

Financial Regulation and AI: A Faustian Bargain?

Christopher Clayton* Antonio Coppola†

July 2025

Abstract

We examine whether and how granular, real-time predictive models should be integrated into central banks’ macroprudential toolkit. First, we develop a tractable framework that formalizes the tradeoff regulators face when choosing between implementing models that forecast systemic risk accurately but have uncertain causal content and models with the opposite profile. We derive the regulator’s optimal policy in a setting in which private portfolios react endogenously to the regulator’s model choice and policy rule. We show that even purely predictive models can generate welfare gains for a regulator, and that predictive precision and knowledge of causal impacts of policy interventions are complementary. Second, we introduce a deep learning architecture tailored to financial holdings data—a graph transformer—and we discuss why it is optimally suited to this problem. The model learns vector embedding representations for both assets and investors by explicitly modeling the relational structure of holdings, and it attains state-of-the-art predictive accuracy in out-of-sample forecasting tasks including trade prediction.

Keywords: Fire Sales, Macroprudential Policy, Artificial Intelligence, Graph Neural Networks, Holdings Data, Embeddings, Stress Testing, Deep Learning, Information Design.

JEL Codes: C4, G1.

*Yale School of Management and NBER; christopher.clayton@yale.edu.

†Stanford University Graduate School of Business and CEPR; acoppola@stanford.edu.

We thank Paul Fontanier, Stefano Giglio, Tarek Hassan, Arvind Krishnamurthy, Matteo Maggiori, Pablo Ottonello, and Jesse Schreger for helpful comments. Financial support from the Stanford GSB Business, Government, and Society (BGS) Initiative is gratefully acknowledged.

1 Introduction

A central concern of macroprudential policy is identifying and mitigating the amplification of shocks through the financial system. Traditional macroprudential frameworks focus on ex ante regulation of well-established sources of financial fragility such as leverage accumulation or maturity mismatches. Yet over the past two decades, the data environment for regulators has changed dramatically. Supervisory filings and rapid data collection now give regulators the ability to observe granular, high-frequency information on investor portfolios. At the same time, predictive technologies such as deep learning have achieved significant gains in out-of-sample performance across domains. These advances raise the question: can real-time, high-dimensional predictive models be used productively in financial regulation and interventions? On the one hand, predictive models can help effectively detect signals of fragility that are not captured by canonical macroprudential metrics: where exactly fire sales may erupt, how crowded trades may unwind, or which asset classes are most exposed to redemption risk. On the other hand, these models are often reduced-form and highly non-linear, with no guarantee that they might recover deep underlying structural forces that are invariant to the regulators' own use of the models.

This paper develops a theoretical and empirical framework to address this question. Our central object of analysis is the regulator's model choice and how it informs intervention. We ask whether regulators should deploy high-performing predictive models when they face uncertainty about the causal consequences of acting on the model's output. In this context, we show that the answer depends on the interaction between the predictive model's informational content as well as forecast accuracy, and the regulator's ability to target interventions and estimates of causal impact of those interventions. Empirically, we introduce and build new predictive architectures using deep learning models tailored to the relational structure of financial holdings data, and we show that these can achieve benchmark-setting predictive performance in the dimensions relevant to the macroprudential task, providing a blueprint for practical implementation by central banks and other regulators.

In our theory, we model a three-period economy with intermediaries and a government regulator. The model builds upon canonical fire sales environments. In the first period, financial intermediaries choose portfolios. In the second, intermediaries face constraints (e.g., pledgeability and collateral constraints) that force them to sell some assets prior to maturity. The key fire sale externalities arise because these assets are sold to second-best users whose productivity depends on the amount purchased. Finally, any assets held to maturity pay off in the third period.

The regulator can intervene in the second period to try to manage fire sales. We allow the regulator to employ a (potentially incomplete) set of liquidation wedges that affect intermediaries' choice of which assets to liquidate. Formally, liquidation wedges take the form of revenue-neutral taxes on selling an asset, although we later show that results extend to applying subsidies for retaining an asset. The key informational friction is that the regulator faces uncertainty over the mapping from intermediaries' asset positions and the policy intervention to realized liquidations and fire sale prices. The regulator has a prior over this mapping. Before undertaking the intervention,

the regulator can choose to deploy a model that delivers a signal about the latent fundamentals (e.g., predicted sales volumes or liquidation discounts). The regulator can then design its intervention based on the model used and the signal produced. Crucially, models can differ in what aspects of the system they inform. A model may provide strong predictive content (about liquidations and prices) but little information about the causal effects of interventions, and vice versa.

We characterize the regulator’s optimal ex post policy rule as a function of its Bayesian posterior over key primitives. The regulator’s optimal rule depends on the product between the predicted causal impact of the policy and the predicted social benefit of the policy. The social benefit of the policy, deriving from raising the fire sale price, is determined by the amount of each type of asset being sold. As a consequence, even a purely predictive model can improve welfare if its forecasts are aligned with dimensions of the system where the regulator has prior knowledge about the causal impact of the policy intervention. For example, if the regulator knows that its intervention will have a strong causal effect on fire sale prices in certain markets, then predictive models that are informative precisely about forced sales in those markets have particular potential to generate welfare gains. This suggests a complementarity between causal knowledge of policy interventions and predictive models that inform the social benefit of intervention.

We then characterize the expected welfare gains associated with the choice of a model by the regulator as well as optimal model choice. The welfare gains depend on the prior expectation of the size of the intervention, and also on the covariance matrix of the policy intervention (assessed from the prior perspective). For purely predictive models that are uninformative as to the causal structure, only this latter term depends on the choice of model. Hence choosing an optimal predictive model amounts to a choice of this covariance matrix. The intuition comes from the law of total variance. The regulator benefits from acquiring precise information on ex-post liquidations to better target its policies. However, not all predictive information is equally valuable. First, predictive precision is more valuable on margins where the regulator has prior knowledge that the policy intervention will have a strong causal impact, reinforcing the complementarity between causal knowledge and the value of predictive information. Second, the value of interventions (and hence of predictive precision) scales with the intervention’s causal impact and with intermediaries’ ex-post costs of portfolio adjustment. The regulator’s optimal model choice thus tends to focus predictive precision on those margins with stronger predicted causal impacts and higher benefits of intervention.

To understand the implications of model deployment on intermediaries’ decision making, we turn to the endogenous ex-ante portfolio response of private agents that anticipate the regulatory intervention. Intermediaries internalize the mapping from their portfolios to expected liquidation costs under the regulator’s anticipated policy, but take as given both what model the regulator will adopt and the intervention and fire sale prices. As a result, the regulator’s model choice affects portfolio selection. We show that the regulator’s model choice and intervention can both discourage intermediaries from holding assets associated with fire sales ex post, but also potentially lead to moral hazard. Intuitively, the ex post intervention has two competing effects on the initial asset

allocation. First, intermediaries directly perceive assets on which the regulator will impose liquidation wedges as more expensive to hold, and so shift portfolios away from these assets. This can create a virtuous effect of the ex-post intervention, potentially substituting for the need to regulate these assets ex ante. However, the countervailing effect is that by using the more informed intervention to raise liquidation prices ex post, the regulator reduces fire sale discounts and so encourages intermediaries to hold more of these assets. This latter effect can be particularly pronounced if the regulator is not able to regulate all holders of certain assets. All of these effects, both virtuous and moral hazard, are limited by the predictability and granularity of the model—that is, the precision with which the regulator can target fire sales in real time.

Further, we explicitly model a regulator’s optimal ex-ante macroprudential intervention in the portfolio holdings of intermediaries. We allow the regulator to employ a set of wedges on asset holdings, modeled as revenue-neutral taxes. The regulator’s optimal tax ex ante starts from the baseline regulation it would apply if it did not intervene ex post, and then makes three adjustments reflecting model choice and the ex post intervention. First, the regulator discourages holdings of assets that make it more costly to acquire precise predictive information ex post. Second, the regulator encourages relative holdings of assets that will cause it to increase the size of its ex-post intervention. Intuitively, the regulator recognizes for these assets that the larger ex-post regulation serves as a substitute for the ex-ante intervention, so the need to intervene ex ante is muted. Finally, the regulator reduces regulation of assets that are expected to be subject to high ex post taxes, since the ex post taxes help discourage ex ante purchases of those assets. Conversely, the regulator increases regulation of assets for which moral hazard is induced by ex post bolstering of asset prices.

Although we illustrate our key insights for policy interventions in the form of liquidation wedges, in practice ex post interventions often resemble “bailout” measures with a subsidy component. We accommodate this policy in an extension in which we assume the regulator ex post imposes an asset retention wedge, formally modeled as a revenue-neutral subsidy for holding an asset to maturity. These subsidies imply a de facto tax on liquidating an asset, and accordingly the optimal intervention and model design remain the same. However, these subsidies introduce an additional moral hazard channel, since assets that are subsidized ex post become more attractive to purchase ex ante. Because this moral hazard effect depends on the expected subsidy size, models that increase predictive precision but do not change the expected tax rate are not subject to it (akin to [Laffont and Tirole 1986](#)). Since in our setup purely predictive models increase precision without changing the expected intervention size, the moral hazard they generate interestingly does not differ relative to the case with liquidation taxes instead of asset retention subsidies.

To empirically assess whether predictive models can in fact deliver useful signals for these macroprudential purposes, we develop a deep learning architecture tailored to the structure of financial holdings data. Much of the modern deep learning toolkit—including models that have been applied in asset pricing contexts—is tailored toward grid and sequence inputs such as images and text. Yet in domains where the data has important graph structure, architectures tailored specifically to

learn from graph structures have achieved breakthrough performance, including for instance when applied to problems of protein folding prediction and drug discovery (Jumper et al. 2021). In our setting, the data naturally form a graph: investors are connected to assets through their positions. We therefore design our empirical deep learning models to explicitly capture and make use of this crucial dimension of the data.

To exploit the relational structure of holdings data, we implement and train a *graph transformer* model, a form of graph neural network (GNN) augmented with an attention mechanism, to learn latent representations of both investors and assets (i.e., investor and asset embeddings into latent vector spaces). The architecture features two properties that are essential for this setting. First, it is permutation-invariant: predictions do not depend on the arbitrary ordering of investors or assets. This is a key distinction with sequence-based architectures (such as sequence transformers which underpin modern large language models) commonly used for text, in which the order of words in a document is essential to their contextual meaning. Second, the graph transformer model is inductive: all model parameters are fully shared across nodes, allowing the model to generalize to new investors or assets without retraining, while at the same time enforcing a strong form of regularization that leads to good generalization out of the model’s training sample. The model is optimally sample-efficient in the sense of requiring minimal degrees of freedom to learn functions of the holdings data that are themselves invariant to arbitrary relabelings of investors and assets—a restriction that reflects the economics of the problem. Intuitively, this occurs because graph-based architectures do not need to use parameters to relearn permutation invariance from the data, but rather they enforce it explicitly through their structure.

