The Inefficient Markets Hypothesis:

A Theory of Casino Capitalism

Christian Hellwig

June 17, 2026

WP 2026-11

https://doi.org/10.21033/wp-2026-11

FEDERAL RESERVE BANK *of* CHICAGO

The Inefficient Markets Hypothesis: A Theory of Casino Capitalism*

Christian Hellwig†

Toulouse School of Economics, Federal Reserve Bank of Chicago and CEPR

June 17, 2026

Abstract

I develop a general equilibrium model of investment and equity markets in which departures from the Efficient Markets Hypothesis generate "Casino Capitalism": time-inconsistent, short-termist firm decisions, firm-level over-investment coupled with inefficient liquidation of positive-NPV projects, and aggregate under-investment with depressed stock prices and distorted inter-temporal savings. These distortions arise even if all agents are fully rational and forward-looking. At the firm level, the wedge between market-implied and fundamental returns on investment causes an endogenous preference reversal that renders the Laissez-faire outcome time-inconsistent and Pareto-inferior. In the aggregate, firm-level attempts to boost stock prices are self-defeating via an aggregate investment wedge as micro distortions translate into macro inefficiency. Restoring firm-level rationality by empowering shareholders can worsen aggregate outcomes and precipitate a complete capital market shutdown.

JEL Codes: G14, G30, G40, D25, E22

Keywords: Capital market imperfections, corporate short-termism, casino capitalism, capital mis-allocation

*This paper was originally presented on the occasion of Nobu Kiyotaki's 70th birthday conference at Princeton University in May 2025; I am immensely grateful to Nobu for decades of scientific and professional advice and for his unwavering support. This paper draws on two earlier papers with Elias Albagli and Aleh Tsyvinski; I thank both for extensive conversations on these and related subjects, while retaining sole responsibility for the content of (and any inaccuracies in) the present manuscript. I am grateful for comments by George-Marios Angeletos, Gadi Barlevy, Kilian Huber, Andreas Schaab, Xavier Vives and multiple seminar and conference audiences. Opinions expressed in this article are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. I acknowledge financial support from the European Research Council (starting grant agreement 263790) and the French National Research Agency (ANR) under the project "ANR23-CE26-0011-01" and the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010).

†christian.hellwig@tse-fr.eu

"*When the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done ... The measure of success attained by Wall Street, regarded as an institution of which the proper social purpose is to direct new investment into the most profitable channels in terms of future yield, cannot be claimed as one of the outstanding triumphs of laissez-faire capitalism.*"

John Maynard Keynes (1936)

1 Introduction

Shareholders as Heroes or Villains? Ever since the Great Depression of the 1930s, Keynes' famous comparison of capital markets to a casino is frequently invoked to describe the excesses and social costs of financial and economic boom-bust cycles. After comparing investment in stock markets to a beauty contest in which success depends more on correctly predicting the other investors' views as summarized by market returns than on fundamental returns these investments generate, Keynes paints shareholders who control firms' investment activities as the villains of Laissez-faire capitalism whose short-sighted pursuit of financial gains misdirects investments away from their most profitable uses at the expense of real wealth creation for overall society.

The Casino Capitalism narrative stands in stark contrast to shareholder value theory (Friedman, 1970; Jensen and Meckling, 1976) and the efficient market hypothesis (Fama, 1970), according to which, in the absence of market power or externalities, well-functioning financial markets induce rational, time-consistent planning by firms, align shareholder value with the social value of firms' investment activities, and direct financial resources to their most productive uses without any need for regulatory oversight. A strong case for Laissez-faire emerges, as policy interventions are at best inconsequential and at worst harmful to economic activity. The efficient markets hypothesis implicitly treats shareholders as the heroes of capitalism, whose pursuit of financial gains helps to maximize the social surplus generated by the firms' investment activities.

The case for shareholder value maximization rests on a parsimonious, internally consistent theory of financial markets with a precise definition of market efficiency as the elimination of all arbitrage opportunities. The no arbitrage principle fully aligns asset prices with their fundamental values and shareholder value with social surplus. This theoretical benchmark lays the burden of proof for policy interventions squarely on the shoulders of proponents of a corrective regulation who must justify its merits by identifying departures from this benchmark they wish to correct.

In this paper, I develop a fully articulated, internally consistent theory for the competing narrative of Casino Capitalism. I consider a dynamic model of firms' investment decisions and show that the efficient markets hypothesis is both sufficient and generically necessary for equilibrium outcomes that are characterized by time-consistent optimal corporate planning and Pareto-efficient allocation of capital in the economy. Outside the efficient markets benchmark, the collective pursuit of financial gains by shareholders instead results in equilibrium outcomes that are aptly summarized by the casino capitalism narrative: shifting corporate objectives or preference reversals, time inconsistency, and firm-level and aggregate outcomes that are strictly Pareto inferior. Importantly, corporate short-termism arises endogenously from capital market imperfections alone, with no need to further appeal to behavioral biases, bounded rationality, agency costs or short-termist incentives of managers. Moreover these inefficiencies can become arbitrarily severe, if firm decisions are sufficiently responsive to expected or realized stock prices, even if capital markets are nearly efficient.

Model and Results. Section 2 introduces a dynamic model of corporate investment and hiring in a shareholder-controlled firm. An initial generation of shareholders make a capital investment in a firm before either selling the firm to a new generation of final shareholders, or liquidating their investment if they are unable to sell the firm at a positive market price. The final shareholders hire labor to generate a stochastic dividend stream.

I characterize firm decisions as the sequential equilibrium of the non-cooperative game between initial and final shareholders and compare equilibrium outcomes when the stock market is efficient with alternative scenarios in which it is not. Under the efficient markets hypothesis, the non-cooperative equilibrium between successive shareholder generations always maximizes total surplus, i.e. the present value of expected future dividend streams net of investment costs, which rationalizes firm decisions as a time-consistent optimal plan. Theorem 1 establishes a generic converse to time-consistent planning under the efficient markets hypothesis: if the stock market departs ever so slightly from the efficient markets benchmark, the sequential equilibrium generically leads to time-inconsistent firm decisions and Pareto-inferior outcomes.

Formally, Theorem 1 shows that with capital market imperfections, shareholder-controlled firms not only fail to maximize total surplus, but the resulting firm decisions cannot be rationalized as time-consistent planning or maximization of *any* common welfarist objective function. Initial shareholders do not internalize how their decisions affect future dividends and thus the welfare of final shareholders, and final shareholders do not internalize how the stock market prices in their expected future hiring decisions and always hire labor efficiently to maximize expected dividends. When capital market imperfections drive a wedge between stock prices and fundamental values, the resulting over- or under-valuation creates an endogenous conflict of interest between initial and final shareholders, which in turn feeds into initial shareholder investment and liquidation incentives.

As a result, initial and final shareholders would both gain if they could somehow agree ex ante on a common plan that increases total surplus by reducing investment distortions or inefficient liquidation, as well as hiring policies that boost share prices by aligning more closely with the policy that maximizes the stock price and thus inital shareholders' expected payoffs.

The resulting corporate strategies then resemble those of a dynamic decision maker whose preferences are anchored on a long-term objective of total surplus, but subject to the "temptation" of deviating to pursue short-term financial gains. When viewed as a unitary dynamic decision-maker, the firm appears to suffer from limited self-control (Strotz, 1955; Gul and Pesendorfer, 2001, 2005), unless by coincidence the technological primitives perfectly align shareholder preferences on maximizing either total surplus or initial shareholder welfare. Kydland and Prescott (1977) emphasize the time inconsistency and Pareto inferiority of discretionary economic policy when economic agents are forward-looking. Theorem 1 shows that capital market imperfections generate a similar form of time-inconsistency inside firms because shareholders do not perfectly internalize how their decisions impact future dividends or prior stock prices.

While the failure of time-consistent corporate planning arises generically for *any* capital market imperfection, Section 3 derives more specific results for a model of stock markets with noisy information aggregation that mirrors the realities of delegated portfolio management, i.e. final shareholders buy stocks through a combination of specialist traders who have noisy private information about the firms' fundamentals (hedge funds), and noise traders (mutual funds), and the share price aggregates information about firm fundamentals with noise. With noisy information aggregation the share price systematically over-reacts to market information, resulting in overly optimistic (resp. pessimistic) assessments of firm fundamentals on the upside (resp. downside), and thus a meaningful departure from the efficient markets benchmark. The over-reaction of the stock price to market information then has real consequences: it causes inefficient liquidation of positive NPV investments when the market grows overly pessimistic about the firms' revenue prospects, and excessive investment ex ante as expected market returns to investment are inflated by both the prospect of overly optimistic market expectations on the upside, and the protection the liquidation option provides against overly pessimistic market expectations on the downside.

Section 4 embeds the analysis of firm decisions with stock market imperfections in a general equilibrium economy with a continuum of firms, in which firm-specific risks are fully diversified in the representative shareholder's portfolios. Theorem 2 shows that capital market imperfections not only cause individual firms to be time-inconsistent and inefficient for the firm itself, but also lead to an allocation of capital across firms that is Pareto inferior, unless no firm is liquidated in equilibrium and capital market imperfections are homogeneous, so that all firms over- or under-invest to the same degree. This is a knife-edge condition that rules out meaningful heterogeneity in risk characteristics, capital and labor intensities across firms or sectors.

In general equilibrium, the firm-level incentives to boost stock prices interact with an endogenous aggregate investment wedge that ties stock prices to corporate earnings. With homogeneous capital market imperfections and no inefficient liquidation, this investment wedge exactly offsets the firm-level incentives to over- or under-invest, so that the equilibrium maximizes firm-level and aggregate surplus. However, with liquidation or heterogeneous capital market imperfections, the equilibrium is both inefficient at the firm level and Pareto inferior in the aggregate: Inefficient liquidations or capital misallocation reduce aggregate earnings, stock prices, and investment. In general equilibrium, initial and final shareholders both lose from lower aggregate earnings, the latter directly, the former indirectly through lower aggregate stock prices. Therefore, everyone would benefit if initial shareholders could somehow agree ex ante on a common plan to avoid capital misallocation and firm liquidation. Moreover, liquidation and heterogeneous capital market imperfections give rise to firm-level inefficiency, i.e. for a given aggregate stock price, initial and final shareholders in almost all firms could gain by agreeing to a plan that increases the firm's total surplus and boosts its share price to transfer a part of these surplus gains to initial shareholders.

Finally, Section 5 discusses implications of market-induced corporate short-termism for capital structure or managerial incentives. For example, giving initial shareholders more control over future corporate strategies (for example by signing managers to long-term incentive contracts) overcomes the preference reversals caused by equity market imperfections. Similarly, debt increases firm value as a commitment device: by diluting final shareholders' hiring incentives (Myers, 1977), debt provides commitment to less hiring, which increases the stock price and reduces liquidation incentives when the firm is under-valued. In partial equilibrium, debt and long-term managerial contracts increase the firm's market value, which benefits initial shareholders. But the impact of these policies on final shareholder payoffs and total surplus is ambiguous and trades off gains from avoiding inefficient liquidation against costs from hiring and investment distortions. In general equilibrium, everyone's welfare moves one-for-one with aggregate dividends. Empowering shareholders can even cause a complete shut-down of the stock market if the resulting hiring distortions are so severe thataggregate corporate earnings are no longer positive.

Related Literature. The analysis draws on ideas from several distinct fields. To my knowledge, they have not been brought together within a single unified framework.

Formally, I combine a canonical model of corporate investment and hiring with a formulation of financial markets that departs from the efficient markets hypothesis. Each of these building blocks is kept deliberately simple so as to focus the analysis on the interaction between the two.

My model of the financial market builds on models of noisy information aggregation (Grossman and Stiglitz, 1980, Hellwig, 1980, Diamond and Verrecchia, 1981), or more specifically on Albagli, Hellwig and Tsyvinski (2024), who characterize equilibrium prices of arbitrary securities in a non-linear noisy rational expectations equilibrium model. I extend the static model of investment and stock markets with noisy information aggregation of Albagli, Hellwig and Tsyvinski (2023) by adding a liquidation option and the subsequent labor choice. The combination of hiring and liquidation introduces a new conflict of interest that creates inefficiencies from the fact that initial and final shareholders are no longer aligned on a common firm objective.

My analysis offers a foundation of behavioral biases in corporate decision-making based on capital market imperfections rather than investor psychology. The investment and liquidation decisions by initial shareholders share similarities with a consumer problem with limited self-control (e.g. Gul and Pesendorfer, 2001), but the connection to hiring decisions by final shareholders generates endogenous preference reversals, time inconsistency, and lack of commitment in firm behavior (e.g. Gul and Pesendorfer, 2005) that result in Pareto-inferior planning. If the firm is viewed as a unitary decision-maker, its behavior can therefore no longer be rationalized as the outcome of a time-consistent optimal plan that maximizes a well-defined firm objective function.

Starting with Kydland and Prescott (1977), the literature on optimal policy emphasizes the value of commitment over discretion when expectations of future policies impact current private sector decisions. I show that the same value of commitment applies to firms whose market value depends on future firm decisions in a manner that is not fully aligned with future earnings. In my model, successive shareholder generations play the same role as successive governments in models of political short-termism due to special-interest politics and electoral turn-over (Alesina and Tabellini, 1990; Battaglini and Coate, 2007), but the conflict of interest arises endogenously from the gap between the firm's market value and its earnings. With special interest politics, government debt has commitment value by constraining future governments' spending decisions (Alesina and Tabellini, 1990; Amador, 2004); I show that the same holds true for corporate debt when current shareholders seek to influence future firm decisions.

Time inconsistency can also arise from aggregating heterogeneous preferences. Jackson and Yariv (2014, 2015) emphasize the impossibility of time-consistent aggregation of heterogeneous time preferences, Davila and Schaab (2025) emphasize preference reversals over efficiency-redistribution tradeoffs in incomplete market environments. Closest to this paper, Caramp, Kozlowski and Teeple

(2024) argue that transaction costs can endogenously generate heterogeneous discount rates in an incomplete market economy in which some agents sell and others hold or purchase stocks at any point in time, resulting in quasi-hyperbolic discounting at the firm level. In my analysis, time-inconsistency instead results from the market's failure to align incentives inter-temporally, but shareholder preferences are fully aligned within each period. Abstracting from heterogeneity of preferences within shareholder generations also facilitates the definition of the firm's objective function in each period and allows me to compare the equilibrium outcome with a well-defined social planner's problem or a Pareto frontier between initial and final shareholder welfare.[1]

The tension between initial and final shareholder preferences introduces novel tradeoffs that shape capital structure or managerial incentives – see Tirole (2006) for a textbook treatment. The model is also consistent with empirical regularities in asset pricing that appear inconsistent with the efficient markets hypothesis, as well as the sensitivity of corporate investment to stock prices. I discuss the related empirical literature in Section 6.

The literature often attributes return anomalies and corporate short-termism to investors' and managers' psychological or belief biases (Stein, 1989; Baker and Wurgler, 2013). I instead interpret such behavior as the rational response of controlling shareholders to capital market imperfections. Bolton, Scheinkman, and Xiong (2006) argue that shareholders may want to distort investments in order to boost short-term stock prices when priors about the firm's prospects are heterogeneous and stock prices contain a speculative component. While the fundamental source of short-termism through departures of prices from fundamental values is similar, the source of market inefficiency is different, and the heterogeneous priors assumption doesn't accommodate efficiency and welfare results as naturally as the present model with common priors and noisy information aggregation.

An alternative explanation attributes corporate short-termism to agency conflicts between shareholders and managers and managerial career concerns (Narayanan, 1985). In these models, a focus on short-term market returns emerges as an efficient incentive contract between shareholders and managers, and corporate short-termism is constrained inefficient only if privately efficient managerial incentives do not internalize the full social returns to investment, such as in the presence of R&D spillovers (Terry, 2023). My analysis instead attributes corporate short-termism to shareholder preferences with imperfect capital markets and shows that shareholder-driven short-termism results in firm-level and aggregate outcomes that are almost never constrained efficient.

[1] The model could be extended to incorporate heterogeneity among initial shareholders by assuming that they are heterogeneous in their ability to trade on private information or exposure to liquidity shocks. Market imperfections then cause conflict of interest not just across but also within shareholder generations.

2 Partial equilibrium: Single firm

Agents, technology, and timing

Consider a single firm operating over four stages. There are two classes of shareholders who are risk-neutral with respect to firm-specific risks: *initial* shareholders, who make the investment decision, and *final* shareholders, who acquire the firm and operate it after observing a public signal about its prospects.

In Stage 1, initial shareholders invest $k \geq 0$ units of capital at increasing, convex cost $C(k)$, with $C'(\cdot) > 0$ and $C''(\cdot) > 0$. The investment is sunk after Stage 1.

In Stage 2, initial shareholders either sell their capital stake to final shareholders at a price $P(z, k)$ that depends on the initial investment k, and a public signal $z \in \mathbb{R}$ about the returns to the investment that becomes available before or while the market opens, or they liquidate the firm, receiving a payoff of 0, if the initial shareholders are unable to sell the firm at a positive price.

In Stage 3, final shareholders hire labor l at an exogenously given competitive wage $w > 0$, with total labor cost $w \cdot l$. In Stage 4, the final shareholders collect the dividend

$$\Pi(\theta, k, l) = A(\theta) k^{\alpha} l^{1-\alpha}, \tag{1}$$

where θ is a fundamental productivity shock, the firm-level TFP $A(\theta)$ is a positive, non-decreasing function of $\theta \in \mathbb{R}$ with $\lim_{\theta \to -\infty} A(\theta) = 0$, and $\alpha \in (0, 1)$ denotes the capital share. Moreover, $A(\theta)\, \lambda^{1-\alpha}$ denotes revenue per unit of capital for a labor intensity $\lambda \equiv l/k$.

The fundamental productivity shock satisfies $\theta \sim \mathcal{N}(0, 1)$. No additional information is available when firms invest at Stage 1, but new information becomes available in the financial market (Stage 2), in the form of the signal z that is normally distributed with distribution $z|\theta \sim \mathcal{N}\left(\theta, \tau^{-1}\right)$, where τ denotes the precision of the signal. This public signal z admits two different interpretations, depending on the micro structure of the market environment: it either represents an exogenous public signal of the fundamentals of the firm, or a sufficient statistic of dispersed investor information that is aggregated through the share price. The latter case imposes additional restrictions: the signal is observable only if the firm is not liquidated (i.e. if there are sufficiently many buyers willing to pay a positive price), and the price function must be invertible along the equilibrium path, so that traders can infer the realized signal from the share price.

