

A Theory of Operational Risk

SULEYMAN BASAK
London Business School and CEPR

ANDREA M. BUFFA
Boston University

February 2016*

Abstract

We study the dynamic decision making of a financial institution in the presence of a novel implementation friction that gives rise to operational risk. We distinguish between internal and external operational risks depending on whether the institution has control over them. Internal operational risk naturally arises in the context of model risk, as the institution exposes itself to operational errors whenever it updates and improves its investment model. In this case, it is no longer optimal to implement the best model available, thus leaving scope for endogenous deviation from it, and hence model sophistication. We show that the optimal exposure to operational risk may well become decreasing in the level of internal operational risk, which in turn makes the exposure to market risk less volatile. We uncover that financial constraints interact with operational risk, whether internal or external, and prompt the institution to always adopt a more sophisticated model. While such constraints are always detrimental when operational risk is internal, they may be beneficial, despite inducing an excessive level of sophistication, when it is external.

JEL Classification: G11, G20, D90, C61.

Keywords: Operational risk, risk management, model sophistication, risk exposure, operational losses, financial constraints.

*Contacts: sbasak@london.edu, buffa@bu.edu. We thank Rui Albuquerque, Patrick Bolton, Jason Donaldson, Denis Gromb, Konstantin Milbradt, Dimitris Papanikolaou, Giorgia Piacentino, Pierre-Olivier Weill, as well as seminar and conference participants at Boston University, USI Lugano, AFA and European Winter Finance for helpful comments and discussions.

Systems enhancements and updates, as well as the requisite training, entail significant costs and create risks associated with implementing new systems and integrating them with existing ones.
Goldman Sachs, 2014 Form 10-K

1 Introduction

Operational risk has always been present, but in the last 20 years, with rapid changes in the financial industry leading to larger and more complex financial institutions, a widespread concern has grown significantly. Jorion (2007) refers to it as the most pernicious form of risk because of its contribution to numerous failures in financial institutions. The Basel Committee on Banking Supervision (2001) defines operational risk as:

*The risk of loss resulting from inadequate or failed internal processes, people, and systems or from external events.*¹

Operational risk is considered internal if the financial institution has control over it, and external if it is due to uncontrollable events such as natural disasters, security breaches, political risk (e.g., see Hull, 2012). Well publicized cases of large operational losses include: Salomon Brothers (\$303 million, 1993), Knight Capital (\$460 million, 2012) and Goldman Sachs (undisclosed amount, 2013) for changes in computing technology and programming errors; Bank of America (\$225 million, 1983), Wells Fargo Bank (\$150 million, 1996) and Freddie Mac (\$207 million, 2001) for systems integration and transaction processing failures; Barings Bank (£830 million, 1995), Société Générale (€4.9 billion, 2008), and UBS (\$2.3 billion, 2011) for rogue trading. Even though large and infrequent operational losses typically make it to the news, operational risk is for the majority part induced by small and frequent operational errors. Not surprisingly, the literature on operational risk is growing rapidly (as discussed below), but it mostly focuses on measurement issues and statistical properties of operational losses. Little is known about the interaction between operational and other risks that financial institutions face, and the choices they make to simultaneously manage the exposure to these risks.

Our goal is to undertake a comprehensive analysis of the decision making of a financial institution in the presence of operational risk. In particular, we study the implications of

¹The Basel Committee on Banking Supervision (2011) also lists eleven Principles (covering governance, the risk management environment and the role of disclosure) to provide guidance to banks on the management of operational risk. The Basel Committee on Banking Supervision (2014) reviews the progress of 60 systemically important banks in implementing the proposed Principles.

operational risk for an institution’s optimal investment decisions within a simple standard dynamic asset allocation framework. To our knowledge, ours is the first attempt to directly embed operational risk into such a familiar framework. The key feature of our work is the presence of a novel implementation friction that gives rise to operational risk. In particular, within our setup, a financial institution makes its decisions based on an investment model, but it has incomplete information on what the true model for investment decisions should be. This induces model risk, and prompts the institution to update and improve its investment model with the arrival of new information over time. When implementing the changes to its model, the financial institution exposes itself to uncertain operational errors (e.g., bugs in programming codes, mistakes in data collection, computer failures, and execution errors), and hence faces an additional source of risk, operational risk. This implies that, if the institution tries to implement a certain model, it will ultimately implement a version of that model which contains operational errors.²

A form of internal operational risk arises when inadequate implementations are more likely to occur at times in which the financial institution tries to implement a very sophisticated model (i.e., a model incorporating larger updates). Operational risk is internal in this case because the institution can control the exposure to this risk by reducing the sophistication of its model. This originates a novel trade-off between model sophistication and operational risk: between a model that is more likely to be profitable, and one that is less likely to contain operational errors. When, instead, an inadequate implementation of the investment model occurs for reasons unrelated to its sophistication, the above trade-off is absent and the corresponding operational risk is external. Our model conveniently nests both cases of internal and external operational risks, and delivers a tractable analysis, admitting closed-form solutions.

Solving for the optimal model sophistication, we demonstrate that the financial institution always adopts the most sophisticated model possible when operational risk is external. This is because compromising the sophistication of its model would not help reduce the exposure

²As disclosed in the 2014 Form 10-K, Goldman Sachs highlights the importance of operational risk with explicit reference to various operational errors: “our businesses ultimately rely on human beings as our greatest resource, and from time-to-time, they make mistakes that are not always caught immediately [...]. These can include calculation errors, mistakes in addressing emails, errors in software development or implementation, or simple errors in judgment. [...] Human errors, even if promptly discovered and remediated, can result in material losses and liabilities for the firm.” Typical cases of operational errors include spreadsheet errors. Recently, JPMorgan reported billions of dollars of losses induced by spreadsheet errors in its VaR model towards the end of January 2012. From the JPMorgan Task Force Report (2013): “[...] further errors were discovered in the Basel II.5 model, including, most significantly, an operational error in the calculation of the relative changes in hazard rates and correlation estimates. Specifically, after subtracting the old rate from the new rate, the spreadsheet divided by their sum instead of their average, as the modeler had intended. This error likely had the effect of muting volatility by a factor of two and of lowering the VaR [...]” “[...] additional operational issues became apparent. For example, the model operated through a series of Excel spreadsheets, which had to be completed manually, by a process of copying and pasting data from one spreadsheet to another.”

to operational risk. In contrast, we show that the institution optimally implements a less sophisticated model when operational risk is internal. This reduces the likelihood of operational errors, and hence operational losses. We find the key determinants of the endogenous model sophistication to be operational and model risks. These risks affect model sophistication in opposite ways. If operational risk is high, the financial institution has a strong incentive to reduce its exposure to this risk by updating its model less over time. Therefore, a high operational risk results in low model sophistication. A financial institution facing high model risk has the opposite incentive. Given the high risk of implementing a model that may differ from the true one, the institution has a strong incentive to incorporate larger upgrades to its model to get closer to the true one. So, high model risk results in high model sophistication.

Given the endogenous choice of model sophistication, we investigate how operational risk affects the exposure of the investment model to this risk. We find that while optimal operational risk exposure is always increasing in operational risk when the risk is external, it may very well be decreasing when this risk is internal. In particular, it is increasing when operational risk is sufficiently low, and it is decreasing otherwise. This result is surprising, as one may expect a financial institution with high operational risk to use a model that is very sensitive to operational errors. This is not the case if operational risk can be controlled by the institution. Indeed, this is because an increase in operational risk has two opposing effects on the optimal operational risk exposure: a positive direct effect due to the higher likelihood of operational errors, and a negative indirect effect due to the optimal reduction in model sophistication. If operational risk is internal, the indirect effect is present and it dominates precisely when this risk is high. In this case, a higher operational risk would result in a lower operational risk exposure.

We also find that operational risk makes the financial institution's exposure to market risk (i.e., investment in risky securities) less volatile when operational risk is internal, and more volatile only when operational risk is external. Notably, in the internal case, a higher likelihood of operational errors is more than offset by a lower optimal model sophistication, which makes the investment model overall less sensitive to market news and operational errors. This offsetting mechanism is absent when operational risk is external. In order to assess how important it is for a financial institution to optimally adjust the sophistication of its model, we quantify the costs of sub-optimal sophistication under operational risk. We demonstrate that these costs are economically significant, and we uncover an asymmetry between the costs of maximal sophistication and unsophistication.

Finally, we study how financial constraints, due to regulation or self-imposed risk limits, affect the financial institution's optimal behavior with operational risk. We find that in the presence of operational risk, whether internal or external, a constrained financial institution always adopts a more sophisticated model than one adopted by an otherwise identical un-

constrained institution. This is because, in the internal case, financial constraints alleviate the trade-off between a more sophisticated model and a model with high operational risk by shielding the institution from large operational errors. In the external case, instead, financial constraints induce a higher model sophistication, prompting the institution to adopt an excessively sophisticated model, because a higher exposure to market news would reduce the relative importance of operational risk (vis-à-vis market risk). At the same time, however, the constraints would shield the excessive exposure to this news. Therefore, operational risk introduces new channels through which financial constraints may influence financial investments. Indeed, by affecting the sophistication of the model, financial constraints alter the institution's exposure to market risk even when the constraints are not binding.

When operational risk is external, we uncover that the financial institution may find financial constraints beneficial. Two opposing effects are present. On one hand, financial constraints are beneficial because they prevent the implementation of extreme market risk exposures, possibly due to large operational errors. On the other hand, they are detrimental as they also prevent valuable market news to be fully incorporated in the investment decisions. We show that when operational risk is external and sufficiently high, the former effect dominates and financial constraints make the institution better off. In contrast, when operational risk is internal, financial constraints are always detrimental since the financial institution can manage the exposure to operational risk more effectively by altering its model sophistication.

Our results have several notable cross-sectional implications. To fix ideas, consider a large investment bank versus a hedge fund. Investment banks are usually more complex institutions, where an accurate implementation of the investment model requires effective communication and coordination among several decentralized divisions (e.g., a research team, a structuring team, and a trading desk). For this reason, it is natural to think of large and complex financial institutions as being more subject to operational errors than hedge funds with a streamlined investment process. However, our findings suggest that, despite being less subject to operational errors, a hedge fund may be overall more exposed to operational risk than a large investment bank. Indeed, the latter reduces its operational risk exposure by adopting a less sophisticated model. Our results also suggest that, given the higher model sophistication, a hedge fund should exhibit more volatile investment decisions. Alternatively, consider two institutions with similar organizational structures (e.g., a large investment bank and a large insurance company). Our model implies that, despite being equally subject to operational errors, their exposures to operational risk may still be different depending on the financial constraints they face. Overall, these implications underscore the importance of having a theory of operational risk to guide in the identification of the determinants of operational losses in the data.

1.1 Related Literature

The importance of operational risk for financial institutions is well recognized and discussed in detail in leading risk management textbooks (e.g., Jorion, 2007; Hull, 2012; Crouhy, Galai and Mark, 2014). However, the difficulties in quantifying operational risk have led most of the recent literature to focus almost exclusively on measurement and estimation issues. Several studies estimate the distribution of operational risk losses using techniques from extreme value theory and data on large and infrequent operational losses (Chavez-Demoulin, Embrechts and Neslehova, 2006; Coleman, 2003; de Fontnouvelle, Rosengren and Jordan, 2006; Ebnother, Vanini, McNeil and Antolinez-Fehr, 2001; Moscadelli, 2004). These estimates, combined with VaR models, are then used to determine regulatory capital against operational losses, as required by the Basel Capital Accord. de Fontnouvelle, Jordan, DeJesus-Ureff and Rosengren (2006) show that the amount of capital held for operational risk may exceed capital for market risk. We view our work as complementary to this literature, as our theory provides a microfoundation of operational losses (and their distribution) by considering the endogenous response of financial institutions to operational risk.

Following the intensity-based framework developed in the reduced form credit risk literature (Jarrow and Turnbull, 1995; Lando, 1998; Duffie and Singleton, 1999; Jarrow and Yu, 2001; Duffie, Eckner, Horel and Saita, 2009), Jarrow (2008) and Chernobai, Jorion and Yu (2011) treat the arrival of operational losses as a conditional Poisson process. Jarrow suggests that both data internal to the firm and market data are needed to estimate the parameters of the operational loss processes. Using a large database of operational losses among U.S. financial institutions, Chernobai, Jorion and Yu find that young and more complex institutions tend to have a higher operational risk exposure, and identify a positive correlation between operational risk and credit risk. As opposed to modeling operational risk in reduced form, we propose a “structural model” of operational risk where operational losses are derived endogenously from the optimal risk exposure of financial institutions.

Franks and Mayer (2001) provide a survey of the European asset management industry, and identify misdealing, settlement problems, and errors in the computation of the asset value as the major sources of operational risk. These findings are confirmed by Biais, Casamatta and Rochet (2003) using a different sample of European fund management companies. Biais, Casamatta and Rochet also propose a theoretical model based on agency frictions in which investors cannot observe the effort that investment companies exert to reduce operational risk. They show that the level of funds’ capital can be useful to provide incentives and hence to reduce operational losses. To the best of our knowledge, this is the only study in which operational risk is not entirely exogenous but rather depends on the optimal actions of financial

institutions. Our model abstracts from delegation frictions, but considers a novel implementation friction that highlights the role of model sophistication as an important determinant of an institution's exposure to operational risk. This implementation friction may arise due to ineffective communication and coordination within a financial institution. Related to this, Vayanos (2003) studies the aggregation of risky positions within a financial institution subject to communication constraints. Biais, Hombert and Weill (2014) derive the optimal trading strategy and equilibrium prices when data collection and aggregation delay the incorporation of relevant information into investment decisions.

In three recent papers Brown, Goetzmann, Liang and Schwarz (2008, 2009, 2012) construct and use a measure of operational risk in the hedge fund industry, called the ω -score, based on funds' mandatory disclosure of past legal and regulatory disputes. Exposure to operational risk, as measured by the ω -score, is associated with subsequent poor fund's returns (Brown, Goetzmann, Liang and Schwarz, 2008), and it may have more predictive power of future fund failure than financial risk (Brown, Goetzmann, Liang and Schwarz, 2009). The ω -score is then refined by integrating a database of operational due diligence reports conducted on behalf of fund investors (Brown, Goetzmann, Liang and Schwarz, 2012). Even though high operational risk could destroy investor value, the authors show that the investors' return-chasing behavior seems to be unaffected by the funds' exposure to this risk.

More broadly, our paper also contributes to the literature on firms' organizational structure (Dessein, 2002; Vayanos, 2003; Dessein and Santos, 2006; Dessein, Galeotti and Santos, 2014). The choice of sophistication of the investment model defines the internal anatomy of a financial firm by requiring, for instance, a specific hierarchical structure, internal due diligence, IT infrastructure, and dedicated quant-desks. Our study is the first to emphasize the connection between operational risk and financial institutions' organizational structure.

The remainder of our paper is organized as follows. Section 2 presents our theory of operational risk, detailing the economic setup and the optimization problem of a financial institution. Section 3 presents our results on the optimal model sophistication, the exposures to market and operational risks, and the costs of sub-optimal sophistication. Section 4 analyzes the optimal behavior of an institution subject to financial constraints and illustrates the beneficial role of constraints in the presence of operational risk. Section 5 concludes. Appendix A contains the proofs, and Appendix B presents results for richer specifications of operational risk.

2 Economy with Operational Risk

The goal in this section is to formally present our theory of operational risk. We consider the simplest possible economic setting in which we incorporate a notion of operational risk into a standard dynamic asset allocation framework. In this setting, we formulate the optimal decision making of a financial institution that accounts for the presence of operational risk.

2.1 Economic Setting

In this section we describe the key ingredients that give rise to an economic role for operational risk in financial markets. A financial institution – our economic agent – relies on a model to make investment decisions; however, uncertainty about this model prompts the institution to change it and improve it over time. The corresponding risk associated with these changes is what we refer to as model risk.

We consider a finite horizon T economy in which market uncertainty is resolved continuously and is driven by a standard Brownian motion w , defined on the probability measure \mathbb{P} . The tradable assets in the economy are a riskless security that pays a constant risk-free rate r , and a risky investment opportunity (e.g., a portfolio of risky loans, a CDO tranche) with return dynamics given by

$$dR_t = (r + \kappa\sigma)dt + \sigma dw_t, \quad (1)$$

where the constants κ and $\sigma > 0$ represent the market price of risk and the volatility of the risky asset, respectively. The drift in (1) is the expected return on the risky investment and is given by the risk-free rate plus a risk premium. In this context, the parameter κ also represents the true model for investment decisions.

The financial institution observes the realizations of the risky return process R_t , but has incomplete information on its dynamics. It deduces σ from the return's quadratic variation, but it must estimate the true model κ via its conditional expectation, rationally updating in a Bayesian fashion. We denote by $\hat{\kappa}_t$ the institution's estimate of the true model κ at time t , given its prior and the realizations of the return on the risky asset. Based on this estimate, the institution's perceived return dynamics are equal to

$$dR_t = (r + \hat{\kappa}_t\sigma)dt + \sigma d\hat{w}_t, \quad (2)$$

where \hat{w} represents the perceived market uncertainty under a new probability measure $\hat{\mathbb{P}}$, and is a Brownian motion satisfying $d\hat{w}_t = dw_t + (\kappa - \hat{\kappa}_t)dt$.