We train the model jointly on two tasks: a masked autoencoder task, in which the model learns to reconstruct partially masked holdings, and a supervised prediction task, in which the model forecasts the cross-sectional pattern of future trades. Both tasks are designed to inform the model about latent economic relationships governing portfolio choice and rebalancing. The training generates embeddings for both assets and investors which are general-purpose and can be used for multiple objectives. We do not train the model explicitly to predict fire sales in a supervised fashion, but rather turn to models which can learn general-purpose embeddings, precisely to avoid in-sample over-fitting and thus avoid a sharp degradation in out-of-sample performance. The training uses quarterly holdings data from Factset, which cover a broad range of institutional investors and assets.

We find that the model performs well on both tasks. It accurately reconstructs positions data with correlations exceeding 90% between predictions and targets, and it achieves strong out-of-sample predictive accuracy of nearly 30% in forecasting trade patterns. Notably, the model’s performance is indeed stable out of sample, due to its inductive structure and shared-parameter design. The high holdings reconstruction fidelity on the autoencoder task should naturally be interpreted in light of the model’s parameter-to-data ratio: with roughly 3.6 million parameters—representing less than 1% of the possible asset-investor pairs in the holdings data—the model extracts economically relevant patterns from high-dimensional data. Further, we show that the model’s performance in

forecasting trading behavior (our second task) remains high even when restricting the sample to the set of active investment managers only and to stress periods, both of which are especially relevant to the macroprudential application we are studying. For example, we show that the model is able to forecast the pattern of asset trades during the market crash of 2020 induced by the Covid pandemic with good accuracy, having been trained only on pre-2020 data.

The predictive success of our deep learning architecture provides empirical support for the potential of model-informed ex post regulation. Moreover, we view our model as a blueprint for real-time regulatory approaches: a framework to transform high-frequency, granular data into relatively low-dimensional representations suitable for intervention design. Although the model is not structurally interpretable in the traditional sense, our theory shows that it can nonetheless play a useful role when paired with prior knowledge of structural effects.

Related Literature. At their heart, our questions are situated within a longstanding intellectual debate concerning whether models that are primarily predictive in nature should be used to inform macroeconomic policy and financial regulation. [Koopmans \(1947\)](#), representative of the viewpoints then prevalent at the Cowles Commission, introduced his landmark “measurement without theory” critique of the earlier empirical studies of business cycles by [Burns and Mitchell \(1946\)](#). These early debates set the stage for the methodological ideas of [Haavelmo \(1944\)](#) and [Friedman \(1953\)](#), and for [Lucas’s \(1976\)](#) eventual critique of policy evaluation carried out without individual-optimization microfoundations. While the qualitative outline of this debate has remained the same, its quantitative content has evolved, as the performance and generalization ability of predictive models has increased sharply. In this context, we show that sufficiently capable predictive models can have a role in the macro-regulatory toolbox, as a complement (rather than a substitute) to traditional structural approaches.

Our analysis relates to four broad areas of the literature. First, we connect to the literature on fire sales, and macroprudential regulation and ex-post interventions. Theoretical contributions include [Bernanke and Gertler \(1986\)](#), [Williamson \(1988\)](#), [Kiyotaki and Moore \(1997\)](#), [Caballero and Krishnamurthy \(2001\)](#), [Lorenzoni \(2008\)](#), [Bianchi \(2011, 2016\)](#), [Stein \(2012\)](#), [Farhi and Werning \(2016\)](#), [Chari and Kehoe \(2016\)](#), [Bianchi and Mendoza \(2018\)](#), [Dávila and Korinek \(2017\)](#), and [Clayton and Schaab \(2022, 2025\)](#). On the empirical side, prior work has focused on particular variables to explain and predict the occurrence and magnitude of fire sales, such as leverage ([Fisher 1933](#); [Kindleberger and Aliber 1978](#); [Brunnermeier and Oehmke 2013](#); [Schularick and Taylor 2012](#); [Adrian and Shin 2014](#); [Krishnamurthy and Muir 2025](#)) and investor composition ([Brainard and Tobin 1968](#); [Coval and Stafford 2007](#); [Haddad, Moreira and Muir 2021](#); [Coppola 2025](#); [Fang, Hardy and Lewis 2025](#)): the empirical contribution of this paper asks whether there is additional value from a more agnostic approach that does not impose an ex-ante focus on particular dimensions of the data.

Second, we connect to the literature on applications of machine learning and deep learning to

finance. Gu, Kelly and Xiu (2020, 2021), Kozak, Nagel and Santosh (2020), Nagel (2021), and Bryzgalova, DeMiguel, Li and Pelger (2023) apply machine learning techniques to the canonical empirical asset pricing problem of valuation in the cross-section of assets. Similarly, Chen, Pelger and Zhu (2024) introduce a deep learning architecture for modeling stock returns. Giglio, Kelly and Xiu (2022) review the intersection of empirical asset pricing and machine learning. Gabaix, Koijen, Richmond and Yogo (2024) use holdings data together with a variety of deep learning sequence models (e.g., Word2Vec and BERT) and recommender systems to estimate asset embeddings, and Gabaix, Koijen, Richmond and Yogo (2025) apply these to explain variation and volatility in credit spreads. Dolphin, Smyth and Dong (2022) and Sarkar (2025) also construct asset embeddings, respectively from return data and textual sources. Methodologically, our paper employs graph neural network architectures (Scarselli, Gori, Tsoi, Hagenbuchner and Monfardini 2008; Hamilton, Ying and Leskovec 2017; Xu, Hu, Leskovec and Jegelka 2018; Wu, Pan, Chen, Long, Zhang and Yu 2020) to learn representations from the relational structure of asset holdings data.¹

Third, we relate to literature studying the deployment of machine learning and large-scale data for economic analysis and policy problems more broadly, including Einav and Levin (2014b,a), Kleinberg, Ludwig, Mullainathan and Obermeyer (2015), Athey (2017, 2018), Mullainathan and Spiess (2017), Kleinberg, Lakkaraju, Leskovec, Ludwig and Mullainathan (2018), Gillis and Spiess (2019), Farboodi and Veldkamp (2020), and Athey and Wager (2021). The question of how non-policy invariant relationships should be exploited by regulators also has earlier conceptual parallels, for instance by Barro and Gordon (1983) in the context of the Phillips curve. Lastly, the model and information design aspect of our theoretical approach connects to work on stress testing (Shapiro and Skeie 2015; Faria-e Castro, Martinez and Philippon 2017; Goldstein and Leitner 2018; Leitner and Williams 2023; Orlov, Zryumov and Skrzypacz 2023; Parlatore and Philippon 2025).

2 A Framework for Regulatory Model Choice

There are N assets. There are I intermediary types (each of equal measure), with a representative intermediary of each type i . There is also a representative arbitrageur. The model has a Beginning-Middle-End structure. In the Beginning, initial asset positions are undertaken. In the Middle, intermediaries may be forced to sell assets prior to maturity to arbitrageurs, who are second best users. In the End, payoffs are distributed and consumption occurs.

Intermediaries. Intermediaries are risk neutral. In the Beginning, intermediary i invests in a vector $q_i = (q_{i1}, \dots, q_{iN})^T$ of assets. If $q_{in} < 0$, then intermediary i has undertaken a negative investment (shorting) asset n .² The cost to intermediary i of producing the asset vector q_i is

¹See also Elliott, Golub and Jackson (2014) and Acemoglu, Ozdaglar and Tahbaz-Salehi (2015) for theory highlighting the importance of the network structure of positions for financial stability.

²We could endow intermediaries with a stock of assets, but for expositional convenience we set that stock to 0.

$C_i(q_i) = p_i^T q_i - \frac{1}{2} q_i^T H_i^q q_i$, where p_{in} is the per-unit cost with $p_i = (p_{i1}, \dots, p_{iN})^T$. The cost component $q_i^T H_i^q q_i$ is a quadratic adjustment/holding cost, where H_i^q is an $N \times N$ matrix. If held to maturity, intermediary i 's holdings of asset n will produce a per-unit return R_{in} in the End, with $R_i = (R_{i1}, \dots, R_{iN})^T$.

In the Middle, intermediary i can sell assets prior to maturity. We denote $\ell_{in} \geq 0$ to be sales by intermediary i at endogenous price γ_n , with $\ell_i = (\ell_{i1}, \dots, \ell_{iN})^T$ and $\gamma = (\gamma_1, \dots, \gamma_N)^T$. Intermediary i faces an adjustment cost $\frac{1}{2} \ell_i^T H_i^\ell \ell_i$ on asset sales, where H_i^ℓ is $N \times N$. Assets will be sold at a discount on their fundamental value ($\gamma < R_i$, see below), and intermediary i faces a set of “rollover constraints” that force asset sales. This set of M constraints is given by

$$A_i^q q_i + \rho_i \leq A_i^\ell \ell_i, \quad (1)$$

where A_i^q, A_i^ℓ are $M \times N$ and ρ_i is $M \times 1$. For example, equation 1 can capture constraints on asset positions (e.g., a limit on debt) or a requirement to raise funds based on asset holdings (e.g., a cost of maintaining the project).³

In the End, intermediary i realizes payoff on assets held to maturity and consumes. Intermediary i 's total payoff (consumption) in the End, inclusive of adjustment costs, is

$$U_i = q_i^T (R_i - p_i) - \ell_i^T (R_i - \gamma) - \frac{1}{2} q_i^T H_i^q q_i - \frac{1}{2} \ell_i^T H_i^\ell \ell_i. \quad (2)$$

Arbitrageurs. A representative arbitrageur is a second-best user of intermediary assets. If the arbitrageur purchases a vector $L = (L_1, \dots, L_N)^T$ of intermediary assets in the Middle, the arbitrageur can use them in a production technology to produce $\mathcal{F}(L) = L^T \bar{\gamma} - \frac{1}{2} L^T \Gamma L$ units of the consumption good in the End, where $\bar{\gamma}$ is $N \times 1$ and Γ is $N \times N$. The representative arbitrageur's payoff is

$$U_A = L^T (\bar{\gamma} - \gamma) - \frac{1}{2} L^T \Gamma L. \quad (3)$$

Information Structure and Timing. Although all model parameters are determined in the Beginning, all agents in the Beginning are uncertain about the true values of the parameters $\Phi = \{A_i^q, A_i^\ell, H_i^\ell, \rho_i, R_i, \bar{\gamma}\}$. That is, agents are uncertain as to the true parameters underlying the rollover constraint (equation 1), the asset return R_i , the arbitrageurs' baseline productivity $\bar{\gamma}$, and the liquidation adjustment cost H_i^ℓ . All agents have a common prior $\Phi \sim \mu_0$ in the Beginning. In the Middle before asset sales and purchases are chosen, the true model parameters become common knowledge of private agents. After becoming common knowledge, intermediaries choose

³We simplify analysis by not having equation 1 depend on the liquidation price γ , as for example in price-dependent collateral constraints. This enables us to maintain a linear-quadratic structure throughout the paper. It is straightforward to extend analysis to include prices in constraints, but the characterization of the ex-ante optimal portfolio would no longer admit a closed-form solution.

asset liquidations and arbitrageurs choose asset purchases.⁴

Market Clearing. Markets must clear for liquidations in the Middle, that is

$$L = \sum_i \ell_i. \quad (4)$$

Competitive Equilibrium. We define a competitive equilibrium as follows.