Payoffs and strategies

Working backwards, I begin from the final shareholders, then characterize the initial shareholders' liquidation decision, and finally their initial investment decision.

Final shareholders' hiring policy. Define $R(z,\lambda) \equiv \mathbb{E}[A(\theta)|z]\,\lambda^{1-\alpha} - w\lambda$ as the *fundamental return* to capital, i.e. the expected dividend per unit of capital, net of labor costs, conditional on the public signal z and labor intensity λ. The final shareholders set $\lambda(z)$ to maximize $R(z,\lambda)$:

$$\lambda(z) \in \arg\max_{\lambda \geq 0} \left\{ \mathbb{E}[A(\theta)|z]\,\lambda^{1-\alpha} - w\lambda \right\}. \tag{2}$$

The optimal labor intensity satisfies the first-order condition

$$(1-\alpha)\,\mathbb{E}[A(\theta)|z]\,\lambda(z)^{-\alpha} = w, \tag{3}$$

which pins down $\lambda^*(z) = \{(1-\alpha)\,\mathbb{E}[A(\theta)|z]/w\}^{1/\alpha}$ as a function of z alone, independently of k. Moreover, final dividends net of labor costs are linear in k and equal to $R(z,\lambda^*(z)) \cdot k$, where

$$R(z,\lambda^*(z)) \equiv \mathbb{E}[A(\theta)|z]\,\lambda^*(z)^{1-\alpha} - w\lambda^*(z) = \alpha \left\{ \left(\frac{1-\alpha}{w}\right)^{1-\alpha} \mathbb{E}[A(\theta)|z] \right\}^{\frac{1}{\alpha}} > 0 \tag{4}$$

represents the equilibrium fundamental return on capital.

Initial shareholders' liquidation decision. Initial shareholders choose a liquidation strategy such that the investment is liquidated whenever its market price falls below 0, i.e. whenever $P(z,k) \leq 0$.[2] Anticipating that the equilibrium share price is linear in k along the equilibrium path, I let $\hat{R}(z,\lambda) \equiv P(z,\lambda;1)$ denote the *market return* to capital for a given signal realization z and expected labor intensity λ, and $P(z,\lambda;k) = \hat{R}(z,\lambda)k$ the share price. For now, I take $\hat{R}(z,\lambda)$ as given; below, I contrast the efficient markets hypothesis (EMH) under which $\hat{R}(z,\lambda) = R(z,\lambda)$ for all (z,λ) with alternative models of equity markets that depart from the EMH. Given the final shareholders' optimal labor intensity $\lambda^*(z)$, the equilibrium share price takes the form $P(z,k) \equiv \hat{R}(z,\lambda^*(z)) \cdot k$.

The initial shareholders' liquidation strategy then divides the set of signal realizations into a *continuation set* $Z \subseteq \mathbb{R}$, such that $\hat{R}(z,\lambda^*(z)) > 0$ if $z \in Z$, and a *liquidation set* $\mathbb{R} \setminus Z$, for

[2] I break ties by assuming the incumbent shareholders liquidate if $\hat{R}(z) = 0$. This tie-breaking rule is consistent with an arbitrarily small, positive return to liquidation.

which $\mathbb{E}\left[\hat{R}(z,\lambda^*(z))|z\in\mathbb{R}\setminus Z\right]\leq 0$. The first condition guarantees that the stock price is strictly positive over the continuation set. The second condition imposes that the expected stock price is non-positive over the liquidation set. When z represents an exogenous public signal about the firm, incumbent shareholders perfectly observe $\hat{R}(z,\lambda^*(z))$ both on and off the equilibrium path and the optimal continuation set is equal to $\bar{Z}\equiv\left\{z\in\mathbb{R}:\hat{R}(z,\lambda^*(z))>0\right\}$.

However, when z represents a sufficient statistic of dispersed information aggregated through the share price, the signal z is publicly observed only if the firm is not liquidated, the equity market is active, and the market return $\hat{R}(z,\lambda^*(z))$ is invertible (i.e., for practical purposes, monotone) for $z\in Z$. This opens the door to a certain form of indeterminacy in equilibrium prices and liquidation outcomes through self-confirming out-of-equilibrium beliefs: if upon seeing a certain price $P>0$, final shareholders always believe that the firm is not worth P, regardless of their prior or private signals, then they will not acquire the firm at that price. But if there is a signal realization $z^{\#}\in\bar{Z}$ for which $P=\hat{R}(z^{\#},\lambda^*(z^{\#}))k$, then signal realization $z^{\#}$ must be part of the liquidation set and shareholders' beliefs at this price remain off-equilibrium. This reasoning explains why the liquidation set can be indeterminate in equilibrium, and the optimality condition for liquidation only requires that the *expected* market return conditional on liquidation cannot be positive. To abstract from the possibility that market inefficiencies arise from this indeterminacy, I restrict attention to the market equilibrium that minimizes liquidation by setting the equilibrium continuation set equal to $\bar{Z}$ and assuming that $\hat{R}(z,\lambda^*(z))$ is invertible for $z\in\bar{Z}$. I verify this invertibility condition for the micro structure of equity markets that I specify below.

Initial shareholders' investment decision. I now move to the initial investment decision. Investment k is chosen to maximize the expected market value of the firm, net of investment costs:

$$\max_{k\geq 0}\left\{\int_{\bar{Z}}\hat{R}\left(z,\lambda^*\left(z\right)\right)\varphi\left(z\right)\,dz\cdot k-C(k)\right\},\tag{5}$$

where $\varphi(\cdot)$ denotes the unconditional density of $z\sim\mathcal{N}(0,1+\tau^{-1})$. The initial shareholders' first-order condition defines the firm's equilibrium investment level k^{LF}:

$$C'(k^{LF})=\mathbb{E}\left(\max\left\{0,\hat{R}\left(z,\lambda^*\left(z\right)\right)\right\}\right).\tag{6}$$

For a given market return $\hat{R}\left(z,\lambda^*\left(z\right)\right)$, the *Laissez-faire equilibrium* of this game is then defined by a hiring strategy $\lambda^*(z)$ for final shareholders that satisfies equation (3), the continuation set $\bar{Z}=\left\{z\in\mathbb{R}:\hat{R}(z,\lambda^*(z))>0\right\}$, and an investment level k^{LF} that satisfies equation (6).

Regularity conditions on returns. I conclude the description of the model set-up with some regularity conditions on the market return function $\hat{R}(z,\lambda)$. These conditions mirror properties that have already been established for fundamental returns $R(z,\lambda)$.

Assumption 1. *$\hat{R}(z,\lambda)$ and $R(z,\lambda)$ are twice continuously differentiable and strictly concave in λ. In addition, $\hat{R}(z,0) = R(z,0) = 0$, $\hat{R}_\lambda(z,0) = R_\lambda(z,0) = \infty$, $\hat{R}$, $\hat{R}_\lambda$, R, and R_λ are all negative for λ sufficiently large, and $\hat{R}$, $\hat{R}_\lambda$, R, and R_λ are strictly increasing in z for given $\lambda > 0$.*

Assumption 1 imposes that the labor input is essential for generating any returns to investment, that the marginal returns to hiring are unbounded at 0, diminishing and negative for sufficiently large λ, and that higher z (positive news about fundamentals) generate both higher market returns to capital and higher marginal returns to labor. It then follows that there exist continuous functions $\lambda^*(z) = \arg\max_{\lambda\geq 0} R(z,\lambda)$ and $\hat{\lambda}(z) = \arg\max_{\lambda\geq 0} \hat{R}(z,\lambda)$ that maximize fundamental and market returns. Moreover, $\lambda^*(z)$ and $\hat{\lambda}(z)$, as well as $R(z,\lambda^*(z))$ and $\hat{R}\left(z,\hat{\lambda}(z)\right)$, by virtue of envelope conditions, are all strictly positive and increasing in z.

Now, strict concavity insures that $\hat{R}(z,\lambda^*(z))$ is positive and strictly increasing in z, whenever $\hat{\lambda}(z) \geq \lambda^*(z)$, implying that values of z for which $\hat{\lambda}(z) \geq \lambda^*(z)$ are always included in the continuation set. However, without additional assumptions, the signs of both $\hat{R}(z,\lambda^*(z))$ and the derivative of $\hat{R}(z,\lambda^*(z))$ w.r.t. z are ambiguous whenever $\hat{\lambda}(z) < \lambda^*(z)$: when the hiring policy exceeds $\hat{\lambda}(z)$, further hires reduce the market returns and may eventually make $\hat{R}(z,\lambda^*(z))$ decreasing in z or even negative. Assuming that $\hat{R}(z,\lambda^*(z))$ is invertible for $z \in \bar{Z}$ in turn guarantees that $\bar{Z}$ must be convex whenever it is non-empty and rules out the possibility that $\hat{R}$ reaches a local extremum on the interior of $\bar{Z}$. Hence $\max\left\{0,\hat{R}(z,\lambda^*(z))\right\}$ must be monotone (increasing or decreasing) in z, and strictly monotone for $z \in \bar{Z}$.

Since $R(z,\lambda^*(z))$ is positive and strictly increasing in z, it is reasonable to focus on cases where the same is true for $\max\left\{0,\hat{R}(z,\lambda^*(z))\right\}$. Below, I confirm that this condition holds in a financial market structure with noisy information aggregation. This in turn implies that $\bar{Z} = \left\{z \in \mathbb{R} : z > z^{LF}\right\}$, where the threshold $z^{LF} > \infty$ either solves $\hat{R}\left(z^{LF},\lambda^*\left(z^{LF}\right)\right) = 0$ or $z^{LF} = -\infty$ if $\hat{R}(z,\lambda^*(z)) > 0$ for all z.

2.1 Time consistency and Pareto efficiency

Interpreting the firm as a unitary decision maker, I define the time-consistency of firm or shareholder preferences and the (intra-firm) Pareto efficiency of allocations $(k, Z, \lambda(\cdot))$ in analogy with the "multiple-selves" representation of time-inconsistent consumers. Time consistency of preferences and Pareto efficiency of allocations are closely related but they are not equivalent, so it will be useful to distinguish between them.

Shareholder preferences are time-inconsistent whenever an optimal dynamic plan for the firm chosen at stage 1 by initial shareholders does not coincide with the continuation plan preferred by final shareholders in stage 3. The initial shareholders' preferred plan $\left(\hat{k}, \hat{Z}, \hat{\lambda}(\cdot)\right)$ sets $C'(\hat{k}) = \mathbb{E}\left(\hat{R}\left(z, \hat{\lambda}(z)\right)\right)$, $\hat{Z} = \mathbb{R}$, and $\hat{\lambda}(\cdot) = \arg\max_{\lambda \geq 0} \hat{R}(z, \lambda)$ to maximize expected market returns, never liquidate the firm, and equalize the marginal cost of capital to the ex ante expected market return on capital. The final shareholders' preferred continuation plan instead always sets $\lambda^*(\cdot) = \arg\max_{\lambda \geq 0} R(z, \lambda)$ to maximize expected fundamental returns. Hence shareholder preferences are *time-inconsistent* whenever $\hat{\lambda}(\cdot) \neq \lambda^*(\cdot)$.

An allocation $(k, Z, \lambda(\cdot))$ is *Pareto inferior*, if there exists an alternative allocation $(k', Z', \lambda'(\cdot))$ that both shareholder groups prefer over $(k, Z, \lambda(\cdot))$, i.e.

$$\int_{Z'} \hat{R}(z, \lambda'(z)) \varphi(z)\, dz \cdot k' - C(k') \geq \int_{Z} \hat{R}(z, \lambda(z)) \varphi(z)\, dz \cdot k - C(k) \tag{7}$$

$$\int_{Z'} \left(R(z, \lambda'(z)) - \hat{R}(z, \lambda'(z))\right) \varphi(z)\, dz \cdot k' \geq \int_{Z} \left(R(z, \lambda(z)) - \hat{R}(z, \lambda(z))\right) \varphi(z)\, dz \cdot k, \tag{8}$$

with one group's preferences being strict. An allocation is Pareto efficient if it is not Pareto inferior. As usual, Pareto efficient allocations can be characterized as the solution to a Paretian social planner's problem that maximizes a weighted sum of initial and final shareholders' expected payoffs. Here the Pareto criterion is applied "intra-firm", i.e. to successive shareholder generations within the same firm. Let $(k^*, \mathbb{R}, \lambda^*(\cdot))$ with $C'(k^*) = \mathbb{E}[R(z)]$ denote the allocation that maximizes expected total surplus $\int_Z R(z, \lambda(z)) \varphi(z)\, dz \cdot k - C(k)$. Obviously, $\left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right)$ and $(k^*, \mathbb{R}, \lambda^*(\cdot))$ are both Pareto efficient.

Proposition 1. *If shareholder preferences are time-consistent then* $\left(k^{LF}, \bar{Z}, \lambda^*(\cdot)\right) = \left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right)$.

Hence time consistency of shareholder preferences is synonymous with the Laissez-faire equilibrium being Pareto efficient and implementing the initial shareholders' preferred allocation.

The efficient markets hypothesis implies time consistency

Consider now the case where the EMH holds and the market return coincides with the fundamental return, or $\hat{R}(z, \lambda) = R(z, \lambda)$ for all (z, λ).[3] Proposition 2 shows that the Laissez-Faire equilibrium under the EMH maximizes expected total surplus.

[3] This condition corresponds to the Efficient Markets Hypothesis in semi-strong form (Fama, 1970).

Proposition 2. *If* $\hat{R}(z,\lambda) = R(z,\lambda)$ *for all* (z,λ)*, then* $(k^{LF}, \bar{Z}, \lambda^*(\cdot)) = \left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right) = (k^*, \mathbb{R}, \lambda^*(\cdot))$.

Under the EMH the market return to capital is equalized to the fundamental return. It follows immediately that shareholder preferences are time-consistent ($\hat{\lambda}(\cdot) = \lambda^*(\cdot)$), which induces the incumbent shareholders both to never liquidate the firm at stage 2 and to make efficient investment decisions in stage 1. In other words, when equity prices are perfectly aligned with fundamental returns, incumbent shareholders who maximize the market value of their investment simultaneously maximize the expected dividends that accrue to final shareholders. The private and social returns to investment and hiring decisions coincide with the market returns, and the allocation that maximizes expected total surplus is decentralized.

Capital market imperfections lead to Pareto-inferior allocations

I conclude this initial section with a generic high-level converse to Propositions 1 and 2: With capital market imperfections, shareholder preferences are generically time-inconsistent and the resulting Laissez-faire equilibrium is generically Pareto inferior. Like Propositions 1 and 2, this result holds for arbitrary return functions $R(z,\lambda)$ and $\hat{R}(z,\lambda)$ that satisfy Assumption 1. It therefore requires neither an explicit micro-foundation for the fundamental returns $R(z,\lambda)$, nor an explicit model of capital market imperfections that are captured by the market return function $\hat{R}(z,\lambda)$. The economically interesting scenario arises when capital market imperfections cause shareholder preferences to be time-inconsistent, i.e. if $\hat{R} \neq R$ and $\hat{\lambda}(\cdot) \neq \lambda^*(\cdot)$.

Theorem 1. *Under Assumption 1, the Laissez-faire equilibrium* $(k^{LF}, \bar{Z}, \lambda^*(\cdot))$ *is Pareto-efficient if and only if* $(k^{LF}, \bar{Z}, \lambda^*(\cdot)) \in \left\{\left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right), (k^*, \mathbb{R}, \lambda^*(\cdot))\right\}$.

Theorem 1 establishes that the Laissez-faire equilibrium is Pareto-efficient if and only if it maximizes either the initial shareholder value or the expected total surplus. Both are knife-edge cases for generic departures from the EMH, as there is no reason to assume that with capital market imperfections shareholder preferences remain time-consistent and final shareholders maximize initial shareholder value, or that initial shareholder incentives exactly line up with maximizing the firms' expected total surplus, a condition that requires $\hat{R}(z, \lambda^*(z)) > 0$ for all z to avoid inefficient liquidation and $\mathbb{E}\left(\hat{R}(z,\lambda^*(z))\right) = \mathbb{E}\left(R(z,\lambda^*(z))\right)$, so that the Laissez-faire investment coincides with k^*. As an immediate corollary, capital market imperfections generically cause shareholder preferences to be time-inconsistent and Laissez-faire equilibria to be Pareto inferior.

By demonstrating the existence of Pareto improvements, Theorem 1 rules out the possibility that the Laissez-faire outcome can be rationalized by any common objective function that depends on initial and final shareholders' expected payoffs. The departure from EMH creates an endogenous preference reversal: incumbent and final shareholder preferences over investment, liquidation and hiring decisions $(k, Z, \lambda(\cdot))$ are not aligned. The final shareholders set $\lambda(\cdot)$ to maximize the fundamental return $R(z,\lambda)$, but do not internalize how this decision affects the market return $\hat{R}(z,\lambda)$, and thus the initial shareholders' incentives to liquidate. The initial shareholders decide to liquidate based on the market return $\hat{R}(z,\lambda)$ and do not internalize the final shareholders' loss of surplus. Hence, the hiring policy $\lambda^*(\cdot)$, coupled with liquidation when $\hat{R}(z,\lambda^*(z)) < 0$, is always Pareto inferior, since with liquidation it is dominated by the hiring policy $\hat{\lambda}(\cdot)$ and no liquidation, which guarantee a positive market return to initial shareholders, and fundamental returns in excess of market returns to final shareholders. The time inconsistency result is immediate when liquidation occurs as part of the Laissez-faire equilibrium, but a similar argument applies even without liquidation, as long as initial and final shareholders disagree about their preferred investment and hiring strategies.

Theorem 1 provides a converse to Proposition 2 by showing that for generic departures from EMH, the Laissez-faire equilibrium is Pareto-efficient *only if* it features no inefficient liquidation and satisfies a highly non-generic incentive alignment condition. In other words, corporate short-termism and Pareto-inferior, seemingly irrational investment and hiring decisions by firms should be viewed as the rule rather than the exception in the presence of capital market imperfections.

Proof of Theorem 1. Since $\left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right)$ and $(k^*, \mathbb{R}, \lambda^*(\cdot))$ are both Pareto-efficient, it suffices to show that $(k^{LF}, \bar{Z}, \lambda^*(\cdot))$ is Pareto inferior whenever $\lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$ (preferences are time-inconsistent) and $(k^{LF}, \bar{Z}) \neq (k^*, \mathbb{R})$ (initial shareholder decisions do not maximize expected total surplus). I establish this conclusion in two steps.