We assume that the financial institution has a normally distributed prior with mean κ_0 and variance ν_0 over the true model κ . Under this assumption, the estimated model is a martingale under the perceived probability measure $\hat{\mathbb{P}}$, and is characterized by the dynamics (Liptser and Shiryaev (2001))

$$d\hat{\kappa}_t = \hat{\nu}_t d\hat{w}_t. \quad (3)$$

The conditional volatility $\hat{\nu}_t$ is given by

$$\hat{\nu}_t = \frac{\nu_0}{1 + \nu_0 t}, \quad (4)$$

where the driving parameter ν_0 is referred to as the *model risk* in what follows.³

2.2 Modeling Operational Risk

Operational risk arises from the inadequate implementation of the models that financial institutions adopt to perform their financial operations. Therefore, our theory builds on a novel implementation friction. According to this friction, an inadequate implementation, which can be caused by different types of errors (e.g., bugs in programming codes, mistakes in data collection and processing), is more likely to occur at times when a financial institution makes changes to its model.

In our framework, as introduced in the previous section, model risk creates the need for a financial institution to update its model with the arrival of new information. In the presence of operational risk, however, operational errors introduce a wedge between the estimated model and the one that is ultimately implemented. In particular, if the financial institution decides to change its model by $d\hat{\kappa}_t$, it will end up implementing a change given by $d\hat{\kappa}_t + \sigma_\epsilon dw_{\epsilon t}$, where $\sigma_\epsilon dw_{\epsilon t}$ captures the operational error. Here, w_ϵ is another standard Brownian motion adapted to the institution's filtration, and represents operational uncertainty. The volatility parameter $\sigma_\epsilon > 0$ is a constant and is our measure of *operational risk*.

³We could in principle start with a stochastic process for the model dynamics satisfying exogenously equation (3). However, the advantage of our formulation is that it provides a microfoundation for the estimated model $\hat{\kappa}_t$, highlighting the underlying determinants. Furthermore, it is consistent with the extant literature on incomplete information and model risk (e.g., Detemple, 1986, 1991; Genotte, 1986; Veronesi, 1999, 2000; Xia, 2001; Cvitanic, Lazrak, Martellini and Zapatero, 2006).

In order for operational risk to play a relevant economic role, we allow the financial institution to potentially manage its exposure to operational errors by controlling the size of the changes it makes to its model. This introduces a new trade-off: by reducing the size of these changes, the institution lowers its exposure to operational errors, but at the same time it compromises on the sophistication of its model. This means that the institution may not select the estimated model, and hence the most sophisticated available, for implementation. Formally, κ_t^* denotes the *implemented model* at time t , with dynamics

$$d\kappa_t^* = \lambda d\hat{\kappa}_t + h(\lambda)\sigma_\epsilon dw_{\epsilon t}. \quad (5)$$

The quantity λ , which is to be determined endogenously, controls the extent to which the financial institution may deviate from the estimated model, and we refer to it as *model sophistication*. Simultaneously, it may affect the sensitivity of the implemented model to operational errors through the function $h(\lambda) \geq 0$, with $h'(\lambda) \geq 0$. The function $h(\lambda)$ simply controls the balance of the trade-off between model sophistication and operational risk, while the quantity $h(\lambda)\sigma_\epsilon$ captures the exposure to operational uncertainty. In practice, the choice of model sophistication involves setting up a specific organizational structure within the institution. Considering the high costs, particularly fixed costs, associated with changes of such organizational structures, we take λ as constant over the horizon considered.

When the institution is able manage the exposure to operational uncertainty, the above formulation captures what is termed as *internal* operational risk, and in this case the function $h(\lambda)$ is strictly increasing. On the contrary, when the institution is not able to control how operational errors affect its model, it is subject to so-called *external* operational risk. A simple way to capture a form of external operational risk within our formulation is by restricting the function $h(\lambda)$ to be flat, $h'(\lambda) = 0$. For the remainder of the paper, we take the function $h(\lambda)$ to be linear, $h(\lambda) = \lambda$, which is not only natural to start with, but it also admits tractability and maintains a simple setting. However, we can also consider a plausible non-linear specification without altering our main conclusions, as discussed in Appendix B.2. Finally, to make the comparison between internal and external operational risks more intuitive, we set $h(\lambda) = 1$ for the case of external operational risk, thus yielding an implemented model with dynamics equal to $d\kappa_t^* = \lambda d\hat{\kappa}_t + \sigma_\epsilon dw_{\epsilon t}$. Our focus is, of course, on internal operational risk; the case of external operational risk enables us to better highlight our main insights.⁴

⁴External operational risk is often modeled as jump risk (e.g., Jarrow, 2008). However, what really matters for our analysis is that a financial institution is not able to affect the exposure to external operational risk. To maintain a simple setting and facilitate the comparison, we consider Brownian uncertainty for both internal and external operational risks.

The implemented model, satisfying (5) with $h(\lambda) = \lambda$, can alternatively be expressed as

$$\kappa_t^* = (1 - \lambda)\kappa_0 + \lambda\hat{\kappa}_t + \lambda\sigma_\epsilon w_{\epsilon t}. \quad (6)$$

This representation shows that the implemented model is given by the sum of two components. The first, $(1 - \lambda)\kappa_0 + \lambda\hat{\kappa}_t$, is the model that the institution aims to implement, and is given by an average of the initial model κ_0 and the estimated model $\hat{\kappa}_t$, weighted by the model sophistication λ . Henceforth, we refer to κ_0 as the *least sophisticated model* and to $\hat{\kappa}_t$ as the *most sophisticated model*, in the sense that there is no better model to be implemented absent operational risk. The second component, $\lambda\sigma_\epsilon w_{\epsilon t}$, captures the operational errors that make the implemented model differ from its target. Operational errors can be further decomposed into operational uncertainty, $w_{\epsilon t}$, and the exposure to it, $\lambda\sigma_\epsilon$. The latter is a key endogenous quantity in our analysis, and we refer to it as *operational risk exposure*. We remark that in practice the true and the most sophisticated models are of course broader and more complex than the simple notion of a single market price of risk. For example, they may involve multiple sources of risk, multiple assets and financial market imperfections, and varying degrees of observability.

Our formulation of the implemented model conveniently nests three special cases in the presence of internal operational risk. When $\sigma_\epsilon = 0$, operational risk is absent, and it serves as our benchmark in the subsequent analysis. When $\lambda = 0$, internal operational risk is fully eliminated by the institution, but it must employ the least sophisticated model for investment decisions, $\kappa_t^* = \kappa_0$. When, on the other extreme, $\lambda = 1$, the institution aims at implementing the most sophisticated model, but it is now exposed to the highest level of operational risk, $\kappa_t^* = \hat{\kappa}_t + \sigma_\epsilon w_{\epsilon t}$.

For a given level of model sophistication λ , the financial institution makes investment decisions based on the implemented model κ_t^* (containing operational errors). Therefore, the institution's effective return dynamics under κ_t^* are equal to

$$dR_t = (r + \kappa_t^* \sigma) dt + \sigma dw_t^*, \quad (7)$$

where w^* is a corresponding Brownian motion under the probability measure \mathbb{P}^* , such that

$$dw_t^* = d\hat{w}_t + (\hat{\kappa}_t - \kappa_t^*) dt. \quad (8)$$

To conclude the description of our setting, the two sources of uncertainty in our economy, market and operational, are naturally taken to be uncorrelated, given the nature of our implementation friction. Although economically less relevant, our analysis can straightforwardly incorporate any correlation between the two sources of uncertainty, whereby our main insights

would remain equally valid.⁵ Moreover, we assume that operational uncertainty does not add information in forming estimates of the true model κ , over and above the one conveyed by the return on the risky asset, $\hat{\mathbb{E}}_s[(\kappa - \hat{\kappa}_t)w_{et}] = 0$ for $s < t$, where $\hat{\mathbb{E}}[\cdot]$ denotes the expectation under the probability measure $\hat{\mathbb{P}}$. Finally, we maintain the assumption that, by the nature of our implementation friction, operational uncertainty cannot be “undone” despite being observed continuously. This preserves the economic realism of a discrete time formulation and the tractability of continuous time.

Remark 1 (A richer specification of operational risk). A richer specification of the dynamics of the implemented model is given by

$$d\kappa_t^* = -\zeta(\kappa_t^* - \hat{\kappa}_t)dt + \lambda d\hat{\kappa}_t + h(\lambda)\sigma_\epsilon dw_{et} \quad (9)$$

where the mean-reverting component, $-\zeta(\kappa_t^* - \hat{\kappa}_t)dt$, makes the implemented model deterministically revert back to the most sophisticated one, and converge to it in the long-run. The parameter $\zeta \geq 0$ is responsible for the speed of mean-reversion. Our formulation in (5) is obtained by “shutting down” this channel, $\zeta = 0$. This implies that the financial institution makes changes to its model only when new information arrives ($d\hat{\kappa}_t \neq 0$). Any mean-reversion component would simply induce the institution to change its model deterministically even in the absence of new information. For instance, when $\zeta = \lambda$, the above specification has the discrete-time interpretation of a financial institution aiming at implementing a model that is given by the λ -weighted average of the most recent implemented model and the most sophisticated one available. Since our main insights remain unchanged in the presence of mean reversion, we leave its treatment to Appendix B.1.

2.3 Optimization Problem

In this section we embed operational risk into the optimal decision making of a financial institution. Because of the implementation friction, this implies that the institution simultaneously chooses the model sophistication and its dynamic investments. Towards that, we proceed by solving the institution’s problem in two steps. First, we fix a level of model sophistication and we solve for the optimal investment as in a standard asset allocation framework. Then, given the optimal dynamic investment decision, we solve for the optimal model sophistication.

⁵If we were to consider correlated sources of uncertainty, the economic interpretation of the operational errors would not be so meaningful. For instance, if the correlation were positive, it would be more likely that positive operational errors would occur when the market went up, and vice-versa for negative correlation. This would imply a systematic bias in the way operational errors are generated.

The financial institution is initially endowed with capital W_0 . It chooses a risky investment process π , where π_t denotes the fraction of capital invested in the risky opportunity at time t . We, henceforth refer to π_t as *market risk exposure*. The institution's capital process W then follows

$$dW_t = W_t(r + \pi_t \sigma \kappa_t^*(\lambda))dt + W_t \pi_t \sigma dw_t^*, \quad (10)$$

where $\kappa_t^*(\lambda)$ is as in (6), and dw_t^* as in (8). The financial institution is guided by a logarithmic objective function over the horizon capital, $\log W_T$.

We next present the two steps we use to solve the institution's problem. The first step is to solve for the optimal dynamic investment for a given model sophistication λ . The optimal market risk exposure $\pi_t^*(\lambda)$ is the solution to the following optimization problem

$$\max_{\pi} \mathbb{E}_0^*[\log W_T] \quad (11)$$

subject to the dynamic budget constraint (10), where $\mathbb{E}^*[\cdot]$ denotes the expectation under the probability measure \mathbb{P}^* . The financial institution faces incomplete markets since there are two sources of uncertainty in the economy, (w^*, w_ϵ) , and only one risky asset available for trading. However, the logarithmic specification admits tractability in solving this problem (see Appendix A).

The second step is to solve for the optimal model sophistication, given the optimal market risk exposure $\pi_t^*(\lambda)$ obtained in the first step. The optimal model sophistication λ^* maximizes the value function $J(\cdot)$ at time 0, under the probability measure $\hat{\mathbb{P}}$,

$$\max_{\lambda} J(\lambda; W_0) = \hat{\mathbb{E}}_0[\log W_T(\pi^*(\lambda))] \quad (12)$$

It is important to emphasize that the financial institution chooses its model sophistication under $\hat{\mathbb{P}}$ since it rationally takes into account what the most sophisticated model is, and importantly how undesirable it is to deviate from it. Therefore, all the models $\kappa_t^*(\lambda)$, each characterized by a different λ , are evaluated against the most sophisticated model $\hat{\kappa}_t$, which is known but not implementable without operational errors.

The optimization problems in the two steps described above can be reasonably interpreted as the decision problems faced by two separate divisions within a financial institution, an executive and an investment division. The executive division decides first the organizational structure of the institution (e.g., hierarchical structure, dedicated quant teams, IT infrastructure, internal due diligence, training programs), hence the level of sophistication λ^* characterizing the model used for investment purposes. Once the organizational structure of the institution is set up,

the investment division is responsible for the model implementation $\kappa_t^*(\lambda^*)$ and the ensuing asset allocation $\pi_t^*(\lambda^*)$. Therefore, our solution procedure entails solving the problem of the investment division first, followed by that of the executive division.

3 Optimal Behavior with Operational Risk

In this section we characterize explicitly the optimal model sophistication and the optimal exposures to market and operational risks. A notable finding is that, while operational risk exposure increases in operational risk when this risk is external, it may actually decrease when operational risk is internal. Moreover, in contrast to the external case, operational risk also decreases the variability of the institution's market risk exposure when it is internal. Finally, we find that not accounting for model sophistication can have sizable losses.

3.1 Model Sophistication and Risk Exposures

The next proposition provides the optimal model sophistication in closed form and presents its key properties.

Proposition 1 (Model sophistication). *In the presence of external operational risk, the financial institution always adopts the most sophisticated model, $\lambda^* = 1$. In the presence of internal operational risk, however, it adopts an under-sophisticated model, $0 < \lambda^* < 1$, with*

$$\lambda^* = \frac{\nu_0 T - \log(1 + \nu_0 T)}{\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2}, \quad (13)$$

Consequently, when operational risk is internal,

- (i) optimal model sophistication is decreasing in operational risk, $\partial\lambda^*/\partial\sigma_\epsilon < 0$;
- (ii) optimal model sophistication is increasing in model risk, $\partial\lambda^*/\partial\nu_0 > 0$.

Proposition 1 reveals that, when operational risk is external, the financial institution aims to implement the most sophisticated model. This is intuitive, as in this case there is no trade-off between model sophistication and operational risk. Instead, when operational risk is internal, the optimal model sophistication trades off a more sophisticated model against higher opera-

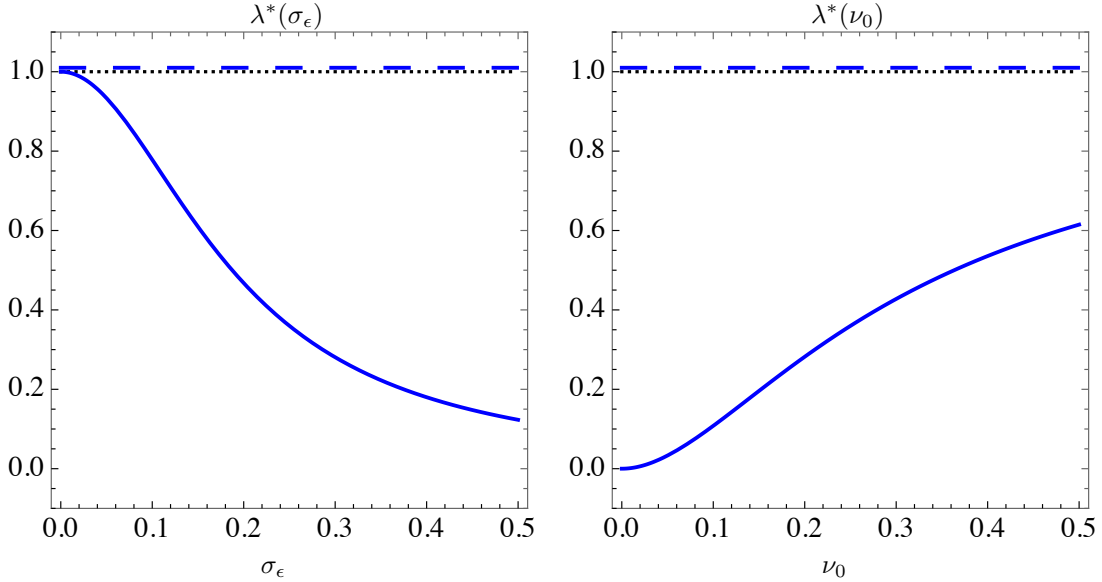


Figure 1: Optimal model sophistication

In this figure we plot the optimal model sophistication λ^* as a function of operational risk σ_ϵ (left panel) and model risk ν_0 (right panel). The dotted line represents the benchmark case of no operational risk, $\sigma_\epsilon = 0$. The solid line refers to the case of internal operational risk, and the dashed line (which coincides with the dotted line) to the case of external operational risk. The parameter values are $\nu_0 = 0.25$, $\sigma_\epsilon = 0.25$, $T = 5$. The plots are typical.

tional risk. To see this more explicitly, one can show that λ^* , which solves the maximization problem defined in (12), also minimizes the following expression

$$\int_0^T \text{vâr}_0[\kappa_t^*(\lambda)] dt - 2 \int_0^T \text{côv}_0[\kappa_t^*(\lambda), \hat{\kappa}_t] dt, \quad (14)$$

where $\text{vâr}[\cdot]$ and $\text{côv}[\cdot]$ denote the variance and covariance under the probability measure $\hat{\mathbb{P}}$. Intuitively, the financial institution wants to achieve two goals. It seeks to implement a model that has low (accumulated) variance, as captured by the first term in (14), but at the same time a model that has a high (accumulated) covariance with the most sophisticated model available, as captured by the second term. The first goal can be achieved by reducing λ , the second by increasing it. This highlights the aforementioned trade-off, which confirms the result in Proposition 1 that the financial institution always chooses a less sophisticated model, $\lambda^* < 1$.