Definition 1. A competitive equilibrium of the model, given true model parameters Φ , is a vector of prices γ and a set of allocations $\{q_i, \ell_i, L\}$ such that:

1. In the Middle, taking as given asset allocations $\{q_i\}$ and model parameters Φ :
 - (a) Intermediary i chooses ℓ_i to maximize utility (equation 2) subject to the rollover constraint (1), taking prices γ as given.
 - (b) Arbitrageurs choose L maximize utility (equation 3), taking prices γ as given.
 - (c) The liquidation markets clear (equation 4).
2. In the Beginning:
 - (a) Intermediary i chooses q_i to maximize expected utility $\mathbb{E}_0[U_i]$, where $\mathbb{E}_0$ denotes the expectation given the prior μ_0 over model parameters Φ .

2.1 Regulator's Model Design and Intervention in the Middle

In the Middle, a regulator is able to intervene in order to try to manage the fire sale price impact of liquidations. Formally, the regulator can impose a vector $\tau_i = (\tau_{i1}, \dots, \tau_{iN})^T$ of revenue-neutral liquidation wedges on each intermediary i , which alter the intermediary's perceived price for selling the asset.⁵ τ_{in} represents a tax on selling asset n , with $\tau_{in} < 0$ being a subsidy for sale. As a result, intermediary i 's payoff in the End is modified to be

$$U_i - (\ell_i^T - \ell_i^{*T})\tau_i \quad (5)$$

⁴Assuming that private agents and the regulator (see below) have a common prior in the Beginning simplifies analysis because it prevents the regulator from learning information about these parameters from inference about private sector beliefs based on the asset allocation. It also eliminates a regulatory incentive based on different beliefs the regulator and agents (e.g., [Fontanier 2025](#).)

Assuming Γ is known to all agents ex ante simplifies analysis by eliminating the ability of a regulator to learn about the price impact of liquidations from observing a signal of γ or ℓ . This will allow us to fully separate predictive and causal channels in our framework. Absent this assumption, predictive models that inform a regulator about γ and ℓ might have even more value ex post because they would also allow the regulator to learn about these structural parameters that determine the causal impact of a policy intervention (even if very little information is learned).

⁵In Section 2.7, we instead assume the regulator must use asset holding subsidies rather than liquidation taxes. The results in the Middle are identical, but the asset holding subsidies introduce additional moral hazard in the Beginning.

where $\ell_i^{*T} \tau_i$ is equilibrium revenue remissions based on the equilibrium asset liquidations ℓ_i^* of intermediary i . Intermediaries take revenue remissions as given. We allow for the possibility that the regulator has potentially incomplete instruments, represented by a regulatory cost $\delta(\tau) = \frac{1}{2} \tau^T \Delta \tau$, where Δ is $NI \times NI$.⁶

Unlike private agents, the regulator does not learn the true model parameters Φ . Instead, the regulator has to use a *model* to make an inference about the true parameters. The regulator's *model* $M \in \mathcal{M}$ formally is a process of drawing a signal s of the parameters. The signal updates the regulator's posterior distribution to $\Phi \sim \mu_p | s, M$. The regulator must choose the wedges τ before private agents move, and so the regulator cannot obtain any new information from observing the market before setting regulation (apart from the signal drawn).

To simplify exposition, we introduce an assumption of matrix symmetry that we maintain throughout the paper.

Assumption 1. *The matrices $H_i^q, H_i^\ell, \Gamma, \Delta$ are symmetric.*

Impact of Regulatory Intervention. To solve the regulator's optimum, we begin by characterizing the impact of the regulator's wedges $\tau = (\tau_1^T, \dots, \tau_I^T)^T$ on the equilibrium in the Middle. Our model is substantially simplified by the information structure: because private agents directly observe Φ in the Middle, for a given vector of asset allocations $q = (q_1^T, \dots, q_I^T)^T$ the equilibrium in the Middle does not depend on the regulator's choice of model M or the realized signal s except through the choice of wedges τ .

The linear-quadratic structure of preferences and the rollover constraint leads to a characterization of the equilibrium as a simple linear system of equations. The following Lemma characterizes this equilibrium.

Lemma 1. *Given asset allocations q , true parameters Φ , and regulatory wedges τ , the equilibrium in the Middle is given by*

$$\ell_i = \bar{\ell}_i + \Lambda_i^q q_i - \Lambda_i^\tau \tau_i - \sum_{j=1}^N \left[\Lambda_j^{q,e} q_j^* - \Lambda_j^{\tau,e} \tau_j^* \right] \quad (6)$$

$$\gamma = \bar{\gamma} - \Gamma \sum_i \ell_i \quad (7)$$

where $\bar{\ell}_i$ ($N \times 1$) and $\Lambda_i^q, \Lambda_i^\tau, \Lambda_i^{q,e}, \Lambda_i^{\tau,e}$ ($N \times N$) are defined in the proof.

The characterization of the equilibrium in Lemma 1 is intuitive. Liquidations start from a benchmark $\bar{\ell}_i$ and increase (for positive Λ_i^q) in asset holdings while decreasing (for positive Λ_i^τ) in the liquidation wedge. Higher liquidations result in a lower liquidation price (for positive Γ) due

⁶One could instead model incompleteness as a restriction $\Delta(\tau) \leq 0$ in which case the regulator's Lagrangian will be $\mathbb{E}[\sum_i U_i | \Gamma, M] - \lambda \Delta(\tau)$ where λ is the Lagrange multiplier. This is akin to the reduced-form cost $\delta(\tau) = \lambda \Delta(\tau)$ except that λ is endogenous.

to more assets having to be absorbed by arbitrageurs (equation 7). These means that liquidations by intermediary i decrease in asset holdings and increase in the liquidation wedges applied to other intermediaries within the same sector and across different sectors (for positive matrix elements), a result of a substitution effect resulting from the increase in equilibrium price. In a model without fire sales, we have $\Lambda_j^{q,e} = \Lambda_j^{\tau,e} = 0$. We distinguish in equation 6 between the asset choices of intermediary i and the wedges applied to intermediary i 's liquidations (denoted with no $*$) as opposed to equilibrium objects (denoted with $*$) that enter because they determine the equilibrium liquidation price.

2.2 Regulator's Optimal Wedges

We solve the regulator's problem by backward induction: first, we characterize the regulator's optimal intervention τ given a choice of model M and signal s . We then characterize the regulator's optimal choice of model M . To simplify analysis, we assume that the regulator places equal weight on all intermediaries but places a welfare weight of zero on arbitrageurs.⁷

Given that liquidation wedges are revenue-neutral, the regulator's optimal choice of τ solves

$$\max_{\tau} \mathbb{E}[\sum_i U_i | s, M] - \delta(\tau),$$

subject to equilibrium determination (Lemma 1).

As preliminaries to the proposition below, we define $\bar{\ell}_i(q) = \bar{\ell}_i + \Lambda_i^q q_i - \sum_i \Lambda_i^{q,e} q_i$ to be the liquidations of intermediary i if there is no regulatory intervention ($\tau = 0$). We define $L(q) = \sum_i \bar{\ell}_i(q)$ to be total liquidations if there is no intervention. The following proposition characterizes optimal liquidation wedges in terms of these objects and model parameters.

Proposition 1. *Given a model M and signal s , the regulator's optimal policy in the Middle is*

$$\tau^* = \mathbb{E} \left[\Xi \middle| s, M \right]^{-1} \mathbb{E} \left[\left(\sum_i \bar{\Lambda}_i^{\tau} \right)^T \Gamma L(q) \middle| s, M \right] \quad (8)$$

where $\Xi = \bar{\Lambda}^{\tau T} + (\sum_i \bar{\Lambda}_i^{\tau})^T \Gamma (\sum_i \bar{\Lambda}_i^{\tau}) + \Delta$, where $\bar{\Lambda}_i^{\tau} = (e_i \otimes I_N)^T \Lambda_i^{\tau} - (\Lambda_1^{\tau,e}, \dots, \Lambda_I^{\tau,e})$ is $N \times NI$ (where e_i is the standard basis vector whose i^{th} element is 1 and $\otimes$ is the Kronecker product), and where $\bar{\Lambda}^{\tau} = (\bar{\Lambda}_1^{\tau T}, \dots, \bar{\Lambda}_I^{\tau T})^T$ is $NI \times NI$.

The optimal wedges of Proposition 1 are familiar from the macroprudential policy literature on fire sales, and intuitively encode an expected marginal cost-marginal benefit trade-off. The marginal

⁷This focuses attention on a distributive externality from shifting wealth between arbitrageurs and intermediaries. This can be incorporated by assuming that arbitrageurs have a high marginal value of wealth in the Beginning but cannot borrow, meaning that Pareto improvements are achieved by raising the liquidation price in the Middle and having a lump sum transfer from intermediaries to arbitrageurs in the Beginning (see e.g., Clayton and Schaab 2025).

cost of the policy is captured in the (conditional expectation of the) inverse matrix Ξ , while the marginal benefit is captured in the expectation.

There are three components to the private and regulatory marginal cost of liquidations, reflecting the distortion of intermediaries' activities away from their private optimum. First, increasing liquidation wedges directly distorts the intermediaries' activities (the first term of Ξ). Second, by using wedges to alter equilibrium prices, the regulator changes the incentives for intermediaries to sell different assets (the second term). Finally, there is the regulatory cost Δ .

The social marginal benefit of regulation arises from the mitigation of the fire sale. This has two components that are central to our analysis. First is the causal effect of the policy intervention on equilibrium liquidation prices, captured by the term $(\sum_i \bar{\Lambda}_i^\tau)^T \Gamma$. The utility consequence of these price changes is proportional to the total value of assets being sold by intermediaries, $L(q)$.

Importantly, this means the social benefit of regulation derives from a productive of the causal impact of the policy intervention, $(\sum_i \bar{\Lambda}_i^\tau)^T \Gamma^T$, and the social value of changing the price, here given by $L(q)$. This means that in designing regulation, there is value not only to identifying causal policy impacts, but also to predicting the magnitude of liquidations $L(q)$. To fully disentangle these two mechanisms, we can assume independence of the predictive and causal components of the system.⁸

Definition 2 (Predictive-Causal Independence). *We have predictive-causal independence if $\{\Lambda_i^\tau\}$ and $\{\bar{\ell}_i, \Lambda_i^q, \Lambda_i^{q,e}\}$ are independent of one another.*

Under predictive-causal independence, we obtain the following trivial corollary to Proposition 1

Corollary 1. *Under predictive-causal independence, the regulator's optimal wedges are*

$$\tau^* = \mathbb{E} \left[\Xi \middle| s, M \right]^{-1} \underbrace{\mathbb{E} \left[\left(\sum_i \bar{\Lambda}_i^\tau \right)^T \Gamma \middle| s, M \right]}_{\text{Causal: Policy Impact on Liquidation Price}} \underbrace{\mathbb{E} \left[L(q) \middle| s, M \right]}_{\text{Predictive: Total Liquidations}} \quad (9)$$

An important implication of Proposition 1 and Corollary 1 is that a model that is purely predictive – that is, that can predict liquidations $L(q)$ – can be valuable to a regulator, provided that the regulator has some prior or posterior knowledge over the causal impact of the policy. The predictive information informs the regulator as to the magnitude and margins for intervention by informing the regulator about the social benefit of the intervention. Corollary 1 implies that, holding fixed the causal impact of the policy, the regulator would design larger-magnitude interventions when the regulator's model predicted that total liquidations would be larger.

