, which proves uniqueness. □

Finally, to complete the proof we show the following result.

Claim 3. *$u(s, s^*) > 0$ for $s < s^*$ and $u(s, s^*) < 0$ for $s > s^*$.*

Proof. From (15), $u(s^*, s^*) = 0$ implies that $\Delta[\tau(., s^*), .]$ changes sign on $[s^* - \varepsilon, s^* + \varepsilon]$. Therefore, by continuity, there exists $\hat{\pi} \in [s^* - \varepsilon, s^* + \varepsilon]$ such that $\Delta[\tau(\hat{\pi}, s^*), \hat{\pi}] = 0$, and hence, $\tau(\hat{\pi}, s^*) = T^*(\hat{\pi}) > 0$.[38] Suppose $\pi < \hat{\pi}$, then $\tau(\pi, s^*) \geq T^*(\hat{\pi})$, and therefore from Lemma 1), $\Delta[\tau(\pi, \hat{s}), \hat{\pi}] \geq 0$. Furthermore, since $\frac{\partial \Delta}{\partial \pi} < 0$, $\Delta[\tau(\pi, \hat{s}), \pi] > \Delta[\tau(\pi, \hat{s}), \hat{\pi}] \geq 0$. Similarly, if $\pi > \hat{\pi}$, then $\Delta[\tau(\pi, \hat{s}), \pi] < 0$. This also shows that $\hat{\pi}$ is uniquely defined.

Suppose $s < s^*$. If $s < \hat{\pi} - \varepsilon$, for any $\pi \in [s - \varepsilon, s + \varepsilon]$, $\Delta[\tau(\pi, s^*), \pi] > 0$, and thus $u(s, s^*) > 0$. If $\hat{\pi} - \varepsilon \leq s < s^*$,

$$\begin{aligned} u(s, s^*) - u(s^*, s^*) &= \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s+\varepsilon} \Delta[\tau(\pi, s^*), \pi] d\pi - \frac{1}{2\varepsilon} \int_{s^*-\varepsilon}^{s^*+\varepsilon} \Delta[\tau(\pi, s^*), \pi] d\pi \\ &= \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s^*-\varepsilon} \Delta[\tau(\pi, s^*), \pi] d\pi - \frac{1}{2\varepsilon} \int_{s+\varepsilon}^{s^*+\varepsilon} \Delta[\tau(\pi, s^*), \pi] d\pi. \end{aligned}$$

$\pi < s^* - \varepsilon$ implies $\pi < \hat{\pi}$, and therefore $\Delta[\tau(\pi, s^*), \pi] > 0$. Thus, $\int_{s-\varepsilon}^{s^*-\varepsilon} \Delta[\tau(\pi, \hat{s}), \pi] d\pi > 0$. $\pi \geq s + \varepsilon$ implies $\pi \geq \hat{\pi}$ and therefore $\Delta[\tau(\pi, s^*), \pi] < 0$. Hence, $\int_{s+\varepsilon}^{s^*+\varepsilon} \Delta[T(\pi, s^*), \pi] d\pi < 0$. Therefore $u(s, s^*) - u(s^*, s^*) = u(s, s^*) > 0$. Symmetrically, if $s > s^*$, $u(s, s^*) < 0$. □

Step 2: Any equilibrium is a threshold equilibrium.

See Online Appendix.

Proof of Proposition 3

The threshold π^* is given by

$$\int_0^1 \Delta[T(q), \pi^*] dq = 0 \Leftrightarrow \int_0^1 p_d[T(q)](\rho \pi^* + \sigma) - p_h[T(q)] \pi^* dq = 0 \tag{16}$$

Total differentiation of (16) with respect to σ yields

$$\begin{aligned} &\int_0^1 \frac{\partial \Delta}{\partial \sigma}[T(q), \pi^*] dq + \int_0^1 \frac{\partial \Delta}{\partial \pi^*}[T(q), \pi^*] \frac{\partial \pi^*}{\partial \sigma} dq = 0 \\ \Leftrightarrow \quad &\int_0^1 p_d[T(q)] dq + \int_0^1 \rho p_d[T(q)] - p_h[T(q)] dq \frac{\partial \pi^*}{\partial \sigma} = 0. \end{aligned} \tag{17}$$

[38] If $T^*(\hat{\pi}) = 0$, then $\hat{\pi} \leq 0$ and $\Delta[T, \hat{\pi}] > 0$ for any $T > 0$, a contradiction.

Rearranging (16)

$$\pi^* \int_0^1 \rho p_d[T(q)] - p_h[T(q)]dq = -\sigma \int_0^1 p_d[T(q)]dq \tag{18}$$

Plugging into (17) yields

$$\frac{\partial \pi^*}{\partial \sigma} = \frac{\pi^*}{\sigma} > 0. \tag{19}$$

Similarly,

$$\begin{aligned} &\int_0^1 \frac{\partial \Delta}{\partial \rho}[T(q), \pi^*]dq + \int_0^1 \frac{\partial \Delta}{\partial \pi^*}[T(q), \pi^*]\frac{\partial \pi^*}{\partial \rho}dq = 0 \\ \Leftrightarrow \quad &\int_0^1 p_d[T(q)]\pi^* dq + \int_0^1 \rho p_d[T(q)] - p_h[T(q)]dq\frac{\partial \pi^*}{\partial \rho} = 0 \end{aligned}$$

Using (18) again,

$$\frac{\partial \pi^*}{\partial \rho} = \frac{(\pi^*)^2}{\sigma} > 0 \tag{20}$$

□

Proof of Proposition 4 part a.

We show here that π^* is a decreasing function of $\frac{\iota}{\lambda}$.

As in (22), let

$$\hat{q}(\iota/\lambda) \equiv \frac{i}{1 - e^{-\iota/\lambda}}.$$

We have

$$U(\pi^*) = \int_0^{\min\{\hat{q}(\iota/\lambda),1\}} \{\rho(\pi^* + \alpha_+)p_d[T(q)] - \pi^* p_h[T(q)]\}dq - \int_{\min\{\hat{q}(\iota/\lambda),1\}}^1 \pi^* p_h[T(q)]dq$$

Note that $U(\pi^*)$ depends on λ and ι both through the boundaries of the integrals and, implicitly, through the functions $p_d(.)$ and $p_h(.)$, that is, the probabilities of execution under each strategy. However, $p_d[T(\hat{q}(\iota/\lambda))] = 0$, and hence, the effect of a marginal change in λ or in ι on the integral boundaries cancels out. As a result, differentiating $U(\pi^*)$ with respect to λ or ι only requires differentiating the integrands.

Let $x \equiv e^{\iota/\lambda}$. Using equations (23), (25) and (26), one obtains

$$\frac{\partial U(\pi^*)}{\partial x} = -\frac{\partial}{\partial x}\left\{[\rho x(\pi^* + \alpha_+) - \pi^*]\int_0^{\min\{\hat{q}(\iota/\lambda),1\}} \frac{q + (1-q)\rho - i}{q + (1-q)\rho x}dq\right\}. \tag{21}$$

It is easy to check that the expression between brackets is increasing in x. Therefore

$$\frac{\partial \pi^*}{\partial x} = -\frac{\frac{\partial U}{\partial x}(\pi^*)}{U'(\pi^*)} < 0.$$

This, in turn, implies that π^* is decreasing in ι and increasing in λ. □

Proof of Proposition 4 part b.

We show here that π^* is an increasing function of i.