Proposition 1 identifies the key determinants of optimal model sophistication with internal operational risk to be model and operational risks. Moreover, the optimal model sophistication inherits the desirable properties of being decreasing in operational risk and increasing in model

risk, as also depicted in Figure 1.⁶ These are intuitive. A financial institution with high operational risk (high σ_ϵ) has an incentive to reduce its model sophistication (low λ) compared to an otherwise identical institution with low operational risk. This is because aiming to implement a more sophisticated model is less important when operational risk is high, given the higher likelihood of operational errors. On the contrary, an institution with high model risk (high ν_0) optimally increases its model sophistication (high λ) compared to an otherwise identical institution with low model risk. In this case, aiming to be closer to the most sophisticated model is more desirable when model risk is high, given the higher risk to implement a model that differs from the true one.

The above results have the following predictions. Consider for instance a large investment bank versus a small hedge fund. An investment bank is usually a more complex institution, where the implementation of the investment model is the result of a more decentralized process among several divisions (e.g., a research team, a quant team, a structuring team, and a trading desk). An accurate implementation, therefore, requires effective communication and coordination among these divisions. For this reason, it is natural to think of large and complex financial institutions as more subject to operational errors. Our findings suggest that large institutions find it optimal to implement a less sophisticated model, which requires lesser changes and updates over time.

We next consider the optimal operational risk exposure, namely the sensitivity of the implemented model to operational uncertainty, and the corresponding distribution of operational losses.

Proposition 2 (Operational risk exposure). *In the presence of external operational risk, the operational risk exposure is always increasing in the level of operational risk σ_ϵ . Instead, with internal operational risk, the optimal operational risk exposure is increasing when operational risk is sufficiently low, $\sigma_\epsilon < \bar{\sigma}_\epsilon$, and it is decreasing otherwise, where the constant $\bar{\sigma}_\epsilon > 0$ is provided in Appendix A.*

In contrast to the case of external operational risk, we find that the operational risk exposure may well become decreasing in operational risk when this risk is internal. This result is surprising, as one may reasonably expect that a financial institution with high operational risk implements a model that is very sensitive to operational uncertainty. This is indeed the

⁶The plots use the parameter values of $\kappa_0 = 0.33$, $\sigma = 0.18$, $\nu_0 \in [0, 0.5]$ with baseline value equal to 0.25, $\sigma_\epsilon \in [0, 0.5]$ with baseline value equal to 0.25, $t = 1$, $T = 5$. We calibrate κ_0 to the average annual market Sharpe ratio, and σ to the annual market standard deviation. We treat realizations of the Sharpe ratio greater than 1.15 as occurring with probability lower than 5%, implying (given the normality assumption) that $\nu_0 \leq 0.5$, and we consider the same range of values for σ_ϵ . We calibrate the horizon T to 5 years as average tenure of organizational structures in financial institutions, and we consider as intermediate date t 1 year.

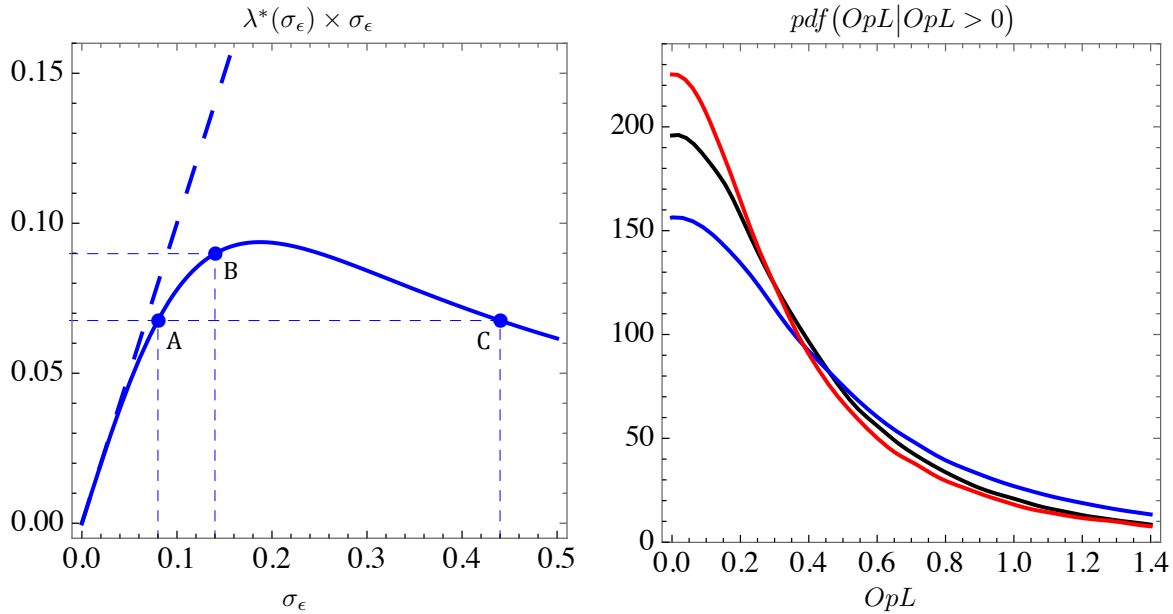


Figure 2: Operational risk exposure and distribution of operational losses

In this figure we plot the optimal operational risk exposure $\lambda^*\sigma_\epsilon$ as a function of operational risk σ_ϵ (left panel), and the conditional distribution of operational losses OpL for different levels of internal operational risk σ_ϵ (right panel). The solid line refers to the case of internal operational risk, and the dashed line to the case of external operational risk. In the right panel, the black line corresponds to $\sigma_\epsilon = 0.1$, the blue line to the baseline case $\sigma_\epsilon = 0.25$, and the red line to $\sigma_\epsilon = 0.5$. The parameter values are as in Figure 1. The plots are typical.

case, but only if the institution is subject to external operational risk, which is a risk that cannot be controlled. When, instead, the institution can affect the exposure to operational risk by changing its organizational structure and hence its model sophistication, the positive effect of operational risk on the sensitivity to operational uncertainty not only is attenuated, but it may even reverse. Moreover, this reversal occurs precisely when operational risk is high. The intuition is as follows. When subject to high internal operational risk, a sophisticated model is particularly undesirable. Therefore, a particularly low model sophistication is optimal, resulting in a model that is overall not very sensitive to operational uncertainty.

Cross-sectionally, the result in Proposition 2 implies that institutions with high internal operational risk may have an optimal operational risk exposure that is lower than the exposure of institutions with low internal operational risk, as depicted by points B and C in the left panel of Figure 2. Moreover, two financial institutions with identical risk exposure, as illustrated by points A and C in the left panel of Figure 2, may well have sharply different model sophistications, as one's operational risk exposure is increasing in operational risk while the other's is decreasing. These results are in sharp contrast with the case of external operational risk. For instance, (reported) risk exposure by itself cannot always be used to infer the institution's

model sophistication and its level of operational risk. Therefore, our findings highlight the importance of distinguishing between internal and external operational risks for the design of regulatory environments, given the different implications on institutions' decision making that these risks may have.

Operational losses, denoted by OpL , are defined as the percentage change in the horizon capital W_T caused exclusively by the materialization of operational errors,

$$OpL = \frac{W_T(\sigma_\epsilon; w_\epsilon = 0) - W_T(\sigma_\epsilon)}{W_T(\sigma_\epsilon)}, \quad (15)$$

and are equal to

$$OpL = \exp \left[\frac{(\lambda^* \sigma_\epsilon)^2}{2} \int_0^T w_{\epsilon t}^2 dt + \lambda^*(1 - \lambda^*) \sigma_\epsilon \int_0^T (\kappa_0 - \hat{\kappa}_t) w_{\epsilon t} dt - \lambda^* \sigma_\epsilon \int_0^T w_{\epsilon t} d\hat{w}_t \right] - 1, \quad (16)$$

as derived in Appendix A.1. Since operational errors may also be beneficial to the financial institution, operational losses can be negative, hence generating operational gains. Even though this is possible, operational losses are positive on average.⁷ The right panel in Figure 2 plots the conditional distribution of operational losses, conditional on the losses being positive $pdf(OpL|OpL > 0)$, for different levels of internal operational risk. In line with the result in Proposition 2, the distribution of operational losses becomes flatter with increasing internal operational risk when this risk is sufficiently small (moving from the black to the blue distribution), and it becomes steeper when this risk is high (from the blue to the red distribution). This finding highlights the endogenous nature of operational losses and their dependence on the optimal choice of model sophistication. Our microfundation for operational errors provides a mechanism that links operational risk to operational losses, and should help the identification of the determinants of these losses in the data.

We conclude this section with the analysis of the optimal market risk exposure.

⁷Note that

$$\mathbb{E}[\log(1 + OpL)] = \mathbb{E} \left[\frac{(\lambda^* \sigma_\epsilon)^2}{2} \int_0^T w_{\epsilon t}^2 dt + \lambda^*(1 - \lambda^*) \sigma_\epsilon \int_0^T (\kappa_0 - \hat{\kappa}_t) w_{\epsilon t} dt - \lambda^* \sigma_\epsilon \int_0^T w_{\epsilon t} d\hat{w}_t \right] = \frac{(\lambda^* \sigma_\epsilon T)^2}{4}.$$

Since by Jensen's inequality $\log(1 + \mathbb{E}[OpL]) > \mathbb{E}[\log(1 + OpL)]$, we can conclude that $\mathbb{E}[OpL] > 0$.

Proposition 3 (Market risk exposure). *In the presence of external operational risk, the optimal market risk exposure is given by $\pi_t^* = \hat{\kappa}_t/\sigma + (\sigma_\epsilon/\sigma)w_{\epsilon t}$. In the presence of internal operational risk, it is instead given by*

$$\pi_t^*(\lambda^*) = \frac{\hat{\kappa}_t}{\sigma} + \frac{(1 - \lambda^*)(\kappa_0 - \hat{\kappa}_t) + \lambda^* \sigma_\epsilon w_{\epsilon t}}{\sigma}, \quad (17)$$

where λ^* is as in (13). Consequently, the variance of the market risk exposure,

$$\text{v}\hat{\text{a}}\text{r}_0[\pi_t^*] = \left(\frac{\lambda^*}{\sigma}\right)^2 \left(\frac{\nu_0^2}{1 + \nu_0 t} + \sigma_\epsilon^2\right) t, \quad (18)$$

- (i) is increasing in operational risk when it is external, $\partial \text{v}\hat{\text{a}}\text{r}_0[\pi_t^*]/\partial \sigma_\epsilon > 0$;
- (ii) is decreasing in operational risk when it is internal, $\partial \text{v}\hat{\text{a}}\text{r}_0[\pi_t^*]/\partial \sigma_\epsilon < 0$.

Operational risk is the risk of implementing a model that contains operational errors. As a consequence, one may naturally think that an institution's exposure to market risk, resulting from the implementation of its model, would be more volatile the higher the level of operational risk is. This is indeed the case in the presence of external operational risk. However, the exact opposite occurs when operational risk is internal. In this case, as Proposition 3 reveals, operational risk always decreases the variability of market risk exposure, implying that this variability is higher when operational risk is absent. The logic behind this counterintuitive result is as follows. An increase in (internal) operational risk σ_ϵ by itself indeed increases the variability of the model and hence of the market risk exposure. However, it simultaneously induces an optimal reduction in model sophistication λ (Proposition 1), which in turn makes the model less sensitive to both market and operational uncertainty. Overall, the latter effect dominates, causing the variability of market risk exposure to be decreasing in the level of internal operational risk. When, instead, operational risk is external this offsetting mechanism is absent.

The results in Proposition 3 have the following cross-sectional implications. Institutions with high internal operational risk always have an optimal market risk exposure that is less volatile than the market exposure of institutions with low internal operational risk. This is depicted by the three points in Figure 3. Moreover, two institutions with different levels of internal operational risk may have identical operational risk exposures (points A and C in Figure 2), while at the same time exhibiting very different variability in their exposure to market risk (points A and C in Figure 3). Therefore, observing an institution with a volatile market exposure does not necessarily imply that this institution is more exposed to operational uncertainty. Indeed, its volatility is mainly driven by fundamental market news that is incorporated into a more sophisticated investment model.

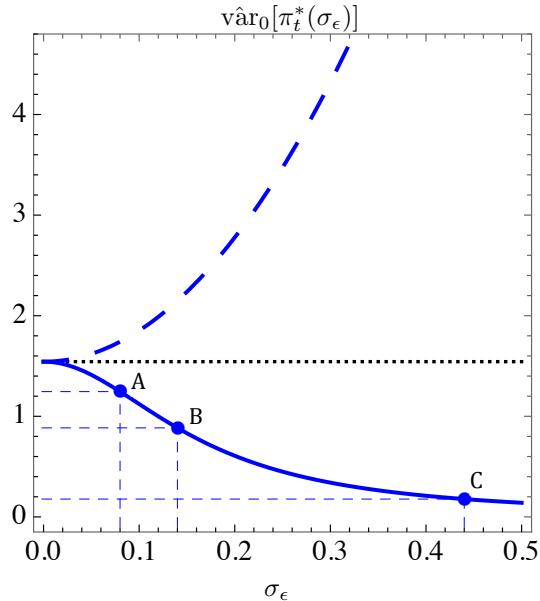


Figure 3: Variability of market risk exposure

In this figure we plot the variance at time 0 of the optimal market risk exposure π_t^* as a function of operational risk σ_ϵ . The dotted line represents the benchmark case of no operational risk, $\sigma_\epsilon = 0$. The solid line refers to the case of internal operational risk, and the dashed line to the case of external operational risk. The parameter values are $\sigma = 0.18$, $t = 1$ and the remainder are as in Figure 1. The plots are typical.

Finally, Equation (18) also reveals that the variance of market risk exposure is increasing in model risk ν_0 . This is the result of two compounding simultaneous forces. First, an increase in model risk directly increases the variability of market risk exposure, as the arrival of more unexpected market news leads to larger changes to the implemented model, and hence to larger changes to the market exposure. At the same time, as shown in Proposition 1, an increase in model risk induces an optimally higher model sophistication, which amplifies the increase in the variability of market risk exposure by making the implemented model more sensitive to both market and operational uncertainty.

3.2 Costs of Sub-Optimal Behavior under Operational Risk

We next analyze the economic significance for a financial institution to be able to manage internal operational risk by optimally adjusting the sophistication of its investment model. In particular, we quantify the costs of adopting the most or the least sophisticated model, instead of the model with optimal sophistication λ^* .

Towards that, we consider a cost measure $\eta_{\lambda=1}$ that represents the percentage of initial capital W_0 that an institution, adopting the most sophisticated model, must additionally have in order to be indifferent to the optimal model sophistication λ^* ,

$$\eta_{\lambda=1} : J(1; W_0(1 + \eta_{\lambda=1})) = J(\lambda^*; W_0). \quad (19)$$

The measure $\eta_{\lambda=1}$ quantifies the cost of maximal sophistication. Similarly, we consider a cost measure $\eta_{\lambda=0}$ that represents the percentage of initial capital W_0 that an institution, adopting the least sophisticated model, must additionally have in order to be indifferent to the optimal model sophistication λ^* ,

$$\eta_{\lambda=0} : J(0; W_0(1 + \eta_{\lambda=0})) = J(\lambda^*; W_0). \quad (20)$$

The measure $\eta_{\lambda=0}$ quantifies the cost of unsophistication.

The next proposition identifies the cost measures of sub-optimal sophistication explicitly in terms of the underlying primitives.

Proposition 4 (Costs of sub-optimal sophistication). *The cost measures for sub-optimally adopting the most and least sophisticated models are given by*

$$\eta_{\lambda=1} = \exp \left[\frac{(1/4)\sigma_\epsilon^4 T^4}{2(\nu_0 T - \log(1 + \nu_0 T)) + \sigma_\epsilon^2 T^2} \right] - 1, \quad (21)$$

$$\eta_{\lambda=0} = \exp \left[\frac{(\nu_0 T - \log(1 + \nu_0 T))^2}{2(\nu_0 T - \log(1 + \nu_0 T)) + \sigma_\epsilon^2 T^2} \right] - 1. \quad (22)$$

Proposition 4 and its corresponding Figure 4 show that the costs of sub-optimal sophistication are driven by the level of operational and model risks. In particular, we find that the cost of maximal sophistication $\eta_{\lambda=1}$ is increasing in operational risk and decreasing in model risk, whereas the cost of unsophistication $\eta_{\lambda=0}$ is decreasing in operational risk and increasing in model risk. As demonstrated in Proposition 1, an increase in operational risk σ_ϵ decreases the optimal model sophistication λ^* , because a higher likelihood of operational errors more than offsets the benefits of sophistication. This in turn makes maximal sophistication more costly, while unsophistication less costly. Instead, an increase in model risk ν_0 increases the optimal model sophistication λ^* , hereby making maximal sophistication less costly and unsophistication more costly. Indeed, a higher level of model risk increases the benefits of sophistication, as the risk of implementing a model that differs from the true one becomes higher.