⁸Naturally absent this assumption, we could perform an analogous decomposition to Corollary 1, but also include the covariance between the two terms.

2.3 Expected Welfare from a Model and Optimal Model Choice

We next ask how a regulator in the Middle would choose a model $M \in \mathcal{M}$ under discretion, taking as given the asset holdings q of intermediaries. Because choice of model in turn impacts intermediaries' optimal portfolio allocations (see Section 2.5), we later consider model choice under commitment.

We begin by characterizing the expected utility to the regulator from a choice M of model, and then characterize optimal model choice. As a preliminary to the proposition below, we define $\theta_i(q) = R_i - (\bar{\gamma} - \Gamma L(q))$ to be the fire sale losses to intermediary i from selling assets prior to maturity in the event of no regulatory intervention. For the presentation in the main text, we focus on the case assuming predictive-causal independence (Definition 2). The proof of Proposition 2 characterizes the general case.

Proposition 2. *Under predictive-causal independence, the regulator's ex-ante expected welfare given positions q and a model M is*

$$V(q, M) = \mathbb{E} \left[q^T (R - p) - \frac{1}{2} q^T H^q q - \ell(q)^T \theta(q) - \frac{1}{2} \ell(q)^T H^\ell \ell(q) \middle| s, M \right] \quad (10)$$

$$+ \frac{1}{2} \mathbb{E}_0[\tau^{*T}] \Psi_0 \mathbb{E}_0[\tau^*] + \frac{1}{2} \text{tr} \left(\Psi_0 \text{cov}_0(\tau^*) \right)$$

where $\Psi_0 = \mathbb{E}_0 \left[2(\sum_i \bar{\Lambda}_i^T)^T \Gamma (\sum_i \bar{\Lambda}_i^T) + \Delta + \bar{\Lambda}^T H^\ell \bar{\Lambda}^T \right]$.

Proposition 2 shows that the regulator's value function over assets q and the model M depends on two sets of terms.

The first line is the regulator's baseline welfare in the absence of intervention. Baseline welfare starts from the return on assets, $q^T (R - p)$, net of losses on assets sold prior to maturity, $-\ell(q)^T \theta(q)$. It then nets out the adjustment costs from the initial portfolio and from liquidations in the Middle. All of these terms are evaluated assuming no intervention, that is $\tau = 0$. As a result, they do not depend on the regulator's choice of model M .

The second line is the welfare gains resulting from the regulator's optimally chosen intervention, which comprises two terms. The first term reflects the expected magnitude of the regulator's intervention, and so depends on $\mathbb{E}_0[\tau^*]$. Because interventions are targeted to the social benefit of intervening (equation 8), this term is quadratic in the expected intervention. It is weighted by the consequences of intervention, reflected by the extent to which interventions move prices (the first term in Ψ_0), the regulatory costs (the second term), and the movement in holding costs from liquidations.

The welfare also depends on the accuracy of the regulator's model and intervention, captured in the second term that depends on the covariance matrix of the policy intervention, $\text{cov}_0(\tau^*)$. The intuition comes from the law of total variance: a perfectly informed regulator would target an intervention based on the true parameters of the system, $\tau^* = \Xi^{-1}(\sum_i \bar{\Lambda}_i^T)^T \Gamma L(q)$. A regulator that

learned no information relative to the prior would have no covariance, so that all benefits would be reflected in the prior beliefs that are used to design the intervention.

In general, both of these terms depend on the choice of model M . The latter term depends on the choice of model because the choice of model determines the covariance matrix of the policy intervention. The former term also depends on the model, because the anticipated choice of model can affect the expected intervention. In particular, even under predictive-causal independence, from Proposition 1 a model that informs about the causal impact of the policy intervention τ will inform about both the costs Ξ of regulation and also about the causal impact $(\sum_i \bar{\Lambda}_i^\tau)^T$ of the policy intervention. We define a *predictive* model as one that is uninformative about the causal structure of the policy intervention.

Definition 3. We say a model M is predictive if $\mathbb{E}[\Xi|s, M] = \mathbb{E}_0[\Xi]$ and $\mathbb{E}[\sum_i \bar{\Lambda}_i^{\tau T} | s, M] = \mathbb{E}_0[\sum_i \bar{\Lambda}_i^{\tau T}]$ for all signals s .

2.4 Optimal Predictive Model Choice

We can now characterize the optimal choice of a model. We maintain simplicity by assuming predictive-causal independence for the main text (Definition 2) and by focusing on predictive models (Definition 3).

We assume that the regulator faces a separable utility cost $\mathcal{C}(M, q)$ from adopting (predictive) model $M \in \mathcal{M}$. By Proposition 2, for a predictive model (Definition 3) the key sufficient statistic of the model from a welfare perspective is the prior covariance matrix over the ex-post policy intervention, $\Sigma_0^\tau = \text{cov}_0(\tau^*)$. We can therefore re-represent costs as $C(\Sigma_0^\tau, q) = \inf_{M \in \mathcal{M} | \text{cov}_0(\tau^*) = \Sigma_0^\tau} \mathcal{C}(M, q)$. We assume that C is differentiable in (Σ_0^τ, q) over an open ball that contains its optimal value.

As a result, for a predictive model we can write the regulator's optimization problem as

$$\max_{\Sigma_0^\tau} \frac{1}{2} \text{tr}(\Psi_0 \Sigma_0^\tau) - C(\Sigma_0^\tau, q).$$

We obtain the following result on the optimal predictive model.

Proposition 3. The regulator's optimal predictive model solves

$$\frac{\partial C(\Sigma_0^\tau, q)}{\partial \Sigma_0^\tau} + \left(\frac{\partial C(\Sigma_0^\tau, q)}{\partial \Sigma_0^\tau} \right)^T = \Psi_0 \quad (11)$$

where $\frac{\partial C(\Sigma_0^\tau, q)}{\partial \Sigma_0^\tau}$ is a square matrix whose ij -th element is $\frac{\partial C(\Sigma_0^\tau, q)}{\partial (\Sigma_0^\tau)_{ij}}$.

Proposition 3 yields an intuitive trade-off on optimal predictive model choice in terms of choice of covariance matrix of the policy intervention. The regulator is willing to pay a higher marginal cost to increase precision on dimensions where Ψ_0 is larger, that is when the policy has greater impact on aggregate liquidations and prices, when the regulatory cost is higher, or when the impact

through holding costs is higher. In this sense, there is a complementarity between predictive power and knowledge of the causal policy impact: the regulator is willing to pay larger costs to acquire precise predictions of aggregate liquidations precisely on dimensions on which the policy impact on liquidations and liquidation prices is anticipated to be largest.

2.5 Privately Optimal Asset Allocations

We now turn to studying how model choice and ex-post intervention shape the ex-ante asset allocations of intermediaries. We begin by studying the private optimum of individual intermediaries, and then study socially optimal interventions.

Intermediary i in the Beginning takes as given the model choice M^* of the regulator and the resulting possible equilibria, and solves

$$\max_{q_i} \mathbb{E}_0 \left[q_i^T (R_i - p_i) - \ell_i^T (R_i + \tau_i^* - \gamma) - \frac{1}{2} q_i^T H_i^q q_i - \frac{1}{2} \ell_i^T H_i^\ell \ell_i \right].$$

Intermediary i 's asset allocation is therefore affected both through the specific intervention τ_i^* anticipated, and also through the liquidation price (which in turn also affects equilibrium liquidations). Intermediary i knows that equilibrium liquidations are determined as in Lemma 1, but only internalizes the effect of its own q_i on its own liquidations (and not on equilibrium liquidation price or equilibrium liquidations, and not on the optimal model choice or intervention).

We next turn to studying the effect of the regulator's model choice on ex-ante asset allocations of intermediaries. Note that Proposition 4 does not rely on predictive-causal independence or a predictive model.

Proposition 4. *The privately optimal asset allocation satisfies*

$$\mathbb{E}_0 \left[H_i^q q_i^* + \Lambda_i^{qT} H_i^\ell \bar{\ell}_i(q^*) \right] = \mathbb{E}_0 \left[R_i - p_i - \Lambda_i^{qT} \theta_i(q^*) \right] - \mathbb{E}_0 \left[\Lambda_i^{qT} \left(\tau_i^* - \left(H_i^\ell \bar{\Lambda}_i^T + \Gamma \left(\sum_i \bar{\Lambda}_i^T \right) \right) \tau^* \right) \right] \quad (12)$$

Proposition 4 expresses the optimal portfolio choice in the form of a marginal cost-marginal benefit trade-off. The left hand side captures the marginal cost of increasing holdings of an asset, which includes both the ex-ante and ex-post adjustment costs of holding more and liquidating more of an asset. This term is evaluated at the no-intervention benchmark, that is as-if we had $\tau^* = 0$.

The right hand side captures the marginal benefit, which is decomposed into a marginal benefit absent the ex-post regulatory intervention (the first term) plus the marginal benefit arising from the impact of intervention (the second term). The first term on the RHS is the baseline expected asset return, $R_i - p_i$, net of costs of liquidations. The liquidation costs are the amount liquidations are changed by increasing holdings, Λ_i^{qT} , times the fire sale loss in liquidation, $\theta_i(q^*)$.

The final term on the RHS captures the impacts of model choice and the ex-post intervention.

This is the only term in equation 12 that depends on model choice and intervention. It captures the cost of holding an asset induced through regulation. The expectation of ex-post policy intervention has two impacts. First, a higher wedge τ_{in} on intermediary i increases the cost of liquidating asset n , discouraging the intermediary from holding portfolios that result in it liquidating that asset ex post. There are also equilibrium costs of the vector of wedges τ through the equilibrium price. In contrast, here a higher wedge τ_{in} encourages the intermediary i to hold more of an asset when the asset price rises as a result, which partially offsets the benefit from raising the price. This is a standard channel of moral hazard.

It is clear that the ex-post intervention affects the optimal asset allocation: a higher expected tax on asset n ex post directly discourages its purchase ex ante, but a higher liquidation price ex post encourages its purchase ex ante. As a result, even a purely predictive model can impact the asset allocation ex ante. In particular, even though under a purely predictive model (Definition 3) the expectation of τ^* is that same as if the regulator ran no model, there is a covariance induced between the tax itself, τ^* , and the impact on liquidations, Λ_i^q , as long as the predictive model is loading at least some on Bayesian inference on the impact of asset holdings on liquidations Λ_i^q . That is to say, focusing on the direct term $\mathbb{E}_0[\Lambda_i^{qT} \tau_i^*]$, for a purely predictive model we can write

$$\mathbb{E}_0 \left[\Lambda_i^{qT} \tau_i^* \right] = \mathbb{E}_0 \left[\Lambda_i^{qT} \right] \mathbb{E}_0 \left[\tau_i \right] + \text{cov}_0 \left(\Lambda_i^{qT}, \mathbb{E}_0[\Xi]^{-1} \mathbb{E}_0 \left[\left(\sum_i \bar{\Lambda}_i^T \right)^T \Gamma \right] \mathbb{E}[L(q)|s, M] \right),$$

where the second term reflects how the true value of Λ_i^q varies with the prediction of liquidations. In a similar fashion, moral hazard can also be exacerbated if the impact of asset allocations on liquidations Λ_i^q is what a key driver of the predictability of ex post liquidations.