First, I show that the Laissez-faire equilibrium must be Pareto inferior if $\bar{Z} \subset \mathbb{R}$ (the firm is liquidated with positive probability). Any Pareto-efficient allocation can be represented as the solution to the decision problem of a Paretian social planner who chooses the initial investment k, the continuation set $Z \subseteq \mathbb{R}$, and the labor intensity $\lambda(\cdot)$, but is subject to the same informational restrictions as initial and final shareholders, and takes as given the market return to capital $\hat{R}(z,\lambda)$. Normalizing the initial shareholders' Pareto weight to 1 and letting $\Gamma \geq 0$ denote the Pareto weight on final shareholders, the social planner solves

$$\max_{k,Z,\lambda(\cdot)} \left\{ \int_Z R^{SP}(z, \lambda(z)) \varphi(z)\, dz \cdot k - C(k) \right\}, \tag{9}$$

where $R^{SP}(z,\lambda) \equiv \Gamma R(z,\lambda) + (1-\Gamma)\, \hat{R}(z,\lambda)$ represents the social planner's expected return on capital, which reduces to a weighted average of the market and fundamental returns, with weights governed by $\Gamma \geq 0$. Under the EMH, $R^{SP}(z,\lambda) = R(z,\lambda) = \hat{R}(z,\lambda)$ regardless of Γ, and the solution to the planner's problem reduces to the surplus-maximizing allocation $(k^*, \mathbb{R}, \lambda^*(\cdot))$. However, when $R(z,\lambda) \neq \hat{R}(z,\lambda)$ and $\Gamma \neq 1$, the social planner's incentives respond to the gap between fundamental and market returns.

The solution to this Paretian social planner's problem implies that it is never Pareto efficient to liquidate the firm when it is under-valued by the market, i.e. for z such that $R(z,\lambda^*(z)) \geq \hat{R}(z,\lambda^*(z))$. However, the Laissez-faire equilibrium only features liquidation for values of z, at which $R(z,\lambda^*(z)) > 0 \geq \hat{R}(z,\lambda^*(z))$, i.e. the firm is strictly under-valued. Therefore if a Laissez-faire equilibrium features a positive probability of liquidation, it cannot be Pareto-efficient.

If $\Gamma \leq 1$, the social planner places can always guarantee a strictly positive return to capital by setting $\lambda(\cdot) = \min\left\{\hat{\lambda}(\cdot), \lambda^*(\cdot)\right\}$. Therefore, rather than protecting incumbent shareholders by liquidating the firm when it is under-valued, the social planner internalizes the initial shareholders' concern for market returns through the hiring policy $\lambda^{SP}(\cdot)$. If $\Gamma > 1$ instead, the planner's returns are always positive if $R(z,\lambda^*(z)) \geq \hat{R}(z,\lambda^*(z))$ and final shareholders are able to acquire the firm at or below fundamental value. If instead $\hat{R}(z,\lambda^*(z)) > R(z,\lambda^*(z))$, the extra cost to final shareholders of purchasing the equity at inflated prices reduces the social planner's return. If $\Gamma - 1$ is sufficiently large, the social planner prefers to liquidate the firm, so as to protect final shareholders from acquiring the firm at inflated prices. However, initial shareholders would never liquidate the firm when it is over-valued in a Laissez-faire equilibrium.

Second, I show that the Laissez-faire equilibrium cannot be Pareto-efficient if $\lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$ and $k^{LF} \neq k^*$, even if the firm is never liquidated. If $k^{LF} \neq k^*$, it is always possible to increase the total surplus by bringing investment closer to k^*. More specifically, setting $k_\varepsilon = k^{LF} + \varepsilon\left(k^* - k^{LF}\right)$ generates a gain of order $\varepsilon > 0$ for final shareholders, at the expense of a reduction in initial shareholders' payoffs that is of order ε^2. In addition, if $\lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$ for $z \in \bar{Z}$, it is always possible to transfer part of these welfare gains to initial shareholders by introducing a small distortion in hiring to boost the share price. Setting $\lambda_\eta(\cdot) = \lambda^*(\cdot) + \eta\left(\hat{\lambda}(\cdot) - \lambda^*(\cdot)\right)$ increases the expected share price, which transfers expected payoffs from final to initial shareholders, of order η, at the expense of a reduction in final shareholder payoffs that is of order η^2. By appropriately scaling the two perturbations, it is then always possible to generate a strict ex ante Pareto improvement.

Conversely, if $(k^{LF}, \bar{Z}) = (k^*, \mathbb{R})$, there is no scope for this kind of "gift exchange" perturbation to yield a Pareto improvement, since the Laissez-faire allocation already maximizes expected total surplus. However, this condition is highly non-generic as it requires $\hat{R}(z,\lambda^*(z)) > 0$ so that the firm is never at risk of liquidation, and $\mathbb{E}\left(\hat{R}\left(z,\lambda^*\left(z\right)\right)\right) = \mathbb{E}\left(R\left(z,\lambda^*\left(z\right)\right)\right)$ so that expected market and fundamental returns to investment coincide ex ante. Both conditions are highly specific to the capital market environment and the structure of market returns, and can typically only hold locally for a very specific set of model primitives.

3 Equity market with noisy information aggregation

So far, I have not attempted to model the sources of capital market imperfections or limits to arbitrage: Theorem 1 shows that shareholder preferences are time-inconsistent and the Laissez-faire equilibrium Pareto-inferior for generic departures from EMH. In this section, focus on a specific source of capital market imperfections which allows me to provide more specific insights into the nature of inefficiencies that are bound to result from corporate short-termism.

To be specific, I depart from the EMH by modeling asset prices as the outcome of noisy information aggregation (NIA) in the spirit of Grossman and Stiglitz (1980) or Hellwig (1980), following the model of Albagli, Hellwig and Tsyvinski (2023, 2024). This model of asset prices provides a specific micro-foundation for a function of market returns $\hat{R}(z,\lambda)$ that is consistent with Assumption 1. It reflects the institutional reality of delegated portfolio management, in which final shareholders acquire equity in a large number of firms (for now, I restrict attention to just one such firm), and delegate their portfolio allocation to two types of traders who purchase equity on their behalf: a unit measure of risk-neutral informed traders, who are indexed by i, and noise traders. I interpret informed traders as hedge funds or specialized investors who have superior information about specific firms and then decide whether to invest in them, up to a limit. The limit on individual exposure captures limits to arbitrage: even when a trader is confident that the firm is under-valued, institutional constraints or capital requirements prevent them from scaling their position without bound. Noise traders instead are interpreted as mutual funds that do not trade actively on firm-specific private information but hold a diversified portfolio across a wide range of different firms, subject to random in and outflows. The randomness in noise trader demand prevents the share price of any individual firm from perfectly revealing the private information of informed traders.

Informed traders observe a private signal $x_i \sim \mathcal{N}(\theta, \beta^{-1})$, which is i.i.d. across traders (conditional on θ), where $\beta > 0$ governs signal precision. After observing x_i, an informed trader submits a price-contingent demand schedule $d_i(\cdot) : \mathbb{R} \to [0,1]$, purchasing the share if and only if the expected fundamental return conditional on their private signal and the price exceeds the price. That is, informed traders cannot short-sell, and can buy at most 1 unit of the shares. An informed trader's strategy is then a function $d\left(x_i, P\right) \in [0,1]$ of the private signal and the price.

Noise traders purchase an exogenous random quantity $\Phi(u)$ of shares, where $u \sim \mathcal{N}(0, \delta^{-1})$ is independent of θ, δ^{-1} is a measure of demand noise, and Φ denotes the standard normal CDF.

The aggregate demand for shares is $D(\theta, P) = \int d(x, P) d\Phi(\sqrt{\beta}(x-\theta)) + \Phi(u)$, where $\Phi(\sqrt{\beta}(x-\theta))$ represents the distribution of private signals x_i conditional on θ. The orders submitted by informed and noise traders are executed at a market-clearing price P such that $D(\theta, P) = 1$.

Let $H(\cdot|x, P)$ denote the traders' posterior cdf of θ, conditional on observing a private signal x, and a market-clearing price P. For given k and $l(z, k)$, a *Noisy Rational Expectations Equilibrium* at stage 2 consists of a demand function $d(x, P)$, a price function $P(\theta, u; k)$, and posterior beliefs $H(\cdot|x, P)$, such that $d(x, P)$ is optimal given the shareholder's beliefs $H(\cdot|x, P)$; $P(\theta, u; k)$ clears the market for all (θ, u) and k; and $H(\cdot|x, P)$ satisfies Bayes' Rule whenever applicable.

3.1 Stage 2: Equilibrium equity price

Fix k and $l(z, k) = \lambda(z) k$. The next proposition characterizes the equilibrium share price in the unique noisy Rational Expectations Equilibrium.

Proposition 3. *Define $z \equiv \theta + 1/\sqrt{\beta} \cdot u$ and $\hat{R}(z, \lambda) \equiv \mathbb{E}[A(\theta)|x = z, z]\lambda^{1-\alpha} - w \cdot \lambda$, and fix an increasing labor schedule $\lambda(z)$. If*

$$\hat{R}(z, \lambda(z)) \equiv \mathbb{E}[A(\theta)|x = z, z]\lambda(z)^{1-\alpha} - w \cdot \lambda(z) \tag{10}$$

is strictly increasing in z whenever $\hat{R}(z, \lambda(z)) > 0$, then the unique equilibrium price under NIA takes the form $P(z; k) = \hat{R}(z, \lambda(z)) \cdot k$ for $z \in \bar{Z}$, and the market signal z satisfies $z|\theta \sim N(\theta, \tau^{-1})$ with precision $\tau = \beta\delta$.

Suppose first that $\lambda > 0$ is fixed and perfectly anticipated by investors in the stock market. Proposition 3 then shows that the unique equilibrium in the asset market has a share price that is equal to $P(z, \lambda; k) = \hat{R}(z, \lambda) k$. Formally, for given λ and k, market-clearing in the asset market requires that P fully reveals the value of $z = \theta + 1/\sqrt{\beta} \cdot u$, or in other words, that z is a sufficient statistic for the information conveyed through the share price. This condition holds because $\hat{R}(z, \lambda)$ is strictly increasing in z. Conditional on θ, z is distributed according to $z \sim \mathcal{N}(\theta, (\beta\delta)^{-1})$, which validates the initial distributional assumption about z with precision $\tau = \beta\delta$.

The key feature of Proposition 3 is the conditional expectation $\mathbb{E}[A(\theta)|x = z, z]$ that is embedded in the market return $\hat{R}(z, \lambda)$. The signal z enters the price formula *twice*: first as the marginal informed trader's private signal $x_i = z$ (with precision β), and second as the public information content of the price itself (with precision $\tau = \beta\delta$). This double-counting is the source of the wedge between $\hat{R}(z,\lambda)$ and the fundamental return $R(z,\lambda) = \mathbb{E}[A(\theta)|z]\lambda^{1-\alpha} - w \cdot \lambda$, or the corresponding share price under the EMH. The share price thus treats the signal z as if it had precision $\beta + \beta\delta$, when in reality its precision is only equal to $\beta\delta$. Hence, the market price is based on an expectation of the marginal return to the investment level k that places a higher weight on the market signal z, relative to its objective information content. Therefore, for high values of z, $\mathbb{E}[A(\theta)|x = z, z] > \mathbb{E}[A(\theta)|z]$ and $\hat{R}(z,\lambda) > R(z,\lambda)$: the market over-values the firm relative to fundamentals. For low values of z, the reverse holds: $\hat{R}(z,\lambda) < R(z,\lambda)$, and the market under-values the firm.

The excess weight the share price attributes to the market signal is fully consistent with Bayesian updating based on private and public information. In the equilibrium representation, the sufficient statistic z represents the private signal of the trader who must be just indifferent between buying or not buying the stock if the market clears, which summarizes the demand for equity shares through noise traders (u) and informed traders (θ). The *identity* of this trader shifts in a systematic way with demand conditions: if informed traders become on average more optimistic (higher θ) or noise trader demand increases (higher u), the marginal informed traders' private signal must also increase to keep the market in equilibrium. To keep this marginal trader indifferent, the market price must increase with z and reveal z publicly to all market participants or outside observers. The expression for $\hat{R}(z,\lambda)$ compounds this second informational effect of z with the market-clearing effect, whereas the fundamental return $R(z,\lambda)$ only contains the informational effect of z.

This logic extends to the case with *information feedback*, i.e. hiring policies $\lambda(z)$ that depend on the information aggregated through the price, with one additional complication: whereas for fixed λ, the function $\hat{R}(z,\lambda)$ is strictly increasing in z and thus invertible, this no longer needs to be the case if λ also varies with z. In particular $\hat{R}(z,\lambda(z))$ may become non-monotone if $\lambda(z)$ exceeds the value $\hat{\lambda}(z)$ that maximizes $\hat{R}(z,\lambda)$, so that $\hat{R}_\lambda(z,\lambda(z)) < 0$. In subsection 3.2, I show that failures of invertibility cannot arise along the equilibrium path, i.e. when $\lambda(z) = \lambda^*(z)$, $\hat{R}(z,\lambda^*(z))$ is strictly increasing in z whenever $\hat{R}(z,\lambda^*(z)) > 0$.

Noisy information aggregation thus causes market returns to depart systematically from fundamental returns, in violation of the EMH. Albagli, Hellwig and Tsyvinski (2024) discuss the asset pricing implications associated with this updating wedge. They show that share price systematically displays a form of "over-confidence" in that the posterior belief embedded in the price attributes too much weight to the market signal. From an ex ante perspective this over-weighting of the sufficient statistic z results in a risk-neutral measure that attributes excess weight on tail risks, on both the upside and the downside. The magnitude of excess weight to tail risks can then be linked to the degree of dispersion in individual return forecasts. As I discuss in Section 6, the model is consistent with several equity return anomalies, such as observed returns to disagreement in equity and bond markets, returns to skewness and interaction effects between both.

3.2 The J-shape of market returns

It remains to identify the precise shape of market returns $\hat{R}(z, \lambda^*(z))$ and confirm that they are invertible w.r.t. z for $z \in \bar{Z}$. I impose the following monotonicity condition:

Assumption 2. $\mathbb{E}[A(\theta)|x = z, z]/\mathbb{E}[A(\theta)|z]$ *is strictly increasing in* z*, with* $\lim_{z \to -\infty} \mathbb{E}[A(\theta)|x = z, z]/\mathbb{E}[A(\theta)|z] = 0$.

The following Lemma then establishes that $\hat{R}(z, \lambda^*(z))$ is "J-shaped":

Lemma 1. *Under Assumption 2,* $\hat{R}(z, \lambda^*(z))$ *is "J-shaped" (Figure 1): there exists a unique* $z^{LF} > -\infty$ *such that* $z \in \bar{Z}$ *and* $\hat{R}(z, \lambda^*(z)) \gtreqless 0$ *if and only if* $z \gtreqless z^{LF}$*. Moreover,* $\hat{R}(z, \lambda^*(z))$ *is strictly increasing in* z *for* $z \in \bar{Z}$ *and* $\lim_{z \to -\infty} \hat{R}(z, \lambda^*(z)) = 0$.

Assumption 2 guarantees that the relative overvaluation of cash flows by the equity market increases continuously with the realized value of z. For sufficiently low values of z, the market return systematically under-estimates the firm's revenue prospects from hiring labor. However, since the expected wage bill is fully priced in, the market return of the firm is eventually negative, meaning that the marginal trader, at low z, believes that the firm's wage bill will exceed its expected revenues –ideally, the marginal trader who prices the equity, and by extension the incumbent shareholders, would prefer to see final shareholders refrain from the hiring they are expected to under-take in the Laissez-faire equilibrium, because it pushes market and fundamental returns in opposite directions.

The market return $\hat{R}(z, \lambda^*(z))$ then inherits the sign pattern described above: it is negative below and positive and increasing above a threshold. This J-shaped pattern is the central feature that distinguishes NIA from the EMH. Moreover, shareholder preferences are time-inconsistent since $\hat{\lambda}(z) = \{(1-\alpha)\mathbb{E}[A(\theta)|x = z, z]/w\}^{1/\alpha}$ differs from $\lambda^*(z) = \{(1-\alpha)\mathbb{E}[A(\theta)|z]/w\}^{1/\alpha}$. Proposition 4 summarizes the departures of the Laissez-faire equilibrium under NIA from the efficient markets benchmark $(k^*, \mathbb{R}, \lambda^*(\cdot))$:

Proposition 4. *Suppose that Assumption 2 holds and equity prices are determined by NIA (equation 10). Then the Laissez-Faire equilibrium* $\left\{k^{LF}, \bar{Z}, \lambda^{LF}(\cdot)\right\}$ *is Pareto-inferior and departs from* $(k^*, \mathbb{R}, \lambda^*(\cdot))$ *as follows:*

(i) The hiring policy maximizes fundamental returns: $\lambda^{LF}(z) = \lambda^*(z)$ *for all* $z > z^{LF}$*.*

(ii) Liquidation is inefficient: the liquidation threshold satisfies $z^{LF} > -\infty$*.*

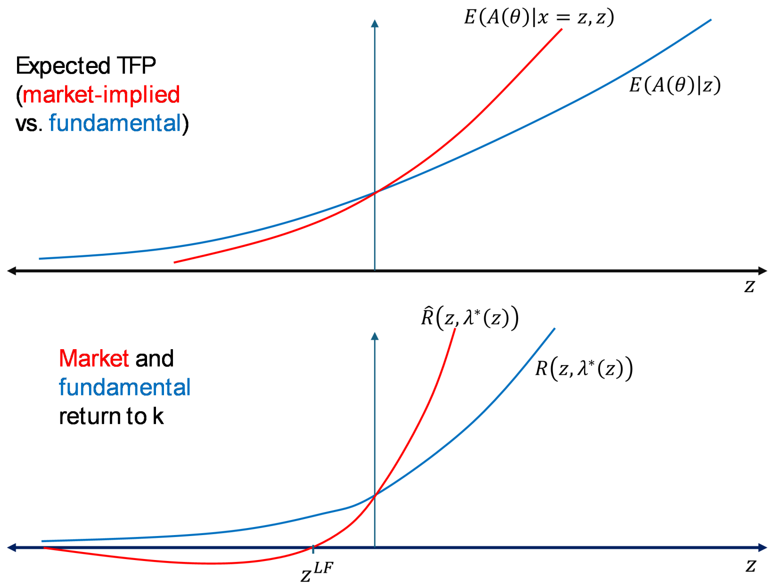

Figure 1:

(iii) Investment is generically inefficient: $k^{LF} \gtreqless k^*$ *if and only if* $\mathbb{E}\left[\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right] \gtreqless \mathbb{E}\left[R\left(z, \lambda^*\left(z\right)\right)\right]$. *Moreover,* $\mathbb{E}\left[\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right] > \mathbb{E}\left[R\left(z, \lambda^*\left(z\right)\right)\right]$, *whenever* $A(\cdot)$ *is symmetric or dominated by upside risk, or* α *is sufficiently low, and* $\lim_{\alpha\to 0} \frac{C'(k^{LF})}{C'(k^*)} = \infty$.