Let $T(q, i)$ be defined as in (2), with the addition of the second argument explicitly recognizing its dependence on i. Let

$$U(\pi) \equiv \int_0^1 \Delta[T(q,i), \pi] dq.$$

Note that π^* solves $U(\pi) = 0$. Note also that $U'(.) < 0$, and from (2), $T(q, i)$ is increasing in i. Consider two cases,

(a) $T_h(i) \geq \iota$

Then, for any $q \in [0, 1)$, $T(q, i) > \iota$, and therefore $\frac{\partial \Delta}{\partial T}[T(q, i), \pi^*] > 0$. This, in turn, implies

$$\frac{\partial U}{\partial i}(\pi^*) = \int_0^1 \frac{\partial \Delta}{\partial T}[T(q,i), \pi^*] \frac{\partial T}{\partial i}(q,i) dq > 0,$$

and finally, by the implicit function theorem,

$$\frac{\partial \pi^*}{\partial i} = -\frac{\frac{\partial U}{\partial i}(\pi^*)}{U'(\pi^*)} > 0.$$

(b) $T_h(i) < \iota$

Let

$$\hat{q}(i) \equiv \frac{i}{p_h(\iota)}. \tag{22}$$

$$\begin{aligned} U(\pi^*) &= \int_0^{\hat{q}(i)} \{\rho(\pi^* + \alpha_+) p_d[T(q,i)] - \pi^* p_h[T(q,i)]\} dq - \int_{\hat{q}(i)}^1 \pi^* p_h[T(q,i)] dq \\ &= \hat{q}(i)[\rho(\pi^* + \alpha_+) - \pi^*] - \int_0^{\hat{q}(i)} \{\rho(\pi^* + \alpha_+)[1 - p_d[T(q,i)]] - \pi^*[1 - p_h[T(q,i)]]\} dq \\ &\quad + \int_{\hat{q}(i)}^1 \pi^*[1 - p_h[T(q,i)]] dq \end{aligned} \tag{23}$$

If $q < \hat{q}(i)$, then $1 - p_d[T(q, i)] = e^{\iota/\lambda}[1 - p_h[T(q, i)]]$. Using $qp_h[T(q, i)] + (1 - q)\rho p_d[T(q, i)] = i$, we get

$$1 - p_d[T(q,i)] = e^{\iota/\lambda}[1 - p_h[T(q,i)]] = e^{\iota/\lambda} \frac{q + (1-q)\rho - i}{q + (1-q)\rho e^{\iota/\lambda}}. \tag{24}$$

Hence, the first integral in (23) becomes

$$[\rho e^{\iota/\lambda}(\pi^* + \alpha_+) - \pi^*] \int_0^{\hat{q}(i)} \frac{q + (1-q)\rho - i}{q + (1-q)\rho e^{\iota/\lambda}} dq. \tag{25}$$

If $q > \hat{q}(i)$, then $p_d[T(q,i)] = 0$. Using this in $qp_h[T(q,i)] + (1-q)\rho p_d[T(q,i)] = i$, we get $1 - p_h[T(q,i)] = \frac{q-i}{q}$. Hence, the second integral in (23) becomes

$$\pi^* \int_{\hat{q}(i)}^1 \frac{q-i}{q} dq \tag{26}$$

Consider the first line of equation (23). $U(\pi^*)$ depends on i both through the boundaries of the integrals (via $\hat{q}(i)$) and through the integrands (via $T(q,i)$). However, since $p_d[T(\hat{q}(i)), i] = 0$, the effect of a marginal change in i that goes through $\hat{q}(i)$ cancels out. Hence, using (24) and (25) to substitute into the second line of (23), we obtain

$$\frac{\partial U}{\partial i}(\pi^*) = [\rho e^{\iota/\lambda}(\pi^* + \alpha_+) - \pi^*] \int_0^{\hat{q}(i)} \frac{1}{q + (1-q)\rho e^{\iota/\lambda}} dq - \pi^* \int_{\hat{q}(i)}^1 \frac{1}{q} dq. \tag{27}$$

Now, if $q > \hat{q}(i)$,

$$\frac{1}{q} = \frac{p_h[(T(q)]}{i}, \tag{28}$$

In addition, rearranging (24),

$$p_d[T(q)] = \frac{q(1 - e^{\iota/\lambda}) + e^{\iota/\lambda} i}{q + (1-q)\rho e^{\iota/\lambda}} \text{ and } p_h[T(q)] = \frac{(1-q)\rho(e^{\iota/\lambda} - 1) + i}{q + (1-q)\rho e^{\iota/\lambda}},$$

which, since $e^{\iota/\lambda} > 1$, implies

$$\frac{p_d[T(q)]}{e^{\iota/\lambda} i} < \frac{1}{q + (1-q)\rho e^{\iota/\lambda}} < \frac{p_h[T(q)]}{i}. \tag{29}$$

Finally, using (27), (28) and (29),

$$i \frac{\partial U}{\partial i}(\pi^*) > \rho(\pi^* + \alpha_+) \int_0^{\hat{q}(i)} p_d[T(q)] dq - \pi^* \int_0^1 p_h[T(q)] dq$$

The RHS of this last inequality is $U(\pi^*) = 0$, and using again the implicit function theorem concludes the proof. □

Proof of Proposition 5

From Proposition 3,

$$\frac{\partial \pi^*}{\partial \sigma} = \frac{\pi^*}{\sigma} \text{ and } \frac{\partial \pi^*}{\partial \rho} = \frac{(\pi^*)^2}{\sigma}. \tag{30}$$

From Proposition 4

$$\frac{\partial \pi^*}{\partial i} > 0 \text{ and } \frac{\partial \pi^*}{\partial \nu} < 0 \tag{31}$$

Keeping in mind that $\pi^* > 0$, combining (30) and (31) delivers the results in Proposition 5.

References

Acharya, V., M. Pagano, and P. Volpin, "Seeking Alpha: Excess Risk Taking and Competition for Managerial Talent," 2013, NBER Working Paper No. 18891.

Acharya, and T. Yorulmazer, "Cash-in-the-Market Pricing and Optimal Resolution of Bank Failures," *Review of Financial Studies* 21 (2008), 2705–42.

Acharya, V., P. Schnabl, and G. Suarez, "Securitization without risk transfer," *Journal of Financial Economics* 107 (2013), 515–536.

Aguerrevere, F., "Real Options, Product Market Competition, and Asset Returns," *Journal of Finance* 64 (2009), 957–83.

Alessandrini, P., A. Presbitero and A. Zazzaro, "Banks, Distances and Firms? Financing Constraints," *Review of Finance* 13 (2008), 261–307.

An, Y., Y. Song and X. Zhang, "The Intermediary Rat Race," 2018, Working Paper,`https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2754235`.

Askenazy, P., D. Thesmar and M. Thoenig "On the Relation Between Organisational Practices and New Technologies: the Role of (Time Based) Competition," *Economic Journal*, 115 (2006), 128–154.

Biais, B., T. Foucault and S. Moinas "Equilibrium Fast Trading," *Journal of Financial Economics*, 116 (2015), 292–313.

Black, S., P. E. Strahan, "Entrepreneurship and Bank Credit Availability," *Journal of Finance*, 57 (2002), 2807–2833.

Bobtcheff, C., J. Bolte and T. Mariotti, "Researcher's Dilemma," *Review of Economic Studies*, 84 (2017), 969-1014.

Bond, P., A. Edmans and I. Goldstein "The Real Effects of Financial Markets," *Annual Review of Financial Economics*, 4 (2012), 339-360.

Bouvard, M., and S. Lee "Agents under time pressure," 2018, Mimeo.

Boyd, J., and G. De Nicoló "The Theory of Bank Risk Taking and Competition Revisited," *Journal of Finance*, 60 (2005), 1329–1343.

Budish, E., P. Cramton and J. Shim "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130 (2015), 1547–1621.

Caballero, R., and R. Pindyck, "Uncertainty, Investment, and Industry Evolution," *International Economic Review* 37 (1996), 641–62.

Carlsson, H., and E. van Damme "Global Games and Equilibrium Selection," *Econometrica*, 61 (1993), 989–1018.

Cetorelli, N., J. McAndrews, and J. Traina "Large and Complex Banks," *Economic Policy Review*, 20 (2014).

Cheng, I., "Corporate Governance Spillovers," 2011, Working Paper, `http://ssrn.com/abstract=1299652`.

Clark, C., "Controlling Risk in a Lightning-Speed Trading Environment," *Chicago Fed Letter*, 272 (2010).

Clark, C., "How to Keep Markets Safe in the Era of High-Speed Trading," *Chicago Fed Letter*, 303 (2012).

Clark, C., and R. Ranjan "How Do Proprietary Trading Firms Control the Risks of High Speed Trading?" 2012, Working Paper, Federal Reserve Bank of Chicago.

Coval, J., J. Jurek, and E. Stafford, "The economics of structured finance," *Journal of Economic Perspectives*, 23 (2009), 3–25.

Dell'Ariccia, G. and R. Marquez, "Information and bank credit allocation," *Journal of Financial Economics*, 72 (2004), 185–214.

Demsetz, R., and P. Strahan, "Diversification, size, and risk at bank holding companies," *Journal of Money, Credit, and Banking*, 72 (1997), 300–313.

Deloitte, *Global risk management survey, 8th edition: Setting a higher bar* (2013).

Dicks, D., "Executive Compensation and the Role for Corporate Governance Regulation," *Review of Financial Studies*, 25 (2012), 1971–2004.

Diamond, D., and P. Dybvig, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91 (1983), 401–419.

Ernst & Young Global Banking and Capital Markets Center, "Remaking financial services: risk management five years after the crisis. A survey of major Financial Institutions," 2013.

Frame, W., and L. White "Empirical Studies of Financial Innovation: Lots of Talk, Little Action?" *Journal of Economic Literature*, 42 (2004), 116-144.