2.6 Socially Optimal Asset Allocation

Finally, we study the constrained efficient asset allocation q of the regulator. Formally, we assume that the regulator can impose a wedge on asset holdings, that is a revenue-neutral tax $t_i = (t_{i1}, \dots, t_{iN})^T$. Because the regulator has a complete set of wedges ex ante, this is equivalent to the regulator directly picking portfolio allocations ex ante.⁹ The following proposition characterizes the optimal wedges $t = (t_1^T, \dots, t_I^T)^T$ when the regulator chooses an optimal predictive model ex post.

⁹It is straight-forward to extend results to incomplete ex-ante instruments, but for brevity we focus on complete instruments.

Proposition 5. *For a predictive model, the social planner's optimal ex-ante portfolio wedges t are*

$$t = \mathbb{E}_0 \left[\left(\sum_j \bar{\Lambda}_j^q \right)^T \Gamma L(q) - \Lambda^{q,eT} \sum_i \left(\theta_i(q) + H_i^\ell \ell_i(q) \right) \right] + \frac{dC(\Sigma_0^\tau, q)}{dq} - \mathbb{E}_0 \left[\left(\sum_i \bar{\Lambda}_i^q \right)^T \Upsilon_0^T \Psi_0 \mathbb{E}_0[\tau^*] - \mathbb{E}_0[\Psi \tau^*] \right] \quad (13)$$

where $\Upsilon_0 = \mathbb{E}_0[\Xi]^{-1} \mathbb{E}_0 \left[\left(\sum_i \bar{\Lambda}_i^\tau \right)^T \Gamma \right]$ and where $\mathbb{E}_0[\Psi \tau]$ is a stacked vector whose i^{th} block is $\mathbb{E}_0 \left[\Lambda_i^{qT} \left(\tau_i - \left(\Gamma \left(\sum_i \bar{\Lambda}_i^\tau \right) + H_i^\ell \bar{\Lambda}_i^\tau \right) \tau \right) \right]$.

The optimal ex-ante asset tax formula is intuitive and is decomposed into two lines. The first line captures the uninternalized effects of intermediaries increasing holdings of an asset in absence of regulatory intervention. First, as q changes, total liquidations change across intermediaries, resulting in changes in the liquidation price and hence changes in revenue proceeds from total liquidations $L(q)$. Second, changes in forced liquidations through changes in the equilibrium price result in losses due to the size of the discount $\theta_i(q)$ and due to changes in marginal holding costs $H_i^\ell \ell_i(q)$. In absence of an ex-post intervention, this first line would capture the standard macroprudential regulation of asset positions.¹⁰

The second line captures the interaction of the ex-ante asset regulation with the ex-post model design and intervention. The first term, dC/dq , reflects how the changing assets held by intermediaries affects the cost of acquiring information through the model ex post. The regulator imposes larger holding taxes on asset positions that make running the model more costly ex post. For example, the regulator might discourage holdings of non-transparent or hard-to-assess assets.

The second term on the second line captures the effect on the expected intervention size (the term $\mathbb{E}_0[\tau^{*T}] \Psi_0 \mathbb{E}_0[\tau^*]$ in equation 10). Intuitively, concentrations into a position q increase total liquidations on the margin by $\sum_i \bar{\Lambda}_i^q$, which leads the regulator ex post to increase the size of the ex-post intervention τ^* accordingly. This increase in ex-post intervention helps to manage the fire sale and so mitigates the impact of the ex-ante increase in asset allocation. This gives a first dimension whereby the ex-post intervention provides a partial substitute for the ex-ante asset tax, and leads the regulator to require potentially smaller interventions ex ante, knowing that fire sales will be managed ex post.

Finally, there are the direct and moral hazard effects on intermediaries, $\mathbb{E}_0[\Psi \tau^*]$, are those in the final term in equation 12 of Proposition 4. First, the prospect of the ex-post tax on liquidations directly discourages holdings of the asset ex ante when holding more promotes liquidations ex post. Second, there is the moral hazard effect: as the liquidation price rises, intermediaries' perceived cost of liquidations $\theta_i(q)$ falls, and so they become more willing to hold more of an asset even if that forces them to liquidate. It is interesting again to observe that for assets for which the direct

¹⁰Since our framework allows negative positions $q_{in} < 0$, we can think of these negative positions as liabilities. As such, it also captures regulation of such liabilities.

effect dominates the indirect price effect, the ex-post tax actually serves as a substitute for ex-ante regulation, leading to a smaller intervention t ex ante.

2.7 Ex-Post Subsidy-Based Interventions

Our baseline analysis assumes that the regulator directly manages liquidations through a wedge/tax on liquidations. In practice, ex-post interventions can also involve “bailout” interventions that subsidize retaining assets rather than taxing selling them (e.g. lender of last resort, debt guarantees). Suppose that rather than taxing liquidations ℓ_i , the regulator instead applies a subsidy on the amount $q_i - \ell_i$ of the asset that is held to maturity. Formally, the payoff to intermediary i in the End is now

$$U_i = U_i + (q_i^T - q_i^{*T})\tau_i - (\ell_i^T - \ell_i^{*T})\tau_i,$$

so that the regulator continues to apply revenue-neutral interventions.¹¹

Because the regulator in the Middle takes the asset allocations as given, the regulator’s optimization problem over both model choice and the ex-post intervention are formally the same as before. Thus, Lemma 1 and Propsitions 1-3 continue to apply. However, the privately optimal asset allocation is affected by the subsidy on asset retention. In particular, in Proposition 4, the expected return on holding asset q_i is raised from R_i to $R_i + \tau_i^*$. Intuitively an ex-post subsidy promotes overinvestment in that asset. This gives rise to a familiar channel of moral hazard.

Relative to the case of a liquidation tax, it is worthwhile to note that the only new term for private intermediaries in their optimization problem is the revenue benefits $q_i^T \mathbb{E}_0[\tau_i^*]$ of their asset position. Therefore, relative to the liquidation tax, the regulator’s model choice only affects intermediaries’ asset allocation ex ante to the extent it changes the expected size of the intervention, $\mathbb{E}_0[\tau_i^*]$. In particular for a predictive model under predictive-causal independence (Definitions 2 and 3), we know that $\mathbb{E}_0[\tau_i^*]$ does not depend on the model used; instead, it is only the covariance matrix of the intervention (its precision) that deepends on the model. It is thus interesting and surprising that purely predictive models that help with prediction but not with causal inference do not exacerbate moral hazard relative to the case of the liquidation tax. In contrast, models that are informative about the causal impact of the policy intervention change expectations of the policy intervention, and are potentially associated with moral hazard. The logic is reminiscent of that of Laffont and Tirole (1986) in the context of regulation of a firm with unobservable effort and uncertain costs.

2.8 Extensions

Better Informed Intermediaries. Our model assumes that in the Beginning, intermediaries and regulators share a common prior over the model parameters. In practice, intermediaries may be

¹¹Since individual intermediaries take revenue remissions as given, there is still moral hazard from the subsidy.

better informed about model parameters. One could extend the model by assuming that intermediaries had a more precise prior than regulators. This would lead a regulator to also infer information about the structural parameters of the economy from observing portfolio holdings directly. One possible method of incorporating is to interpret $C(M, q)$ as a cost of the inference the regulator is undertaking, and so interpret the model as both a processing of information (including from such Bayesian inference). One possible advantage of the ex-post intervention would be its ability to react to information discovered from observing intermediaries' choices, which might reduce risks of moral hazard or imperfectly calibrated regulation ex ante. That is, better-informed intermediaries who saw a regulator was under-regulating an asset n ex ante would also know that the regulator would discover the mistake ex post, and so intervene more strongly upon it. Exploring the effects of differential information between the regulator and intermediaries is an interesting direction.

Dynamic Learning. We embed a one-shot learning problem. One could extend our framework to embed dynamic information acquisition by assuming that our baseline Beginning-Middle-End model was a stage game played at each date $t = 0, 1, \dots$. In this environment, one could think of there being a true distribution μ^* from which model parameters are drawn each period, so that the regulator and intermediaries learn about this true distribution over time. The regulator and intermediaries would carry information forward at each date, and so the regulator would consider both how a model acquired information on the current crisis and also how it informed about the underlying variables. We conjecture that this would make predictive models relatively myopically useful: they would give potentially substantial information for intervening in the current crisis, but relatively little information about the underlying structural parameters, and so might be of limited use in updating beliefs about the true distribution of parameters. This could face the regulator with an interesting dynamic trade-off between myopically acquiring more predictive information, and trying to uncover the true structural parameters that would also be useful for designing regulation and interventions during the next crisis.

Commitment vs. Discretion in Model Choice. We have assumed the regulator chooses the model and ex-post policy intervention with discretion, after asset allocations are chosen. As highlighted by the results in Section 2.5, this leads to a potential time consistency problem in model choice since the model choice affects ex-ante asset allocations. This time consistency problem will likely be more pronounced the more incomplete the regulator's ability to regulate portfolio positions is ex ante, for example if the regulator cannot regulate all intermediaries or all assets. In our ongoing work, we are exploring the implications of commitment versus discretion for optimal model design and welfare gains from use of a predictive model.

3 Graph Representation Learning for Holdings Data

We now move to an empirical examination of whether high-dimensional forecasting models can in fact be used successfully for macroprudential applications. This analysis establishes several points. First, it provides a blueprint for the practical implementation of deep learning models by central banks and other financial supervisors. Second, it demonstrates that such models do in fact have significant predictive power—which is a crucial motivating fact for the theory. Third, motivated by the theoretical setting, it lays out several key design principles that deep learning architectures should follow when applied to financial holdings data.

GNN Architectures and Holdings Data. We begin by introducing a deep learning architecture tailored to holdings data, and we discuss why it is optimally suited to this setting. Much of the modern deep learning toolkit is optimized for grid inputs such as images and sequence inputs such as text. Indeed, architectures that have proved successful are those that exploit the particular structure of the data they are modeling: examples include convolutional neural networks for images and other grid-structured data, and text transformers—which underpin modern large language models—for textual data. Financial holdings data does not fit neatly into either of these categories. Instead, its defining feature is that it has very rich *relational structure*: the data can be naturally thought of as representing a graph connecting investors and assets, with the information relevant to the learning task contained in the graph’s edges, i.e. the positions connecting investors to assets.