Hence the Laissez-Faire equilibrium under NIA leads to inefficient (over-)investment and excessive liquidation. For sufficiently low z it becomes impossible to sell the firm at a positive price. Initial shareholders then prefer to liquidate their investment, even though the fundamental returns to the investment remain strictly positive. Theorem 1 further implies that the equilibrium with NIA must be Pareto inferior. In addition, the gap between market and fundamental returns, along with the liquidation option, leads to a departure of ex ante investment incentives from the social optimum. The initial shareholders' first-order condition for investment can be written as

$$\frac{C'(k^{LF})}{C'(k^*)} = \frac{\mathbb{E}\left[\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right]}{\mathbb{E}[R(z, \lambda^*\left(z\right))]} > \frac{\mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right]}{\mathbb{E}[R(z, \lambda^*\left(z\right))]}. \tag{11}$$

Therefore, the equilibrium features over-investment, whenever $\mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right] \geq \mathbb{E}[R(z, \lambda^*\left(z\right))]$. The difference between $\mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right]$ and $\mathbb{E}[R(z, \lambda^*\left(z\right))]$ in turn depends on the shape of $A\left(\cdot\right)$ and $\lambda^{LF}\left(z\right)$, as well as the labor share $1-\alpha$. Notice first that k^{LF} converges to k^*as $\alpha \to 1$ if $A\left(\cdot\right)$ is symmetric. As $\alpha \to 1$, $\mathbb{E}\left[\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right] \to \mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right]$, as the incentive to over-hire disappears as the labor share goes to zero and therefore there is also no longer any incentive to liquidate, and $\mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right] = \mathbb{E}[R(z, \lambda^*\left(z\right))]$ if the return function is symmetric as over-valuation of returns on the upside is exactly offset by undervaluation of returns on the downside. For $\alpha < 1$, the information feedback and the liquidation option together imply that $\mathbb{E}\left[\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right] > \mathbb{E}\left[\hat{R}(z, \lambda^*\left(z\right))\right] > \mathbb{E}[R(z, \lambda^*\left(z\right))]$, which causes the firm to over-invest, even with symmetric $A\left(\cdot\right)$. This over-investment is further reinforced when $A(\cdot)$ shifts towards upside risk. When $A(\cdot)$ instead is dominated by downside risk, return asymmetry and information feedback push in opposite directions, and the Laissez-Faire equilibrium features under-investment, if the exogenous downside risk in $A\left(\cdot\right)$ more than offsets the endogenous upside risk stemming from both the liquidation option and the information feedback.[4]

Lastly, the endogenous upside risk from the liquidation option and the information feedback becomes arbitrarily strong when the labor share is close to 1. In the limit as $\alpha \to 0$, the firm is liquidated whenever it is under-valued, but the over-valuation on the upside becomes so extreme that it induces arbitrarily large over-investment relative to the efficient markets benchmark. Even small departures from the EMH can therefore cause arbitrarily large investment distortions, if market returns are sufficiently sensitive to the firm's expected hiring policy.

3.3 Discussion of model assumptions and implications

Several features of the model deserve further discussion.

Investment and liquidation. To understand the mechanism more clearly, it is instructive to decompose the market return:

$$\hat{R}\left(z, \lambda^*(z)\right) = R(z, \lambda^*(z)) \left\{ \frac{\mathbb{E}[A(\theta)|x = z, z]}{\mathbb{E}[A(\theta)|z]} - (1 - \alpha) \right\}. \tag{12}$$

The term $\mathbb{E}[A(\theta)|x = z, z]/\mathbb{E}[A(\theta)|z]$ represents the extent to which the marginal informed trader over- or under-estimates the firm's expected marginal revenue to labor. The term $1 - \alpha$ represents the compensation to labor. Assumption 2 insures that for low z, the first term is small (the marginal trader is pessimistic) relative to the second, hence $\hat{R}\left(z, \lambda^*(z)\right) < 0$.

When $\hat{R}\left(z, \lambda^*(z)\right) < 0$, incumbent shareholders prefer to liquidate, even if $R(z, \lambda^*(z)) > 0$. Liquidation is inefficient: positive-NPV projects are abandoned only because the market does not adequately compensate shareholders for the labor that would need to be employed after the sale. This inefficient liquidation region guarantees that the Laissez-faire equilibrium is Pareto inferior.

[4] Albagli, Hellwig and Tsyvinski (2023, 2024) discuss at length the implications of NIA and payoff asymmetry for asset returns and investment incentives in a version of the model without the stage 3 labor decision. Here the same insights apply but asymmetries are partly endogenous to the value of the liquidation option and the information feedback through labor choices. The overall return may still be dominated by upside risk and yield over-investment even if $A\left(\cdot\right)$ has more downside uncertainty, if these endogenous sources of upside risk are sufficiently strong.

The investment first-order condition yields $k^{LF} \gtreqless k^*$ if and only if $\mathbb{E}[\max\left\{0, \hat{R}\left(z, \lambda^*(z)\right)\right\}] \gtreqless \mathbb{E}\left[R\left(z, \lambda^*(z)\right)\right]$. As discussed in connection with Proposition 4, it is possible to have cases with $k^{LF} = k^*$, if exogenous downside risks from the technology exactly offset endogenous upside risks from the liquidation option and the information feedback. In that case, inefficient liquidation is the sole source of inefficiency. But these cases require knife-edge conditions on the fundamentals.

Inefficiencies in liquidation or investment (or both), i.e. having $\left(k^{LF}, z^{LF}\right) \neq (k^*, \mathbb{R})$, are necessary and sufficient to establish the preference reversal argument underlying Theorem 1. With NIA, the liquidation option guarantees that the Laissez-faire equilibrium is Pareto inferior even in the absence of conditions guaranteeing inefficient investment for all primitives.

Lack of commitment and preference reversal over hiring. Final shareholders condition on the public signal z and choose labor to maximize the fundamental return: $\lambda^{LF}(\cdot) = \lambda^*(\cdot)$. However, efficient hiring ex post is not synonymous with maximizing the stock price, which requires setting $\hat{\lambda}(\cdot)$. If final shareholders could commit upfront to a hiring policy that maximizes the market returns for $z \leq z^{LF}$, then the latter remains strictly positive and the firm is never liquidated. However, this alternative strategy is not credible: ex post the final shareholders gain from setting $\lambda^{LF}(\cdot) = \lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$, resulting in strictly better firm performance but lowering the market returns below 0. Anticipating that a hiring policy that preserves positive market returns ex ante is not credible, initial shareholders then strictly prefer to liquidate when $z \leq z^{LF}$.

The fact that the firm's expected marginal revenue to labor is not correctly priced into market returns is the source of time inconsistency of shareholder preferences and key for the Pareto inefficiency result: if we had $\lambda^*(\cdot) = \hat{\lambda}(\cdot)$, the hiring decision remains time-consistent, the initial shareholders never liquidate and the firm decisions maximize initial shareholder value. It is only because of the preference reversal over hiring that there is potential for disagreement over investment and liquidation decisions. The fact that final shareholders can't commit in advance to a hiring policy that maximizes the stock price is key to showing that the endogenous preference reversal causes Laissez-faire outcomes to be Pareto-inferior. Kydland and Prescott (1977) emphasize that discretionary economic policy is generally inconsistent with rational planning when economic agents make forward-looking decisions. Here, the market imperfection leads to a similar failure of efficient, time-consistent planning by firms when shareholders do not fully internalize the impact of their decisions on future dividends or past stock prices.

The Nesting of the EMH, scalability of labor, and the magnitude of distortions. The model nests the EMH as a limiting case. As the noise in the information aggregation process vanishes — formally, as $\delta \to \infty$ (noise trader demand variance goes to zero) — the price signal becomes perfectly informative and $\hat{R}(z, \lambda^*(z)) \to R(z, \lambda^*(z))$, recovering the EMH. The same applies if $\beta \to 0$. In this limit, all the distortions identified above disappear: investment converges to the first-best, liquidation vanishes, and aggregate stock prices reflect fundamental values.

However, the convergence to efficiency as $\delta \to \infty$ is not uniform in α. For any $\alpha < 1$, the inefficiency vanishes as $\delta \to \infty$. But for any $\delta < \infty$, the inefficiency is bounded away from zero regardless of how small the noise is. Moreover, the magnitude of the inefficiency is governed by α. As $\alpha \to 1$, the labor cost term in the decomposition (12) goes to zero, removing the scope for inefficient liquidation. Conversely, as $\alpha \to 0$, the labor share approaches one, the liquidation option is most valuable, and the investment distortion becomes arbitrarily large.

This interaction between NIA and scalability of labor is what makes the model's implications quantitatively significant: if the sensitivity of hiring to expected fundamental returns, α^{-1}, is sufficiently large, even a minor capital market imperfections can cause large economic inefficiencies.

The role of information feedback I assumed that final shareholders hire labor based on the information content of the share price, which is summarized by z. The preference reversal stems from the fact that the marginal shareholder overweighs z relative to the final shareholders; this source of disagreement is at the core of the discrepancy between fundamentals and market valuations. The information feedback, i.e. the fact that final shareholders condition their decisions on z, is key for the time inconsistency result: if instead the final shareholders perfectly observe θ before making their hiring decision, their hiring choice, as a function of θ, satisfies $\lambda^*(\theta) = \{(1-\alpha)\, A(\theta)/w\}^{1/\alpha}$ and maximizes both the fundamental return $R(z, \lambda^*(\cdot)) = \alpha \left(\frac{1-\alpha}{w}\right)^{1-\alpha} \mathbb{E}\left[A(\theta)^{1/\alpha}|z\right]$ and the market return $\hat{R}\,(z, \lambda^*(\cdot)) = \alpha \left(\frac{1-\alpha}{w}\right)^{1-\alpha} \mathbb{E}\left[A(\theta)^{1/\alpha}|x = z, z\right]$. As a result, the firm remains mis-priced, but shareholder preferences remain time-consistent and market returns remain strictly positive, ruling out inefficient liquidation and coordinating shareholders on maximizing initial shareholder value –akin to the the discussion of the $\lambda^*(\cdot) = \hat{\lambda}\,(\cdot)$ condition in Proposition 1.

This last observation, however, is knife-edge: as long as θ remains uncertain and final shareholders use the market signal z to update their beliefs about the firm's marginal product of labor, noisy information aggregation breaks the time consistency of shareholder preferences with regards to the firm's hiring decision, causing the Laissez-faire equilibrium to become Pareto inferior.

The role of noisy information aggregation. To simplify the analysis, I assumed that preferences are homogeneous within shareholder generations, and I incorporated informed and noise trading on the side of final shareholders. The homogeneous preference assumption abstracts from heterogeneity and conflicts of interests within shareholder generations.[5] Moreover, placing informed and noise trading on the same side allows me to abstract from the differences in returns these two groups experience, since the excess returns from trading on private information exactly cancel out the trading losses from stochastic liquidity shocks, from the perspective of final shareholders. It is possible, but beyond the scope of the present analysis, to allow initial shareholders to sell because of stochastic liquidity needs or private information, or to allow for heterogeneity in access to information or exposure to liquidity shocks. This would alter the perceived market returns for different types of traders and introduce additional conflicts of interest within generations, which would only reinforce the overarching conclusion that capital market imperfections become a source of preference reversals and conflict of interest between shareholders.

Assumptions about firm decisions and financial market imperfections. I focused on a stylized model of initial investment and subsequent hiring decisions to capture the notion of sequential decisions within the same firm. The basic principles, however, apply to many other decision margins: initial project choice (i.e. in which of several technologies to invest), investment duration, capital structure, or any other dynamic firm decision that affects the risk embedded in the firm's capital. Similarly, I focused on hiring as the final shareholders' decision margin, but this interpretation should not be taken too literally.

Theorem 1 showed that capital market imperfections cause generic preference reversals and Pareto-inferior equilibrium outcomes, provided that the reduced form fundamental and market returns to investment that describe the interaction between initial and final shareholders share the properties highlighted by the theorem. Preference reversals arise generically whenever the financial market imperfection introduces a mis-alignment between the initial and final shareholders' preferred investment, liquidation and hiring decisions, in which case there exists a possibility for a Pareto improvement via a "gift exchange" perturbation that marginally tweaks shareholder decisions in the direction favored by the other party.

Rationality and the source of time-inconsistency. Finally, a crucial feature of the model is that all agents are fully rational and forward-looking. Initial shareholders correctly anticipate the equilibrium stock price and the final shareholders' hiring policy. Final shareholders maximize their expected profits conditional on all available information. Informed traders correctly process their private signals and respond optimally to prices. The inefficiency that arises is not due to any

[5] See Jackson and Yariv (2015), Caramp et al. (2024), and Davila and Schaab (2025) for results on the role of static heterogeneity for time-inconsistent preference aggregation.

agent making a mistake or failing to process information correctly; it is a consequence of the market equilibrium itself failing to aggregate information in a way that aligns private incentives with social surplus. As the comparison of Proposition 2 and Theorem 1 makes clear, departures from EMH are the *only* source of time-inconsistent shareholder preferences and Pareto inferior allocations.

This distinguishes the present theory from accounts of corporate short-termism that emphasize bounded rationality or conflicts of interest between shareholders and managers. The model does not assume that shareholders are irrationally exuberant, that they extrapolate past returns, or that they fail to discount the future correctly. Nor does corporate short-termism arise for incentive reasons as the outcome of a bilaterally efficient incentive contract between the shareholders and CEO of the company, against the backdrop of fully efficient financial markets.

Instead, it is the capital market imperfections that cause the mis-alignment of valuations and incentives across shareholder generations and between private and social returns to investment and hiring. The preference reversal between initial and final shareholders coupled with a lack of commitment to future strategies leave the firm in a position in which at each stage its controlling shareholders maximize short-term market returns without internalizing the impact of their decisions on future dividends or prior stock prices –the corporate analogue of temptation-driven consumer preferences (Gul and Pesendorfer 2001, 2005).

The Gul-Pesendorfer interpretation of time-inconsistency differs from the Kydland-Prescott interpretation primarily in the definition of the policy objective. In the Gul-Pesendorfer view, capital market imperfections cause initial shareholders to deviate from the firm's long-term objective, which is given by its total surplus. In the Kydland-Prescott view, the firm maximizes initial shareholder value, and time-inconsistency emerges from the lack of commitment to future firm policies. In some cases these two views have similar implications for the time inconsistency problem of the firm (i.e. with regards to inefficient liquidation), but in other cases they have opposite implications (i.e. with regards to investment and hiring policies). Both perspectives on the time consistency problem are useful: The Kydland-Prescott view represents a corporate finance perspective, in which initial shareholders look to implement a long-term strategy that maximizes the firm's stock market value. The Gul-Pesendorfer view instead represents a regulator who looks to maximize total surplus. From this regulator's perspective, if equity markets cannot be counted upon to align shareholder preferences with total surplus, then shareholders cannot be counted upon to implement socially desirable outcomes. Quite the contrary: whereas the controlling shareholders act as the social planner's allies under the EMH, they become her primary opponent once the EMH no longer holds.

4 General equilibrium with a continuum of firms

I now extend the analysis to a general equilibrium setting with a continuum of firms. This extension is important for two reasons. First, it allows me to examine when the firm-level distortions identified in the partial equilibrium analysis survive if stock prices are determined endogenously as part of a general equilibrium. Second, it reveals a new layer of inefficiency: an aggregate investment wedge through which each firm's attempt to maximize its market return is self-defeating in the aggregate.

The general equilibrium model builds on Albagli, Hellwig and Tsyvinski (2023). There is a measure 1 continuum of firms indexed by $i \in \mathcal{I} \equiv [0,1]$. There are representative incumbent and final shareholders who hold diversified portfolios across all firms. Firms are ex ante identical w.r.t. firm-specific uncertainty: the fundamental shocks θ_i and information signals z_i are i.i.d. across firms. There is no aggregate uncertainty. There may be ex ante differences across firms, i.e. cash-flows $\Pi_i(\theta, k, l)$, investment costs $C_i(k)$, share prices $P_i(z, \lambda, k_i)$ or equivalently fundamental and market returns $R_i(z, \lambda)$ and $\hat{R}_i(z, \lambda)$ may all vary across i, as do the investment, liquidation and hiring policies $\left(k_i, \bar{Z}_i, \lambda_i(\cdot)\right)$. Assumption 1 holds for all i; in addition I assume that $\frac{d \ln C_i(k)}{d \ln k} = \kappa > 1$ for all $i \in \mathcal{I}$ in order to link aggregate investment to aggregate dividends. To simplify the exposition, I assume that $\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*(z)\right)\right] / \mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$ is either constant across all i (homogeneous capital market imperfections) or strictly increasing (heterogeneous capital market imperfections).[6]

4.1 Set-up

Representative final shareholders. The representative final shareholder's preferences over aggregate consumption D and the aggregate share price are given by

$$u(sD) - sP, \tag{13}$$

where $D = \int_{i \in \mathcal{I}} k_i \int_{Z_i} R_i(z_i, \lambda_i) \varphi(z) dz\, di$ is aggregate dividends, $P = \int_{i \in \mathcal{I}} \int_{Z_i} P_i(z_i, \lambda_i, k_i) \varphi(z) dz\, di$ is the aggregate stock price and $s \in \mathbb{R}$ the share of the aggregate equity portfolio they purchase (in equilibrium, $s = 1$) . I assume that the utility function u is strictly increasing and concave with $\chi(D) \equiv -Du''(D)/u'(D) < 1$. The representative final shareholder's first-order condition for s results in the following first-order condition for the aggregate stock price:

$$P = Du'(D) \tag{14}$$

[6] Since it is without loss of generality to rank-order firms by their value of $\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*(z)\right)\right] / \mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$, this simplification only assumes that heterogeneity of capital market imperfections is generic in that w.p. 1, no two firms have the same value of $\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*(z)\right)\right] / \mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$.