Froot, K., D. Scharfstein and J. Stein "Risk Management: Coordinating Corporate Investment and Financing Policies," *Journal of Finance*, 48 (1993), 1629–1658.

Garleanu, N., and L. Pedersen "Liquidity and Risk Management," *American Economic Review*, 97 (2007), 193-197.

Goldstein, I., and A. Pauzner "Demand–Deposit Contracts and the Probability of Bank Runs," *Journal of Finance*, 60 (2005), 1293–1327.

Grenadier, S., "Option Exercise Games: An Application to the Equilibrium Investment Strategies of Firms," *Review of Financial Studies*, 15 (2002), 691–721.

Gromb, D., and D. Vayanos "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, 66 (2002), 361–407.

Grossman, S., and M. Miller "Liquidity and Market Structure," *Journal of Finance*, 43 (1988), 617–633.

Hellmann T., K. Murdock and J. Stieglitz, “Liberalization, Moral Hazard in Banking, and Prudential Regulation: Are Capital Requirements Enough?” *American Economic Review*, 90 (2000), 147–165.

Keeley, M., “Deposit Insurance, Risk, and Market Power in Banking,” *American Economic Review*, 80 (1990), 1183–1200.

Kendall, C., “The time cost of information in financial markets,” *Journal of Economic Theory*, 176 (2018), 118–157.

Kirilenko, A., and A. Lo, “Moore’s Law versus Murphy’s Law: Algorithmic Trading and Its Discontents,” *Journal of Economic Perspectives*, 27 (2013), 51–72.

Kondor, P., “Risk in Dynamic Arbitrage: The Price Effects of Convergence Trading,” *Journal of Finance*, 64 (2009), 631–655.

Kroszner, R. S., and P. E. Strahan, “What drives deregulation? Economics and politics of the relaxation of bank branching restrictions,” *Quarterly Journal of Economics*, 114 (1999), 1437–1467.

Landier, A., D. Sraer and D. Thesmar “Financial Risk Management: when does independence fail?” *American Economic Review*, 99 (2013), 454–458.

Leahy, J., “Investment in Competitive Equilibrium: The Optimality of Myopic Behavior,” *Quarterly Journal of Economics*, 108 (1993), 1105–1133.

Maćkowiak, B., and M. Wiederholt “Information Processing and Limited Liability,” *American Economic Review*, 102 (2012), 30–34.

Martinez-Miera, D., and R. Repullo “Does Competition Reduce the Risk of Bank Failure?” *Review of Financial Studies*, 23 (2010), 3638–3664.

Mehta, A., M. Neukirchen, S. Pfetsch and T. Poppensieker “Managing market risk: Today and tomorrow,” 2012, McKinsey Working Papers on Risk No 32.

Morris S., and H.S. Shin "Global Games: Theory and Applications," *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, (2003).

Pagnotta, E., and T. Philippon "Competing on speed," *Econometrica*, 86 (2018), 1067–1115.

Parlour, C., and U. Rajan "Competition in Loan Contracts," *American Economic Review*, 91 (2001), 1311–1328.

Petersen, M., and R. Rajan "Does Distance still Matter? The Information Revolution in Small Business Lending," *Journal of Finance*, 57 (2002), 2533–2570.

Raff, K., "Information Externalities in Corporate Governance," 2011, Working Paper, `http://ssrn.com/abstract=1787130`.

Rajan, R. "Has Finance Made the World Riskier?"*European Financial Management*, 12 (2006), 499–533.

Rajan, U., A. Seru and V. Vig, "The failure of models that predict failure: Distance, incentives, and defaults,"*Journal of Financial Economics*, 115 (2015), 237–260.

Ruckes, M. "Bank Competition and Credit Standards,"*Review of Financial Studies*, 17 (2004), 1073–1102.

Rampini, A., and S. Viswanathan "Collateral, Risk Management, and the Distribution of Debt Capacity," *Journal of Finance*, 65 (2010), 2293–2322.

Shleifer, A., and R. Vishny "The Limits of Arbitrage," *Journal of Finance*, 52 (1997), 35–55.

Smith, C., and R. Stulz "The Determinants of Firm's Hedging Policies," *Journal of Financial and Quantitative Analysis*, 20 (1985), 391–406.

Stulz, R. "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis*, 19 (1984), 127–140.

Stulz, R. "Risk management failures: What are they and when do they happen?," *Journal of Applied Corporate Finance*, 20 (2008), 39–48.

Stulz, R. "Governance, Risk Management, and Risk-Taking in Banks," 2014, ECGI Working Paper No 427/2014, http://ssrn.com/abstract=2457947.

Tobin, J. "A Proposal for International Monetary Reform," *Eastern Economic Journal*, 4 (1978), 153–159.

c Zhu, C. "Big Data as a Governance Mechanism," 2018, Working Paper, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3164624.

Online Appendix: Additional proofs and extensions

Proof of Proposition 2, Step 2: Any equilibrium is a threshold equilibrium.

As mentioned in main text, our setup has the same properties (one-sided strategic complementarities) as Goldstein and Pauzner (2005). The proof follows their strategy, and we provide here a simplified demonstration that restricts attention to symmetric pure-strategy equilibria but follows the same steps as the complete proof. We refer the reader to the aforementioned paper for a proof that allows for any possible strategy.

Suppose all traders play an equilibrium strategy that maps their signal s_k into a trading behaviour that can be hasty or deliberate. Given this strategy, for each realization of π, a mass $q(\pi)$ of traders are hasty, which maps one-to-one into a deadline $\tau(\pi)$ at which the opportunity is exhausted. [The additional complexity in the proof in Goldstein and Pauzner (2005) comes from the possibility that $q(\pi)$ could be random when allowing a larger set of strategies.] Hence, while we do not make this dependence explicit to save on notation, the function $\tau(.)$ depends on equilibrium strategies and fully captures for each trader the impact of other traders' strategy on his expected payoff. Specifically,

$$u[s, \tau(.)] \equiv \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s+\varepsilon} \Delta[\tau(\pi), \pi] d\pi$$

denote the net benefit of being deliberate for a trader with a signal s in this equilibrium.

Let s_A denote the signal below which traders are always deliberate, that is

$$s_A \equiv \inf\{s : u[s, \tau(.)] \leq 0\}.$$

The existence of dominance regions guarantees the existence of s_A. Note also that $u[s, \tau(.)]$ is continuous in s, which implies $u[s_A, \tau(.)] = 0$.

Suppose that traders do not follow a threshold strategy. Then, there exists signals $s > s_A$ such that $u[s, \tau(.)] \geq 0$. Let s_B be their infimum:

$$s_B \equiv \inf\{s > s_A : u[s, \tau(.)] \geq 0\}.$$

By continuity again, $u[s_B, \tau(.)] = 0$, and therefore $u[s_A, \tau(.)] = u[s_B, \tau(.)]$, that is,

$$\frac{1}{2\varepsilon} \int_{s_A-\varepsilon}^{s_A+\varepsilon} \Delta[\tau(\pi), \pi] d\pi = \frac{1}{2\varepsilon} \int_{s_B-\varepsilon}^{s_B+\varepsilon} \Delta[\tau(\pi), \pi] d\pi. \tag{OA.1}$$

By definition, for any $s < s_A$, $u[s, \tau(.)] > 0$ and for any $s \in (s_A, s_B)$, $u[s, \tau(.)] < 0$. For $s > s_B$, the sign of $u[s, \tau(.)]$ is indeterminate. The proof consists in showing that (OA.1) cannot hold.

Let $\overline{s}_A \equiv \min\{s_A + \varepsilon, s_B - \varepsilon\}$ and $\underline{s}_B \equiv \max\{s_A + \varepsilon\, s_B - \varepsilon\}$. Cancelling out the (potentially empty) region

$[s_B - \varepsilon, s_A + \varepsilon]$ in (OA.1), one obtains

$$\overline{u}[s_A, \tau(.)] \equiv \frac{1}{2\varepsilon}\int_{s_A-\varepsilon}^{\overline{s}_A} \Delta[\tau(\pi), \pi]d\pi = \frac{1}{2\varepsilon}\int_{\underline{s}_B}^{s_B+\varepsilon} \Delta[\tau(\pi), \pi]d\pi \equiv \overline{u}[s_B, \tau(.)]. \tag{OA.2}$$

Note that the two integrals have the same length: $\overline{s}_A - s_A + \varepsilon = s_B + \varepsilon - \underline{s}_B \equiv d$.