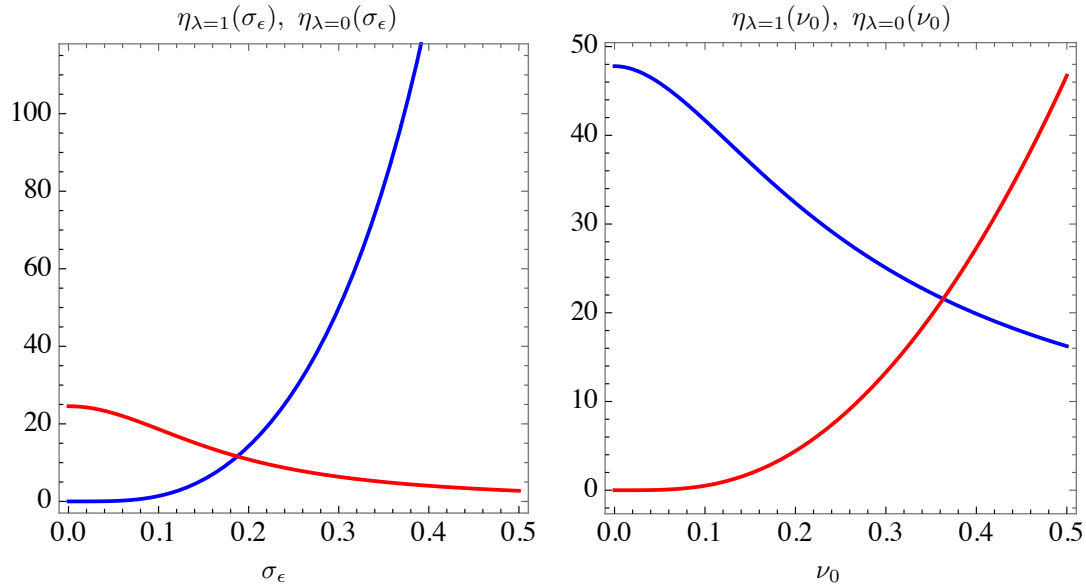


Figure 4: Costs of sub-optimal model sophistication

In this figure we plot the cost measures of sub-optimal model sophistication $\eta_{\lambda=1}$ and $\eta_{\lambda=0}$ in percentages, as a function of operational risk σ_ϵ (left panel) and model risk ν_0 (right panel). The blue line represents the cost of maximal sophistication $\eta_{\lambda=1}$, and the red line represents the cost of unsophistication $\eta_{\lambda=0}$. The parameter values are as in Figure 1.

Importantly, our model allows us to assess quantitatively the significance of a financial institution’s optimal behavior under operational risk. The left panel in Figure 4 illustrates that when operational risk is severe, a financial institution may need twice as much capital to compensate for the potential losses associated with an maximal sophisticated model. When instead operational risk is mild, it may need an amount of additional capital in the range of 20% to compensate for the potential losses associated with an unsophisticated model. These large magnitudes reveal the economic importance of optimally managing internal operational risk. Moreover, these magnitudes suggest that maximal sophistication may be more detrimental than unsophistication, highlighting a notable asymmetry between the costs of sub-optimal sophistication with respect to operational risk. Finally, the right panel in Figure 4 shows that, for a financial institution with an intermediate level of operational risk, the cost of maximal sophistication decreases from roughly 50% to 15% for increasing levels of model risk, whereas the cost of unsophistication increases from 0 to almost 50%. This confirms the economic relevance of optimal model sophistication for a wide range of model risks.

4 Operational Risk and Financial Constraints

In this section we study how financial constraints interact with operational risk and thereby influence the institutions' choice of model sophistication. Our main finding is that in the presence of operational risk, whether internal or external, a constrained financial institution optimally adopts a more sophisticated model as compared to an unconstrained one. Moreover, we uncover that an institution subject to external operational risk implements an excessively sophisticated model and may find financial constraints beneficial. In contrast, an institution subject to internal operational would always find those constraints detrimental.

4.1 Optimization Problem with Constraints

We consider an economic environment in which a financial institution is subject to financial constraints either due to regulation or self-imposed risk limits. More specifically, we focus on upper- and lower-bound constraints on the proportion of capital invested in the risky opportunity. Let \mathcal{C} denote the set of feasible investment choices,

$$\pi_t \in \mathcal{C} \equiv \{\pi \in \mathbb{R} : \underline{\alpha} \leq \pi \leq \bar{\beta}\} \quad (23)$$

with $\bar{\beta} > \underline{\alpha}$. This set features some of the most common constraints financial institutions face. Examples of lower-bound constraints $\pi_t \geq \underline{\alpha}$ include short-selling constraints and concentration constraints. Examples of upper-bound constraints $\pi_t \leq \bar{\beta}$ include leverage constraints, borrowing constraints, and margin requirements. We focus our analysis on the case in which the least sophisticated model κ_0 implies an asset allocation, π_0 , that does not violate the financial constraints, $\underline{\alpha} \leq \pi_0 \leq \bar{\beta}$. Indeed, this allows us to isolate the effects of possible future violations of the financial constraints on the optimal model sophistication. Under our maintained assumption, these effects are not confounded by the fact that the least sophisticated model may be constrained to begin with, admitting a clearer comparison across constrained environments (see footnote 9).

The optimization problem faced by the financial institution follows the steps presented in section 2.3. The optimal market risk exposure for a given model sophistication $\pi_t^c(\lambda)$ maximizes the objective $\mathbb{E}_0^*[\log W_T]$, subject to the dynamic budget constraint (10) and the financial constraints $\pi_t \in \mathcal{C}$. The optimal model sophistication λ^c is then obtained by solving

$$\max_{\lambda} J^c(\lambda; W_0) = \hat{\mathbb{E}}_0[\log W_T(\pi^c(\lambda))], \quad (24)$$

where $J^c(\cdot)$ denotes the value function in the presence of financial constraints.

4.2 Optimal Behavior with Constraints and Operational Risk

The optimal model sophistication in the presence of financial constraints cannot be obtained in closed-form given the highly non-linear nature of the problem. However, the next proposition provides a characterization for it and reports a key property.

Proposition 5 (Model sophistication with financial constraints). *In the presence of operational risk, internal or external, financial constraints \mathcal{C} increase the optimal model sophistication,*

$$\lambda^c \geq \lambda^*, \quad (25)$$

where λ^* is the unconstrained model sophistication in (13). When operational risk is internal, the financial institution still adopts an under-sophisticated model, $0 < \lambda^c < 1$, which solves the equation

$$\lambda^c = 1 - \frac{\int_0^T \left[\Omega \left(\frac{\sigma_{\underline{\alpha}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) - \Omega \left(\frac{\sigma_{\bar{\beta}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) \right] \lambda^c \sigma_\epsilon^2 t \, dt}{\int_0^T \left[\Omega \left(\frac{\sigma_{\underline{\alpha}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) - \Omega \left(\frac{\sigma_{\bar{\beta}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) \right] \nu_0 \hat{\nu}_t t \, dt}, \quad (26)$$

where the function $\Omega(x) = n(x)x + \int_x^\infty n(s)ds$ and $n(\cdot)$ represents the probability density function of a standard normal distribution. Instead, when operational risk is external, the financial institution adopts an excessively sophisticated model, $\lambda^c > 1$, which solves the equation

$$\lambda^c = 1 + \frac{\int_0^T \left[n \left(\frac{\sigma_{\bar{\beta}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) \frac{\sigma_{\bar{\beta}} - \kappa_0}{(\lambda^{c^2} \nu_0 \hat{\nu}_t + \sigma_\epsilon^2 t)^{\frac{3}{2}}} - n \left(\frac{\sigma_{\underline{\alpha}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) \frac{\sigma_{\underline{\alpha}} - \kappa_0}{(\lambda^{c^2} \nu_0 \hat{\nu}_t + \sigma_\epsilon^2 t)^{\frac{3}{2}}} \right] (\nu_0 \hat{\nu}_t t) \sigma_\epsilon^2 t \, dt}{\int_0^T \left[\Omega \left(\frac{\sigma_{\underline{\alpha}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) - \Omega \left(\frac{\sigma_{\bar{\beta}} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) \right] \nu_0 \hat{\nu}_t t \, dt}. \quad (27)$$

Absent operational risk, the most sophisticated model is always implemented, even in the presence of financial constraints: $\lambda^c = \lambda^* = 1$. This implies that the implemented model is not affected by financial constraints when operational risk is not present.⁸ In contrast, when operational risk is present, financial constraints always affect the degree of model sophistication that financial institutions optimally choose. This result underscores the fact that operational

⁸Setting $\sigma_\epsilon = 0$ in both Equation (26) and Equation (27), it is immediate to see that $\lambda^c = 1$.

risk, whether internal or external, interacts with a constrained environment that limits the investment decisions of financial institutions.

Proposition 5 reveals that a financial institution that is subject to financial constraints always adopts a more sophisticated model, as compared to an otherwise identical institution that faces no constraints. However, while the implemented model remains under-sophisticated when operational risk is internal, it becomes excessively sophisticated when operational risk is external. We highlight the different economic mechanisms behind these results in the next paragraphs.

Consider the materialization of an operational error in the implemented model. A positive operational error may induce an over-exposure to market risk, while a negative operational error may induce an under-exposure to market risk. When operational risk is internal, a financial institution can reduce the likelihood of these events by adopting a less sophisticated model. However, a less sophisticated model delivers investment decisions that are less accurate. As discussed in the previous section, the optimal model sophistication trades off these two opposing effects. The presence of financial constraints of the form in (23), say leverage or short-selling constraints, would limit the over- or under-exposure to market risk when large realizations of operational uncertainty w_ϵ occur. As a consequence, the trade-off between a more sophisticated model and a model with high internal operational risk gets alleviated, thus inducing a constrained financial institution to increase its model sophistication.⁹

Even though the implemented model of a constrained financial institution subject to internal operational risk incorporates more market news, $\lambda^c \geq \lambda^*$, it still under-reacts to this news when compared to the most sophisticated model, $\lambda^c < 1$. In contrast, with external operational risk, financial constraints induce a financial institution to optimally over-react to market news through the choice of a model which is excessively sophisticated, $\lambda^c > 1$. The intuition for this finding is as follows. Despite the exposure to external operational risk being constant, and hence not affected by the choice of model sophistication, the financial institution can reduce the relative contribution of operational uncertainty to the total variability of the investment

⁹ The knife-edge case of $\lambda^c = \lambda^*$ obtains only when one of the two constraints is initially exactly binding ($\pi_0 = \underline{\alpha}$ or $\pi_0 = \bar{\beta}$) and the other is absent ($\bar{\beta} = \infty$ or $\underline{\alpha} = -\infty$). Moreover, if our maintained assumption $\underline{\alpha} \leq \kappa_0/\sigma \leq \bar{\beta}$ is violated, the analysis become less transparent as the mechanism highlighted above gets confounded by different initial conditions characterizing constrained and unconstrained financial institutions. In this case, the trade-off between model sophistication and operational risk may not be alleviated and model sophistication may be reduced simply because the scope for sophistication is initially constrained. The comparison between a constrained and an unconstrained institution becomes less informative. Indeed, we would not be able to attribute the change in model sophistication solely to the implemented model possibly violating financial constraints in the future, as it could also be attributed to the use of a constrained model to begin with.

model by increasing its sophistication. In fact, the ratio between the conditional variance of the implemented model driven only by operational uncertainty and its total variance,

$$\frac{\sigma_\epsilon^2}{\lambda^2 \nu_0 \hat{\nu}_t + \sigma_\epsilon^2},$$

is decreasing in λ . This alone creates an incentive for the institution to increase the sophistication of its model, even to the extent of making it overly sophisticated ($\lambda > 1$). As shown in Proposition 1, when financial constraints are absent, this incentive is perfectly offset by the implicit cost of adopting a model that over-reacts to market news, thus rendering the most sophisticated model ($\lambda = 1$) optimal. In contrast, when a financial institution is subject to financial constraints, the implicit cost of adopting a model that over-reacts to market news becomes less severe. Indeed, the presence of financial constraints would limit the over- or under-exposure to market risk when large realizations of market uncertainty w occur. As a consequence, a constrained financial institution optimally increases λ above 1, thus implementing an excessively sophisticated model.

Therefore, financial constraints affect model sophistication through different channels. Effectively, they shield financial institutions from large operational errors when operational risk is internal, and they shield them from an excessive exposure to market news when operational risk is external.

The desirable properties of the optimal model sophistication, discussed in Proposition 1 for the case of unconstrained financial institutions, hold true also for the case of constrained institutions: λ^c is decreasing in operational risk, and increasing in model risk. Figure 5 depicts the (percentage) difference in model sophistication between a constrained and an otherwise identical unconstrained financial institution, $(\lambda^c - \lambda^*)/\lambda^*$, as a function of financial constraints, and operational risk. The left and right panels show that when financial constraints are tightened ($\underline{\alpha}$ increases, $\bar{\beta}$ decreases), the difference in model sophistication tends to increase. This is the case for both internal (top panels) and external (bottom panels) operational risks, in line with the economic intuition aforementioned. However, when the constraints are very tight, the difference in model sophistication, although remaining positive, may also decrease. This occurs because a very tight constraint also limits the advantage of adopting a more sophisticated model (internal operational risk), and increases the disadvantage of adopting an overly sophisticated model (external operational risk). Indeed, a model that is very responsive to market uncertainty (i.e., to positive or negative market news) is less valuable when the market risk exposure is very constrained. The top-right panel shows that an increase in internal operational risk first widens and then narrows the difference in model sophistication. When internal operational risk is particularly high, a sophisticated model is particularly undesirable,

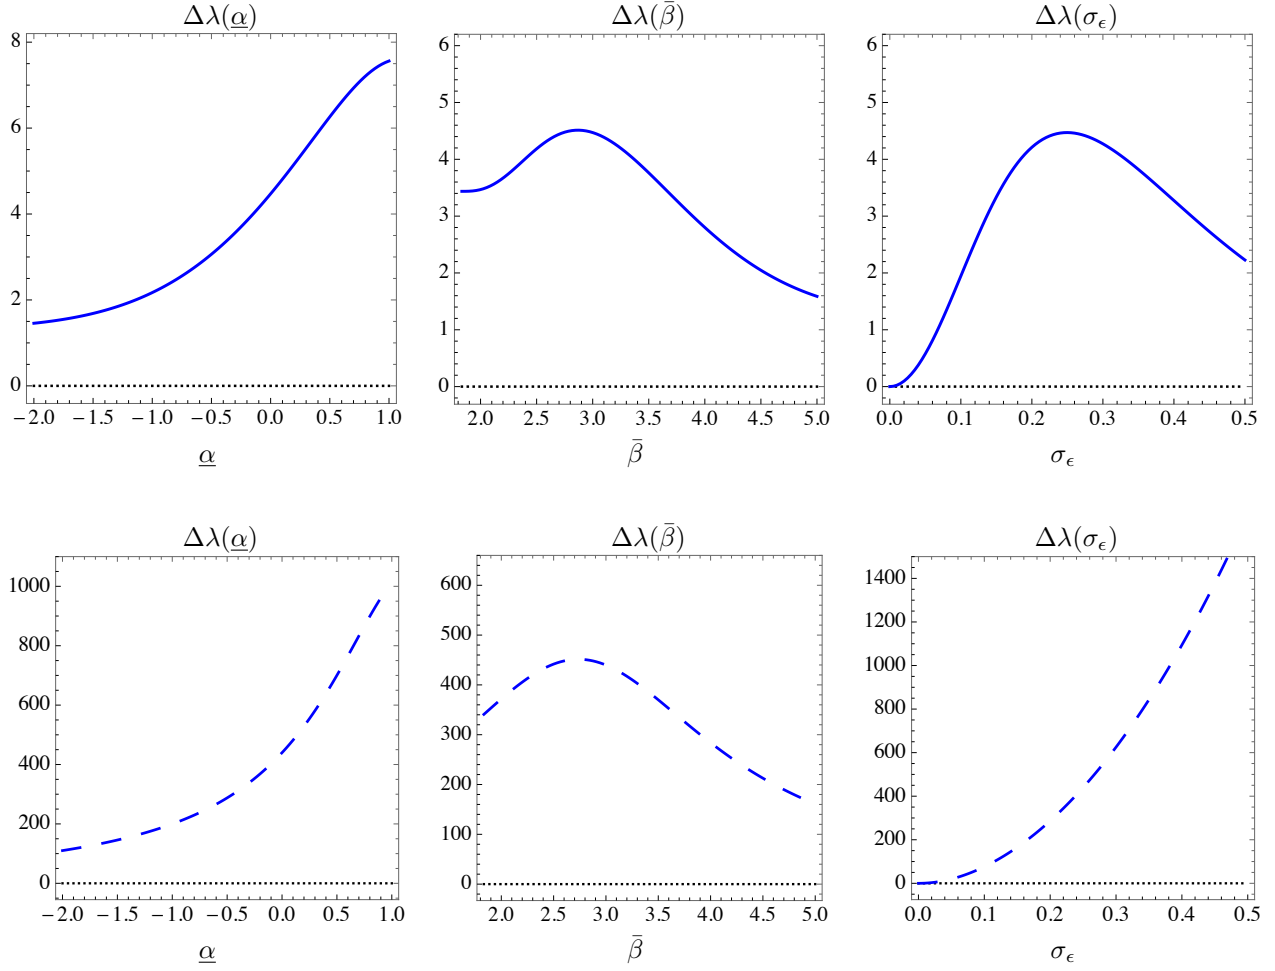


Figure 5: Change in model sophistication with financial constraints

In this figure we plot the change in model sophistication $\Delta\lambda \equiv (\lambda^c - \lambda^*)/\lambda^*$ in percentages induced by the financial constraints $\underline{\alpha} \leq \pi_t \leq \bar{\beta}$, as a function of the lower-bound constraint $\underline{\alpha}$ (left panels), the upper-bound constraint $\bar{\beta}$ (central panels), and operational risk σ_ϵ (right panels). The dotted line represents the benchmark case of no operational risk ($\sigma_\epsilon = 0$). The solid line refers to the case of internal operational risk (top panels), and the dashed line to the case of external operational risk (bottom panels). The parameter values are $\kappa_0 = 0.33$, $\underline{\alpha} = 0$, $\bar{\beta} = 3$, and the remainder are as in Figure 3. The plots are typical.

and the shielding role of financial constraints against operational risk becomes less relevant. In contrast, bottom-right panel shows that an increase in external operational risk always widens the difference in model sophistication. When external operational risk increases, the incentive to reduce the relative contribution of operational uncertainty to the total variability of the implemented model becomes stronger.