In other words, the representative final shareholder discounts cash flows at the aggregate intertemporal marginal rate of substitution $u'(D)$. Since $\chi(D) < 1$, P is increasing in D.

As before, final shareholders acquire the portfolio of firms through informed traders, or hedge funds, and noise traders, or mutual funds, capturing in a simplified way the reality of delegated portfolio management. For each firm, there is a unit measure of risk-neutral informed traders who can each purchase up to one share in the firm, or alternatively lend their funds to another informed trader at a market-determined interest rate Q. In addition, a random fraction of equity shares is purchased by noise traders or mutual funds. I interpret the random component of flows at the firm level as resulting mechanically from the flow of orders into and out of a collection of mutual funds who are all diversified across many firms, and who do not trade on private information.

The market for each firm is modeled as in Section 3; distributional assumptions at the level of each individual firm are identical to before and not repeated here. I assume that the final shareholders allocate their resources between mutual funds and hedge funds in a manner that covers the aggregate equity purchases of hedge funds and mutual funds, which are completely deterministic, given that firm-level uncertainty is fully diversified for both types of investors. The lending market between informed traders allows to reallocate funds towards informed investors that are the most inclined to purchasing equity shares, with the more pessimistic informed investors acting as lenders. This formulation embeds the model of equity markets at the firm-level into a broader general-equilibrium model with a continuum of firms and no aggregate uncertainty.

Laissez-faire equilibrium. Following the same characterization steps as in Section 3, the market-clearing equity price for each firm is characterized by the indifference condition for the marginal informed investor, which yields $P_i(z,k) = Q^{-1} \cdot \hat{R}_i(z, \lambda_i(z)) \cdot k$. Aggregating across all firms, the aggregate stock price P and aggregate dividends D are given by

$$P = \frac{1}{Q} \int_{i \in \mathcal{I}} k_i \int_{Z_i} \hat{R}_i(z, \lambda_i(z)) \varphi(z) dz di \quad \text{and} \quad D = \int_{i \in \mathcal{I}} k_i \int_{Z_i} R_i(z, \lambda_i(z)) \varphi(z) dz di. \tag{15}$$

Combining equations (14) and (15) leads to the following expression for the market-clearing lending rate Q as a function of the continuation set $\bar{Z}$ and the firm-level stock market imperfections:

$$Q = \frac{1}{u'(D)} \frac{\int_{i \in \mathcal{I}} k_i \int_{Z_i} \hat{R}_i(z, \lambda_i(z)) \varphi(z) dz di}{\int_{i \in \mathcal{I}} k_i \int_{Z_i} R_i(z, \lambda_i(z)) \varphi(z) dz di}. \tag{16}$$

A *Laissez-faire equilibrium* consists of $\left\{k_i^{LF}, \bar{Z}_i, \lambda_i^{LF}(\cdot); D^{LF}, P^{LF}, Q^{LF}\right\}$ with hiring policy $\lambda_i^{LF}(z) = \lambda_i^*(z) \equiv \arg\max_{\lambda \geq 0} R_i(z, \lambda)$ for $z \in \bar{Z}_i$, continuation set $\bar{Z}_i = \left\{z \in \mathbb{R} : \hat{R}_i(z, \lambda_i^*(z)) > 0\right\}$, investment satisfying $C_i'(k_i^{LF}) = \mathbb{E}\left[\max\left\{0, \hat{R}_i\left(z, \lambda_i^*(z)\right)\right\}\right]/Q^{LF}$ and $\left\{Q^{LF}, P^{LF}, D^{LF}\right\}$ satisfying equations (15) and (16). This equilibrium definition augments the firm-level optimality conditions for hiring, liquidation and investment, the latter adjusted by the market lending rate, with the representative final shareholder's Euler equation and the aggregation equations for stock prices, dividends, and the lending rate Q.

Efficient Markets Allocation

In general equilibrium, aggregate surplus is given by the equally weighted sum of the representative initial and final shareholder's payoffs $u(D) - \int_i C_i(k_i)\, di$, where $D = \int_i k_i \int_{\bar{Z}_i} R_i(z, \lambda_i(z))\varphi(z) dz di$. As in Section 2, the surplus-maximizing hiring policy sets $\lambda_i(\cdot) = \lambda_i^*(\cdot)$, the firm is never liquidated ($Z_i = \mathbb{R}$), and investment satisfies $C_i'(k_i^*) = u'(D^*)\mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$, with aggregate dividends given by $D^* = \int_{i\in\mathcal{I}} k_i^* \mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right] di$.

As in Section 2, the Laissez-faire equilibrium maximizes aggregate surplus, whenever the EMH is satisfied at the firm level, or $\hat{R}_i(z, \lambda) = R_i(z, \lambda)$ for all $i \in \mathcal{I}$ and (z, λ): If $\lambda_i^{LF}(\cdot) = \lambda_i^*(\cdot)$, $\hat{R}_i\left(z, \lambda_i^*(z)\right) = R_i\left(z, \lambda_i^*(z)\right) > 0$ and the firm is never liquidated. Substituting $\hat{R}_i(z, \lambda) = R_i(z, \lambda)$ into equation (16) and the optimality condition for investment yields $Q^{LF} = 1/u'\left(D^{LF}\right)$ and $C'(k_i^{LF}) = u'\left(D^{LF}\right)\mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$, and therefore $k_i^{LF} = k_i^*$ and $D^{LF} = D^*$, so that the market discounts returns at the same rate as the representative final shareholder and investment maximizes aggregate surplus. Hence, under the EMH the Laissez-faire equilibrium is equal to $\left\{k_i^{LF}, \bar{Z}_i, \lambda_i^{LF}(\cdot), D^{LF}\right\} = \left\{k_i^*, \mathbb{R}, \lambda_i^*(\cdot), D^*\right\}$, with $Q^{LF} = 1/u'(D^*)$ and $P^{LF} = D^* u'(D^*)$.

4.2 Time consistency, intra-firm and aggregate Pareto efficiency

In general equilibrium, capital market imperfections can cause inefficiencies both at the level of individual firms, i.e. preventing initial and final shareholders from coordinating on a time-consistent optimal plan, and at the aggregate level, resulting in inefficient aggregate savings and inefficient cross-sectional allocation of investment. As before, shareholder preferences for firm i are time-inconsistent whenever $\lambda_i^*(z) \neq \hat{\lambda}_i(z) \equiv \arg\max_{\lambda\geq 0} \hat{R}_i(z, \lambda)$ for some $z \in \mathbb{R}$. A plan $(k_i, Z_i, \lambda_i(\cdot))$ for firm i is *intra-firm Pareto-efficient* if there exists no alternative plan $(k_i', Z_i', \lambda_i'(\cdot)) \neq (k_i, Z_i, \lambda_i(\cdot))$ that is weakly preferred by initial and final shareholders in firm i, for given aggregate variables (D, P, Q), and strictly preferred by one group of shareholders:

$$\int_{Z_i'} Q^{-1}\hat{R}_i\left(z, \lambda_i'(z)\right)\varphi(z)\, dz \cdot k_i' - C_i(k_i') \geq \int_{Z_i} Q^{-1}\hat{R}_i\left(z, \lambda_i(z)\right)\varphi(z)\, dz \cdot k_i - C_i(k_i) \tag{17}$$

$$\int_{Z_i'} \left(u'(D)\, R\left(z, \lambda_i'(z)\right) - Q^{-1} \hat{R}_i\left(z, \lambda_i'(z)\right)\right) \varphi(z)\ dz \cdot k_i' \geq \int_{Z_i} \left(u'(D)\, R\left(z, \lambda_i(z)\right) - Q^{-1} \hat{R}_i\left(z, \lambda_i(z)\right)\right) \varphi(z)\ dz \cdot k_i \tag{18}$$

with one inequality being strict. The results of Sections 2 and 3 directly carry over to intra-firm Pareto inefficiency in the general equilibrium model after discounting fundamental returns by $u'(D)$ and market returns by Q^{-1}. If shareholder preferences in firm i are time-consistent then the firm's equilibrium strategy $\left(k_i^{LF}, \bar{Z}_i\right)$ maximizes the initial shareholders' expected payoffs and is equal to $\left(\hat{k}_i, \mathbb{R}, \lambda_i^*(\cdot)\right)$, where $C_i\left(\hat{k}_i\right) = Q^{-1}\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*(z)\right)\right]$, whereas if shareholder preferences in firm i are time-inconsistent, the firm's equilibrium strategy $\left(k_i^{LF}, \bar{Z}_i, \lambda_i^*(\cdot)\right)$ is intra-firm Pareto inferior unless $\left(k_i^{LF}, \bar{Z}_i\right) = \left(k_i^{\dagger}, \mathbb{R}\right)$ where $C_i'\left(k_i^{\dagger}\right) = u'(D)\,\mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right]$, or equivalently $u'(D)\,\mathbb{E}\left[R_i\left(z, \lambda_i^*(z)\right)\right] = Q^{-1}\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*(z)\right)\right]$. This knife-edge condition requires that the aggregate wedge between the market discount factor Q^{-1} and the representative final shareholders' inter-temporal marginal rate of substitution $u'(D)$ exactly offsets the gap between expected market and fundamental returns. With heterogeneous capital market imperfections, intra-firm Pareto efficiency holds only for a measure 0 of firms.

An allocation $\left(k_i, Z_i, \lambda_i(\cdot)\right)_{i\in\mathcal{I}}$ for all firms $i \in \mathcal{I}$ is *intra-firm Pareto-efficient* if the plans $\left(k_i, Z_i, \lambda_i(\cdot)\right)$ are intra-firm Pareto-efficient for all $i \in \mathcal{I}$. An allocation $\left(k_i, Z_i, \lambda_i(\cdot)\right)_{i\in\mathcal{I}}$ is *aggregate Pareto-efficient* if there exists no alternative allocation $\left(k_i', Z_i', \lambda_i'(\cdot)\right)_{i\in\mathcal{I}}$ that is weakly preferred by the representative initial and final shareholder, i.e.

$$D'u'(D') - \int_{i\in\mathcal{I}} C_i\left(k_i'\right) di \geq Du'(D) - \int_{i\in\mathcal{I}} C_i\left(k_i\right) di \quad \text{and} \quad D' \geq D \tag{19}$$

with one strict inequality, where D and D' are derived from $\left(k_i, Z_i, \lambda_i(\cdot)\right)_{i\in\mathcal{I}}$ and $\left(k_i', Z_i', \lambda_i'(\cdot)\right)_{i\in\mathcal{I}}$ using equation (15). This definition exploits the fact that initial and final shareholders only care about the aggregate stock price, not about the prices in individual firms. Capital market imperfections at the firm level thus affect aggregate Pareto efficiency only indirectly, through their impact on investment and aggregate dividends. The key distinction between intra-firm and aggregate Pareto efficiency is that the former takes the aggregate stock price, dividend and interest rate (D, P, Q) as given, while the latter internalizes how the decisions of each firm impact these aggregate conditions.

Two observations follow immediately. First, aggregate Pareto efficiency is invariant to firm-level departures from the EMH, i.e. the precise form of the market returns $\hat{R}_i(z, \lambda)$. In other words, a Paretian social planner only cares about aggregate earnings and the aggregate stock price, not the extent to which individual firms are over- or under-valued. Second, because P and $u(D) - P$ are both increasing in D, the social planner's problem can be decomposed into a choice of liquidation set Z_i and hiring policy $\lambda_i(\cdot)$ that maximize each firm's fundamental return $R_i \equiv \int_{Z_i} R_i(z, \lambda_i(z)) \varphi(z)\, dz$, a Pareto efficient allocation of investment k_i for given fundamental returns R_i and aggregate investment expenditure $\mathcal{C} = \int_{i\in\mathcal{I}} C_i(k_i)\, di$, and an efficient choice of $\mathcal{C}$ that trades off between the initial shareholders' cost of investment $\mathcal{C}$ and initial and final shareholder benefits through higher aggregate dividends D.

Therefore $(Z_i, \lambda_i(\cdot))_{i\in\mathcal{I}} = (\mathbb{R}, \lambda_i^*(\cdot))_{i\in\mathcal{I}}$ for any aggregate Pareto-efficient allocation, i.e. hiring policies are undistorted and no firm is ever liquidated. In addition, the allocation of investment across firms is aggregate Pareto-efficient for a given distribution of fundamental returns R_i whenever $C_i'(k_i)/R_i$ is equalized across $i \in \mathcal{I}$ and bounded below by $u'(D)(1-\chi(D))$, where $D = \int_{i\in\mathcal{I}} k_i R_i di$. When $C_i'(k_i)/R_i = u'(D)(1-\chi(D))$ for all $i \in \mathcal{I}$, the allocation maximizes the representative initial shareholder's payoffs. With $C_i'(k_i)/R_i = u'(D)$ for all $i \in \mathcal{I}$, the allocation maximizes aggregate total surplus, and for even higher values of $C_i'(k_i)/R_i$, the aggregate Pareto-efficient allocation places even higher weight on final shareholders who gain from larger aggregate investment and corporate earnings.

4.3 Firm-level imperfections cause aggregate inefficiency

Theorem 2 provides necessary and sufficient conditions for intra-firm and aggregate Pareto inefficiency in general equilibrium.

Theorem 2. *Under Assumption 1 and if $\lambda_i^*(\cdot) \neq \hat{\lambda}_i(\cdot)$ for all $i \in \mathcal{I}$, the following are equivalent:*

1. *The Laissez-faire equilibrium is aggregate Pareto efficient.*
2. *The Laissez-faire equilibrium is intra-firm Pareto efficient.*
3. *$\bar{Z}_i = 0$ for all $i \in \mathcal{I}$, i.e. no firm is liquidated in equilibrium, and $\mathbb{E}\left[\hat{R}_i(z, \lambda_i^*(z))\right] / \mathbb{E}\left[R_i(z, \lambda_i^*(z))\right]$ is constant across $i \in \mathcal{I}$, i.e. capital market imperfections are homogeneous.*

Theorem 2 shows that the necessary and sufficient conditions for intra-firm and aggregate Pareto efficiency coincide and they remain very special. Besides requiring no inefficient liquidation, they also require homogeneous capital market imperfections. The latter condition allows $\mathbb{E}\left[\hat{R}_i(z, \lambda_i^*(z))\right]$ to differ from $\mathbb{E}\left[R_i(z, \lambda_i^*(z))\right]$ as long as their ratio is the same for all firms – for example it holds if all firms are ex ante identical.

Consider first the case in which both of these conditions are satisfied. In this case, the equilibrium is aggregate Pareto efficient and maximizes total surplus: it features efficient hiring and liquidation choices, allocates capital efficiently across firms, and the scale of investment efficiently trades off the marginal investment costs against the marginal corporate earnings increases. Moreover, the interest rate Q adjusts so that the Laissez-faire equilibrium is also intra-firm Pareto efficient, i.e.

for all firms the wedge between Q and $1/u'(D)$ exactly offsets the wedge between market and fundamental returns: the aggregate investment wedge between $1/Q$ and the representative final shareholder's intertemporal marginal rate of substitution $u'(D)$ exactly counter-balances the firm-level investment distortion in each firm's investment decision, so that the Laissez-faire equilibrium is both intra-firm and aggregate Pareto efficient.

With liquidation or heterogeneous capital market imperfections, there is scope for aggregate Pareto improvements by undoing inefficient liquidation or capital misallocation. These increases in aggregate returns generate a strict aggregate Pareto improvement without the need for hiring distortions to transfer surplus gains from final to initial shareholders. Moreover, inefficient liquidation can never be intra-firm Pareto efficient, and heterogeneous capital market imperfections imply that intra-firm Pareto efficiency fails for (almost) all firms in the cross-section: the endogenous adjustment of Q to over- or under-valuation of stocks only corrects the average incentive to over- or under-invest, but cannot undo the capital misallocation.

As in Albagli, Hellwig and Tsyvinski (2023), the firm-level investment incentives for given Q interact with aggregate investment efficiency through the investment wedge $Qu'(D)$: initial shareholders collectively do not internalize the impact of their decisions on the market interest rate Q, but each firm takes Q as given and maximizes its own market value. The aggregate stock price instead is tied to aggregate dividends through the final shareholders' Euler equation. Hence each firm's attempt to boost its own share price is self-defeating in general equilibrium. Without liquidation and cross-sectional misallocation of capital, the aggregate investment wedge exactly offsets the common incentive to over- or under-invest, restoring firm-level and aggregate surplus maximization. But if firm-level attempts to boost stock prices instead cause inefficiencies, they reduce aggregate dividends and stock prices below the surplus-maximizing level.

The next Proposition summarizes the implications of heterogeneous capital market imperfections and/or inefficient liquidation for general equilibrium allocations. Recall that the model with noisy information aggregation always features liquidation, so Theorem 2 implies that the equilibrium is intra-firm and aggregate Pareto inferior.

Proposition 5. *Suppose that Assumption 1 holds. If the Laissez-faire equilibrium is Pareto-inferior, it features aggregate under-investment, i.e.* $\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right)di < \int_{i\in\mathcal{I}} C_i\left(k_i^*\right)di$*, depressed stock prices, i.e.* $P^{LF} < D^* \cdot u'(D^*)$*, and aggregate savings distortions, i.e.* $Q^{LF} \gtreqless 1/u'(D^{LF})$ *if and only if* $\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} \hat{R}_i(z,\lambda_i^*(z))\varphi(z)dzdi \gtreqless \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i(z,\lambda_i^*(z))\varphi(z)dzdi$*.*

The most striking aspect of Proposition 5 is the reversal of the investment distortion between partial and general equilibrium. In partial equilibrium, noisy information aggregation generates excess investment if there is sufficient upside risk: incumbent shareholders over-invest to take advantage of the over-valuation for high-z realizations. In general equilibrium, the opposite holds: the aggregate investment expenditures are below the first-best value.[7] When all firms simultaneously over-invest relative to a given market discount factor $1/Q$ and liquidate with positive probability, aggregate dividends D^{LF} fall below the first-best level D^*. This depresses the representative final shareholder's willingness to pay for equity, reducing the market-clearing price and the market discount rate $1/Q$. In turn, lower stock prices discourage investment at the firm level, and the general equilibrium settles at an inefficiently low level of aggregate investment, earnings and stock prices. With homogeneous capital market imperfections and inefficient liquidation, this under-investment also extends to the firm level, i.e. $k_i^{LF} < k_i^*$ for all $i \in \mathcal{I}$, but with heterogeneous capital market imperfections, some firms with very high $\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*\left(z\right)\right)\right]/\mathbb{E}\left[R_i\left(z, \lambda_i^*\left(z\right)\right)\right]$ may over-invest at the same time as the aggregate investment falls below the level that maximizes aggregate surplus.