Notice next that since Δ is monotonically decreasing in π, so is the function

$$v_A[\pi, \tau(.)] \equiv \frac{1}{2\varepsilon}\int_{s_A-\varepsilon}^{\overline{s}_A} \Delta[\tau(\tilde{\pi}), \pi]d\tilde{\pi},$$

and $v_A[\overline{s}_A, \tau(.)] < u[s_A, \tau(.)] < v_A[s_A - \varepsilon, \tau(.)]$. Therefore, there exists $\pi_A \in [s_A - \varepsilon, \overline{s}_A]$, such that $v_A[\pi_A, \tau(.)] = \overline{u}[s_A, \tau(.)]$. Similarly, there exists $\pi_B \in [\underline{s}_B, s_B + \varepsilon]$, such that

$$v_B[\pi_B, \tau(.)] \equiv \frac{1}{2\varepsilon}\int_{\underline{s}_B}^{s_B+\varepsilon} \Delta[\tau(\pi), \pi_B]d\pi = \overline{u}[s_B, \tau(.)].$$

Using again the strict monotonicity of Δ in π and $\pi_A < \pi_B$, we get that

$$\overline{u}[s_B, \tau(.)] = v_B[\pi_B, \tau(.)] < v_B[\pi_A, \tau(.)]. \tag{OA.3}$$

The end of the proof consists in showing that $v_B[\pi_A, \tau(.)] \leq v_A[\pi_A, \tau(.)] = \overline{u}[s_A, \tau(.)]$ which, together with (OA.3) contradicts (OA.2) and hence, (OA.1).

Let $\overleftarrow{\tau}(\pi) \equiv \tau(\overline{s}_A + s_A - \varepsilon - \pi)$. $\overleftarrow{\tau}(\pi)$ is the mirror image of $\tau(\pi)$ over $[s_A - \varepsilon, \overline{s}_A]$, that is, when π increases from $s_A - \varepsilon$ to $\overline{s}_A$, $\tau(\pi)$ follows the same path as $\overleftarrow{\tau}(\pi)$ when π decreases from $\overline{s}_A$ to $s_A - \varepsilon$. Hence,

$$v_A[\pi_A, \tau(.)] = \frac{1}{2\varepsilon}\int_{s_A-\varepsilon}^{\overline{s}_A} \Delta[\tau(\pi), \pi_A]d\pi = \frac{1}{2\varepsilon}\int_{s_A-\varepsilon}^{\overline{s}_A} \Delta[\overleftarrow{\tau}(\pi), \pi_A]d\pi$$

Claim 4. *$\overleftarrow{\tau}(.)$ is strictly increasing on $[s_A - \varepsilon, \overline{s}_A]$, and $\tau(.)$ is weakly increasing on $[\underline{s}_B, s_B + \varepsilon]$. Furthermore, $\overleftarrow{\tau}(.)$ increases at a faster rate on $[s_A - \varepsilon, \overline{s}_A]$ than $\tau(.)$ on $[\underline{s}_B, s_B + \varepsilon]$. That is, if $(\pi_a, \pi_b) \in [s_A - \varepsilon, \overline{s}_A] \times [\underline{s}_B, s_B + \varepsilon]$ and $\overleftarrow{\tau}(\pi_a) = \tau(\pi_b)$, then $\overleftarrow{\tau}(\pi_a + \hat{d}) \geq \tau(\pi_b + d)$ for $\hat{d} > 0$.*

Proof. Suppose that $\pi = s_A - \varepsilon$, then agents receive signals in $[s_A - 2\varepsilon, s_A]$ and therefore for all agents (except at s_A), $u[s, \tau(.)] > 0$. It follows that almost all agents are deliberate and therefore $\tau(s_A - \varepsilon) = \overleftarrow{\tau}(\overline{s}_A) = T_d$. Suppose that π increases by $\hat{d}$, then agents with signals in $[s_A - 2\varepsilon, s_A - 2\varepsilon + \hat{d})$ are replaced one for one with agents with signals in $(s_A, s_A + \hat{d}]$. That is, one substitutes agents for whom $u[s, \tau(.)] > 0$ with agents for whom $u[s, \tau(.)] < 0$. As a result $\tau(.)$ (resp. $\overleftarrow{\tau}(.)$) decreases (resp. increases) at the fastest possible rate. Symmetrically, $\tau(.)$ increases on $[\underline{s}_B, s_B + \varepsilon]$, but at a (weakly) slower rate: there can be values of s in $[s_B, s_B + \varepsilon]$ such that $u[s, \tau(.)] < 0$ in which case, as π increases, one substitute hasty agents with other hasty agents, leaving $\tau(.)$ unchanged. □

Claim 5. *For any* $\hat{d} \in [0, d]$, $\tau(\underline{s}_B + \delta) \leq \overleftarrow{\tau}(s_A - \varepsilon + \delta)$.

Proof. Notice that $(s_A, \overline{s}_A]$ and $[\underline{s}_B, s_B)$ have the same measure and $u[s, \tau(.)]$ is always strictly negative on these two segments. By contrast, $[s_A - \varepsilon, s_A)$ and $(s_B, s_B + \varepsilon]$ have the same measure but while $u[s, \tau(.)]$ is always strictly positive on the first segment, it can change signs on the second one. This implies $\tau(\underline{s}_B) \leq \tau(\overline{s}_A) = \overleftarrow{\tau}(s_A - \varepsilon)$. Claim 4 completes the proof. □

Note that if $\Delta(., \pi)$ was monotonically increasing (that is, under global strategic complementarities), Claim 5 would directly imply that $v_A[\pi_A, \tau(.)] \geq v_B[\pi_A, \tau(.)]$.

Claim 6. $v_A[\pi_A, \tau(.)] \geq 0$.

Proof. Note first that the monotonicity result in Claim 4 can be extended: $\tau(.)$ is weakly decreasing on $[\overline{s}_A, s_A + \varepsilon]$. Indeed, as π increases in this interval, one substitutes deliberate traders with deliberate or hasty traders (the latter can be deliberate if $s_A + \varepsilon > s_B$). Let $\hat{\pi}_A \equiv \inf\{\pi \in [s_A - \varepsilon, s_A + \varepsilon] : \Delta[\tau(\pi_A), \pi_A] \leq 0\}$, which is well defined since $u[s_A, \tau(.)] = 0$. Using the single-crossing property of Δ together with the monotonicity of $\tau(.)$ on $[s_A - \varepsilon, s_A + \varepsilon]$ and the fact that $\frac{\partial \Delta}{\partial \pi} < 0$, we get that for any $\pi \in (\hat{\pi}_A, s_A + \varepsilon]$, $\Delta[\tau(\pi), \pi] < 0$ and for any $\pi \in [s_A - \varepsilon, \hat{\pi}_A)$, $\Delta[\tau(\pi), \pi] > 0$. In words the integrand in $u[s_A, \tau(.)]$ is positive below a threshold and negative above it, which together with the fact that

$$u[s_A, \tau(.)] = \int_{s_A - \varepsilon}^{s_A + \varepsilon} \Delta[\tau(\pi), \pi] d\pi = 0$$

implies that

$$\overline{u}[s_A, \tau(.)] = \int_{s_A - \varepsilon}^{\overline{s}_A} \Delta[\tau(\pi), \pi] d\pi \geq 0,$$

which is equivalent to $v_A[\pi_A, \tau(.)] \geq 0$. □

We want to show

$$v_A[\pi_A, \tau(.)] - v_B[\pi_A, \tau(.)] = \int_{s_A - \varepsilon}^{\overline{s}_A} \Delta[\overleftarrow{\tau}(\pi), \pi_A] d\pi - \int_{\underline{s}_B}^{s_B + \varepsilon} \Delta[\tau(\pi), \pi_A] d\pi \geq 0 \tag{OA.4}$$

Suppose that $\tau(s_B + \varepsilon) < T^*(\pi_A)$. Then, since $\tau(.)$ is weakly increasing on $[\underline{s}_B, s_B + \varepsilon]$, for any $\pi \in [\underline{s}_B, s_B + \varepsilon]$, $\tau(\pi) < T^*(\pi_A)$ and therefore $\Delta[\tau(\pi), \pi_A] < 0$. Hence, $v_B[\pi_A, \tau(.)] < 0$, and since, from Claim 6, $v_A[\pi_A, \tau(.)] \geq 0$, (OA.4) holds.

Suppose that $\tau(\underline{s}_B) \geq T^*(\pi_A)$.[OA1] Then $\Delta(., \pi_A)$ strictly increasing for $T > T^*(\pi_A)$ and Claim 5 imply that (OA.4) holds.

[OA1] Note that we implicitly assume $\pi_A > 0$. If $\pi_A \leq 0$, then $\Delta(., \pi_A)$ is monotonically increasing and the result is immediate.