The next proposition presents the effects of financial constraints on the optimal market risk exposure.

Proposition 6 (Market risk exposure with financial constraints). *In the presence of financial constraints \mathcal{C} , the optimal market risk exposure is*

$$\pi_t^c = \pi_t^*(\lambda^c) + \left[(\underline{\alpha} - \pi_t^*(\lambda^c))^+ - (\pi_t^*(\lambda^c) - \bar{\beta})^+ \right], \quad (28)$$

where $\pi_t^*(\cdot)$ is the unconstrained market risk exposure in (17). Consequently, when the constrained market risk exposure is not binding, $\underline{\alpha} < \pi_t^c < \bar{\beta}$, it coincides with the unconstrained market risk exposure absent operational risk, $\pi_t^c = \pi_t^*$, and it always differs in the presence of operational risk,

$$\pi_t^c - \pi_t^* = (\lambda^c - \lambda^*) \left(\frac{\hat{\kappa}_t - \kappa_0 + \sigma_\epsilon w_{et}}{\sigma} \right) \neq 0, \quad (29)$$

where λ^* is as in (13) and λ^c solves (26) when operational risk is internal and solves (27) when external.

Equation (28) reveals that the constrained market risk exposure π_t^c is given by the unconstrained market risk exposure evaluated at the constrained model sophistication λ^c , plus an additional term that arises due to the presence of financial constraints. Such constraints affect the optimal market risk exposure not only directly by imposing lower and/or upper limits, but also indirectly by changing the sophistication of the model that a constrained institution optimally adopts. This implies that, as reported in (29), even when financial constraints are not binding, the market risk exposure differs from the unconstrained one. This highlights a new channel through which financial constraints affect financial investments.

4.3 Benefits of Financial Constraints

We next study how financial constraints affect the profitability of a financial institution. In particular, we analyze whether financial constraints can be beneficial when operational risk is present.

We consider a benefit-cost measure of financial constraints η^c that represents the percentage of initial capital W_0 that an institution without financial constraints must additionally have in order to be indifferent to the constraints \mathcal{C} :

$$\eta^c : \quad J^c(\lambda^c; W_0) = J(\lambda^*; W_0(1 + \eta^c)). \quad (30)$$

When this measure is negative ($\eta^c < 0$) financial constraints are detrimental, and when positive ($\eta^c > 0$) they are beneficial. Even though we have an analytical representation for the benefit-

cost measure (as reported in Appendix A.2) we find it more helpful to present our results with plots.

Figure 6 depicts the benefits-cost measure of financial constraints. The left panel reveals that financial constraints are costly when operational risk is absent (dotted line) or when it is internal (solid line). Indeed one would expect the constraints to typically be detrimental ($\eta^c < 0$) since they restrict the set of possible investment choices (i.e., market risk exposures). In stark contrast, the left panel also reveals that financial constraints may be beneficial ($\eta^c > 0$) when operational risk is external (dashed line), particularly for higher levels of operational risk. This occurs because of two opposing effects. First, financial constraints may be beneficial as they prevent the implementation of extreme market risk exposures that are possibly due to large realizations of operational uncertainty. For this reason, constraints help reduce the exposure to operational risk. At the same time, financial constraints may be detrimental as they also prevent valuable market news to be fully incorporated in the investment choices. When operational risk is external and sufficiently high, the former effect dominates and financial constraints make a financial institution better off. In contrast, when operational risk is internal, the exposure to operational risk is managed more efficiently by altering the sophistication of the investment model, and when it is absent, there is no operational risk to be managed. Therefore, in these cases financial constraints are always detrimental.¹⁰

Notably, the largest costs of financial constraints are always associated with the case of no operational risk, implying that financial constraints are less costly when operational risk is present. This is because, absent operational risk, an institution’s market risk exposure is more responsive to market news (since the most sophisticated model is always implemented), and hence it is more likely that the financial constraints are binding. Moreover, the shielding role of constraints against operational risk are absent. The center and left panels in Figure 6 depict the effects of tighter financial constraints. When operational risk is absent or is internal, the constraints become increasingly costlier as $\underline{\alpha}$ increases (left panel), or $\bar{\beta}$ decreases (right panel). With external operational risk, instead, the benefits of financial constraints increase first and, once the constraints become very tight, they decrease. Moreover, these benefits are economically significant, as unveiled by all the plots in Figure 6.

Our findings also provide novel insights on the economic role of self-imposed risk limits when financial institutions are subject to forms of operational risk that cannot be directly managed internally. Consider for instance an error made by mis-typing a trade, also known as a “fat finger” mistake. This represents a type of external operational risk, since reducing

¹⁰In the presence of external operational risk, financial constraints are also beneficial ($\eta^c > 0$) for lower levels of model risk. The same aforementioned intuition holds, as low model risk makes the second opposing effect less important.

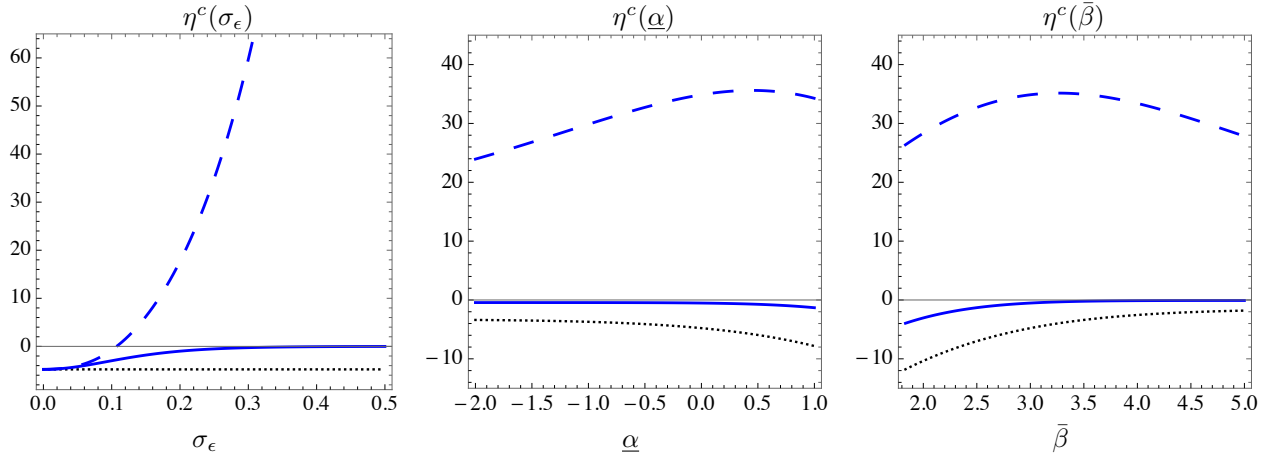


Figure 6: Benefits of financial constraints

In this figure we plot the benefit-cost measure of financial constraints η^c in percentages, as a function of operational risk σ_ϵ (left panel), the lower-bound constraint $\underline{\alpha}$ (center panel), and the upper-bound constraint $\bar{\beta}$ (right panel). The dotted line represents the benchmark case of no operational risk ($\sigma_\epsilon = 0$). The solid line refers to the case of internal operational risk, and the dashed line to the case of external operational risk. The parameter values are as in Figure 5. The plots are typical.

the sophistication of the model would not reduce the exposure to these errors. However, the presence of risk limits, either self-imposed or due to regulation, would contain the damage of these mistakes.

5 Conclusions

In this paper we study the decision making of a financial institution in the presence of a novel implementation friction that gives rise to operational risk. This is a first step towards incorporating operational risk into a dynamic asset allocation framework. Our setting emerges rich in implications. When operational risk is internal, a financial institution optimally adopts a less sophisticated investment model in order to reduce the likelihood and the severity of operational losses. Because of the endogenous nature of model sophistication, the exposure of a financial institution to operational risk becomes decreasing in the level of operational risk precisely when this risk is high. As a consequence, operational risk also makes the financial institution's exposure to market risk less volatile. Furthermore, financial constraints that institutions may face due to regulatory or self-imposed risk limits, interact with both internal and external operational risk, and ultimately affect the optimal choice of model sophistication. In particular, a constrained financial institution has the incentive to increase the sophistication of its invest-

ment model, as compared to an otherwise identical unconstrained institution. However, while the implemented model remains under-sophisticated if operational risk is internal, it becomes excessively sophisticated if it is external. When model sophistication is not effective in reducing the exposure to operational risk, i.e., when operational risk is external, financial constraints prove to be beneficial.

It would be of interest to undertake our analysis in several other directions. One extension would be to allow for an inter-temporal choice of model sophistication. In this setting, a financial institution would initially choose an organizational structure taking into account the possibility to change this structure in the future. This would allow us to examine the dynamics of model sophistication. Another direction for further investigation would be to study the implications of operational risk on equilibrium asset prices. Interestingly, even if the operational errors affecting the investment decisions of different financial institutions were to cancel each other out, operational risk would still indirectly affect equilibrium prices through the institutions' optimal choices of model sophistication. Moreover, endogenizing asset prices would also offer the opportunity to perform a normative analysis with emphasis on policy implications.

Appendix A: Proofs

Proof of Proposition 1. As discussed in Section 2.3, the financial institution's optimization problem can be solved in two steps. The optimal model sophistication λ^* is obtained in the second step.

Step 1. The first step is to solve for the optimal market risk exposure π_t^* for a given model sophistication λ . This entails solving the optimization problem in (11) subject to the dynamic budget constraint (10). From the dynamic budget constraint, we obtain the institution's capital at time T ,

$$W_T = W_0 \exp \left[\int_0^T \left(r + \pi_t \sigma \kappa_t^*(\lambda) - \frac{1}{2} \pi_t^2 \sigma^2 \right) dt + \int_0^T \pi_t \sigma d\hat{w}_t^* \right]. \quad (\text{A.1})$$

Consequently, the optimization problem in (11) subject to (10) becomes

$$\{\pi_t^*\}_0^T \in \arg \max_{\pi} \mathbb{E}_0^* \left[\int_0^T \left(\pi_t \sigma \kappa_t^*(\lambda) - \frac{1}{2} \pi_t^2 \sigma^2 \right) dt \right]. \quad (\text{A.2})$$

where $\kappa_t^*(\lambda)$ is as in (6). For a given realization of $(w_t^*, w_{\epsilon t})$, we can solve the optimization point-wise, thus obtaining

$$\pi_t^*(\lambda) = \frac{\kappa_t^*(\lambda)}{\sigma} = \frac{\hat{\kappa}_t}{\sigma} + \frac{(1 - \lambda)(\kappa_0 - \hat{\kappa}_t) + \lambda \sigma_{\epsilon} w_{\epsilon t}}{\sigma}. \quad (\text{A.3})$$

Step 2. The second step is to solve for the optimal model sophistication λ^* , given the optimal market risk exposure $\pi_t^*(\lambda)$ in (A.3). This entails solving the optimization problem in (12). Substituting (A.3) into (A.1) and evaluating the institution's horizon capital under the probability measure $\hat{\mathbb{P}}$, making use of (8), we obtain

$$W_T = W_0 \exp \left[\int_0^T \left(r + \kappa_t^*(\lambda) \hat{\kappa}_t - \frac{1}{2} \kappa_t^*(\lambda)^2 \right) dt + \int_0^T \kappa_t^*(\lambda) d\hat{w}_t \right]. \quad (\text{A.4})$$

It follows that the value function in (12) is equal to

$$J(\lambda; W_0) = \log W_0 + rT + \hat{\mathbb{E}}_0 \left[\int_0^T \left(\kappa_t^*(\lambda) \hat{\kappa}_t - \frac{1}{2} \kappa_t^*(\lambda)^2 \right) dt + \int_0^T \kappa_t^*(\lambda) d\hat{w}_t \right], \quad (\text{A.5})$$

which can be further simplified to

$$J(\lambda; W_0) = \log W_0 + \left(r + \frac{1}{2} \kappa_0^2 \right) T + \int_0^T \text{c}\hat{\text{ov}}_0[\kappa_t^*(\lambda), \hat{\kappa}_t] dt - \frac{1}{2} \int_0^T \text{v}\hat{\text{ar}}_0[\kappa_t^*(\lambda)] dt. \quad (\text{A.6})$$

When operational risk is internal, κ_t^* is as in (6), the covariance between the implemented and the most sophisticated models and the variance of the implemented model are given by

$$\text{côv}_0[\kappa_t^*(\lambda), \hat{\kappa}_t] = \lambda \left(\frac{\nu_0^2 t}{1 + \nu_0 t} \right), \quad (\text{A.7})$$

$$\text{vâr}_0[\kappa_t^*(\lambda)] = \lambda^2 \left(\frac{\nu_0^2 t}{1 + \nu_0 t} + \sigma_\epsilon^2 t \right), \quad (\text{A.8})$$

respectively. Consequently,

$$\lambda^* \in \arg \max_{\lambda} \lambda \int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} dt - \frac{\lambda^2}{2} \int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} + \sigma_\epsilon^2 t dt, \quad (\text{A.9})$$

yielding

$$\lambda^* = \frac{\int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} dt}{\int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} + \sigma_\epsilon^2 t dt}, \quad (\text{A.10})$$

which satisfies $0 < \lambda^c < 1$ for any non-degenerate levels of model and operational risks. Since

$$\int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} dt = \nu_0 T - \log(1 + \nu_0 T), \quad (\text{A.11})$$

(13) obtains. The partial derivatives of λ^* in (13) with respect to σ_ϵ and ν_0 are equal to

$$\frac{\partial \lambda^*}{\partial \sigma_\epsilon} = - \frac{(\nu_0 T - \log(1 + \nu_0 T)) \sigma_\epsilon T^2}{(\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2)^2}, \quad (\text{A.12})$$

$$\frac{\partial \lambda^*}{\partial \nu_0} = \frac{\nu_0 \sigma_\epsilon^2 T^4}{2(1 + \nu_0 T) (\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2)^2}, \quad (\text{A.13})$$

respectively, where the first one is always negative and the second one always positive.

When operational risk is external, $\kappa_t^* = (1 - \lambda)\kappa_0 + \lambda\hat{\kappa}_t + \sigma_\epsilon w_{\epsilon t}$, the covariance between the implemented and the most sophisticated models remains as in (A.7), whereas the variance of the implemented model becomes equal to

$$\text{vâr}_0[\kappa_t^*(\lambda)] = \lambda^2 \left(\frac{\nu_0^2 t}{1 + \nu_0 t} \right) + \sigma_\epsilon^2 t. \quad (\text{A.14})$$

Consequently,

$$\lambda^* \in \arg \max_{\lambda} \left(\lambda - \frac{\lambda^2}{2} \right) \int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} dt - \frac{1}{2} \int_0^T \sigma_\epsilon^2 t dt, \quad (\text{A.15})$$

yielding $\lambda^* = 1$. □

Proof of Proposition 2. When operational risk is internal, the optimal operational risk exposure equals

$$\lambda^* \sigma_\epsilon = \frac{(\nu_0 T - \log(1 + \nu_0 T)) \sigma_\epsilon}{\nu_0 T - \log(1 + \nu_0 T) + (1/2) \sigma_\epsilon^2 T^2}. \quad (\text{A.16})$$

The partial derivative of $\lambda^* \sigma_\epsilon$ with respect to σ_ϵ ,

$$\frac{\partial \lambda^* \sigma_\epsilon}{\partial \sigma_\epsilon} = \frac{(\nu_0 T - \log(1 + \nu_0 T))}{(\nu_0 T - \log(1 + \nu_0 T) + (1/2) \sigma_\epsilon^2 T^2)^2} \left(\nu_0 T - \log(1 + \nu_0 T) - \frac{\sigma_\epsilon^2 T^2}{2} \right), \quad (\text{A.17})$$

is positive for $\sigma_\epsilon < \bar{\sigma}_\epsilon$, and negative otherwise, where

$$\bar{\sigma}_\epsilon = \frac{\sqrt{2(\nu_0 T - \log(1 + \nu_0 T))}}{T} \quad (\text{A.18})$$

is obtained by setting the partial derivative in (A.17) equal to 0.

When operational risk is external, the optimal operational risk exposure equals σ_ϵ , and hence it is always increasing in σ_ϵ . \square

Proof of Proposition 3. We obtain the optimal market risk exposure for the cases of internal and external operational risk by substituting (13) and $\lambda^* = 1$ into (A.3), respectively.