The last part of Proposition 5 identifies the aggregate investment wedge. Equation (16) shows that $1/Q^{LF}$ can be decomposed into the representative final shareholder's discount factor $u'(D^{LF})$ and the ratio $\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} \hat{R}_i(z, \lambda_i^*(z))\varphi(z)dzdi / \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dzdi$, which represents the aggregate price-earnings ratio multiplied by the market interest rate Q^{LF}. This second term enters as a wedge into the investment first-order condition and is typically strictly larger than 1. Since

$$\frac{\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} \hat{R}_i(z, \lambda_i^*(z))\varphi(z)dzdi}{\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dzdi} \geq \frac{\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*\left(z\right)\right)\right]di}{\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left[R_i\left(z, \lambda_i^*\left(z\right)\right)\right]di}$$

with strict inequality whenever $\bar{Z}_i$ is non-empty for a positive measure of firms, this wedge reduces investment whenever returns are symmetric or dominated by upside risk and $\alpha < 1$. The wedge only increases investment only if $\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*\left(z\right)\right)\right]di < \int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left[R_i\left(z, \lambda_i^*\left(z\right)\right)\right]di$ and firms on average have an incentive to under-invest for given Q. In that case, the aggregate investment wedge does not completely offset the firm-level incentives to over-invest unless there is no liquidation or cross-sectional mis-allocation of capital.

To summarize, in general equilibrium, corporate short-termism induced by capital market imperfections never boosts aggregate investment, dividends and stock prices above the level that maximizes total surplus, and it depresses aggregate investment, dividends and stock prices below the surplus-maximizing level if it induces firms to engage in investment and liquidation strategies that are privately efficient but socially wasteful.

[7]The condition that and $\frac{d\ln C_i(k)}{d\ln k} = \kappa > 1$ for all i is only used to show that efficiency losses from capital market imperfections translate into aggregate under-investment. Without this condition, it may be possible that capital market imperfections reduce aggregate dividends and aggregate stock prices, but increase aggregate investment, if firms with high $\mathbb{E}\left[\hat{R}_i\left(z, \lambda_i^*\left(z\right)\right)\right]/\mathbb{E}\left[R_i\left(z, \lambda_i^*\left(z\right)\right)\right]$ also have the highest sensitivity of investment to stock prices.

5 Implications for corporate finance

The failure of time-consistent planning by firms that was highlighted by Theorems 1 and 2 has implications for corporate finance. In this section, I discuss two policies through which initial shareholders may attempt to overcome the time-consistency problem: capital structure and long-term managerial contracts that maximize initial shareholder value.

Under the EMH, long-term managerial contracts maximize firm value and implement Pareto-efficient investment and hiring. With capital market imperfections, they restore time consistency on the initial shareholders' preferred policy $\left(\hat{k}, \mathbb{R}, \hat{\lambda}\left(\cdot\right)\right)$. However, while overcoming inefficient liquidation by cutting back on hiring on the downside, they also create extra distortions by encouraging over-hiring or "empire-building" to boost stock prices on the upside. Similarly under the EMH, the Modigliani-Miller theorem shows that debt does not increase the overall firm value. Rather, debt reduces firm value if it dilutes final shareholders' hiring incentives (Myers, 1977). However, with departures from EMH, debt instead *increases* firm value as a commitment device, i.e. by making it credible for final shareholders to hire below the ex post efficient level, debt allows the firm to avoid inefficient liquidation. In a nutshell, the capital market imperfection turns debt overhang from a negative to a positive for the firm.

In partial equilibrium, long-term managerial contracts and debt always increase investment and benefit initial shareholders. These policies also benefit (resp. hurt) final shareholders and they are Pareto-improving (resp. Pareto inferior) in general equilibrium if the earnings gains from avoiding inefficient liquidation are larger (resp. smaller) than the earnings losses from more severe hiring and investment distortions. This tradeoff is again shaped by the relative size of upside and downside risks and the elasticities of hiring and investment to realized and expected stock market returns. Shareholder value maximization can even cause a complete collapse of the stock market if these additional hiring distortions become so severe that aggregate corporate earnings turn negative.

These two examples illustrate that conventional wisdom based on EMH reasoning need not hold in the presence of capital market imperfections. Policies that maximize firm value and allocate capital efficiently under EMH (i.e. shareholder value maximization) can amplify distortions and cause severe inefficiencies when EMH fails, while policies that are viewed as inefficient (i.e. debt overhang) under EMH increase firm value with capital market imperfections.

Empowering shareholders may worsen market outcomes

Consider a policy that gives initial shareholders full control over the firm's operations, including the hiring policy $\lambda\left(\cdot\right)$ that is implemented after the firm is sold to final shareholders. This policy follows the logic of maximizing "Shareholder Value", i.e. the market value of the firm to initial shareholders, and implements the ex post hiring policy $\hat{\lambda}(z) = \arg\max_{\lambda \geq 0} \hat{R}(z, \lambda)$ that maximizes the expected market return. Such a hiring policy restores time-consistency, allows the firm to avoid inefficient liquidation and improves investment incentives. To implement this policy, the initial shareholders need to hire a CEO whose incentives are tied to the stock price performance in a way that maximizes shareholder value, with protections that make it too costly for final shareholders to renegotiate the CEO's contract and implement a different hiring policy after acquiring the firm.

Proposition 6. *Suppose $\lambda(\cdot) = \hat{\lambda}(\cdot)$ and Assumption 1 is satisfied. Then in equilibrium, firm decisions are time-consistent and there is no liquidation.*

1. In partial equilibrium, investment and share prices are higher but firm performance (fundamental returns) is worse than in the Laissez-faire equilibrium, for $z \in \bar{Z}$. Initial shareholders always gain relative to the Laissez-faire equilibrium. Social surplus and final shareholder payoffs decrease whenever $\mathbb{E}\left(\max\left\{0, \hat{R}(z, \lambda^*(z))\right\}\right) \geq \int_{\bar{Z}} R(z, \lambda^*(z))\varphi(z)\,dz \geq \mathbb{E}\left(R\left(z, \hat{\lambda}(z)\right)\right)$.

2. In general equilibrium, $D^{LF} \gtreqless \hat{D}$, $P^{LF} \gtreqless \hat{P}$ and $\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right) di < \int_{i\in\mathcal{I}} C_i\left(\hat{k}_i\right) di$ and the Laissez-faire equilibrium Pareto-dominates (resp. is dominated by) shareholder value maximization, if and only if $\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dzdi \gtreqless \int_{i\in\mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i(z)\right)\right) di$.

If $\int_{i\in\mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i(z)\right)\right) di \leq 0$, the unique equilibrium features a complete market shutdown: $\hat{k} = \hat{D} = \hat{P} = 0$. This scenario arises if α_i is sufficiently small and $A_i(\cdot)$ has sufficient upside risk (is unbounded above) for a positive measure of $i \in \mathcal{I}$.

By giving initial shareholders exactly what they want —full control over present and future firm operations— shareholder value maximization restores time consistency. But it does so at the cost of distorting the firm's hiring decisions to inflate stock prices rather than to maximize firm performance. These hiring distortions allow the firm to avoid inefficient liquidation when the firm is under-valued, which benefits both initial and final shareholders. But they can also be very costly in states of the world where the firm is over-valued and not at risk of liquidation, since initial shareholders take advantage by encouraging "empire-building": excessive hiring to boost short-term stock prices, while leaving final shareholders with higher labor costs and lower revenues ex post. Therefore, whether shareholder value maximization is Pareto improving really comes down to comparing the efficiency gains from avoiding inefficient liquidation on the downside to the efficiency losses from hiring and investment distortions or empire-building on the upside. The condition $\mathbb{E}\left(\max\left\{0, \hat{R}(z, \lambda^*(z))\right\}\right) \geq \int_{\bar{Z}} R(z, \lambda^*(z))\varphi(z)\,dz \geq \mathbb{E}\left(R\left(z, \hat{\lambda}(z)\right)\right)$ guarantees that the losses from hiring and investment distortions always dominate the gains from avoiding inefficient liquidation, so that shareholder value maximization unambiguously lowers total surplus and final shareholders' expected payoffs.

In general equilibrium, all shareholders gain from higher aggregate corporate earnings since the aggregate stock price transmits the change in aggregate earnings from final to initial shareholders. Whether shareholder value maximization is preferable to the Laissez-faire equilibrium then comes down to comparing aggregate earnings under these two scenarios. The shift towards shareholder value maximization avoids inefficient liquidation, which increases aggregate earnings and welfare, but creates costly hiring distortions, and may increase or reduce the mis-allocation of capital, depending on which firms are most prone to over-invest in response to higher stock market returns.[8] Shareholder value maximization is always strictly Pareto inferior to the Laissez-faire equilibrium if the latter features no inefficient liquidation and homogeneous capital market imperfections.

With noisy information aggregation, these tradeoffs depend once again on the asymmetry between upside and downside risks and the labor share $1-\alpha$. If there is sufficient upside risk, i.e. $A(\cdot)$ has sufficient variation at the top, and α is sufficiently close to 0, then $R\left(z, \hat{\lambda}(z)\right) < 0$ for high enough values of z. Moreover, this threshold decreases to $\hat{z}$ as $\alpha \to 0$, and $R\left(z, \hat{\lambda}(z)\right)$ becomes arbitrarily negative for $z > \hat{z}$. It follows that $\mathbb{E}\left(R\left(z, \hat{\lambda}(z)\right)\right) < 0$ for sufficiently low α, i.e. under shareholder value maximization, expected earnings are negative.

While some firms may have negative expected earnings in equilibrium, the adverse effects of shareholder value maximization are especially striking, if aggregate corporate earnings are negative. If $\int_{i \in \mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i(z)\right)\right) di \leq 0$, final shareholders would simply refuse to invest in stocks, and the unique equilibrium features no investment and no equity trade. In other words, the equity market completely ceases to play its role of channeling investments to their most productive use. The market shut-down stems from a commitment problem among initial shareholders: Initial shareholders would want to commit collectively to hiring policies that guarantee positive cash flows to final shareholders. But such policies would not be credible since initial shareholders in each firm would have an incentive to deviate and engage in a hiring policy that inflates their own firm's market value, while ultimately generating losses for final shareholders. Without a commitment to hiring policies that protect the interests of final shareholders, there does not exist an equilibrium with positive aggregate dividends and positive share prices. But if initial shareholders do not expect to be able to transfer their investment to final shareholders at positive prices, the only equilibrium features no initial investment, liquidation and a complete shutdown of equity markets.

[8] I conjecture that shareholder value maximization worsens capital misallocation, if shareholders of firms that have the most to gain from boosting stock prices by distorting labor also have the most to gain from increasing their investment ex ante, thus simultaneously increasing their investment while reducing their earnings under shareholder value maximization.

The commitment value of debt

Consider now an extension of the baseline model of Section 2, in which initial shareholders are given the option, at the market stage, to issue some corporate debt $b > 0$ per unit of capital invested in the firm, in addition to selling their equity in the market. Debt is also sold through the market. Debt changes the the return structure for final shareholders, whose realized cash flows take the form $\max\{0, \Pi(\theta, 1, \lambda) - b\} - w\lambda$, i.e. the final shareholders earn the residual cash flows after paying bond holders. Hence, given b and λ, there exists a threshold $\bar{\theta}(\lambda, b)$ such that $\Pi(\theta, 1, \lambda) \gtreqless b$ if and only if $\theta \gtreqless \bar{\theta}(\lambda, b)$, and the firm stays solvent only if $\theta \geq \bar{\theta}(\lambda, b)$. To isolate the commitment channel, I abstract from other channels that impact the role of debt, such as tax shields or bankruptcy costs.

It follows that the expected fundamental return to the firm $R(z, \lambda) = R^E(z, \lambda; b) + R^B(z, \lambda; b)$ is unchanged, but debt affects the split between returns on debt $R^B(z, \lambda; b) = \mathbb{E}[\min\{b, \Pi(\theta, 1, \lambda)\}|z]$ and equity $R^E(z, \lambda; b) = \mathbb{E}[\max\{0, \Pi(\theta, 1, \lambda) - b\}|z] - w\lambda$. Specifically, the returns to debt satisfy $R^B(z, \lambda; 0) = 0$, R^B and R^B_λ are increasing in b and converge to R and R_λ as $b \to \infty$ (hence R^E and R^E_λ are positive but decreasing in b and converge to 0 as $d \to \infty$). Let $\lambda^*(z, b) = \arg\max_{\lambda \geq 0} R^E(z, \lambda; b)$ denote the hiring policy that maximizes the fundamental return to equity, for a given level of debt $b > 0$. It follows that $\lambda^*(z, b)$ is increasing in z, but decreasing in b and $\lambda^*(z, b) \in [0, \lambda^*(z)]$: debt dilutes final shareholders' incentives to hire.[9]

Debt and equity are subject to the same capital market imperfections. Under NIA, the market returns to debt and equity satisfy $\hat{R}^B(z, \lambda; b) = \mathbb{E}[\min\{\Pi(\theta, 1, \lambda), b\}|x = z, z]$ and $\hat{R}^E(z, \lambda; b) = \mathbb{E}[\max\{0, \Pi(\theta, 1, \lambda) - b\}|x = z, z] - wl$, where informational parameters in both markets are the same and supply shocks are perfectly correlated. These conditions imply that $\hat{R}(z, \lambda) = \hat{R}^B(z, \lambda; b) + \hat{R}^E(z, \lambda; b)$, i.e. the capital structure does not affect the market value of the firm, for given choices of λ.[10] The initial shareholders can condition the amount of debt that is sold in the market as a function of equity and bond prices, or equivalently, as a function of the market signal z.

Under the EMH, i.e. if $\hat{R}^B(z, \lambda; b) = R^B(z, \lambda; b)$ and $\hat{R}^E(z, \lambda; b) = R^E(z, \lambda; b)$ for all (z, λ, b), issuing debt is costly for the firm as it inefficiently reduces hiring, hence the surplus-maximizing

[9] This formulation assumes, for simplicity, that final shareholders are fully liable for paying labor. The results would essentially be unchanged if labor was paid out of final cash flows and treated as another form of debt, provided that w represents the expected payroll cost per unit of labor: the key is that under any such specification, final shareholders do not internalize the return to λ that goes to their bond holders in the form of lower default risk.

[10] Departures from EMH and the introduction of a separate market for corporate debt could also impact firm value through other channels, for example by increasing the available information with additional market signals. Alternatively, if capital market imperfections are not symmetric between the two markets, then $\hat{R}^B(z, \lambda; b) + \hat{R}^E(z, \lambda; b)$ varies with b –this could arise for example if debt was priced efficiently, i.e. $\hat{R}^B(z, \lambda; b) = R^B(z, \lambda; b)$ or lending was contracted with an outside lender, like a bank, rather than through the market, while equity remains subject to information frictions –see Albagli, Hellwig and Tsyvinski (2024) and Banerjee, Breon-Drish and Smith (2025) for further discussion. I abstract from these additional channels by focusing on an environment, in which, conditional on its investment and hiring policy, the market value of the firm is invariant to its capital structure.

capital structure sets $b^*(z) = 0$ for all z, so the firm maximizes total surplus in the Laissez-faire equilibrium as an equity-only firm, investing k^* and setting the hiring policy $\lambda^*(z)$.

This result no longer applies with departures from the EMH. Initial shareholders can implement any hiring policy $\lambda(\cdot)$ such that $\lambda(z) \in [0, \lambda^*(z)]$ for all $z \in \mathbb{R}$ by choosing the appropriate level of debt –hence they can reduce hiring below the first-best level, but not increase it beyond that level. The initial shareholders' preferred hiring policy sets $\lambda^B(z) = \min\left\{\hat{\lambda}(z), \lambda^*(z)\right\}$. Hence, the firm hires fewer workers than the equity-only benchmark when initial shareholders would want to cut back on hiring ($\hat{\lambda}(z) < \lambda^*(z)$), but the firm doesn't issue any debt and the hiring policy is unchanged at $\lambda^*(z)$ when initial shareholders would strictly prefer to increase hiring ($\hat{\lambda}(z) \geq \lambda^*(z)$). Furthermore, $\hat{R}\left(z, \lambda^B(z)\right) > 0$ for all z, so the firm is never liquidated. Let $\left\{k^B, \bar{Z}^B, \lambda^B(\cdot), b(\cdot)\right\}$ denote the Laissez-faire equilibrium with debt, indexed by i and augmented by $\left\{D^B, P^B, Q^B\right\}$ in the general equilibrium extension, where $D^B = \int_{i\in\mathcal{I}} k_i^B \mathbb{E}\left[R_i\left(z, \lambda^B(z)\right)\right] di$ denotes aggregate corporate earnings and $P^B = D^B u'\left(D^B\right)$ the aggregate value of stock and bond markets combined.

Proposition 7. *Suppose Assumption 1 is satisfied, and suppose that there exist $z'$and z'' such that $\lambda^*(z')/\hat{\lambda}(z') > 1 > \lambda^*(z'')/\hat{\lambda}(z'')$. With debt, initial shareholders implement $\lambda^B(\cdot) = \min\left\{\hat{\lambda}(\cdot), \lambda^*(\cdot)\right\}$ and there is no liquidation, but the equilibrium with debt remains intra-firm and aggregate Pareto-inferior.*

1. In partial equilibrium, investment k^B and stock prices are higher for z, but firm performance (fundamental returns) is worse than in the equity only equilibrium for $z \in \bar{Z}$, and strictly so if $\hat{\lambda}(z) < \lambda^*(z)$. Initial shareholders always gain relative to the Laissez-faire equilibrium. Social surplus and final shareholder payoffs decrease whenever $\mathbb{E}\left(\max\left\{0, \hat{R}(z, \lambda^*(z))\right\}\right) \geq \int_{\bar{Z}} R(z, \lambda^*(z))\varphi(z)\,dz \geq \mathbb{E}\left(R\left(z, \min\left\{\hat{\lambda}(z), \lambda^*(z)\right\}\right)\right)$.