Graph neural networks (GNNs) are a class of deep learning models that specifically model and exploits relational structure in the data (Scarselli et al. 2008; Hamilton et al. 2017; Wu et al. 2020). By iteratively propagating and aggregating information along the edges that connect graph nodes (in this case, assets and investors), GNNs learn embeddings for each node in the graph—i.e., representations of the nodes in a latent vector space that effectively capture the characteristics relevant to the tasks the model is trained against. In practice, in our implementation an investor’s embedding evolves based on the embeddings of the assets they hold, and those asset embeddings, in turn, adjust in light of the embeddings of the investors that include them. Through iterated rounds of this neighbor-aggregation mechanism (also known as *message passing*), information flows along the network of positions, allowing the model to learn explicitly from the relational structure of the holdings data.

GNNs have state-of-the-art empirical success in domains where relational structure is paramount. The protein-folding model AlphaFold, for example, uses a graph-based architecture to encode the three-dimensional interactions among amino acids, and it has dramatically advanced the field of protein structure prediction (Jumper et al. 2021). Traffic-forecasting systems used for products such as Google Maps represent road networks as graphs and learn congestion and travel-time patterns directly from the topology of streets and highways (Derrow-Pinion et al. 2021). Similarly, in frontier drug discovery models, molecules are treated as graphs of atoms and bonds, and GNNs have revolutionized the prediction of chemical properties and binding affinities (Jiang et al. 2021).

Each of these breakthroughs rests on the fundamental ability of GNNs to natively handle and learn from graph-structured data.

In the context of asset holdings data, GNNs have two key properties that other deep architectures lack: permutation invariance and inductive learning. First, the models are *permutation invariant*, meaning that the learning process and estimation results do not depend on any arbitrary relabeling or reshuffling of investor and asset identifiers. The architecture attains permutation invariance because aggregation operates over unordered sets of neighbors, and by doing so they respect a fundamental symmetry of the problem.

Second, the GNN architecture features *inductive learning*. Crucially, all of the model’s parameters are shared, in the sense that there are no parameters specific to particular assets or investors. Therefore the same learned representation rules apply across all nodes and edges: given a graph $\mathcal{G}$, regardless of the number and identity of the nodes, the trained GNN architecture is able to construct asset and investor embeddings simply from the graph structure and the node characteristics. Inductive generalization thus follows naturally from parameter sharing: the model can immediately generate embeddings and forecasts for new, unseen investors or assets without any retraining. This feature is not only valuable for real-time regulatory applications, but also it enforces a strong form of regularization upon the model, preventing overfitting of the training data and leading to good generalization to unseen, out-of-sample data.

This combination of relational representation learning, proven empirical performance, permutation invariance, and inductive design makes GNNs a particularly well suited deep learning architecture for the holdings data setting and for macroprudential applications. Our specific implementation is a *graph transformer*, which incorporates an attention mechanism in the basic GNN architecture, as we discuss below.

Core Architecture Specification. To formally specify the GNN graph transformer architecture, we start by modeling the holdings data as a bipartite graph—meaning a graph with two distinct types of nodes, and whose edges only connect nodes of the two different types. The holdings data forms a bipartite graph in which investors connect to assets via position edges. Formally, we let a given cross-section of the data be represented as $\mathcal{G} = (\mathcal{I}, \mathcal{A}, \mathcal{E})$, where $\mathcal{I} = \{1, \dots, I\}$ indexes investors, $\mathcal{A} = \{1, \dots, N\}$ is the set of assets, and $\mathcal{E}$ is the set of edges (i.e., positions). A certain position exists in the graph if the corresponding investor-asset pair is in the set of edges: whenever investor i holds asset a , letting w_{ia} be the size of the position, we have that $(i, a, w_{ia}) \in \mathcal{E}$.¹² We write $\mathcal{V} = \mathcal{I} \cup \mathcal{A}$ as the set of all nodes, both assets and investors. We let $N(v)$ be the neighborhood set for a given node v . This is the set of all nodes that are connected to v : for assets, $N(v)$ corresponds to the investors holding the asset, while for investors this corresponds to the set of assets in the investor’s portfolio.

¹²The positions can alternatively be written as $w_{ai} = w_{ia}$. We allow for both notations so as to keep the rest of the formal architecture specification symmetric.

Each node $v \in \mathcal{V}$ has an associated set of characteristics x_v . The characteristics can be either numerical (e.g., the total amount outstanding of a given security) or categorical (e.g., the type of institutional investor). For categorical features, these are pre-embedded to a vector space of dimension d_c using maps that are learned jointly with the rest of the model’s parameters. The node characteristics vector x_v concatenates over all individual features, both scalar-valued ones and vector-valued categorical ones.

The graph transformer learns node embeddings (both asset embeddings and investor embeddings) at various hierarchical levels of information aggregation. The first-level embeddings, which we denote as $h_v^{(0)}$, are obtained simply by embedding the node characteristics vectors x_v into a d_h -dimensional hidden space via a learnable map ϕ :

$$h_v^{(0)} = \phi(x_v). \quad (14)$$

Next, the architecture passes these layer-zero node embeddings through several successive layers of message passing which aggregate information over the graph: this message passing stage is the core of the GNN model and is what allows the model to learn from the data’s relational structure. To increase the expressivity and learning capability of the model, we integrate attention mechanisms in each of the GNN message passing layers, which allow the architecture to learn from the data which positions should be given more or less weight in any given information aggregation steps. This incorporation of attention layers is what makes our architecture a graph transformer.

We perform L distinct layers of message passing. The message passing at layer ℓ unfolds in several distinct steps. First, the prior-layer node embeddings are passed through a feed-forward layer $M^{(\ell)} : \mathbb{R}^{d_h} \rightarrow \mathbb{R}^{d_h}$ to construct *node messages*:¹³


$$M_v^\ell = M^{(\ell)}(h_v^{(\ell-1)}). \quad (15)$$

Next, the attention mechanism computes aggregation weights that dictate how the individual node messages will be weighted when passing information over the graph. We allow for S_A distinct and independent attention heads. Each attention head $s = 1, \dots, S_A$ forms attention weights using learnable projection matrices $W_{qs}, W_{ks} \in \mathbb{R}^{d_h \times d_h}$ which project prior-layer embeddings into query and key spaces, respectively—akin to how text transformers form attention values for the interactions of tokens in a text sequence (Vaswani et al. 2017). The attention weight $\alpha_{vu}^{(\ell)}$ indicates the importance of messages from each node u in the neighborhood $N(v)$ of node v at layer ℓ , and it is formed by averaging over the individual attention heads:

$$\alpha_{vu}^{(\ell)} = \frac{1}{S_A} \sum_{s=1}^{S_A} \frac{\exp\{(W_{qs} h_v^{(\ell-1)})^T (W_{ks} h_u^{(\ell-1)})\}}{\sum_{u' \in N(v)} \exp\{(W_{qs} h_v^{(\ell-1)})^T (W_{ks} h_{u'}^{(\ell-1)})\}}. \quad (16)$$

¹³All feed-forward layers in our architecture use GELU non-linearities (Hendrycks and Gimpel 2016). These help avoid dead gradients during training.

Figure 1: **Model architecture**



Notes: We visualize the architecture of our graph transformer model diagrammatically.

Having established attention weights $\alpha_{vu}^{(\ell)}$, the message passing algorithm then constructs an aggregated message $m_v^{(\ell)}$ for each node v by averaging over the messages received from each of its neighbors $u \in \mathcal{N}(v)$, according to:

$$m_v^{(\ell)} = \sum_{u \in \mathcal{N}(v)} w_{vu} \alpha_{vu}^{(\ell)} M_u^{(\ell)}, \quad (17)$$

The node embeddings at the next hierarchical stage ℓ are then updated using another feed-forward layer $U^{(\ell)} : \mathbb{R}^{2d_h} \rightarrow \mathbb{R}^{d_h}$ which maps the current embeddings and the current aggregated messages into the next-stage representations:

$$h_v^{(\ell)} = U^{(\ell)}(h_v^{(\ell-1)}, m_v^{(\ell)}). \quad (18)$$

After L rounds of message-passing, a readout feed-forward layer $\rho : \mathbb{R}^{d_h} \rightarrow \mathbb{R}^{d_e}$ produces final asset and investor embeddings:

$$e_v = \rho(h_v^{(L)}). \quad (19)$$

When referring to node embeddings in subsequent sections, unless otherwise specified we refer to these final representations e_v . Because all parameters are shared across nodes and layers, this architecture is both permutation invariant and inductive, allowing predictions on unseen investors or assets without retraining.

Prediction Heads. To train and leverage the learned embeddings e_v , we attach two task-specific heads to the graph transformer. The first task uses a masked autoencoder (MAE) objective, where we randomly mask a subset of edges (including non-existent edges drawn at random, which we represent using $w_v = 0$) and ask the model to predict whether the edge exists ($w_v \neq 0$) and the associated position size w_v . The second task uses a supervised trade prediction objective, where we ask the model to engage in a pure forecasting task, using embeddings computed using time t holdings information to predict the cross-sectional pattern of investor trades in the future, between time t and time $t + 1$.

For the masked autoencoder (MAE) objective, we randomly mask a subset of edges and predict masked weights $\hat{w}_{ia} = f_{\text{AE}}(e_i, e_a)$ using a feed-forward head layer f_{AE} which takes as inputs the embeddings for the given masked investor-asset pair (i, a) . The autoencoder objective minimizes a mean squared error loss defined over the divergence between the true edge weights w_v and the predicted edge weights $\hat{w}_v$:

$$\mathcal{L}_{\text{AE}} = \sum_{(i,a) \in \mathcal{V}_{\text{masked}}} (w_{ia} - \hat{w}_{ia})^2, \quad (20)$$

where $\mathcal{V}_{\text{masked}}$ denotes the set of masked edges.

For trade prediction, we construct the targets by first defining the percent changes in holdings

for a given position in asset a by investor i between time t and time $t + 1$ (in practice, quarters) as:

$$\Delta\%q_{ia,t} = \frac{q_{ia,t} - q_{ia,t-1}}{q_{ia,t-1}}. \quad (21)$$

We strip away all common movement in trades for a particular asset a , such as for example movement induced by changes in the asset’s valuation which affect all investors. To do this, we construct cross-sectional z-scores of the trades $\Delta\%q_{ia,t}$, which subtract the average percentage position change for asset a in the given time period ($\overline{\Delta\%q_{a,t}}$) and divide by the standard deviation of the same position changes ($\sigma(\Delta\%q_{a,t})$), placing all assets in all time periods on the same scale. The cross-sectional trade z-scores are thus defined as:

$$y_{ia,t} = \frac{\Delta\%q_{ia,t} - \overline{\Delta\%q_{a,t}}}{\sigma(\Delta\%q_{a,t})}. \quad (22)$$

The cross-sectional trade patterns captured by $y_{ia,t}$ are the targets for the supervised prediction head. Specifically, we construct predicted trades $\hat{y}_{ia} = f_{\text{TP}}(e_i, e_a)$ using a feed-forward layer head f_{TP} which acts on the relevant pair of asset and investor embeddings.¹⁴ The supervised objective is again defined over the mean squared error between the actual trades and the predicted trades:

$$\mathcal{L}_{\text{TP}} = \sum_{(i,a) \in \mathcal{V}} (y_{ia} - \hat{y}_{ia})^2. \quad (23)$$

Model Training. To recap, the model contains several learnable components, all of which are parameterized using a high-dimensional set of parameters. The trainable components include the feature map ϕ , the pre-embedding functions for categorical characteristics, the message feed-forward layers $M^{(\ell)}$, the attention mechanism projection matrices W_{qs} and W_{ks} , the node update function $U^{(\ell)}$, the embeddings projection layer ρ , and the task-specific prediction heads f_{AE} and f_{TP} . The model architecture is summarized visually in Figure 1. We collect the set of parameters in all these learnable components in the vector Θ .