Finally, suppose that $\tau(\underline{s}_B) < T^*(\pi_A) \leq \tau(s_B + \varepsilon)$. Let $\overleftarrow{\pi}(T) \equiv \overleftarrow{\tau}^{-1}(T)$ and

$$\pi(T) \equiv \begin{cases} \underline{s}_B & \text{if } T = \tau(\underline{s}_B) \\ \max\{\pi \in (\underline{s}_B, s_B + \varepsilon] : \tau(\pi) = T\} & \text{if } \tau(\underline{s}_B) < T \leq \tau(s_B + \varepsilon) \end{cases}$$

In words, $\overleftarrow{\pi}(.)$ and $\pi(.)$ are inverse functions of $\overleftarrow{\tau}(.)$ and $\tau(.)$. Therefore from Claim 4, $\overleftarrow{\pi}(.)$ and $\pi(.)$ are increasing and $\pi(.)$ increases faster than $\overleftarrow{\pi}(.)$. Using this notation and Claim 5, rewrite (OA.4):

$$v_A[\pi_A, \tau(.)] - v_B[\pi_A, \tau(.)] = \int_{s_A - \varepsilon}^{\overline{s}_A} \Delta[\overleftarrow{\tau}(\pi), \pi_A] d\pi - \int_{\underline{s}_B}^{s_B + \varepsilon} \Delta[\tau(\pi), \pi_A] d\pi$$

$$= \quad - \int_{\tau(\underline{s}_B)}^{\tau(s_A - \varepsilon)} \Delta[T, \pi_A] d\pi(T) \tag{OA.5}$$

$$+ \int_{\tau(s_A - \varepsilon)}^{T^*(\pi_A)} \Delta[T, \pi_A] d\overleftarrow{\pi}(T) - \int_{\tau(s_A - \varepsilon)}^{T^*(\pi_A)} \Delta[T, \pi_A] d\pi(T) \tag{OA.6}$$

$$+ \int_{\overleftarrow{\pi}[T^*(\pi_A)]}^{\overline{s}_A} \Delta[\overleftarrow{\tau}(\pi), \pi_A] d\pi - \int_{\pi[T^*(\pi_A)]}^{s_B + \varepsilon} \Delta[\tau(\pi), \pi_A] d\pi. \tag{OA.7}$$

This equation is a decomposition of the two integrals in (OA.4) along the interval $[\tau(\underline{s}_B), \tau(\overline{s}_A)]$. Note that this decomposition assumes $\tau(s_A - \varepsilon) \leq T^*(\pi_A)$, the demonstration would be a fortiori true if $\tau(s_A - \varepsilon) > T^*(\pi_A)$.

At the bottom, (OA.5) corresponds to the part of the integral in v_B with values of T below the lowest T in v_A, that is, $\tau(s_A - \varepsilon)$. These values are below $T^*(\pi_A)$, therefore this part of v_B is negative and (OA.5) is strictly positive. In the interval $[\tau(s_A - \varepsilon), T^*(\pi_A))$ integrands in both v_A and v_B are negative as T is still below $T^*(\pi_A)$. However, $d\overleftarrow{\pi}(T) \leq d\pi(T)$, implies that (OA.6) is positive (Intuitively, v_B visits any negative values of Δ that v_A takes but "stays longer" at each of them.) Finally, (OA.7) corresponds to values of T above T^*. Note first that $\tau(s_A - \varepsilon) \geq \tau(\underline{s}_B)$ (Claim 5) and $\overleftarrow{\pi}(.)$ increasing more slowly than $\pi(.)$ imply $\overleftarrow{\pi}[T^*(\pi_A)] - s_A + \varepsilon \leq \pi[T^*(\pi_A)] - \underline{s}_B$, and therefore, $\overline{s}_A - \overleftarrow{\pi}[T^*(\pi_A)] \geq s_B - \varepsilon - \pi[T^*(\pi_A)]$. (i.e., the LHS integral takes a larger range of values of π than the RHS integral.) In addition, from Claim 4, for $\delta > 0$, $\overleftarrow{\tau}\{\overleftarrow{\pi}[T^*(\pi_A)] + \delta\} \geq \tau\{\pi[T^*(\pi_A)] + \delta\} \geq 0$, and $\Delta(., \pi_A)$ is strictly increasing for $T \geq 0$, (Lemma 1). Hence, (OA.7) is positive (intuitively, $\overleftarrow{\tau}(.)$ shifts more weight towards high values of Δ than $\tau(.)$.)

This eventually shows $v_B[\pi_A, \tau(.)] \leq v_A[\pi_A, \tau(.)] = \overline{u}[s_A, \tau(.)]$, which together with $\overline{u}[s_B, \tau(.)] = v_B[\pi_B, \tau(.)] < v_B[\pi_A, \tau(.)]$ shows $\overline{u}[s_B, \tau(.)] < \overline{u}[s_A, \tau(.)]$, a contradiction. □

Exogenous deadline

We present here the extension with an exogenous deadline informally discussed in Section 4.1. Suppose that the trading opportunity disappears at the first of these two times: (1) an endogenous deadline T at which the mass of traders who have executed the trade reaches I and the opportunity is depleted; (2) an exogenous deadline X, exponentially distributed with intensity $1/\chi$.

Under the hasty strategy, the implementation probability becomes

$$\begin{aligned} p_h(T) &= \int_0^T e^{-\frac{t}{\chi}} e^{-\frac{t}{\lambda}} \frac{1}{\lambda} dt \\ &= \frac{\chi}{\lambda+\chi}\left[1-e^{-\left(\frac{1}{\lambda}+\frac{1}{\chi}\right)T}\right]. \end{aligned}$$

Under the hasty strategy, the implementation probability is 0 if $T \leq \iota$ and otherwise,

$$\begin{aligned} p_d(T) &= \int_0^{T-\iota} e^{-\frac{t+\iota}{\chi}} e^{-\frac{t}{\lambda}} \frac{1}{\lambda} dt \\ &= \frac{\chi e^{-\frac{\iota}{\chi}}}{\lambda+\chi}\left[1-e^{-\left(\frac{1}{\lambda}+\frac{1}{\chi}\right)(T-\iota)}\right]. \end{aligned}$$

The definition of the net (private) benefit of risk management is unchanged:

$$\Delta(T) = p_d(T)\rho(\pi+\alpha_+) - p_h(T)\pi.$$

Note also that if $T > \iota$,

$$\frac{\partial \Delta}{\partial T} = \frac{1}{\lambda} e^{-\left(\frac{1}{\lambda}+\frac{1}{\chi}\right)T}\left[e^{\frac{\iota}{\lambda}}\rho(\pi+\alpha_+)-\pi\right] > 0.$$

Since the shape of $\Delta(.)$ is unchanged, Lemma 1 still holds, and there exists a unique threshold T^* such that a trader chooses risk management if and only if $T > T^*$, and T^* is an increasing function of π. T^* can be expressed as a function of π. In particular, if $\pi < 0$, it is optimal to be deliberate, and $T^* = 0$. Conversely,

$$\lim_{T\to+\infty} \Delta(T) = \frac{\chi}{\lambda+\chi}\left[\pi - \rho e^{-\frac{\iota}{\chi}}(\pi+\alpha_+)\right], \tag{OA.8}$$

and hence, if the expression between brackets is positive, then RM is never profitable, even when T is arbitrarily large. That is, rearranging (OA.8), if $\pi > -\frac{1-\rho}{e^{\iota/\chi}-\rho}\alpha_-$, the hasty strategy is dominant and $T^* = +\infty$. In between these two bounds, T^* is a strictly increasing function of π.

The bounds $\underline{\pi}'$ and $\overline{\pi}'$ can be defined in the same way as in the main text, and one shows that the deliberate equilibrium exists iff $\pi < \overline{\pi}'$, while the hasty equilibrium exists if $\pi > \overline{\pi}'$. It is however no longer true that the deliberate strategy is collectively optimal for traders. Indeed, the difference in aggregate surplus between every trader being deliberate and every trader being hasty is

$$\begin{aligned} &\left[\int_\iota^{T_d} \frac{1}{\chi} e^{-\frac{t}{\chi}}\left(1-e^{-(t-\iota)/\lambda}\right)dt + e^{-\frac{T_d}{\chi}} I\right]\rho(\pi+\alpha_+) \\ &-\left[\int_0^{T_h} \frac{1}{\chi} e^{-\frac{t}{\chi}}\left(1-e^{-t/\lambda}\right)dt + e^{-\frac{T_h}{\chi}} I\right]\pi \end{aligned} \tag{OA.9}$$

For every realization of π, there exists ι sufficiently large relative χ that (OA.9) becomes negative. That is, while in the baseline case, the aggregate profits when all traders are deliberate arbitrarily was independent of ι, it can now be made arbitrarily small by taking ι sufficiently large.