2. In general equilibrium, $D^{LF} \gtreqless D^B$, $P^{LF} \gtreqless P^B$ and $\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right) di \gtreqless \int_{i\in\mathcal{I}} C_i\left(k_i^B\right) di$ and the equity-only equilibrium Pareto-dominates (resp. is dominated by) the corporate debt equilibrium, if and only if $\int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dzdi \gtreqless \int_{i\in\mathcal{I}} k_i^B \mathbb{E}\left[R_i\left(z, \min\left\{\hat{\lambda}_i(z), \lambda_i^*(z)\right\}\right)\right] di$.

Proposition 7 shows that issuing debt unambiguously increases the initial shareholders' expected payoffs by reducing hiring when $\hat{\lambda}(z) < \lambda^*(z)$ and avoiding inefficient liquidation. This boosts investment incentives in partial equilibrium. By diluting hiring incentives ex post, debt restores commitment to a hiring policy below the ex post efficient level. This in turn increases the interim market value at which the firm can be sold and allows initial shareholders to avoid inefficient liquidation for realizations of z at which expected market returns for an equity-only firm are negative. Whether debt also raises social surplus and final shareholder welfare then comes down to a tradeoff between surplus gains from avoiding inefficient liquidation if $z \notin \bar{Z}$, and losses from hiring distortions for $z \in \bar{Z}$ and more severe ex ante over-investment when $\mathbb{E}\left(\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right) \geq \int_{\bar{Z}} R\left(z, \lambda^*\left(z\right)\right)\varphi(z)\,dz$. Debt reduces total surplus when the extra hiring distortions and more sever over-investment outweigh the gains from avoiding inefficient liquidation ($\int_{\bar{Z}} R\left(z, \lambda^*\left(z\right)\right)\varphi(z)\,dz \geq \mathbb{E}\left(R\left(z, \lambda^B\left(z\right)\right)\right)$). But debt can benefit both initial and final shareholders if losses from inefficient liquidation or under-investment at the equity-only equilibrium are sufficiently large to dominate losses from additional hiring distortions.

In general equilibrium, debt is socially beneficial if and only if it increases aggregate corporate earnings, or if the gains from avoiding costly liquidation outweigh the costs of additional labor distortions. With noisy information aggregation, this occurs if α is sufficiently small and $A\left(\cdot\right)$ has sufficient downside risk.

In summary, issuing debt credibly implements hiring policies that avoid losses from inefficient liquidation, and thereby overcomes the final shareholder's lack of commitment. Notice that debt remains an inherently imperfect solution to the time-consistency issue: by aligning shareholder incentives on stock prices, it limits the firm's downside risks. But it does so at the cost of introducing new distortions into hiring when the firm is not at risk of liquidation. As a consequence, even if issuing debt leads to a Pareto improvement over the equity-only equilibrium, it does not restore full Pareto efficiency whenever $k^B \neq k^*$. In partial equilibrium, it is always possible to further increase both initial and final shareholders' expected payoffs with a common perturbation of investment in the direction of k^* and hiring in the direction of $\hat{\lambda}\left(\cdot\right)$ along the same lines as in Theorem 1. In general equilibrium, aggregate dividends can be increased (benefiting both initial and final shareholders) by setting a hiring policy $\lambda_i^*\left(\cdot\right)$ without liquidation.

6 Empirical evidence

The model was not designed to match specific empirical facts, but to highlight the channels through which equity market imperfections provide a generic source of time inconsistency in firm decisions. The model's implications for stock prices and firm decisions with noisy information aggregation nevertheless resonate with several stylized facts about asset returns and corporate investment.

Asset pricing

The central prediction of the noisy information aggregation model is that the stock price responds too strongly to the information that is aggregated through the market. When z is high (resp. low), the marginal informed trader becomes overly optimistic (resp. pessimistic) about the firm's prospects. From an ex ante perspective, this excess sensitivity to market information translates into a U-shaped pricing kernel that places too much weight on extremely high or low realizations of fundamentals. This pattern is verified empirically as the variance premium puzzle (Carr and Wu, 2009) and consistent with direct evidence on non-monotone pricing kernels (Jackwerth, 2000; Christoffersen, Heston and Jacobs, 2013; Audrino, Huitema and Ludwig, 2022).

Excess weight on tail risks implies that return asymmetry has first-order effects on equity prices and returns: there are negative returns to skewness in stock markets as positively skewed securities are over-priced on average and generate lower returns than negatively skewed securities that are under-priced. Boyer, Mitton and Vorkink (2010) and Conrad, Dittmar and Ghysels (2013) confirm this negative relationship at the firm level and document significant variation in returns with skewness of up to 8% p.a. after controlling for other return factors.

Albagli, Hellwig and Tsyvinski (2024) show that measures of belief dispersion or forecast disagreement can serve as a sufficient statistic for excess weight on tail risks, since the latter derives from investors' private information. The model thus predicts that securities with upside risk experience more over-pricing and lower returns when there is higher disagreement. Diether, Malloy and Scherbina (2002) confirm this prediction by showing that equity returns are substantially lower for firms with high disagreement, with return differences of up to 7% p.a. after controlling for other factors; on the other hand Guntay and Hackbarth (2010) document increasing returns to disagreement in bond markets which are characterized by downside risk. Finally, the model suggests increasing differences in returns, i.e. more upside risk leads to more negative returns to disagreement, while higher disagreement makes the returns to skewness more negative. I am not aware of direct evidence for this last implication but to the extent that we can proxy for skewness with value-vs.-growth (Zhang 2013), Yu (2011) documents that the value premium is increasing with disagreement and the returns to disagreement are decreasing with the book-to-market ratio, consistent with the increasing difference property.

In summary, the core model of asset prices with noisy information aggregation that I used to model departures from the EMH rationalizes several well-known empirical patterns regarding returns in equity markets. The investment part of the model discusses how these return anomalies impact firm decisions and performance when shareholders care about stock prices.

Corporate investment

Starting from Tobin (1969), an ample literature posits and documents sensitivity of firm investment to stock prices, both in the aggregate (Barro, 1990; Blanchard, Rhee and Summers, 1993) and at the firm level (Morck, Shleifer and Vishny, 1990). One channel for stock price sensitivity of investment is that firm decisions respond to information that is aggregated through equity markets because this information is perceived as valuable to firm decisions and complementary to other information sources to which the firm has access, i.e. the information feedback (Chen, Goldstein and Jiang, 2007) that is also at the heart of the model. Alternatively, investment may respond to non-fundamental components of stock prices through financing constraints, sentiments or managerial incentives that are tied to stock prices, and a substantial literature has sought to identify the investment response to non-fundamental price movements (Baker, Stein and Wurgler, 2003). For example, firms may adjust their strategies to cater to market sentiments (Polk and Sapienza, 2009). Such investment responses to non-fundamental price fluctuations imply that abnormally high investment is often followed by abnormally low earnings and returns (Hoberg and Phillips, 2010). Amihud and Levy (2022) document that stock illiquidity negatively impacts investment using proxies that correlate with the price impact of trade and are known to raise the firm's cost of capital.

The present model combines elements of all three channels: an information feedback channel since information aggregated through the market is valuable for firm decisions, a catering channel since initial shareholder incentives depend on market returns that differ systematically from fundamental returns, and a liquidity channel that captures the excess sensitivity of stock prices to market information, as well as liquidity shocks from noise trading.[11]

Firm investment can also be directly linked to skewness and forecast dispersion in equity markets, providing suggestive evidence for the mechanism linking investment to stock prices in the model. Gebhardt, Lee and Swaminathan (2001) show that analyst forecast dispersion is negatively related to firms' cost of capital, consistent with the model's implication that a larger departure from the EMH, as proxied by forecast dispersion, raises stock prices and investment and lowers the firm's implicit cost of capital. Bloom, Guvenen and Salgado (2025) document that, controlling for first- and second moments, skewness of firm fundamentals has strong predictive power for corporate investment, employment, output and earnings. They show that recessions are characterized by a shift towards more firm-level downside risk, and this shift adversely impacts firms' investment and hiring decisions. Such a relation is difficult to reconcile with standard risk-return tradeoffs that focus on first and second moments, but arises naturally when variation in third moments induces variation in market relative to fundamental returns, and market returns feed into investment and hiring decisions. A shift towards more downside risk then amplifies the tension between initial and

[11] While I am not aware of comparable results for the impact of stock prices on employment at the firm level, Belo, Lin and Bazdresch (2014) document that firms with higher employment growth experience lower subsequent stock market returns. The model accounts for this empirical pattern since higher z simultaneously leads to more hiring and more over-valuation of equity and thus lower subsequent returns.

final shareholder preferences, resulting in more inefficient liquidation and amplifying the aggregate investment wedge in general equilibrium, when downside risk is large.

A separate literature provides direct evidence of corporate short-termism in publicly traded firms. Graham, Harvey and Rajgopal (2005) document that a large share of managers prefers to forego positive-NPV investments to meet short-term earnings targets. Terry (2023) provides evidence suggestive of active earnings management: earnings announcements bunch at the analysts' prior earnings target, firms that just meet the target subsequently have lower R&D growth and patenting, and firms that just miss their earnings target experience lower stock returns and their executives receive lower compensation and are more likely to leave the firm. Asker, Farre-Mensa and Ljungqvist (2014) document that publicly traded firms have consistently lower investment rates than otherwise identical privately held firms, consistent with the hypothesis that equity markets put pressure on firms and managers to boost short-term earnings and stock prices at the expense of longer-term investments. While these papers typically rationalize corporate short-termism through contracting frictions and managerial career concerns (Narayanan 1985), the present paper proposes an alternative theory based on time-inconsistency resulting from financial market imperfections.

Gormsen and Huber (2025a) directly measure firms' perceived cost of capital from transcripts of corporate conference calls. They document the existence of a "cost of capital wedge" between the perceived and actual cost of capital, and show that there is substantial dispersion in the firms' perceived cost of capital which appears to be unrelated to firm's actual financing costs but more closely tied to external investor sentiment or analyst's return expectations.

Their stylized facts map nicely to general equilibrium model in Section 4.[12] Using the first-order condition for investment, the firms' model-implied perceived cost of capital can be represented as

$$\frac{\int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dz}{C_i'(k_i^{LF})} = Q^{LF} \cdot \Delta_i, \text{ where } \Delta_i \equiv \frac{\int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dz}{\int_{\bar{Z}_i} \hat{R}_i(z, \lambda_i^*(z))\varphi(z)dz}, \tag{20}$$

combining the marginal trader's expected return Q^{LF} with a firm-specific adjustment Δ_i that captures variation in expected value-weighted returns $Q^{LF} \cdot \frac{R_i(z,\lambda_i^*(z))}{\hat{R}_i(z,\lambda_i^*(z))}$ across firms. Gormsen and Huber (2025a) show that firms' perceived cost of capital is correlated with leverage and book-to-market ratios, which is consistent with the model implication that Δ_i is lower for firms that have more upside risk. Moreover, the firms' perceived cost of capital co-varies with external analyst's return expectations, which is consistent with the definition of Δ_i as an expected value-weighted

[12] In addition, (Gormsen and Huber 2025b) directly measure firms' discount rates and document the presence of a "discount rate wedge" between the firms' discount rate and the perceived cost of capital. The present model does not feature the analogue of the discount rate wedge, which appears more naturally in models with heterogeneous discount rates, such as Caramp, Kozlowski and Teeple (2024).

return on equity and suggests that firms factor market expectations into their investment decisions.

In addition, Gormsen and Huber (2025a) show that perceived costs of capital only display weak co-movement with their estimates of actual financing costs. They regress proxies for actual cost of capital (based on realized returns on equity and debt, as well as expected returns based on CAPM and Fama-French) on the perceived cost of capital and find pass-through coefficients of 0.5 and below. When instead the actual cost of capital is estimated using analysts' expected return on equity at the firm level, the estimated pass-through coefficients are much higher and no longer statistically significantly different from 1.

Without aggregate uncertainty, the model-implied cost of capital is equal to the aggregate stock market return $1/u'\left(D^{LF}\right)$, which can be represented as

$$1/u'\left(D^{LF}\right) = Q^{LF} \cdot \Delta, \text{ where } \Delta \equiv \frac{\int_i k_i \int_{\bar{Z}_i} R_i(z, \lambda_i^*(z))\varphi(z)dzdi}{\int_i k_i \int_{\bar{Z}_i} \hat{R}_i(z, \lambda_i^*(z))\varphi(z)dzdi}. \tag{21}$$

Hence, the firm's perceived cost of capital $Q^{LF} \cdot \Delta_i$ differs from their actual cost of capital $1/u'\left(D^{LF}\right)$ by a firm-specific factor Δ_i/Δ that captures the cross-sectional variation in value-weighted returns, relative to the overall market average. Therefore, regressing noisy measures of the firms' financial cost of capital $\ln\left[Q^{LF} \cdot \Delta\right]$ on their perceived cost of capital $\ln\left[Q^{LF} \cdot \Delta_i\right]$ naturally leads to attenuation bias due to the fact that $\ln\left[Q^{LF} \cdot \Delta_i\right]$ is mechanically correlated with the residual that includes $-\ln\left(\Delta_i/\Delta\right)$. On the other hand, if I interpret external analysts' return expectations as a noisy proxy for the value-weighted expectation $Q^{LF} \cdot \Delta_i$ of firms' equity returns and use it to proxy for the firms financial cost of capital, then regressing analyst-based measures of the financial cost of capital on the firms' perceived cost of capital will not suffer from the same attenuation bias and lead to a regression coefficient much closer to 1. Hence the general equilibrium model presented in this paper rationalizes the empirical patterns documented in Gormsen and Huber (2025a) regarding the cost of capital wedge between the firms' perceived and actual cost of capital and its co-movement with analyst or market expectations.

7 Conclusion

I develop a theory of Casino Capitalism in which generic departures from the efficient markets hypothesis generate the full range of inefficiencies that critics of shareholder value theory have long identified. Formally, I show that time-consistent rational planning by firms emerges as an equilibrium outcome if and only if incentives of successive shareholder generations are perfectly aligned. This incentive alignment condition is always satisfied with efficient markets, but almost always violated with even small departures from EMH. The resulting tension between initial and final shareholder preferences further alters fundamental trade-offs that shape decisions about capital structure or managerial incentives within firms. I establish these results both for a broad range of capital market imperfections that can be modeled as violations of the EMH, and for a simple model of noisy information aggregation in equity markets that delivers specific predictions that are consistent with cross-sectional evidence on stock returns and corporate investment.

At the firm level, the market imperfection creates an endogenous preference reversal between initial shareholders whose objective is to maximize market returns on investment, and final shareholders who want to maximize corporate earnings. The resulting time inconsistency leads to a Laissez-faire equilibrium that is strictly Pareto inferior for both initial and final shareholders, featuring excess investment combined with inefficient liquidation of positive-NPV projects. In general equilibrium, investment decisions that are privately optimal for initial shareholders are collectively inefficient as individual attempts to boost firm-level stock prices are self-defeating, depressing aggregate dividends, stock prices, and investment through an aggregate investment wedge.

The model was intentionally stylized and meant to isolate the key forces driving time inconsistency result. But the main conceptual result linking endogenous preference reversals among shareholders to departures from the EMH relies on a limited set of assumptions that hold across a broad range of models incorporating market imperfections. It is therefore likely to apply well beyond the noisy information aggregation model, even if capital markets operate close to but not exactly in accordance with the efficient markets hypothesis.

The main inefficiency results in this paper cast doubt on the general prescription that maximizing shareholder value is synonymous with maximizing social surplus. If one accepts the premise that arbitrage is costly and real world stock markets do not reach the theoretical ideal of perfect efficiency, then one should view efficient, time-consistent planning by firms as the exception rather than the rule. Moreover, even a small departure from the EMH can generate arbitrarily large distortions when investment and hiring decisions are sufficiently sensitive to stock prices. The race between market (im-)perfection and scalability determines the quantitative significance of the distortions, and this race can be won decisively by scalability even when the stock market is close to efficient. Further development of the empirical and policy implications of this fundamental insight for understanding firm decisions is left to future research.

References

[1] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2023). "Imperfect Financial Markets and Investment Inefficiencies," American Economic Review, 113 (9), 2323-2354.

[2] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2024). "Information aggregation with asymmetric asset payoffs," Journal of Finance, 79 (4), 2715-2758.

[3] Alesina, Alberto, and Guido Tabellini (1990). "A Positive Theory of Fiscal Deficits and Government Debt," Review of Economic Studies, 57 (3), 403-414.

[4] Amador, Manuel (2004). "A Political Economy Model of Sovereign Debt Repayment," working paper, Stanford University.

[5] Amihud, Yakov, and Shai Levi (2023). "The Effect of Stock Liquidity on the Firm's Investment and Production," Review of Financial Studies, 36 (3), 1094-1147.

[6] Asker, John, Joan Farre-Mensa and Alexander Ljungqvist (2014). "Corporate Investment and Stock Market Listing: A Puzzle?" Review of Financial Studies, 28 (2), 342-390.

[7] Audrino, Francesco, Robert Huitema, and Markus Ludwig (2021). "An Empirical Implementation of the Ross Recovery Theorem as a Prediction Device," Journal of Financial Econometrics, 19 (2), 291-312.

[8] Baker, Malcolm, Jeremy Stein, and Jeffrey Wurgler (2003). "When Does the Market Matter? Stock Prices and the Investment of Equity-Dependent Firms," Quarterly Journal of Economics, 118 (3), 969-1005.

[9] Baker, Malcolm, and Jeffrey Wurgler (2013). "Behavioral Corporate Finance: an Updated Survey," in: Handbook of Economics and Finance (Constantinides, Harris, Stulz, eds.), Elsevier, Vol. 2 A, 357-424.

[10] Banerjee, Snehal, Bradyn Breon-Drish and Kevin Smith (2025). "Asymmetric Information, Disagreement, and the Valuation of Debt and Equity," Journal of Financial Economics, 165, 103995.

[11] Barro, Robert (1990). "The Stock Market and Investment," Review of Financial Studies 3 (1), 115-131.

[12] Battaglini, Marco, and Stephen Coate (2007). "Inefficiency in Legislative Policy-Making: A Dynamic Analysis," American Economic Review, 97 (1), 118-149.

[13] Belo, Federico, Xiaoji Lin, and Santiago Bazdresch (2014). "Labor Hiring, Investment, and Stock Return Predictability in the Cross Section," Journal of Political Economy, 122 (1), 129-177.

[14] Benmelech, Efraim, Eugene Kandel, and Pietro Veronesi (2010). "Share-Based Compensation and CEO (Dis)Incentives," Quarterly Journal of Economics, 125 (4), 1769-1820.