When operational risk is internal, the variance of the optimal market risk exposure is equal to

$$\hat{\text{var}}_0[\pi_t^*] = \left(\frac{\nu_0 T - \log(1 + \nu_0 T)}{\sigma (\nu_0 T - \log(1 + \nu_0 T) + (1/2) \sigma_\epsilon^2 T^2)} \right)^2 \left(\frac{\nu_0^2}{1 + \nu_0 t} + \sigma_\epsilon^2 \right) t. \quad (\text{A.19})$$

Since

$$\frac{\partial \hat{\text{var}}_0[\pi_t^*]}{\partial \sigma_\epsilon} = - \frac{\left(\sigma_\epsilon^2 T^2 + \log(1 + \nu_0 T) - \frac{\nu_0 T}{1 + \nu_0 t} \right) \lambda^{*2} \sigma_\epsilon t}{\sigma^2 (\nu_0 T - \log(1 + \nu_0 T) + (1/2) \sigma_\epsilon^2 T^2)}, \quad (\text{A.20})$$

and $\log(1 + x) - x/(1 + x) > 0$ for any $x > 0$, the variance of the optimal market risk exposure is always decreasing in σ_ϵ .

When operational risk is external, the variance of the optimal market risk exposure is equal to

$$\hat{\text{var}}_0[\pi_t^*] = \frac{1}{\sigma^2} \left(\frac{\nu_0^2}{1 + \nu_0 t} + \sigma_\epsilon^2 \right) t, \quad (\text{A.21})$$

which is increasing in σ_ϵ . \square

Proof of Proposition 4. The value function associated with the optimal model sophistication $J(\lambda^*; W_0)$ in the presence of internal operational risk is obtained by substituting (13) into (A.6):

$$J(\lambda^*; W_0) = \log W_0 + \left(r + \frac{1}{2}\kappa_0^2 \right) T + \frac{(\nu_0 T - \log(1 + \nu_0 T))^2}{2(\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2)}. \quad (\text{A.22})$$

The value functions associated with the cases of maximal sophistication and unsophistication are equal to

$$J(1; W_0) = \log W_0 + \left(r + \frac{1}{2}\kappa_0^2 \right) T + \frac{1}{2} \left(\nu_0 T - \log(1 + \nu_0 T) - \frac{\sigma_\epsilon^2 T^2}{2} \right), \quad (\text{A.23})$$

$$J(0; W_0) = \log W_0 + \left(r + \frac{1}{2}\kappa_0^2 \right) T, \quad (\text{A.24})$$

respectively. Given the definitions in (19) and (20), the costs of maximal sophistication and unsophistication are given by

$$\eta_{\lambda=1} = \exp [J(\lambda^*; W_0) - J(1; W_0)] - 1, \quad (\text{A.25})$$

$$\eta_{\lambda=0} = \exp [J(\lambda^*; W_0) - J(0; W_0)] - 1, \quad (\text{A.26})$$

thus yielding (21) and (22), respectively. \square

Lemma A.1. Consider the standardized normal bivariate distribution of probability density

$$\varphi(x, y; \rho) = \gamma(2\pi)^{-1} \exp \left\{ -\frac{\gamma^2}{2} (x^2 - 2\rho xy + y^2) \right\}$$

where ρ is the coefficient of correlation between x and y , and $\gamma \equiv (1 - \rho^2)^{-1/2}$. Consider the truncation $x \geq a$, where a is a constant. Then, the following moments hold true:

$$\mathbb{E}(x \mathbf{1}_{\{x \geq a\}}) = n(a) \quad (\text{A.27})$$

$$\mathbb{E}(y \mathbf{1}_{\{x \geq a\}}) = \rho n(a) \quad (\text{A.28})$$

$$\mathbb{E}(x^2 \mathbf{1}_{\{x \geq a\}}) = n(a) a + 1 - N(a) \quad (\text{A.29})$$

$$\mathbb{E}(y^2 \mathbf{1}_{\{x \geq a\}}) = \rho^2 n(a) a + 1 - N(a) \quad (\text{A.30})$$

$$\mathbb{E}(xy \mathbf{1}_{\{x \geq a\}}) = \rho [n(a) a + 1 - N(a)] \quad (\text{A.31})$$

where

$$n(a) = (2\pi)^{-1/2} e^{-(1/2)a^2}, \quad N(a) = \int_{-\infty}^a n(s) ds. \quad (\text{A.32})$$

Proof. See Weiler (1959). \square

Lemma A.2. Consider two sets of deterministic time-weights ω_{1t} and ω_{2t} , defined over the time interval $[0, T]$, that satisfy the following conditions:

(i) The inter-temporal sum of the time-weights over the time interval $[0, T]$ is equal to 1:

$$\int_0^T \omega_{it} dt = 1 \quad \text{for } i = 1, 2. \quad (\text{A.33})$$

(ii) The initial time-weights are 0: $\omega_{i0} = 0$ for $i = 1, 2$.

(iii) The time-weights are increasing and concave in time:

$$\frac{\partial \omega_{it}}{\partial t} > 0, \quad \frac{\partial^2 \omega_{it}}{\partial t^2} < 0 \quad \forall t \in [0, T] \quad \text{and for } i = 1, 2. \quad (\text{A.34})$$

(iv) Without loss of generality,

$$\left. \frac{\partial \omega_{1t}}{\partial t} \right|_{t=0} > \left. \frac{\partial \omega_{2t}}{\partial t} \right|_{t=0}. \quad (\text{A.35})$$

Consider a deterministic function of time $f(t)$. If $f(t)$ is decreasing in time,

$$\frac{\partial f(t)}{\partial t} \leq 0 \quad \forall t \in [0, T], \quad (\text{A.36})$$

then

$$\int_0^T f(t) \omega_{1t} dt \geq \int_0^T f(t) \omega_{2t} dt. \quad (\text{A.37})$$

Proof. Properties (i), (ii) and (iii) imply that the time-weights ω_{1t} and ω_{2t} cross only once in the time interval $[0, T]$. Adding property (iv) implies that $\omega_{1T} < \omega_{2T}$. Therefore, it immediately follows that

$$\int_0^t \omega_{1s} ds \geq \int_0^t \omega_{2s} ds \quad \forall t \in [0, T] \quad (\text{A.38})$$

Integrating by parts, we obtain that

$$\int_0^T f(t) (\omega_{1t} - \omega_{2t}) dt = f(T) \int_0^T (\omega_{1t} - \omega_{2t}) - \int_0^T \frac{\partial f(t)}{\partial t} \left[\int_0^t (\omega_{1s} - \omega_{2s}) ds \right] dt \geq 0. \quad (\text{A.39})$$

The first term on the RHS of (A.39) is equal to zero because of (A.33), and the second term is (weakly) positive because of (A.36) and (A.38). \square

Proof of Proposition 5. We follow closely the two steps in Proposition 1 to solve the optimization problem of a constrained financial institution.

Step 1. We first determine the optimal market risk exposure π_t^c (for a given model sophistication λ) by solving the optimization problem in (11) subject to the dynamic budget constraint (10) and the financial constraints in (23). The logarithmic preferences enables us to solve the dynamic optimization problem with constraints as a point-wise constrained problem. We could alternatively use the methodology of Cvitanic and Karatzas (1992), who study a class of portfolio choice problems with constraints and general preferences. The institution's capital at time T is as in (A.1). Consequently, the optimization problem in (11) subject to (10) and (23) becomes

$$\{\pi_t^c\}_0^T \in \arg \max_{\pi \in \mathcal{C}} \mathbb{E}_0^* \left[\int_0^T \left(\pi_t \sigma \kappa_t^*(\lambda) - \frac{1}{2} \pi_t^2 \sigma^2 \right) dt \right]. \quad (\text{A.40})$$

where $\kappa_t^*(\lambda)$ is as in (6). For a given realization of (w_t^*, w_{et}) , we can solve the optimization point-wise, thus obtaining

$$\pi_t^c(\lambda) = \frac{\kappa_t^*(\lambda)}{\sigma} + \left[\left(\underline{\alpha} - \frac{\kappa_t^*(\lambda)}{\sigma} \right)^+ - \left(\frac{\kappa_t^*(\lambda)}{\sigma} - \bar{\beta} \right)^+ \right], \quad (\text{A.41})$$

where we adopt the notation $x^+ \equiv \max\{x, 0\}$.

Step 2. We now determine the optimal model sophistication λ^c with financial constraints, given the optimal market risk exposure $\pi_t^c(\lambda)$ in (A.41), by solving the optimization problem in (24). Substituting (A.41) into (A.1) and evaluating the institution's horizon capital under the probability measure $\hat{\mathbb{P}}$, making use of (8), we obtain

$$W_T = W_0 \exp \left[\int_0^T \left(r + \sigma \pi_t^c(\lambda) \hat{\kappa}_t - \frac{\sigma^2}{2} \pi_t^c(\lambda)^2 \right) dt + \int_0^T \sigma \pi_t^c(\lambda) d\hat{w}_t \right]. \quad (\text{A.42})$$

It follows that the value function in (24) is equal to

$$J^c(\lambda; W_0) = \log W_0 + rT + \int_0^T \sigma \hat{\mathbb{E}}_0[\pi_t^c(\lambda) \hat{\kappa}_t] dt - \frac{1}{2} \int_0^T \sigma^2 \hat{\mathbb{E}}_0[\pi_t^c(\lambda)^2] dt. \quad (\text{A.43})$$

To make use of Lemma A.1, we conveniently write $\sigma \pi_t^c(\lambda)$ as

$$\sigma \pi_t^c(\lambda) = \kappa_t^*(\lambda) \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}} - \kappa_t^*(\lambda) \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \beta\}} + \alpha - \alpha \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}} + \beta \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}, \quad (\text{A.44})$$

where $\alpha \equiv \sigma \underline{\alpha}$ and $\beta \equiv \sigma \bar{\beta}$. Consequently, the moments in (A.43) are equal to

$$\begin{aligned} \sigma \hat{\mathbb{E}}_0[\pi_t^c(\lambda) \hat{\kappa}_t] &= \hat{\mathbb{E}}_0[\kappa_t^*(\lambda) \hat{\kappa}_t \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] - \hat{\mathbb{E}}_0[\kappa_t^*(\lambda) \hat{\kappa}_t \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}] + \alpha \kappa_0 \\ &\quad - \alpha \hat{\mathbb{E}}_0[\hat{\kappa}_t \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] + \beta \hat{\mathbb{E}}_0[\hat{\kappa}_t \mathbb{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}], \end{aligned} \quad (\text{A.45})$$

$$\begin{aligned}\sigma^2 \hat{\mathbb{E}}_0[\pi_t^c(\lambda)^2] &= \hat{\mathbb{E}}_0[\kappa_t^*(\lambda)^2 \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] - \hat{\mathbb{E}}_0[\kappa_t^*(\lambda)^2 \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}] \\ &\quad + \alpha^2 \hat{\mathbb{P}}_0[\kappa_t^*(\lambda) < \alpha] + \beta^2 \hat{\mathbb{P}}_0[\kappa_t^*(\lambda) \geq \beta].\end{aligned}\tag{A.46}$$

Since both the most sophisticated and the implemented models at time t are (unconditionally) normally distributed, we can express them as

$$\hat{\kappa}_t = \kappa_0 + \hat{\sigma}_t \hat{z}_t,\tag{A.47}$$

$$\kappa_t^* = \kappa_0 + \sigma_t^*(\lambda) z_t^*,\tag{A.48}$$

where

$$\hat{\sigma}_t = \sqrt{\nu_0 \hat{\nu}_t t},\tag{A.49}$$

$$\sigma_t^*(\lambda) = \sqrt{\lambda^2 \nu_0 \hat{\nu}_t t + h(\lambda)^2 \sigma_\epsilon^2 t},\tag{A.50}$$

and $\hat{\nu}_t$ is as in (4). Note that our cases of internal and external operational risk are nested in the above specification for $h(\lambda) = \lambda$ and $h(\lambda) = 1$, respectively. (z_t^*, \hat{z}_t) are distributed according to a standardized normal bivariate distribution for any t , with coefficient of correlation $\varrho_t = \lambda \nu_0 \hat{\nu}_t t / \hat{\sigma}_t \sigma_t^*(\lambda)$. The conditions $\kappa_t^*(\lambda) \geq \alpha$ and $\kappa_t^*(\lambda) \geq \beta$ imply that $z_t^* \geq a_t(\lambda)$ and $z_t^* \geq b_t(\lambda)$, respectively, where

$$a_t(\lambda) \equiv \frac{\alpha - \kappa_0}{\sigma_t^*(\lambda)}, \quad b_t(\lambda) \equiv \frac{\beta - \kappa_0}{\sigma_t^*(\lambda)}.\tag{A.51}$$

We use the results in Lemma A.1 to evaluate the expectations and probabilities in (A.45) and (A.46):

$$\hat{\mathbb{E}}_0[\hat{\kappa}_t \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] = (1 - N(a_t(\lambda))) \kappa_0 + n(a_t(\lambda)) \varrho_t \hat{\sigma}_t,\tag{A.52}$$

$$\hat{\mathbb{E}}_0[\hat{\kappa}_t \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}] = (1 - N(b_t(\lambda))) \kappa_0 + n(b_t(\lambda)) \varrho_t \hat{\sigma}_t,\tag{A.53}$$

$$\begin{aligned}\hat{\mathbb{E}}_0[\kappa_t^*(\lambda) \hat{\kappa}_t \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] &= (1 - N(a_t(\lambda))) (\kappa_0^2 + \varrho_t \hat{\sigma}_t \sigma_t^*(\lambda)) \\ &\quad + n(a_t(\lambda)) (a_t(\lambda) \varrho_t \hat{\sigma}_t \sigma_t^*(\lambda) + \kappa_0 \varrho_t \hat{\sigma}_t + \kappa_0 \sigma_t^*(\lambda)),\end{aligned}\tag{A.54}$$

$$\begin{aligned}\hat{\mathbb{E}}_0[\kappa_t^*(\lambda) \hat{\kappa}_t \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}] &= (1 - N(b_t(\lambda))) (\kappa_0^2 + \varrho_t \hat{\sigma}_t \sigma_t^*(\lambda)) \\ &\quad + n(b_t(\lambda)) (b_t(\lambda) \varrho_t \hat{\sigma}_t \sigma_t^*(\lambda) + \kappa_0 \varrho_t \hat{\sigma}_t + \kappa_0 \sigma_t^*(\lambda)),\end{aligned}\tag{A.55}$$

$$\begin{aligned}\hat{\mathbb{E}}_0[\kappa_t^*(\lambda)^2 \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \alpha\}}] &= (1 - N(a_t(\lambda))) (\kappa_0^2 + \sigma_t^*(\lambda)^2) \\ &\quad + n(a_t(\lambda)) (a_t(\lambda) \sigma_t^*(\lambda)^2 + 2\kappa_0 \sigma_t^*(\lambda)),\end{aligned}\tag{A.56}$$

$$\begin{aligned}\hat{\mathbb{E}}_0[\kappa_t^*(\lambda)^2 \mathbf{1}_{\{\kappa_t^*(\lambda) \geq \beta\}}] &= (1 - N(b_t(\lambda))) (\kappa_0^2 + \sigma_t^*(\lambda)^2) \\ &\quad + n(b_t(\lambda)) (b_t(\lambda) \sigma_t^*(\lambda)^2 + 2\kappa_0 \sigma_t^*(\lambda)),\end{aligned}\tag{A.57}$$

$$\hat{\mathbb{P}}_0[\kappa_t^*(\lambda) < \alpha] = N(a_t(\lambda)),\tag{A.58}$$

$$\hat{\mathbb{P}}_0[\kappa_t^*(\lambda) \geq \beta] = 1 - N(b_t(\lambda)),\tag{A.59}$$

where $n(\cdot)$ and $N(\cdot)$ represent the probability density function and cumulative distribution function of a standard normal distribution, respectively, and are defined in (A.32). Substituting (A.52)-(A.59) into (A.43), using (A.45) and (A.46), we obtain

$$J^c(\lambda; W_0) = \log W_0 + \left(r + \frac{1}{2}\kappa_0^2\right)T - \frac{1}{2} \int_0^T N(a_t(\lambda))(\alpha - \kappa_0)^2 + (1 - N(b_t(\lambda)))(\beta - \kappa_0)^2 dt \\ + \int_0^T [N(b_t(\lambda)) - N(a_t(\lambda))] \lambda \nu_0 \hat{\nu}_t - \frac{1}{2} [\Omega(a_t(\lambda)) - \Omega(b_t(\lambda))] \sigma_t^*(\lambda)^2 dt, \quad (\text{A.60})$$

where

$$\Omega(x) \equiv n(x)x + 1 - N(x). \quad (\text{A.61})$$

The optimal model sophistication in the presence of financial constraints λ^c maximizes the value function $J^c(\lambda; W_0)$ in (A.60). Exploiting the properties of standard normal distributions,

$$N'(x) = n(x), \quad n'(x) = -n(x)x, \quad (\text{A.62})$$

we obtain the following first-order condition:

$$\int_0^T \left\{ [N(b_t(\lambda^c)) - N(a_t(\lambda^c))] + [n(a_t(\lambda^c))a_t(\lambda^c) - n(b_t(\lambda^c))b_t(\lambda^c)] \frac{\lambda^c}{\sigma_t^*(\lambda^c)} \frac{\partial \sigma_t^*(\lambda^c)}{\partial \lambda^c} \right\} \nu_0 \hat{\nu}_t dt \\ - \int_0^T [\Omega(a_t(\lambda^c)) - \Omega(b_t(\lambda^c))] \sigma_t^*(\lambda^c) \frac{\partial \sigma_t^*(\lambda^c)}{\partial \lambda^c} dt = 0. \quad (\text{A.63})$$

We next specialize the proof for two cases of internal and external operational risk.