However, the coordination failure remains. Indeed, when both equilibria co-exist, the deliberate one is Pareto-optimal. To see this, notice that a trader's profit in the deliberate equilibrium has to be higher under the deliberate strategy than under the hasty strategy keeping the (endogenous) deadline unchanged, that is, equal to T_d. The profit from a hasty strategy is, in turn, higher under the deliberate deadline T_d than under the hasty deadline T_h. It follows that every trader is better off in the deliberate equilibrium than in the hasty one.

Decision making in continuous time

We consider here a variation of the baseline model in Section 2 where the trader can adjust his risk management decision "on the fly" as the trading game unfold. This extension is informally discussed in Section 4.4. We show that, as in the original version of the model, strategic complementarities generate socially inefficient risk management decisions. Interestingly, in this specification, a unique pure-strategy equilibrium obtains even without perturbing the model with dispersed information.

We introduce two modifications to the original model. First, traders, instead of being committed to a risk management decision, can execute the trade at any point in time after they have located the trading opportunity. That is, a trader k who identifies the potential trade and engages into risk management can decide to execute the trade before the procedure is completed and produces information on the private value α_k. Second, we assume that the risk management time frame, instead of being deterministic (equal to ι) follows an exponential distribution with intensity $\frac{1}{\iota}$. This second change introduces continuity through time in the expected payoff from risk management: at any t, conditional on having located the trading opportunity before t, trader k has a probability $\frac{1}{\iota}dt$ of learning α_k between t and $t + dt$.

Given these assumptions, a trader who finds the investment opportunity at time t enters into a continuous decision process: at any subsequent time, he needs to decide wether to execute the trade or extend the risk management process for another "small" period dt in the hope of learning his private value α_k. This decision process stops either with the trader executing the trade before learning his private value, or with the trader learning α_k and optimizing his trading decision, or with the trading opportunity disappearing before the trader had a chance to implement his strategy. To simplify the exposition, we assume that traders do not observe trades by other traders. In other words, we allow traders' strategy to depend only on time.[OA2] Since, conditional on learning α_k, trader k has a dominant option, a strategy only needs to specify the investigation period following the discovery, after which he

[OA2] One can in fact show that allowing for the trader's strategy to depend on the mass of executed up to time t does not change the results.

executes the trade if he has not learnt α_k. Note that a trader's decision to execute a trade without learning α_k does not depend on the point in time at which the opportunity was located. Note also that if a trader is willing to execute a trade at time t without knowing α_k, he will a fortiori be willing to do so at any subsequent time. Hence, trader's k strategy can be captured in a single variable τ_k. If trader k finds the investment opportunity before τ_k but does not learn α_k before τ_k, he executes the trade at τ_k. If trader k finds the investment opportunity after τ_k, he executes the trade immediately.

T_h is defined as in the previous sections, but because the investigation time is now random, the definition of T_d (the time at which the trading opportunity is depleted if traders never execute without learning α_k) changes. Specifically, T_d solves

$$\rho\left(1-\frac{\frac{1}{\iota}e^{-\lambda T_d}-\lambda e^{-\frac{1}{\iota}T_d}}{\frac{1}{\iota}-\lambda}\right)=i.$$

Note that, instead of being a binary decision as in the original specification, risk management is now a continuous variable: the higher τ_k, the more likely it is that trader k makes an informed trading decision. When $\tau_k = T_d$, trader k is fully deliberate: he never trades without knowing α_k. We show the following result.

Proposition OA.1. *There is a unique pure-strategy equilibrium. In this equilibrium $\tau_k = \tau^\star$ for all k, with $T_h < \tau^\star < T_d$ and $\tau^\star$ is a decreasing function of i.*

Proof. Notice first that if $t < T_h$, investigating is a dominating strategy: even if every firm follows the hasty strategy, the investment opportunity will not be fully exhausted until T_h. Hence we can delete all strategies $\tau_j < T_h$.

Given that any firm that finds the opportunity before T_h investigates until T_h, at T_h, the fraction of firms that have found the investment opportunity and do not know their α_k is

$$m(T_h)=1-e^{-T_h/\lambda}-\left(1-\frac{\frac{1}{\iota}e^{-T_h/\lambda}-\lambda^{-1}e^{-\frac{1}{\iota}T_h}}{\frac{1}{\iota}-\lambda^{-1}}\right)$$

while the size i of the investment opportunity is

$$i(T_h)=i-\rho\left(1-\frac{\frac{1}{\iota}e^{-T_h/\lambda}-\lambda^{-1}e^{-\frac{1}{\iota}T_h}}{\frac{1}{\iota}-\lambda^{-1}}\right).$$

Note that $m(T_h) < i(T_h)$. Intuitively, this reflects the fact that we derived T_h assuming that all firms would follow the hasty strategy, while firms will in fact follow a strategy where they investigate at least until T_h (if they do not learn α_k before T_h). We move now to the next round of deletion of dominated strategy.

Suppose that all firms play the hasty strategy from T_h on. Then the mass of investment between T_h (included) and $t > T_h$ is

$$m(T_h) + e^{-T_h/\lambda} - e^{-t/\lambda}.$$

Hence, it is a dominating strategy to choose $\tau \geq T_h^1$ where T_h^1 solves

$$m(T_h) + e^{-T_h/\lambda} - e^{-T_h^1/\lambda} = i(T_h) \Leftrightarrow 1 - e^{-(T_h^1 - T_h)/\lambda} = e^{T_h/\lambda}[i(T_h) - m(T_h)]$$

By continuing to iterate, we obtain an increasing sequence $(T_h^n)_n$ which converges. Hence the only strategies that survives iterated deletion of strictly dominated strategies are such that $\tau_j \geq T_h^\infty$, where $\tau_j \geq T_h^\infty$ is solution to $i(T) - m(T) = 0$, that is,

$$1 - e^{-T_h^\infty/\lambda} - \rho\left(1 - \frac{\frac{1}{\iota}e^{-T_h^\infty/\lambda} - \lambda^{-1}e^{-\frac{1}{\iota}T_h^\infty}}{\frac{1}{\iota} - \lambda^{-1}}\right) = s_0. \tag{OA.10}$$

(OA.10) has a unique solution if $i < \rho$, and one can show $T_h < T_h^\infty < T_d$.

This concludes the first part of the proof.

Conjecture an equilibrium where a strictly positive mass of firms choose to investigate at $t = T_h^\infty$. Let i_t denote the size of the investment opportunity at t. Let p_t denote the ratio of the size of the investment opportunity to the mass of firms that have found the opportunity but have not invested yet. By definition of T_h^∞, $p_t < 1$ if $t > T_h$ and $p_t = 1$ if $t = T_h^\infty$.

Finally, let T denote the time at which the investment opportunity is depleted. In equilibrium, the following condition must be true

$$T = \max_k \tau_k \tag{OA.11}$$

Suppose indeed that there exists τ_j such that $\tau_j > T$ (intuitively, firms keep deliberating until after the investment opportunity is depleted). For dt "small", the net benefit of waiting at $T - dt$ is

$$\frac{1}{\iota}dt\frac{\pi + 1}{2} - \pi,$$

which becomes negative for dt small enough. Hence, $\tau_k > T$ cannot be an equilibrium strategy. Thus, for all j, $\tau_k \leq T$, which implies that all firms that have found the opportunity before T invest at T at the latest. Since $p_t \leq 1$ for $t \geq T_h^\infty$, it implies that the investment opportunity is fully depleted exactly at $\max_k \tau_k$.

The net benefit of waiting at $T - dt$ for the firm that plays T is

$$\frac{1}{\iota}dt\frac{\pi + 1}{2} + (1 - \frac{1}{\iota}dt)\pi p_T - \pi,$$

which becomes negative for dt small enough unless $p_T = 1$. Thus the only possible equilibrium is $\tau_k = T_h^\infty$ for all

k. □

Proposition OA.1 is consistent with the conclusions of the original model. First, externalities between traders create a coordination failure. While it would be optimal for them too coordinate on a fully deliberate strategy, $\tau_k = T_d$, traders sometimes execute trades without learning α_k. Second, competitive pressure intensifies this inefficiency: when i goes down and the preemption motive becomes more stringent, traders spend less time deliberating in equilibrium. Third, in spite of externalities, there can still be some risk management, that is, $\tau_k > 0$ in equilibrium. To understand this, consider the case of a trader who locates the trading opportunity at the very beginning of the game ($t = 0$). Even if he anticipates that every other trader will execute the trade as soon as he finds it, he knows that the investment opportunity cannot be fully depleted before T_h. In other words, it is a strictly dominant strategy to be deliberate between 0 and T_h, and one can delete strategies τ_k smaller than T_h. But now, given that a fraction of traders who found the opportunity before T_h have realized that their private value was low by T_h and exited the market without trading, at T_h, the mass of traders who have identified the trading opportunity and have not executed the trade yet is strictly smaller than the size of the current investment opportunity. Hence, it is again a dominant strategy to deliberate a little longer. Reiterating this deletion of strictly dominated strategies, one can construct a series of thresholds that converge to $\tau^\star$. At $\tau^\star$, the mass of traders who have identified the trading opportunity and have not executed the trade yet is exactly equal to the size of the current size of the investment opportunity. The end of the proof consists in showing that there cannot be an equilibrium strategy strictly higher than this threshold.