The model is trained end-to-end by minimizing a joint loss which combines the mean squared error losses from the two training tasks:

$$\min_{\Theta} \mathcal{L}(\Theta) = \mathcal{L}_{\text{AE}} + \kappa \mathcal{L}_{\text{TP}}, \quad (24)$$

where $\kappa > 0$ determines the relative weight of the two training tasks. We optimize the model parameters Θ using the Adam optimizer (Kingma 2014).

¹⁴While for compactness we are not carrying through time subscripts on the embeddings e_v , naturally for the trade prediction task we use embeddings estimated using the graph $\mathcal{G}$ at time $t - 1$ to construct the predictions for time t , $y_{ia,t}$.

Table 1: **Model hyperparameters**

Hidden dimension (d_h)	256
Embedding dimension (d_e)	128
Layers (L)	3
Attention heads	4
Dropout rate	0.1
Learning rate	10^{-2}
Weight decay	10^{-5}
Loss weight (κ)	1

Notes: We list the hyperparameters used for our graph transformer architecture and for training.

Optimality of Graph-Based Architectures for Holdings Data. Before moving on the empirical implementation, we discuss more precisely the sense in which message-passing, graph-based architectures are optimal in the context of holdings-based problems. To do this, we lay out a few definitions. Let $W \in \mathbb{R}^{I \times N}$ be the full holdings matrix with entries w_{ia} , and let $f: \mathbb{R}^{I \times N} \rightarrow \mathbb{R}^d$ be a functional acting on the graph $\mathcal{G}$ represented by W . A continuous graph functional is permutation-invariant if, for all permutation matrices $P_1 \in \mathbb{R}^{I \times I}$, $P_2 \in \mathbb{R}^{N \times N}$, it satisfies $f(P_1 W P_2^T) = f(W)$. Informally, permutation invariance means that if we were to arbitrarily relabel columns and rows of the holdings matrix (i.e., investors and assets), the output of f would remain the same. All regulatory or prediction targets we care about (future trading patterns, systemic risk scores, etc.) are assumed to lie in this family, reflecting the economics of the problem.

A well-known idea in the literature on graph deep learning is that in order to represent a permutation-invariant mapping with the fewest parameters, one should enforce permutation invariance via shared (message-passing) parameters (Zaheer et al. 2017, Maron et al. 2018, Xu et al. 2018). This is the sense in which the models are optimally sample-efficient. For illustration, compare two classes of models. First, consider the class of GNNs which implement permutation-invariant message-passing layers as described above, with shared parameters. Second, consider the class of sequence or grid networks that act on the flattened matrix $\text{vec}(W)$ under a fixed but arbitrary ordering of rows and columns (this class includes recurrent neural networks, convolutional neural networks, sequence transformers, and so on). Intuitively, GNNs can approximate permutation-invariant graph functionals without carrying superfluous degrees of freedom that sequence/grid networks would have to “use up” to relearn permutation invariance from data. Message-passing GNNs are efficient because the architecture is itself permutation-invariant, avoiding the need to use additional parameters to learn and enforce it.

4 Empirical Implementation

We train our deep learning model on quarterly institutional holdings from Factset, starting in 2005Q1. The training-validation split is done at the level of quarters: we sort quarters into training and validation sets. The holdings data gives us observations of the positions graph $\mathcal{G}$ for each quarter’s cross-section, and we also use it to construct the standardized trade indicators $y_{ia,t}$. We do not include any quarters following 2019Q3 in the training set, so that we can use the Covid crisis of 2020 as a particularly strict test episode, in the sense that the Covid quarters are not just out of the training sample, but also the model is only trained with data prior to the start of the crisis, mimicking the way in which the model would be deployed in an actual regulatory scenario.

We construct the node feature vectors x_v using reference information from Factset as well as from the Global Capital Allocation Project (GCAP) security master file (Coppola et al. 2021, Coppola 2025). For assets, the feature vectors x_v include asset class, currency of denomination, amount outstanding, number of holders, average position size, standard deviation of position size, as well as bond sub-class and coupon for debt securities. For investors, they include institution type (such as open-end mutual funds, ETFs, separate accounts, etc.), manager style (including flags for active vs. passive portfolio management and strategy types), total AUM, number of positions, average position size, and standard deviation of position size. In principle, our architecture also allows for the use of global features that vary over time but not across nodes: these can be introduced as vectors which enter message-passing in the same way for all nodes in a given time period. In ongoing work, we are integrating global features and assessing the impact on the model’s performance.¹⁵ Similarly, we are exploring the use of price data both as predictive features and targets for the model.


Hyperparameters are chosen as in Table 1. In particular, we set the hidden dimension to $d_h = 256$ and the final embeddings dimensionality to $d_e = 128$. We use $L = 3$ layers of message-passing: we do not increase L beyond this number to prevent over-smoothing problem, whereby the node embeddings would converge to similar values: intuitively, over-smoothing occurs for higher numbers of message-passing iterations since each consecutive iteration increases the *receptive field* of each node, i.e. the set of nodes that the final-layer embeddings attend to, and for high L values the receptive fields of all nodes converge to the largest possible field, which is the set of all nodes in the graph. We allow for four attention heads and we give equal weight to the two training tasks by setting $\kappa = 1$. Altogether, the model has a total of 3,640,465 parameters.

We introduce dropout during training for additional regularization, with a dropout rate of 0.1. The Adam optimizers uses a starting learning rate of 10^{-2} with progressive weight decay. We train the model using a compute node with four NVIDIA H100 GPUs. We train for up to 500 epochs, with early stopping based on the loss on the validation sample.

Out-of-sample, the MAE head achieves a correlation above 0.90 between reconstructed and true positions, indicating that the GNN captures structural regularities in holdings. As mentioned in the

¹⁵The weight to be placed on global features can be made learnable by the model. Global features may include time series measures such as aggregate credit spreads and other macro series.

Figure 2: **Performance metrics**



Notes: We plot the correlations between the trained model’s predictions and the targets for the two tasks (masked autoencoder and cross-sectional trade prediction). The blue bars show performance on the training set, while the red bars show out-of-sample performance on the validation set.

introduction, the autoencoder prediction is best interpreted in the context of the model’s parameter-to-data ratio: in this case, the model’s roughly 3.6 million parameters represent less than 1% of possible investor-asset pairs in the data. The trade-prediction head yields an average correlation of just under 0.30 between predicted and realized trade indicators, with minimal degradation from training to validation sets. Figure 2 plots the correlation between the trained model’s predictions and the targets for both tasks: the blue bars show the performance on the training sample, while the red bars show the out-of-sample performance on the validation set. The fully inductive design, sharing parameters across nodes and layers, ensures stable performance on unseen investors and assets, such that the model performs very similarly out-of-sample as it does on the training data. In ongoing work, we are performing a descriptive analysis of the asset and investor embeddings produced by the model, so as to provide greater interpretability of the model’s predictions.

A natural question is whether the predictive ability of the model comes primarily from relatively more mechanical aspects of the data, such as by correctly assessing the trades of passive index-tracking investors. In Figure 3a, we show that this is not the case by reporting the out-of-sample performance of the model on a sub-sample consisting only of active investors. We also show performance on a sub-sample that only includes open-end mutual funds, as these are the investor category with the highest degree of coverage within the Factset holdings data. In both cases, the forecasting performance of the model is quantitatively similar to that on the full validation sample.

An additional possible concern is that the model’s predictive ability may be concentrated in calm periods rather than the market stress episodes where macroprudential interventions are most

relevant. To rule this out, in Figure 3b we also show the performance on the trade prediction task separately for each quarter in the sample, as a time series: since the train-validation split occurs at the level of quarters, this time series naturally combines both training and validation data. The shaded gray areas correspond to market stress periods, defined as those when the St. Louis Fed Financial Stress Index is above 1.5. The predictive performance is consistently high both in stress periods and in non-stress periods.¹⁶ In particular, the model’s performance during the Covid crisis yields a particularly stringent test, since none of the quarters following 2019Q3 are included in the training set: the model is only trained with pre-crisis data, precisely as it would be deployed in an actual regulatory scenario, and nonetheless it displays high accuracy in predicting the patterns of trading during the course of the crisis quarters. Altogether, these results demonstrate the efficacy of graph-based predictive models for real-time macroprudential surveillance.


5 Conclusion

This paper develops a theoretical and empirical framework for understanding the role of high-dimensional predictive models in financial regulation. We formalize the tradeoffs regulators face when deploying models that deliver precise forecasts but limited causal insight, and we characterize when and how such models can improve welfare. We introduce a graph-based deep learning architecture tailored to holdings data, which we use to learn representations of assets and investors which achieve state-of-the-art results in forecasting trading patterns with minimal out-of-sample performance loss. Our empirical analysis demonstrates that real-time prediction of portfolio dynamics is feasible and provides a blueprint for practical implementation. While predictive models are not substitutes for structural content or causal inference, our results suggest that they can meaningfully complement structural knowledge, particularly when regulators have the ability to target ex post interventions.


¹⁶The FRED ticker for the St. Louis Fed Financial Stress Index is STLFSI4. We also note that the model’s performance exhibits a slight downtrend occurring between 2012 and 2017, stabilizing by the end of the period.

Figure 3: **Performance metrics: heterogeneity**

(a) Validation set performance, by subsample



(b) Performance for trade prediction, by date



Notes: Panel A shows the out-of-sample correlation on the validation set between the trained model's predictions and the targets, for both tasks (trade prediction and masked autoencoder). We show this in the subsamples consisting of open-end mutual funds only and of active funds only. Panel B shows the correlation between trade predictions and targets on the full training and validation data, by quarter. Shaded gray areas correspond to periods where the St. Louis Fed Financial Stress Index is above 1.5.