Internal operational risk. When operational risk is internal, given our linear specification $h(\lambda) = \lambda$, the standard deviation of the implemented model in (A.50) becomes equal to

$$\sigma_t^*(\lambda) = \sqrt{\lambda^2(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}, \quad (\text{A.64})$$

with partial derivative with respect to model sophistication

$$\frac{\partial \sigma_t^*(\lambda)}{\partial \lambda} = \frac{\sigma_t^*(\lambda)}{\lambda}. \quad (\text{A.65})$$

Substituting (A.64) and (A.65) into (A.63), we obtain

$$\int_0^T [\Omega(a_t(\lambda^c)) - \Omega(b_t(\lambda^c))] \{ (1 - \lambda^c) \nu_0 \hat{\nu}_t - \lambda^c \sigma_\epsilon^2 t \} dt = 0, \quad (\text{A.66})$$

which coincides with (26) after rearranging terms and using the definitions in (A.51).

We next show that $\lambda^c \geq \lambda^*$. To this purpose, let $f(t)$ denote a function of time t , defined as

$$f(t) \equiv \Omega \left(\frac{\sigma \underline{\alpha} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right) - \Omega \left(\frac{\sigma \bar{\beta} - \kappa_0}{\sqrt{\lambda^{c^2}(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2)t}} \right), \quad (\text{A.67})$$

where $\Omega(\cdot)$ is as in (A.61). It follows from the first-order condition (A.66) that the optimal model sophistication with internal operational risk can be expressed as

$$\lambda^c = \frac{\int_0^T f(t) \nu_0 \hat{\nu}_t t dt}{\int_0^T f(t) (\nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t dt}, \quad (\text{A.68})$$

where the above expression is not an explicit solution for λ^c as the function $f(t)$ depends on λ^c . However, since $f(t) > 0$, we can conclude that $0 < \lambda^c < 1$ for any non-degenerate levels of model and operational risks. Absent financial constraints, corresponding to the nested case $\underline{\alpha} = -\infty$ and $\bar{\beta} = +\infty$, the function $f(t) = 1$, thus yielding an explicit solution for λ^* ,

$$\lambda^* = \frac{\int_0^T \nu_0 \hat{\nu}_t t dt}{\int_0^T (\nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t dt}, \quad (\text{A.69})$$

which is as in (A.10). Given (A.68) and (A.69), the following inequality

$$\int_0^T f(t) \frac{\nu_0 \hat{\nu}_t t}{\int_0^T \nu_0 \hat{\nu}_s s ds} dt \geq \int_0^T f(t) \frac{(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t}{\int_0^T (\nu_0 \hat{\nu}_s + \sigma_\epsilon^2) s ds} dt. \quad (\text{A.70})$$

implies $\lambda^c \geq \lambda^*$. We make use of Lemma A.2 to show that the sufficient conditions for (A.70) to hold are satisfied. Let us define,

$$\omega_{1t} \equiv \frac{\nu_0 \hat{\nu}_t t}{\int_0^T \nu_0 \hat{\nu}_s s ds}, \quad \omega_{2t} \equiv \frac{(\nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t}{\int_0^T (\nu_0 \hat{\nu}_s + \sigma_\epsilon^2) s ds}. \quad (\text{A.71})$$

It is immediate to see that the time-weights ω_{1t} and ω_{2t} in (A.71) satisfy conditions (i) and (ii) in Lemma A.2. The following partial derivatives,

$$\frac{\partial \omega_{1t}}{\partial t} = \frac{1}{\nu_0 T - \log(1 + \nu_0 T)} \left(\frac{\nu_0^2}{(1 + \nu_0 t)^2} \right) > 0, \quad (\text{A.72})$$

$$\frac{\partial \omega_{2t}}{\partial t} = \frac{1}{\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2} \left(\frac{\nu_0^2}{(1 + \nu_0 t)^2} + \sigma_\epsilon^2 \right) > 0, \quad (\text{A.73})$$

$$\frac{\partial^2 \omega_{1t}}{\partial t^2} = -\frac{2}{\nu_0 T - \log(1 + \nu_0 T)} \left(\frac{\nu_0^3}{(1 + \nu_0 t)^3} \right) < 0, \quad (\text{A.74})$$

$$\frac{\partial^2 \omega_{2t}}{\partial t^2} = -\frac{2}{\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2} \left(\frac{\nu_0^3}{(1 + \nu_0 t)^3} \right) < 0, \quad (\text{A.75})$$

confirm that condition (iii) is satisfied. Evaluating (A.72) and (A.73) at $t = 0$, we can deduce that also condition (iv) is satisfied, given that $\log(1 + x) - x + x^2/2 > 0$ for any $x > 0$. We conclude this part of the proof by showing that $f(t)$ is decreasing:

$$\begin{aligned} \frac{\partial f(t)}{\partial t} &= \frac{\partial \Omega(a_t)}{\partial a_t} \frac{\partial a_t}{\partial \sigma_t^*} \frac{\partial \sigma_t^*}{\partial t} - \frac{\partial \Omega(b_t)}{\partial b_t} \frac{\partial b_t}{\partial \sigma_t^*} \frac{\partial \sigma_t^*}{\partial t} \\ &= \frac{\partial \sigma_t^*}{\partial t} \left(\frac{\partial \Omega(a_t)}{\partial a_t} \frac{(\kappa_0 - \alpha)}{\sigma_t^{*2}} + \frac{\partial \Omega(b_t)}{\partial b_t} \frac{(\beta - \kappa_0)}{\sigma_t^{*2}} \right) < 0 \end{aligned} \quad (\text{A.76})$$

since $\Omega'(x) = -n(x)x^2 < 0$ for any $x \in \mathbb{R}$, $\partial \sigma_t^*/\partial t$ is positive given the sign of the partial derivative in (A.73), and $\alpha \leq \kappa_0 \leq \beta$ under our maintained assumption.

External operational risk. When operational risk is external, $h(\lambda) = 1$, the standard deviation of the implemented model in (A.50) becomes equal to

$$\sigma_t^*(\lambda) = \sqrt{\lambda^2 \nu_0 \hat{\nu}_t t + \sigma_\epsilon^2 t}, \quad (\text{A.77})$$

with partial derivative with respect to model sophistication

$$\frac{\partial \sigma_t^*(\lambda)}{\partial \lambda} = \frac{\sigma_t^*(\lambda)}{\lambda} - \frac{\sigma_\epsilon^2 t}{\lambda \sigma_t^*(\lambda)}. \quad (\text{A.78})$$

Substituting (A.77) and (A.78) into (A.63), we obtain

$$\begin{aligned} &\int_0^T [\Omega(a_t(\lambda^c)) - \Omega(b_t(\lambda^c))] (1 - \lambda^c) \nu_0 \hat{\nu}_t t dt \\ &+ \int_0^T \frac{[n(b_t(\lambda^c))b_t(\lambda^c) - n(a_t(\lambda^c))a_t(\lambda^c)]}{(\lambda^{c2} \nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t} (\nu_0 \hat{\nu}_t t) \sigma_\epsilon^2 t dt = 0, \end{aligned} \quad (\text{A.79})$$

which coincides with (27) after rearranging terms and using the definitions in (A.51). Since the second term on the RHS of (27) is always positive for any non-degenerate levels of model and operational risks, we can conclude that $\lambda^c > 1$. Since by Proposition 1 $\lambda^* = 1$ in the presence of external operational risk, it follows that $\lambda^c > \lambda^*$. \square

Proof of Proposition 6. We obtain the optimal market risk exposure in the presence of financial constraints for the cases of external and internal operational risk by substituting λ^c , obtained from (27) and (26), into (A.41), respectively. When the constrained market risk exposure is not binding, $\underline{\alpha} < \pi_t^c(\lambda^c) < \bar{\beta}$, the difference between the constrained and unconstrained market risk exposure is given by

$$\pi_t^c(\lambda^c) - \pi_t^*(\lambda^*) = \pi_t^c(\lambda^c) - \pi_t^*(\lambda^*), \quad (\text{A.80})$$

which is different from zero in the presence of operational risk since, from Proposition 5, $\lambda^c \geq \lambda^*$, and it is equal to zero absent operational risk, since $\lambda^c = \lambda^* = 1$. \square

A.1 Operational Losses

Given the definition of operational losses in (15) for the case of internal operational risk, and the institution's horizon capital in (A.4),

$$W_T(\sigma_\epsilon; w_\epsilon = 0) = W_0 \exp \left[\int_0^T \left(r + (\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0))\hat{\kappa}_t - \frac{1}{2}(\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0))^2 \right) dt + \int_0^T (\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0)) d\hat{w}_t \right], \quad (\text{A.81})$$

$$W_T(\sigma_\epsilon) = W_0 \exp \left[\int_0^T \left(r + (\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0 + \sigma_\epsilon w_{\epsilon t}))\hat{\kappa}_t - \frac{1}{2}(\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0 + \sigma_\epsilon w_{\epsilon t}))^2 \right) dt + \int_0^T (\kappa_0 + \lambda^*(\hat{\kappa}_t - \kappa_0 + \sigma_\epsilon w_{\epsilon t})) d\hat{w}_t \right]. \quad (\text{A.82})$$

Taking the ratio between (A.81) and (A.82) yields

$$\exp \left[\frac{(\lambda^* \sigma_\epsilon)^2}{2} \int_0^T w_{\epsilon t}^2 dt + \lambda^*(1 - \lambda^*) \sigma_\epsilon \int_0^T (\kappa_0 - \hat{\kappa}_t) w_{\epsilon t} dt - \lambda^* \sigma_\epsilon \int_0^T w_{\epsilon t} d\hat{w}_t \right]. \quad (\text{A.83})$$

A.2 Benefit-Cost Measure of Financial Constraints

Given the definition in (30), the benefit-cost measure of financial constraints with internal operational risk is equal to

$$\begin{aligned} \eta^c &= \exp [J^c(\lambda^c; W_0) - J(\lambda^*, W_0)] - 1 \quad (\text{A.84}) \\ &= \exp \left[-\frac{1}{2} \int_0^T N(a_t(\lambda^c))(\alpha - \kappa_0)^2 + (1 - N(b_t(\lambda^c)))(\beta - \kappa_0)^2 dt \right. \\ &\quad + \int_0^T [N(b_t(\lambda^c)) - N(a_t(\lambda^c))] \lambda^c \nu_0 \hat{\nu}_t - \frac{1}{2} [\Omega(a_t(\lambda^c)) - \Omega(b_t(\lambda^c))] \lambda^{c2} (\nu_0 \hat{\nu}_t + \sigma_\epsilon^2) t dt \\ &\quad \left. - \frac{1}{2} \left(\frac{(\nu_0 T - \log(1 + \nu_0 T))^2}{\nu_0 T - \log(1 + \nu_0 T) + (1/2)\sigma_\epsilon^2 T^2} \right) \right] - 1, \quad (\text{A.85}) \end{aligned}$$

where λ^* is as in (13), and λ^c solves (26).

The benefit-cost measure of financial constraints with external operational risk is equal to

$$\eta^c = \exp [J^c(\lambda^c; W_0) - J(1, W_0)] - 1 \quad (\text{A.86})$$

$$\begin{aligned} &= \exp \left[-\frac{1}{2} \int_0^T N(a_t(\lambda^c))(\alpha - \kappa_0)^2 + (1 - N(b_t(\lambda^c)))(\beta - \kappa_0)^2 dt \right. \\ &\quad + \int_0^T [N(b_t(\lambda^c)) - N(a_t(\lambda^c))] \lambda^c \nu_0 \hat{\nu}_t t - \frac{1}{2} [\Omega(a_t(\lambda^c)) - \Omega(b_t(\lambda^c))] (\lambda^{c^2} \nu_0 \hat{\nu}_t t + \sigma_\epsilon^2 t) dt \\ &\quad \left. - \frac{1}{2} \left(\nu_0 T - \log(1 + \nu_0 T) - \frac{\sigma_\epsilon^2 T^2}{2} \right) \right] - 1, \end{aligned} \quad (\text{A.87})$$

where λ^c solves (27).

Appendix B: Richer Specifications of Operational Risk

B.1 Mean-Reversion in Implemented Model Dynamics

In this section we consider our analysis in a richer framework with the implemented model featuring deterministic mean-reversion, and investigate the ensuing optimal model sophistication and risk exposures. As discussed in Remark 1, we consider $\zeta = \lambda$ which captures the case of a financial institution aiming at implementing a model that is given by the λ -weighted average of the most recent implemented model and the most sophisticated one available. For simplicity, consider the discrete time formulation,

$$\kappa_{t+dt}^* = (1 - \lambda)\kappa_t^* + \lambda\hat{\kappa}_{t+dt} + \lambda\sigma_\epsilon(w_{et+dt} - w_{et}). \quad (\text{B.1})$$

Subtracting κ_t^* from both sides and taking limit $dt \rightarrow 0$, we obtain the following stochastic differential equation:

$$d\kappa_t^* = -\lambda(\kappa_t^* - \hat{\kappa}_t)dt + \lambda d\hat{\kappa}_t + \lambda\sigma_\epsilon dw_{et} \quad (\text{B.2})$$

Proposition B.1 (Model sophistication with mean-reversion). *When the implemented model features a determining mean-reversion towards the most sophisticated model, the optimal model sophistication with external operational risk is given by $\lambda^* = 1$, while with internal operational risk it solves the following equation:*

$$\begin{aligned} &4(1 - \lambda)\lambda^2((1 - \lambda)(1 + \nu_0 T) + \nu_0) \left(E_i \left(2\lambda T + \frac{2\lambda}{\nu_0} \right) - E_i \left(\frac{2\lambda}{\nu} \right) \right) \\ &+ \nu_0 e^{\frac{2\lambda}{\nu_0}} \left(\nu_0 + \lambda \left(2\lambda^2(1 + \nu_0 T) - \lambda(\nu_0 + 4(1 + \nu_0 T) + \sigma_\epsilon^2 T) + 2(1 + \nu_0 T) \right) \right. \\ &\left. + e^{2\lambda T} \left(\lambda(\lambda(4 - 2\lambda + \nu_0 + \sigma_\epsilon^2 T) - 2) - \nu_0 \right) \right) = 0, \end{aligned} \quad (\text{B.3})$$

where $E_i(\cdot)$ denotes the exponential integral function, $E_i(x) = -\int_{-x}^{\infty} e^{-t}/t dt$.