Position Limits

We explore here a variation of the baseline model mentioned in Section 2.1, where a positive signal from risk management allows a trader to trade more units of the underlying asset. Specifically, assume each trader is restricted to trading one unit without risk management, but under active risk management, if the private value is found to be positive, i.e., $\alpha = \alpha^+$, the trader is allowed to trade $n > 1$ units. In what follows, we trace out the implications of this model change.

Strategies remain monotonic. For a given deadline T, the payoff from the hasty strategy remains the same as in the baseline model, $p_h(T)\pi$. However, the payoff from the deliberate strategy is now multiplied by n and becomes $p_d(T)n\rho(\pi + \alpha_+)$. Let $\Delta(T, n)$ denote the net value of risk management in this setting, so it is $\Delta(T, 1)$ for the baseline model. Then,

$$\begin{aligned}\Delta(T,n) &\equiv p_d(T)n\rho(\pi+\alpha_+) - p_h(T)\pi \\ &= \Delta(T,1) + p_d(T)(n-1)\rho(\pi+\alpha_+). \qquad \text{(OA.12)}\end{aligned}$$

In comparison to the baseline model, there is one more term $p_d(T)(n-1)\rho(\pi+\alpha_+)$, which captures the additional benefit of risk management in this setting is that the trader is allowed to trade $n-1$ more units if the private value turns out to be positive and the strategy can be implemented. Recall from the baseline model that $\Delta(T,1)$ is U-shaped and reaches its minimum at $T=\iota$, at and to the left of which point the payoff from the deliberate strategy is 0. Now note that the additional term in $(OA.12)$ also equals 0 for all $T \leq \iota$ and is strictly increasing in T for $T \geq \iota$, for all $n>1$. Hence, like $\Delta(T,1)$, $\Delta(T,n)$ is negative for $T \leq \iota$ and strictly increasing for $T \geq \iota$, converging to $n\rho(\pi+\alpha_+)-\pi>0$ in the limit $T \to \infty$. As a result, strategies remain monotonic in T for all $n>1$: a trader plays the deliberate strategy iff T exceeds a unique threshold T^* (which decreases in n). In aggregate, this means that the fraction of hasty traders $q(T,n)$ monotonically decreases in T for all $n \geq 1$, as in the baseline model.

Time pressure may decrease in the fraction of hasty traders for large n. As in the baseline model, let q denote the fraction of hasty traders. Time T by which the trading opportunity is exhausted now satisfies

$$qp_h(T)+(1-q)\rho p_d(T)n=i. \tag{OA.13}$$

The function $T(q,n)$, defined by equation (OA.13), can be decreasing in q, in which case the equilibrium analysis is essentially the same as in the baseline model. For example, consider speed parameters such that $p_h(T)-p_d(T)$ is large and $p_d(T)$ is close to zero for all $T(q,n)$ with $q \in [0,1]$.[OA3] It can then be that $\rho p_d(T)n < p_h(T)$ for all $T=T(q,n)$ with $q \in [0,1]$, which in turn would imply that $T(q,n)$ is decreasing in q. However, unlike in the baseline model, $T(q,n)$ need not always be decreasing. Depending on parameters, it can (in part or everywhere) be increasing in q for large enough n. This reduces the scope for equilibrium multiplicity. In particular, if $T(q,n)$ monotonically increases in $q \in [0,1]$, the equilibrium is always unique, as we discuss below.

Equilibrium may be unique, but still exhibits scope for coordination failure. For parameter constellations under which $T(q,n)$ is increasing for all $q \in [0,1]$, any intersection between $q(T,n)$ and $T(q,n)$ with $q \in [0,1]$ is always unique. To see, this, let $T^{-1}(q,n)$ be the inverse of $q(T,n)$, which is monotonically decreasing for $q \in [0,1]$. Then, $T(q,n)-T^{-1}(q,n)$ is monotonically increasing and hence assumes the value zero at most once for $q \in (0,1)$. If it does, the associated fraction $q^* \in (0,1)$ denotes the equilibrium fraction of traders that play the hasty strategy; otherwise, either the hasty or the deliberate strategy is strictly dominant (see Figure 3 below). At any rate, however, the equilibrium is unique. Having said that, the scope for coordination failure remains: Both corner equilibria where $q^*=0$ and interior equilibria where $q^* \in (0,1)$ involve the play of hasty strategies even though aggregate trader welfare is highest for $q=1$. Moreover, since traders' payoffs are symmetric in all equilibria – either because they play the same strategy (corner equilibria) or are indifferent between them (interior equilibria)

[OA3]This requires that internal speed $1/\iota$ is sufficiently smaller than external speed $1/\lambda$. Intuitively, these are environments in which markets are fast and risk management is very disruptive.

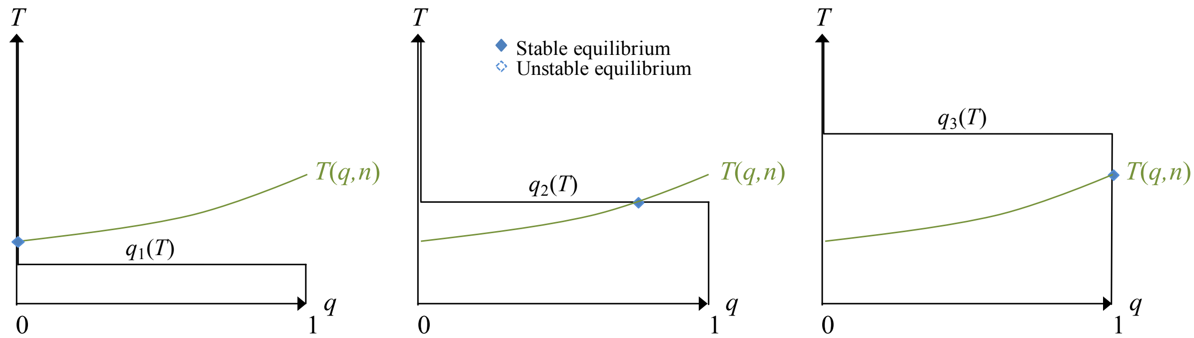

Figure 3: Equilibrium uniqueness when $T(q, n)$ is increasing in q. $T(q, n)$ depicts how time pressure depends on the firms' risk management choices, while the step function $q(T)$ depicts the firms' optimal risk management choice as a function of time pressure for various levels of the common value π (i.e., the ex ante value of a trade). Higher $q(T)$ correspond to higher values of π. Intersections between $q(T)$ and $T(q)$ constitute equilibria. In the first two graphs, the equilibrium is constrained inefficient.

– the equilibrium outcomes with $q < 1$ are Pareto-dominated by the payoffs under $q = 1$.

Is the assumed trade restriction privately optimal? It should be noted that in the above setting, the exogenous restriction of trading to one unit in the absence of risk management, although the ex ante value of each additional unit is positive, is suboptimal from the firm's point of view; even without risk management, the firm would want the trader to trade all n units. A more coherent assumption for "position limits" that may be relaxed pending risk management approval would be that n_s units can be traded without approval because this amount of trading does not raise any critical concerns, but trading n_r additional units may require approval because it raises risk concerns to a point where the firm may want to check its exposure prior to trading. All firms with or without risk management would allow their traders to trade n_s units. If the *ex ante* value of the n_r additional units is negative (without further information), then no firm would permit those trades without risk management approval (case A). If the *ex ante* value of the additional units is positive, firms face a choice: a firm can still require approval in order to ascertain that the trades are attractive, i.e., have positive private value (case B), or it may discard the approval process in favor of fast execution, effectively relaxing the position limits imposed on traders (case C). Our model is a reduced-form representation of the above with n_s and n_r, respectively, normalized to 0 and 1. The three cases are captured by $\pi < 0$ (case A), $\pi > 0$ with the deliberate strategy (case B), and $\pi > 0$ with the hasty strategy (case C). As our model illustrates, the interesting strategic interactions concern whether firms, when $\pi > 0$, give traders' discretion to trade the n_r additional units or require approval – or put differently, whether they relax traders' position limits when market conditions are good. Our results show that, with respect to these interactions, the deadline T will monotonically decrease in the fraction of firms that dispense with the approval requirement, generating a positive feedback loop and thus strategic complementarity.