References

- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “Systemic risk and stability in financial networks,” *American Economic Review*, 2015, *105* (2), 564–608.
- Adrian, Tobias and Hyun Song Shin**, “Procyclical leverage and value-at-risk,” *The Review of Financial Studies*, 2014, *27* (2), 373–403.
- Athey, Susan**, “Beyond prediction: Using big data for policy problems,” *Science*, 2017, *355* (6324), 483–485.
- Athey, Susan**, “The impact of machine learning on economics,” in “The economics of artificial intelligence: An agenda,” University of Chicago Press, 2018, pp. 507–547.
- Athey, Susan and Stefan Wager**, “Policy learning with observational data,” *Econometrica*, 2021, *89* (1), 133–161.
- Barro, Robert J and David B Gordon**, “A positive theory of monetary policy in a natural rate model,” *Journal of political economy*, 1983, *91* (4), 589–610.
- Bernanke, Ben S and Mark Gertler**, “Agency costs, collateral, and business fluctuations,” 1986.
- Bianchi, Javier**, “Overborrowing and Systemic Externalities in the Business Cycle,” *American Economic Review*, December 2011, *101* (7), 3400–3426.
- Bianchi, Javier**, “Efficient Bailouts?,” *American Economic Review*, December 2016, *106* (12), 3607–59.
- Bianchi, Javier and Enrique G. Mendoza**, “Optimal Time-Consistent Macroprudential Policy,” *Journal of Political Economy*, 2018, *126* (2), 588–634.
- Brainard, William C and James Tobin**, “Pitfalls in financial model building,” *The American economic review*, 1968, *58* (2), 99–122.
- Brunnermeier, Markus K and Martin Oehmke**, “Bubbles, financial crises, and systemic risk,” *Handbook of the Economics of Finance*, 2013, *2*, 1221–1288.
- Bryzgalova, Svetlana, Victor DeMiguel, Sicong Li, and Markus Pelger**, “Asset-pricing factors with economic targets,” *Available at SSRN*, 2023, *4344837*.
- Burns, Arthur F and Wesley C Mitchell**, *Measuring business cycles*, National bureau of economic research, 1946.

- Caballero, Ricardo J. and Arvind Krishnamurthy**, “International and domestic collateral constraints in a model of emerging market crises,” *Journal of Monetary Economics*, December 2001, *48* (3), 513–548.
- Chari, V. V. and Patrick J. Kehoe**, “Bailouts, Time Inconsistency, and Optimal Regulation: A Macroeconomic View,” *American Economic Review*, September 2016, *106* (9), 2458–93.
- Chen, Luyang, Markus Pelger, and Jason Zhu**, “Deep learning in asset pricing,” *Management Science*, 2024, *70* (2), 714–750.
- Clayton, Christopher and Andreas Schaab**, “Multinational Banks and Financial Stability,” *The Quarterly Journal of Economics*, 01 2022, *137* (3), 1681–1736.
- Clayton, Christopher and Andreas Schaab**, “Bail-Ins, Optimal Regulation, and Crisis Resolution,” *The Review of Financial Studies*, 03 2025, p. hhaf002.
- Coppola, Antonio**, “In safe hands: The financial and real impact of investor composition over the credit cycle,” *The Review of Financial Studies*, 2025, p. hhaf017.
- Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger**, “Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens,” *The Quarterly Journal of Economics*, 2021, *136* (3), 1499–1556.
- Coval, Joshua and Erik Stafford**, “Asset fire sales (and purchases) in equity markets,” *Journal of Financial Economics*, 2007, *86* (2), 479–512.
- Derrow-Pinion, Austin, Jennifer She, David Wong, Oliver Lange, Todd Hester, Luis Perez, Marc Nunkesser, Seongjae Lee, Xueying Guo, Brett Wiltshire et al.**, “Eta prediction with graph neural networks in google maps,” in “Proceedings of the 30th ACM international conference on information & knowledge management” 2021, pp. 3767–3776.
- Dolphin, Rian, Barry Smyth, and Ruihai Dong**, “Stock embeddings: Learning distributed representations for financial assets,” *arXiv preprint arXiv:2202.08968*, 2022.
- Dávila, Eduardo and Anton Korinek**, “Pecuniary Externalities in Economies with Financial Frictions,” *The Review of Economic Studies*, 02 2017, *85* (1), 352–395.
- Einav, Liran and Jonathan Levin**, “The data revolution and economic analysis,” *Innovation Policy and the Economy*, 2014, *14* (1), 1–24.
- Einav, Liran and Jonathan Levin**, “Economics in the age of big data,” *Science*, 2014, *346* (6210),

1243089.

- Elliott, Matthew, Benjamin Golub, and Matthew O Jackson**, “Financial networks and contagion,” *American Economic Review*, 2014, *104* (10), 3115–3153.
- Fang, Xiang, Bryan Hardy, and Karen K Lewis**, “Who holds sovereign debt and why it matters,” *The Review of Financial Studies*, 2025, p. hhaf031.
- Farboodi, Maryam and Laura Veldkamp**, “Long-run growth of financial data technology,” *American Economic Review*, 2020, *110* (8), 2485–2523.
- Farhi, Emmanuel and Iván Werning**, “A theory of macroprudential policies in the presence of nominal rigidities,” *Econometrica*, 2016, *84* (5), 1645–1704.
- Faria-e Castro, Miguel, Joseba Martinez, and Thomas Philippon**, “Runs versus lemons: Information disclosure and fiscal capacity,” *The Review of Economic Studies*, 2017, *84* (4), 1683–1707.
- Fisher, Irving**, “The debt-deflation theory of great depressions,” *Econometrica: Journal of the Econometric Society*, 1933, pp. 337–357.
- Fontanier, Paul**, “Optimal policy for behavioral financial crises,” *Journal of Financial Economics*, 2025, *166*, 104005.
- Friedman, Milton**, “The methodology of positive economics,” *Essays In Positive Economics*, 1953.
- Gabaix, Xavier, Ralph SJ Koijen, Robert Richmond, and Motohiro Yogo**, “Asset embeddings,” *Available at SSRN 4507511*, 2024.
- Gabaix, Xavier, Ralph SJ Koijen, Robert Richmond, and Motohiro Yogo**, “Upgrading Credit Pricing and Risk Assessment through Embeddings,” *Available at SSRN*, 2025.
- Giglio, Stefano, Bryan Kelly, and Dacheng Xiu**, “Factor models, machine learning, and asset pricing,” *Annual Review of Financial Economics*, 2022, *14* (1), 337–368.
- Gillis, Talia B and Jann L Spiess**, “Big data and discrimination,” *The University of Chicago Law Review*, 2019, *86* (2), 459–488.
- Goldstein, Itay and Yaron Leitner**, “Stress tests and information disclosure,” *Journal of Economic Theory*, 2018, *177*, 34–69.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu**, “Empirical asset pricing via machine learning,”

- The Review of Financial Studies*, 2020, *33* (5), 2223–2273.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu**, “Autoencoder asset pricing models,” *Journal of Econometrics*, 2021, *222* (1), 429–450.
- Haavelmo, Trygve**, “The probability approach in econometrics,” *Econometrica: Journal of the Econometric Society*, 1944, pp. iii–115.
- Haddad, Valentin, Alan Moreira, and Tyler Muir**, “When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed’s response,” *The Review of Financial Studies*, 2021, *34* (11), 5309–5351.
- Hamilton, Will, Zhitao Ying, and Jure Leskovec**, “Inductive representation learning on large graphs,” *Advances in neural information processing systems*, 2017, *30*.
- Hendrycks, Dan and Kevin Gimpel**, “Gaussian error linear units (gelus),” *arXiv preprint arXiv:1606.08415*, 2016.
- Jiang, Dejun, Zhenxing Wu, Chang-Yu Hsieh, Guangyong Chen, Ben Liao, Zhe Wang, Chao Shen, Dongsheng Cao, Jian Wu, and Tingjun Hou**, “Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models,” *Journal of cheminformatics*, 2021, *13*, 1–23.
- Jumper, John, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko et al.**, “Highly accurate protein structure prediction with AlphaFold,” *nature*, 2021, *596* (7873), 583–589.
- Kindleberger, Charles and Robert Aliber**, *Manias, panics and crashes: a history of financial crises*, Springer, 1978.
- Kingma, Diederik P**, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- Kiyotaki, Nobuhiro and John Moore**, “Credit cycles,” *Journal of political economy*, 1997, *105* (2), 211–248.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan**, “Human decisions and machine predictions,” *The quarterly journal of economics*, 2018, *133* (1), 237–293.

- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer**, “Prediction policy problems,” *American Economic Review*, 2015, *105* (5), 491–495.
- Koopmans, Tjalling C**, “Measurement without theory,” *The Review of Economics and Statistics*, 1947, *29* (3), 161–172.
- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh**, “Shrinking the cross-section,” *Journal of Financial Economics*, 2020, *135* (2), 271–292.
- Krishnamurthy, Arvind and Tyler Muir**, “How credit cycles across a financial crisis,” *The Journal of Finance*, 2025, *80* (3), 1339–1378.
- Laffont, Jean-Jacques and Jean Tirole**, “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 1986, *94* (3), 614–641.
- Leitner, Yaron and Basil Williams**, “Model secrecy and stress tests,” *The Journal of Finance*, 2023, *78* (2), 1055–1095.
- Lorenzoni, Guido**, “Inefficient Credit Booms,” *The Review of Economic Studies*, 07 2008, *75* (3), 809–833.
- Lucas, Robert E**, “Econometric policy evaluation: A critique,” in “Carnegie-Rochester conference series on public policy,” Vol. 1 North-Holland 1976, pp. 19–46.
- Maron, Haggai, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman**, “Invariant and equivariant graph networks,” *arXiv preprint arXiv:1812.09902*, 2018.
- Mullainathan, Sendhil and Jann Spiess**, “Machine learning: an applied econometric approach,” *Journal of Economic Perspectives*, 2017, *31* (2), 87–106.
- Nagel, Stefan**, *Machine learning in asset pricing*, Princeton University Press, 2021.
- Orlov, Dmitry, Pavel Zryumov, and Andrzej Skrzypacz**, “The design of macroprudential stress tests,” *The Review of Financial Studies*, 2023, *36* (11), 4460–4501.
- Parlatore, Cecilia and Thomas Philippon**, “Designing stress scenarios,” *The Journal of Finance*, 2025, *80* (2), 833–873.
- Sarkar, Suproteem K**, “Economic representations,” 2025.
- Scarselli, Franco, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini**, “The graph neural network model,” *IEEE transactions on neural networks*, 2008,

20 (1), 61–80.

Schularick, Moritz and Alan M Taylor, “Credit booms gone bust: monetary policy, leverage cycles, and financial crises, 1870–2008,” *American Economic Review*, 2012, 102 (2), 1029–1061.

Shapiro, Joel and David Skeie, “Information management in banking crises,” *The Review of Financial Studies*, 2015, 28 (8), 2322–2363.

Stein, Jeremy C., “Monetary Policy as Financial Stability Regulation*,” *The Quarterly Journal of Economics*, 01 2012, 127 (1), 57–95.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, 2017, 30.

Williamson, Oliver E, “Corporate finance and corporate governance,” *The journal of finance*, 1988, 43 (3), 567–591.

Wu, Zonghan, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu, “A comprehensive survey on graph neural networks,” *IEEE transactions on neural networks and learning systems*, 2020, 32 (1), 4–24.

Xu, Keyulu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka, “How powerful are graph neural networks?,” *arXiv preprint arXiv:1810.00826*, 2018.

Zaheer, Manzil, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola, “Deep sets,” *Advances in neural information processing systems*, 2017, 30.