Proof. Let θ_t denote the difference between the implemented and the most sophisticate models, $\theta_t(\lambda) \equiv \kappa_t^*(\lambda) - \hat{\kappa}_t$. Given the dynamics in (3) and (B.2), it follows that

$$d\theta_t(\lambda) = -\lambda\theta_t(\lambda)dt - (1 - \lambda)\hat{\nu}_t d\hat{w}_t + \lambda\sigma_\epsilon dw_{\epsilon t}. \quad (\text{B.4})$$

Let consider the function $g(\theta_t, t) = \theta_t(\lambda)e^{\lambda t}$. By Ito's lemma,

$$dg(\theta_t, t) = (1 - \lambda)e^{\lambda t}\hat{\nu}_t d\hat{w}_t + \lambda e^{\lambda t}\sigma_\epsilon dw_{\epsilon t}. \quad (\text{B.5})$$

Since $\theta_0 = 0$, integrating both sides of (B.5) between 0 and t , we obtain

$$\theta_t(\lambda) = (1 - \lambda) \int_0^t e^{-\lambda(t-s)} \hat{\nu}_s d\hat{w}_s + \lambda \int_0^t e^{-\lambda(t-s)} \sigma_\epsilon dw_{\epsilon t}. \quad (\text{B.6})$$

Note that, since $\hat{\nu}_0[\hat{\kappa}_t]$ does not depend on λ , the optimal model sophistication λ^* maximizing (A.6) also minimizes the following,

$$\lambda^* \in \arg \min_{\lambda} \int_0^T \hat{\text{var}}_0[\theta_t(\lambda)] dt, \quad (\text{B.7})$$

where

$$\hat{\text{var}}_0[\theta_t(\lambda)] = (1 - \lambda)^2 \int_0^t e^{-2\lambda(t-s)} \left(\frac{\nu_0}{1 + \nu_0 t} \right)^2 ds + \lambda^2 \sigma_\epsilon^2 \int_0^t e^{-2\lambda(t-s)} ds. \quad (\text{B.8})$$

Applying Fubini's theorem,

$$\int_0^T \int_0^t e^{-2\lambda(t-s)} \left(\frac{\nu_0}{1 + \nu_0 t} \right)^2 ds dt \quad (\text{B.9})$$

$$= \int_0^T \left(\frac{\nu_0}{1 + \nu_0 s} \right)^2 \int_s^T e^{-2\lambda(t-s)} dt ds \quad (\text{B.10})$$

$$= \int_0^T \left(\frac{\nu_0}{1 + \nu_0 s} \right)^2 \left(\frac{1 - e^{-2\lambda(T-s)}}{2\lambda} \right) ds \quad (\text{B.11})$$

$$= \frac{e^{-2\lambda(T+\frac{1}{\nu_0})}}{2\lambda} \left(e^{\frac{2\lambda}{\nu_0}} (e^{2T\lambda} - 1) \nu_0 - 2\lambda E_i \left(2\lambda T + \frac{2\lambda}{\nu_0} \right) + 2\lambda E_i \left(\frac{2\lambda}{\nu_0} \right) \right), \quad (\text{B.12})$$

and

$$\int_0^T \int_0^t e^{-2\lambda(t-s)} ds dt = \frac{e^{-2\lambda T} - 1 + 2\lambda T}{4\lambda^2}. \quad (\text{B.13})$$

Therefore, λ^* minimizes

$$\frac{(1 - \lambda)^2 e^{-2\lambda(T+\frac{1}{\nu_0})}}{2\lambda} \left(e^{\frac{2\lambda}{\nu_0}} (e^{2T\lambda} - 1) \nu_0 - 2\lambda E_i \left(\frac{2\lambda}{\nu_0} + 2\lambda T \right) + 2\lambda E_i \left(\frac{2\lambda}{\nu_0} \right) \right) + \lambda^2 \sigma_\epsilon^2 \left(\frac{e^{-2\lambda T} - 1 + 2\lambda T}{4\lambda^2} \right). \quad (\text{B.14})$$

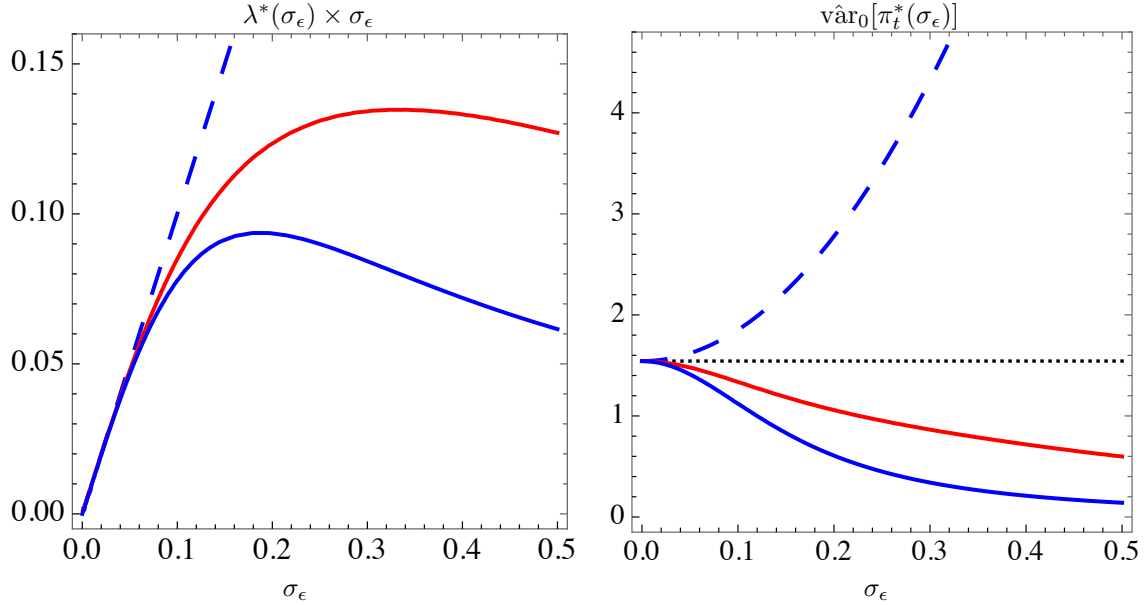


Figure B.1: Operational risk exposure and variability of market risk exposure with mean-reversion

In this figure we plot the optimal operational risk exposure $\lambda^* \sigma_\epsilon$ (left panel), and the variance at time 0 of the optimal market risk exposure π_t^* (right panel), as a function of operational risk σ_ϵ . The dotted line represents the benchmark case of no operational risk, $\sigma_\epsilon = 0$. The dashed line refers to the case of external operational risk, the solid blue line to the case of internal operational risk without mean-reversion, the solid red line to the case of internal operational risk with mean-reversion. The parameter values are as in Figure 3. The plots are typical.

Taking the first-order condition with respect to λ , and simplifying, we obtain (B.3). \square

Figure B.1 shows that our results on the optimal operational and market risk exposures, in Proposition 2 and Proposition 3 respectively, are valid for the case of mean-reversion in the implemented model dynamics. The red line in both panels refers to the case of internal operational risk with mean-reversion.

B.2 Non-Linear Operational Risk Trade-Off

In this section we generalize our framework to incorporate a non-linear operational risk trade-off function $h(\lambda)$ given in (5). The implemented model is given by

$$\kappa_t^* = (1 - \lambda)\kappa_0 + \lambda\hat{\kappa}_t + h(\lambda)\sigma_\epsilon w_{et}. \quad (\text{B.15})$$

To keep the maximization problem well defined, we maintain that the second derivative of $h(\lambda)$ is not too negative:

$$h''(\lambda) > -\frac{h'(\lambda)^2}{h(\lambda)} \quad (\text{B.16})$$

Proposition B.2 (Model sophistication with non-linear operational risk). *When the operational risk trade-off is captured by a non-linear function $h(\lambda)$, the optimal model sophistication solves the following equation:*

$$(1 - \lambda^*) (\nu_0 T - \log(1 + \nu_0 T)) - h(\lambda^*) h'(\lambda^*) \frac{\sigma_\epsilon^2 T^2}{2} = 0. \quad (\text{B.17})$$

Proof. Following the steps in Proposition 1, the covariance between the implemented and the most sophisticated models is as in (A.7), whereas the variance of the implemented model is now equal to

$$\text{var}_0[\kappa_t^*(\lambda)] = \lambda^2 \left(\frac{\nu_0^2 t}{1 + \nu_0 t} \right) + h(\lambda)^2 \sigma_\epsilon^2 t. \quad (\text{B.18})$$

Therefore,

$$\lambda^* \in \arg \max_{\lambda} \left(\lambda - \frac{\lambda^2}{2} \right) \int_0^T \frac{\nu_0^2 t}{1 + \nu_0 t} dt - \frac{h(\lambda)^2}{2} \int_0^T \sigma_\epsilon^2 t dt, \quad (\text{B.19})$$

yielding the first-order condition (B.17). It follows from condition (B.16) that the second-order condition of the optimization problem is satisfied. \square

We next present two examples of non-linear operational risk trade-off, which allow for analytic solutions of the optimal model sophistication λ^* . Example 1 considers a concave trade-off, Example 2 a convex one.

Example 1 (Concave operational risk trade-off). Let us consider the function form $h(\lambda) = \lambda^{\frac{3}{4}}$. Since $h''(\lambda) < 0$, the operational risk trade-off is concave in model sophistication. This implies that, starting from the most sophisticated model, a larger reduction in model sophistication (compared to the linear case) is needed to significantly reduce the exposure to operational risk. The optimal model sophistication can be solved explicitly and is given by

$$\lambda^* = 1 - \frac{3}{32} \left(\sqrt{64 + 9\sigma_\epsilon^4 X^2} - 3\sigma_\epsilon^2 X \right) \sigma_\epsilon^2 X, \quad (\text{B.20})$$

where

$$X \equiv \frac{T^2}{2(\nu_0 T - \log(1 + \nu_0 T))}. \quad (\text{B.21})$$

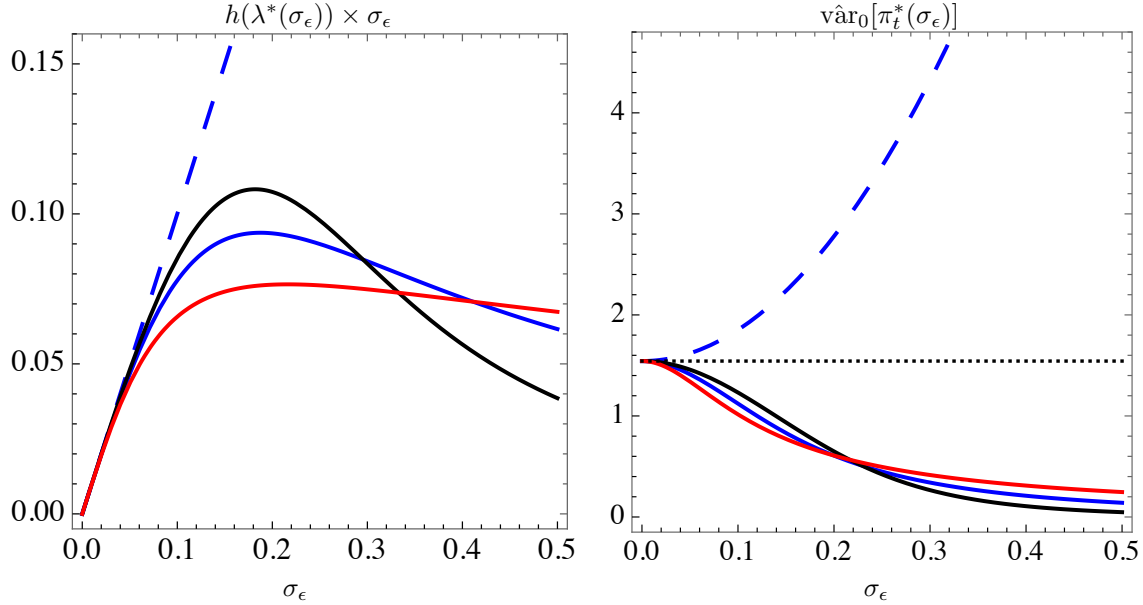


Figure B.2: Operational risk exposure and variability of market risk exposure with non-linear trade-off

In this figure we plot the optimal operational risk exposure $h(\lambda^*)\sigma_\epsilon$ (left panel), and the variance at time 0 of the optimal market risk exposure π_t^* (right panel), as a function of operational risk σ_ϵ . The dotted line represents the benchmark case of no operational risk, $\sigma_\epsilon = 0$. The dashed line refers to the case of external operational risk and the solid lines to the cases of internal operational risk with different operational risk trade-offs. The blue line corresponds to a linear trade-off ($h(\lambda) = \lambda$), the black line to a concave trade-off ($h(\lambda) = \lambda^{3/4}$), the red line to a convex trade-off ($h(\lambda) = \lambda^{3/2}$). The parameter values are as in Figure 3. The plots are typical.

Example 2 (Convex operational risk trade-off). Let us consider the function form $h(\lambda) = \lambda^{3/2}$. Since $h''(\lambda) > 0$, the operational risk trade-off is convex in model sophistication. This implies that, starting from the least sophisticated model, a larger increase in model sophistication (compared to the linear case) is needed to significantly increase the exposure to operational risk. The optimal model sophistication can be solved explicitly and is given by

$$\lambda^* = \frac{2}{1 + \sqrt{1 + 6\sigma_\epsilon^2 X}}, \quad (\text{B.22})$$

where X is as in (B.21).

Our results in Proposition 2 and Propostion 3, as well as the corresponding cross-sectional implications, remain equally valid, as shown in Figure B.2.

References

- Basel Committee on Banking Supervision, 2001, Operational Risk, Working paper Bank for International Settlement.
- Basel Committee on Banking Supervision, 2014, Review of the Principles for the Sound Management of Operational Risk, Working paper Bank for International Settlement.
- Biais, B., C. Casamatta, and J.C. Rochet, 2003, Operational Risk and Capital Requirements in the European Investment Fund Industry, Working paper 239 IDEI.
- Biais, B., J. Hombert, and P.-O. Weill, 2014, Equilibrium Pricing and Trading Volume under Preference Uncertainty, *Review of Economic Studies*, 1–37.
- Brown, S., W. Goetzmann, B. Liang, and C. Schwarz, 2008, Mandatory Disclosure and Operational Risk: Evidence from Hedge Fund Registration, *Journal of Finance* 63(6), 2785–2815.
- Brown, S., W. Goetzmann, B. Liang, and C. Schwarz, 2009, Estimating Operational Risk for Hedge Funds: The ω -Score, *Financial Analysts Journal* 5(1), 43–53.
- Brown, S., W. Goetzmann, B. Liang, and C. Schwarz, 2012, Trust and Delegation, *Journal of Financial Economics* 103(2), 221–234.
- Chavez-Demoulin, V., P. Embrechts, and J. Neslehova, 2006, Quantitative Models for Operational Risk: Extremes, Dependence and Aggregation, *Journal of Banking and Finance* 30, 2635–2658.
- Chernobai, A., P. Jorion, and F. Yu, 2011, The Determinants of Operational Risk in U.S. Financial Institutions, *Journal of Financial and Quantitative Analysis* 46(6), 1683–1725.
- Coleman, R., 2003, Op risk Modelling for Extremes, *Operational Risk* 4(1), 6–9.
- Crouhy, M., D. Galai, and R. Mark, 2014, The Essentials of Risk Management, McGraw-Hill, New York.
- Cvitanic, J., and I. Karatzas, 1992, Convex Duality in Constrained Portfolio Optimization, *Annals of Applied Probability* 2(4), 767–818.
- Cvitanic, J., A. Lazrak, L. Martellini, and F. Zapatero, 2006, Dynamic Portfolio Choice with Parameter Uncertainty and the Economic Value of Analysts Recommendations, *Review of Financial Studies* 19(4), 1113–1156.
- de Fontnouvelle, P., V. DeJesus-Rueff, J.S. Jordan, and E.S. Rosengren, 2006, Capital and Risk: New Evidence on Implications of Large Operational Losses, *Journal of Money, Credit, and Banking* 38, 1819–1846.
- de Fontnouvelle, P., E.S. Rosengren, and J.S. Jordan, 2006, Implications of Alternative Operational Risk Modeling Techniques, In *The Risk of Financial Institutions*, M. Carey, and R.M. Stulz (Eds.), University of Chicago Press, Chicago, 475–512.

- Dessein, W., 2002, Authority and Communication in Organizations, *Review of Economic Studies* 69, 811–838.
- Dessein, W., A. Galeotti, and T. Santos, 2014, Rational Inattention and Organizational Focus, Working paper Columbia Business School.
- Dessein, W., and T. Santos, 2006, Adaptive Organizations, *Journal of Political Economy* 114(5), 956–995.
- Detemple, J., 1986, Asset Pricing in a Production Economy with Incomplete Information, *Journal of Finance* 41(2), 383–391.
- Detemple, J., 1991, Further Results on Asset Pricing with Incomplete Information, *Journal of Economic Dynamics and Control* 15(3), 425–454.
- Duffie, D., A. Eckner, G. Horel, and L. Saita, 2009, Frailty Correlated Default, *Journal of Finance* 64, 2089–2123.
- Duffie, D., and K. J. Singleton, 1999, Modeling Term Structures of Defaultable Bonds, *Review of Financial Studies* 12, 687–720.
- Ebnother, S., P. Vanini, A. McNeil, and P. Antolinez-Fehr, 2001, Modelling Operational Risk, Working paper ETH Zurich.
- Franks, J.R., and C. Mayer, 2001, Risks and Regulation in European Asset Management: Is There a Role for Capital Requirements?, Research report OXERA.
- Genotte, G., 1986, Optimal Portfolio Choice Under Incomplete Information, *Journal of Finance* 41(3), 733–746.
- Hull, J.C., 2012, Risk Management in Financial Institutions, Wiley, New Jersey.
- Jarrow, R.A., 2008, Operational Risk, *Journal of Banking and Finance* 32, 870–879.
- Jarrow, R.A., and S.M. Turnbull, 1995, Pricing Derivatives on Financial Securities Subject to Credit Risk, *Journal of Finance* 50, 53–85.
- Jarrow, R.A., and F. Yu, 2001, Counterparty and the Pricing of Defaultable Securities, *Journal of Finance* 56, 1765–1799.
- Jorion, P., 2007, Value at Risk: The New Benchmark for Managing Financial Risk, McGraw-Hill, New York.
- JPMorgan Chase & Co., 2013, Report of JPMorgan Chase & Co. Management Task Force Regarding 2012 CIO Losses.
- Lando, D., 1998, On Cox Processes and Credit Risky Securities, *Review of Derivatives Research* 2, 99–120.
- Liptser, R., and A.N. Shiryaev, 2001, Statistics of Random Processes, Springer-Verlag, Berlin-Heidelberg.

- Moscadelli, M., 2004, The Modelling of Operational Risk: Experience with the Analysis of the Data Collected by the Basel Committee, Technical report 517 Banca d'Italia.
- Vayanos, D., 2003, The Decentralization of Information Processing in the Presence of Interactions, *Review of Economic Studies* 70, 667–695.
- Veronesi, P., 1999, Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model, *Review of Financial Studies* 12(5), 975–1007.
- Veronesi, P., 2000, How Does Information Quality Affect Stock Returns?, *Journal of Finance* 55(2), 807–837.
- Weiler, H., 1959, Means and Standard Deviation of a Truncated Normal Bivariate Distribution, *Australian Journal of Statistics* 1, 73–81.
- Xia, Y., 2001, Learning about Predictability: the Effect of Parameter Uncertainty on Dynamic Asset Allocation, *Journal of Finance* 56, 205–246.