Speed investments

In our model, the technological speed parameters λ and ι are exogenous. Here, we briefly discuss some of the basic effects that endogenous investments in speed would have on the analysis, with a focus on establishing that the concerns about coordination failure would remain.

As shown in Section 4.2, strategic complementarities also emerge if risk management involves fixed costs rather than latency, and an increase in either type of cost makes coordination failure more likely. Now, consider situations in which these costs are inversely related. In particular, suppose firms can shorten risk management latency from ι to $\hat{\iota} < \iota$ by paying an upfront cost $k > 0$. Intuitively, this should make a hasty equilibrium more difficult to sustain.

To see this, focus on the parameter constellations for which the hasty equilibrium and the deliberate equilibrium coexist given a delay of ι:

$$\Delta(T_h) < 0 < \Delta(T_d). \tag{OA.14}$$

Consider the impact of the investment k at both "ends" of (OA.14), beginning with the left inequality, that is, the viability of the hasty equilibrium. Let $\Delta^{\hat{\iota}}(.)$ denote the net value of risk management to a trader who has made the investment to shorten the delay to $\hat{\iota}$. This investment is profitable if the payoff, net of k, from being deliberate with a risk-management delay of $\hat{\iota}$ is higher than the payoff from being hasty. That is, if $\Delta^{\hat{\iota}}(T_h) - k > 0$, the hasty equilibrium unravels: At least one trader strictly gains from making the investment and activating risk management, and as deviations extend the deadline (i.e., reduce time pressure), everyone deviates. While this validates the intuition that allowing for investments in risk-management speed favours deliberate strategies, it does not mean that there is no coordination failure. Indeed, these investments are deadweight costs, too, given a deliberate equilibrium equilibrium exists also in their absence. This new inefficiency is even more apparent at the other end of (OA.14). If $\Delta^{\hat{\iota}}(T_d) - k > \Delta(T_d)$, the deliberate equilibrium *without* investment collapses, and a less efficient deliberate equilibrium *with* investment emerges instead. Essentially, rather than vanish, the coordination problem morphs from a "race to the bottom" (lack of risk management) to an "arms race" (speed investments).[OA4]

Going further, arms race and inefficient risk management may compound each other once we allow for investments in external (search) speed λ^{-1}, as well as in risk-management speed. In general, the question of which type of speed firms would invest in has no simple answer. As a start, this depends on the cost and "speed-per-dollar" returns on the two types of investments, as determined by technological and institutional factors. To complicate matters further, these investments are complements when the firm plays the deliberate strategy. However, introducing the

[OA4]This seems loosely related to the insight in Hellwig and Veldkamp (2009) that information acquisition choices (here, risk management) inherit their strategic dependencies from the *underlying* game. In our framework, the key externalities originate from preemptive competition: If others are fast, individuals' incentives to be fast increase. This feedback loop is then reflected in the information strategies, be it as a race to the bottom or as an arms race, yet in either case driven by the objective of becoming faster.

possibility of investments in external speed has two clear-cut effects that bolster the existence of hasty equilibria.

First, suppose the firm in the previous example can alternatively deploy k to raise its search speed from λ^{-1} to $\hat{\lambda}^{-1}$. That is, the firm chooses between $(\lambda, \hat{\iota})$ and $(\hat{\lambda}, \iota)$. Further consider the case $\Delta^{\hat{\iota}}(T_h) - k > 0$ where, barring investments in λ, the hasty equilibrium unravels. By (1), this condition can be written

$$p_d^{\hat{\iota}}(T_h)\rho(\pi + \alpha_+) > p_h^{\hat{\iota}}(T_h)\pi + k \tag{OA.15}$$

where $p_s^{\hat{\iota}}(.)$ denotes the implementation probability under strategy s given the speed parameters $(\lambda, \hat{\iota})$, and T_h is the deadline if everyone is hasty given those parameters. Suppose a firm anticipates all other firms to invest in $(\hat{\lambda}, \iota)$ and be hasty with the resultant deadline $T_h^{\hat{\lambda}}$. Conditional on investing k, it is a best response to also pick $(\hat{\lambda}, \iota)$ and be hasty if and only if

$$p_d^{\hat{\iota}}(T_h^{\hat{\lambda}})\rho(\pi + \alpha_+) < p_h^{\hat{\lambda}}(T_h^{\hat{\lambda}})\pi. \tag{OA.16}$$

Observe that $p_h^{\hat{\lambda}}(T_h^{\hat{\lambda}}) = p_h^{\hat{\iota}}(T_h) = i$ given either probability involves everyone being hasty with the same search speed. In addition, $p_d^{\hat{\iota}}(T_h^{\hat{\lambda}}) < p_d^{\hat{\iota}}(T_h)$ because $T_h^{\hat{\lambda}} < T_h$. It follows that for k sufficiently small, i.e., when the return to investing in either risk-management or search speed is high, (OA.15) can hold even if (OA.16) is true. In that case, the hasty equilibrium reemerges, now coupled with investment in external speed.[OA5] In other words, once firms can pick which type of speed to invest in, the strategic complementarities span risk management and investment: Firms that believe others to invest in external speed in order to be hasty are inclined to follow suit.

The second effect that favors hastiness is that investments in internal speed only pay off when the deliberate strategy is played, whereas investments in external speed improve both strategies; in the example, $p_d^{\hat{\lambda}}(T) > p_d(T)$ and $p_h^{\hat{\lambda}}(T) > p_h(T)$ whereas $p_d^{\hat{\iota}}(T) > p_d(T)$ and $p_h^{\hat{\iota}}(T) = p_h(T)$, for any T. Whenever there is ambiguity about which strategy will be played in equilibrium – due to equilibrium multiplicity in the baseline model or due to fundamental uncertainty in the global games version – this biases investments toward external speed. This bias reflects that the search for opportunities is the *primary* activity, while reviewing them is *secondary* to successful search. Put another way, if investments in search or risk-management speed had the same effect on the overall speed of the deliberate strategy, i.e., if $p_d^{\hat{\lambda}}(T) = p_d^{\hat{\iota}}(T)$, then the intrinsic primacy of search would make it *weakly dominant* to pick $(\hat{\lambda}, \iota)$, that is, to maximize external speed, and *strictly dominant* for any positive probability, however small, that the hasty strategy is played. (Notably, this is independent of (OA.14), (OA.15), or (OA.16).) The robust set of equilibria would hence be that under speed parameters $(\hat{\lambda}, \iota)$, which maximally skew the speed ratio $\frac{\iota}{\hat{\lambda}}$ toward generating risk management failures. In this case, the arms race *fuels* the race to the bottom.

[OA5]Small k also satisfy that the firm does not prefer being hasty without the investment: $p_h^{\hat{\lambda}}(T_h^{\hat{\lambda}})\pi - k > p_h(T_h^{\hat{\lambda}})\pi$. Being deliberate without the investment is suboptimal: $\Delta(T_h^{\hat{\lambda}}) < \Delta(T_h)$ because $T_h^{\hat{\lambda}} < T_h$, and $\Delta(T_h)$ is negative by (OA.14). Last, it is also suboptimal for the firm to be deliberate with an investment in $(\hat{\lambda}, \iota)$ instead: $\Delta^{\hat{\lambda}}(T_h^{\hat{\lambda}}) < \Delta(T_h)$, since either side involves everyone being hasty with uniform speed, $p_h^{\hat{\lambda}}(T_h^{\hat{\lambda}}) = p_h(T_h)$, but the speed ratio is higher on the left side, $p_h^{\hat{\lambda}}(T^{\hat{\lambda}}) - p_d^{\hat{\lambda}}(T^{\hat{\lambda}}) > p_h(T_h) - p_d(T_h)$, which makes risk management even less attractive.

To summarize the central insights from the above discussion, when firms can invest in speed, coordination failures can also be manifested in terms of over-investment, and more, there is an intrinsic advantage of investing in external rather than internal speed, which further stacks the market environment against risk management. Normatively, this means that not only may there be too much investment in speed overall (as already noted by Budish, Cramton and Shim (2015) and Biais, Foucault and Moinas (2015)), but conditional thereupon, also that firms might have insufficient incentives to channel the investment into *internal* (risk management) speed.