[15] Blanchard, Olivier, Changyong Rhee, and Lawrence Summers (1993). "The Stock Market, Profit and Investment," Quarterly Journal of Economics, 108 (1), 115-136.

[16] Bolton, Patrick, Jose Scheinkman, and Wei Xiong (2006). "Executive Compensation and Short-Termist Behaviour in Speculative Markets," Review of Economic Studies, 73 (3), 577-610.

[17] Boyer, Brian, Todd Mitton, and Keith Vorkink (2010). "Expected Idiosyncratic Skewness," Review of Financial Studies, 23 (1), 169-202.

[18] Caramp, Nicolas, Julian Kozlowski and Keisuke Teeple (2024). "Liquidity and Investment in General Equilibrium," working paper, UC Davis, St. Louis Fed and University of Waterloo.

[19] Carr, Peter, and Liuren Wu (2009). "Variance Risk Premiums," Review of Financial Studies 22 (3), 1311-1341.

[20] Chen, Qi, Itay Goldstein, and Wei Jiang (2007). "Price Informativeness and Investment Sensitivity to Stock Price." Review of Financial Studies, 20(3), 619-650.

[21] Christoffersen, Peter, Steven Heston and Kris Jacobs (2013). "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel," Review of Financial Studies 26 (8), 1963-2006.

[22] Conrad, Jennifer, Robert Dittmar, and Eric Ghysels (2013). "Ex Ante Skewness and Expected Stock Returns," Journal of Finance, 68 (1), 85-124.

[23] Davila, Eduardo, and Andreas Schaab (2025). "Time Inconsistency with Heterogeneous Agents," working paper, Yale University and UC Berkeley.

[24] Diamond, Douglas and Robert Verrecchia (1981). "Information Aggregation in a Noisy Rational Expectations Economy," Journal of Financial Economics, 9 (3), 221-235.

[25] Diether, Karl, Christopher Malloy, and Anna Scherbina (2002). "Differences of Opinion and the Cross Section of Stock Returns," Journal of Finance, 57 (5), 2113-2141.

[26] Fama (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work," Journal of Finance, 25 (2), 383-417.

[27] Friedman (1970). "The Social Responsibility of Business is to increase its Profits," New York Times, September 13, 1970, 32-33.

[28] Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan (2001). "Toward an Implied Cost of Capital," Journal of Accounting Research, 39(1), 135-176.

[29] Graham, John R., Campbell R. Harvey, and Shiva Rajgopal (2005). "The Economic Implications of Corporate Financial Reporting," Journal of Accounting and Economics, 40(1-3), 3-73.

[30] Gormsen, Niels J. and Kilian Huber (2025a). "Firms' Perceived Cost of Capital", working paper, University of Chicago.

[31] Gormsen, Niels J. and Kilian Huber (2025a). "Corporate Discount Rates," American Economic Review, 115 (6), 2001-2049.

[32] Grossman, Sanford J. and Joseph E. Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets," American Economic Review, 70 (3), 393-408.

[33] Güntay, Levent, and Dirk Hackbarth (2010). "Corporate Bond Credit Spreads and Forecast Dispersion." Journal of Banking and Finance, 34 (10), 2328-2345.

[34] Gul, Faruk, and Wolfgang Pesendorfer (2001). "Temptation and Self-Control." Econometrica, 69 (6), 1403-1435.

[35] Gul, Faruk, and Wolfgang Pesendorfer (2005). "The Revealed Preference Theory of Changing Tastes." Review of Economic Studies, 72 (2), 429-448.

[36] Jensen and Meckling (1976). "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," Journal of Financial Economics, 3 (4), 305-360.

[37] Hellwig, Martin (1980). "On the Aggregation of Information in Competitive Markets," Journal of Economic Theory, 22, 477-498.

[38] Hoberg, Gerard, and Gordon Phillips (2010). "Real and Financial Industry Booms and Busts." Journal of Finance 65(1), 45-86.

[39] Holmstrom, Bengt, and Jean Tirole (1993). "Market Liquidity and Performance Monitoring," Journal of Political Economy, 101(4), 678-709.

[40] Jackson and Yariv (2014). "Present Bias and Collective Dynamic Choice in the Lab," American Economic Review, 104 (12), 4184-4202.

[41] Jackson and Yariv (2015). "Collective Dynamic Choice: The Necessity of Time Inconsistency," American Economic Journal: Microeconomics, 7 (4), 150-178.

[42] Jackwerth, Jens (2000). "Recovering Risk Aversion from Option Prices and Realized Returns" Review of Financial Studies, 13 (2), 433–451.

[43] Keynes, John M. (1936). "The General Theory of Employment, Interest and Money," McMillan, London.

[44] Kydland, Finn, and Edward C. Prescott (1977). "Rules Rather than Discretion: The Inconsistency of Optimal Plans," Journal of Political Economy, 85 (3), 473-491.

[45] Morck, Randall, Andrei Shleifer, and Robert Vishny (1990). "The Stock Market and Investment: is the Market a Sideshow?" Brookings Papers on Economic Activity, 1990 (2), 157-215.

[46] Myers, Stewart (1977). "Determinants of Corporate Borrowing," Journal of Financial Economics, 5 (2), 147–175.

[47] Narayanan (1985). "Managerial Incentives for Short-Term Results," Journal of Finance, 40 (5), 1469-1484.

[48] Polk, Christopher and Paola Sapienza (2009). "The Share Market and Corporate Investment: A Test of Catering Theory." Review of Financial Studies, 22 (1), 187-217.

[49] Salgado, Sergio, Nicholas Bloom and Fatih Guvenen (2025). "Skewed Business Cycles," working paper, University of Pennsylvania, Stanford University and University of Minnesota.

[50] Stein, Jeremy (1988). "Takeover Threats and Managerial Myopia," Journal of Political Economy, 96 (1), 61-80.

[51] Stein, Jeremy (1989). "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," Quarterly Journal of Economics, 104 (4), 655-669.

[52] Strotz, Robert H. (1955). "Myopia and Inconsistency in Dynamic Utility Maximization," Review of Economic Studies, 23 (3), 165-180.

[53] Terry, Stephen (2023). "The Macro Impact of Short-termism," Econometrica 91 (5), 1881-1912.

[54] Tirole, Jean (2006). "Theory of Corporate Finance," Princeton Economic Press.

[55] Tobin (1969). "A General Equilibrium Approach to Monetary Theory," Journal of Money, Credit and Banking, 1 (1), 15-29.

[56] Yu, Jialin (2011). "Disagreement and Return Predictability of Stock Portfolios," Journal of Financial Economics, 99 (1), 162-183.

[57] Zhang, Xiao-Jun (2013). "Book-to-Market Ratio and Skewness of Stock Returns." The Accounting Review, 88 (6), 2213-2240.

8 Appendix: Proofs

Proof of Proposition 1:

If $\lambda^*(\cdot) = \hat{\lambda}(\cdot)$, then $\hat{R}(z, \lambda^*(z)) = \hat{R}(z, \hat{\lambda}(z)) > 0$ and $\bar{Z} = \mathbb{R}$, so the liquidation option is never exercised. The first-order condition for investment (6) for investment reduces to $C'(k^{LF}) = \mathbb{E}\left[\hat{R}(z, \hat{\lambda}(z))\right]$, and therefore $k^{LF} = \hat{k}$.

Proof of Proposition 2:

The labor intensity is efficient by (3). Moreover, when $\hat{R}(z, \lambda^*(z)) = R(z, \lambda^*(z)) > 0$, $\bar{Z} = \mathbb{R}$ and the liquidation option is never exercised in a Laissez-faire equilibrium. Finally, first-order condition (6) for investment reduces to $C'(k^{LF}) = \mathbb{E}\left[R(z, \lambda^*(z))\right]$, and therefore $k^{LF} = k^*$.

Proof of Theorem 1:

Since $\left(\hat{k}, \mathbb{R}, \hat{\lambda}(\cdot)\right)$ and $(k^*, \mathbb{R}, \lambda^*(\cdot))$ are both Pareto-efficient, it suffices to show that $(k^{LF}, \bar{Z}, \lambda^*(\cdot))$ is Pareto inferior whenever $\lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$ and $(k^{LF}, \bar{Z}) \neq (k^*, \mathbb{R})$.

I first show that the Laissez-faire equilibrium must be Pareto inferior if $\bar{Z} \subset \mathbb{R}$. Any Pareto efficient allocation can be represented as the solution to the Pareto planner problem

$$\max_{k, Z, \lambda(\cdot)} \left\{ \int_Z R^{SP}(z, \lambda(z)) \varphi(z)\, dz \cdot k - C(k) \right\},$$

where $R^{SP}(z,\lambda) \equiv \Gamma R(z,\lambda) + (1-\Gamma)\hat{R}(z,\lambda)$. If $\Gamma \le 1$, $R_{\lambda}^{SP}(z,\lambda) > 0$ for any z and λ close to zero. Therefore $\lambda^{SP}(z) > 0$ and $R^{SP}(z) = \max_{\lambda \ge 0} R^{SP}(z,\lambda) > 0$, so that liquidation cannot be optimal. If $\Gamma > 1$,

$$R^{SP}(z,\lambda^{SP}(z)) \ge R^{SP}(z,\lambda^*(z)) = R(z,\lambda^*(z)) + (\Gamma-1)\left(R(z,\lambda^*(z)) - \hat{R}(z,\lambda^*(z))\right) > 0$$

whenever $R(z,\lambda^*(z)) \ge \hat{R}(z,\lambda^*(z))$, where the last inequality uses $R(z,\lambda^*(z)) > 0$. At the same time, if $\hat{R}(z,\lambda^*(z)) \ge R(z,\lambda^*(z)) > 0$, the firm is not liquidated in equilibrium. Therefore, if $\bar{Z} \subset \mathbb{R}$ (the firm is liquidated with positive probability in equilibrium), the Laissez-faire equilibrium cannot be Pareto efficient.

Second, I show that the Laissez-faire equilibrium cannot be Pareto-efficient if $\lambda^*(\cdot) \neq \hat{\lambda}(\cdot)$ and $k^{LF} \neq k^*$, even if $\bar{Z} = \mathbb{R}$. If $\mathbb{E}\left(\hat{R}(z,\lambda^*(z))\right) \neq \mathbb{E}\left(R(z,\lambda^*(z))\right)$ and $\lambda^*(z) \neq \hat{\lambda}(z)$ with positive probability, consider a perturbation that sets $k_\varepsilon = k^{LF} + \varepsilon\left(k^* - k^{LF}\right)$ and $\lambda_\eta(z) = \lambda^*(z) + \eta\left(\hat{\lambda}(z) - \lambda^*(z)\right)$. The perturbation changes final shareholders' expected payoffs by

$$\mathbb{E}\left(R(z,\lambda_\eta(z)) - \hat{R}(z,\lambda_\eta(z))\right)k_\varepsilon - \left\{\mathbb{E}\left(R(z,\lambda^*(z))\right) - \mathbb{E}\left(\hat{R}(z,\lambda^*(z))\right)\right\}k^{LF}$$

$$= \varepsilon\left(k^* - k^{LF}\right)\left(C'(k^*) - C'\left(k^{LF}\right)\right) - \eta k^{LF}\mathbb{E}\left(\hat{R}_\lambda(z,\lambda^*(z))\left(\hat{\lambda}(z) - \lambda^*(z)\right)\right) + o\left(\eta^2 + \varepsilon\eta + \varepsilon^2\right)$$

and incumbent shareholders' expected payoffs by

$$\mathbb{E}\left(\hat{R}(z,\lambda_\eta(z))\right)k_\varepsilon - \mathbb{E}\left(\hat{R}(z,\lambda^*(z))\right)k^{LF} - \left(C(k_\varepsilon) - C\left(k^{LF}\right)\right)$$

$$= \eta k^{LF}\mathbb{E}\left(\hat{R}_\lambda(z,\lambda^*(z))\left(\hat{\lambda}(z) - \lambda^*(z)\right)\right) + o\left(\eta^2 + \varepsilon\eta + \varepsilon^2\right).$$

Since $\hat{R}_\lambda(z,\lambda^*(z)) \gtreqless 0$ if and only if $\hat{\lambda}(z) \gtreqless \lambda^*(z)$, it follows that $\hat{R}_\lambda(z,\lambda^*(z))\left(\hat{\lambda}(z) - \lambda^*(z)\right) > 0$ unless $\hat{\lambda}(z) = \lambda^*(z)$, and therefore $k^{LF}\mathbb{E}\left(\hat{R}_\lambda(z,\lambda^*(z))\left(\hat{\lambda}(z) - \lambda^*(z)\right)\right) > 0$. Moreover, $k^* \gtreqless k^{LF}$ if and only if $C'(k^*) \gtreqless C'\left(k^{LF}\right)$ implies $\left(k^* - k^{LF}\right)\left(C'(k^*) - C'\left(k^{LF}\right)\right) > 0$ unless $k^* = k^{LF}$. By setting $\varepsilon\left(k^* - k^{LF}\right)\left(C'(k^*) - C'\left(k^{LF}\right)\right) = 2\eta k^{LF}\mathbb{E}\left(\hat{R}_\lambda(z,\lambda^*(z))\left(\hat{\lambda}(z) - \lambda^*(z)\right)\right)$ and $\eta > 0$ sufficiently small so that higher-order terms can be ignored, the perturbation increases both incumbent and final shareholders' expected payoffs. Hence the Laissez-faire equilibrium cannot be Pareto-efficient.

Proof of Proposition 3:

Fix a price function P and the associated signal $z(P)$ that the price reveals. An informed trader with signal x_i purchases the share if and only if $\mathbb{E}[A(\theta)|x_i, z(P)]\lambda(z(P))^{1-\alpha} - w\lambda(z(P)) \ge P/k$, i.e., if $x_i \geq z(P)$ by the monotone likelihood ratio property and the definition of $\hat{R}$. Market clearing requires $1 = \Phi(u) + 1 - \Phi(\sqrt{\beta}(z(P) - \theta))$, which gives $z(P) = \theta + (1/\sqrt{\beta})u$. Substituting back into the informed trader's indifference condition yields $P/k = \mathbb{E}[A(\theta)|x = z, z]\lambda(z)^{1-\alpha} - w\lambda(z)$, and therefore the unique equilibrium price function must satisfy $P(z,k) = \hat{R}(z, \lambda(z)) \cdot k$, for realizations $z \in \bar{Z}$ at which there is no liquidation. Moreover, monotonicity guarantees that the price function is invertible w.r.t. z, and hence P is informationally equivalent to z, for given investment k.

Proof of Lemma 1:

Write $\hat{R}(z, \lambda^*(z))$ as

$$\hat{R}(z, \lambda^*(z)) = R(z, \lambda^*(z)) \left\{ 1 + \alpha^{-1} \left(\frac{\mathbb{E}[A(\theta)|x = z, z]}{\mathbb{E}[A(\theta)|z]} - 1 \right) \right\}.$$

Hence there exists a unique z^{LF}, such that $\hat{R}(z, \lambda^*(z)) \gtreqless 0$ iff $z \gtreqless z^{LF}$, where z_1 satisfies $(1-\alpha)\mathbb{E}[A(\theta)|z^{LF}] = \mathbb{E}[A(\theta)|x = z^{LF}, z^{LF}]$. Next, consider $z \geq z^{LF}$: since $R(z, \lambda^*(z))$ and $\hat{R}(z, \lambda^*(z)) / R(z, \lambda^*(z))$ are both positive and strictly increasing for $z > z^{LF}$, it follows that $\hat{R}(z, \lambda^*(z))$ must be positive and strictly increasing for $z > z^{LF}$. Lastly, since $\lim_{z \to -\infty} R(z, \lambda^*(z)) = 0$, it follows that $0 \geq \lim_{z \to -\infty} \hat{R}(z, \lambda^*(z)) \geq \lim_{z \to -\infty} R(z, \lambda^*(z)) (1 - \alpha^{-1}) = 0$, i.e. $\hat{R}(z, \lambda^*(z))$ converges to 0 "from below" as $z \to -\infty$.

Proof of Proposition 4:

Part (i) follows from the fact that the final shareholders' first-order condition for labor coincides with the planner's FOC. For part (ii), Lemma 1 shows that $\hat{R}(z, \lambda^*(z)) < 0$ for z sufficiently low, hence there must exists a liquidation threshold $z^{LF} > -\infty$ at which $\hat{R}\left(z^{LF}, \lambda^*\left(z^{LF}\right)\right) = 0$. For part (iii), note that $k^{LF} \gtreqless k^*$ if and only if $\mathbb{E}\left[\max\left\{0, \hat{R}(z, \lambda^*(z))\right\}\right] \gtreqless \mathbb{E}[R(z, \lambda^*(z))]$ and $\mathbb{E}\left[\max\left\{0, \hat{R}(z, \lambda^*(z))\right\}\right] > \mathbb{E}\left[\hat{R}(z, \lambda^*(z))\right]$, so that $k^{LF} > k^*$ whenever $\mathbb{E}\left[\hat{R}(z, \lambda^*(z))\right] \geq \mathbb{E}[R(z, \lambda^*(z))]$. But we can write

$$\begin{aligned} \mathbb{E}\left[\hat{R}(z, \lambda^*(z))\right] - \mathbb{E}[R(z, \lambda^*(z))] \quad = \quad & \lambda^*(\hat{z})^{1-\alpha} \int (A(\xi\theta) - A(\theta))\, d\Phi(\theta) \\ & + \mathbb{E}\left\{ (\mathbb{E}[A(\theta)|x = z, z] - \mathbb{E}[A(\theta)|z]) \left(\lambda^*(z)^{1-\alpha} - \lambda^*(\hat{z})^{1-\alpha} \right) \right\} \end{aligned}$$

where $\xi = \left(1 + \frac{\beta}{(1+\beta+\beta\delta)^2}\left(1 + \delta^{-1}\right)\right)^{1/2} > 1$ only depends on the informational parameters (β, δ) and $\hat{z}$ satisfies $\mathbb{E}[A(\theta)|x = \hat{z}, \hat{z}] = \mathbb{E}[A(\theta)|\hat{z}]$. The first term in this expression captures asymmetry in A and is positive/zero/negative whenever A is dominated by upside risk/symmetric/dominated by downside risk. The second term is strictly positive whenever $\alpha < 1$, but converges to 0 as $\alpha \to 1$.

Hence $\mathbb{E}\left[\max\left\{0,\hat{R}\left(z,\lambda^*\left(z\right)\right)\right\}\right]>\mathbb{E}\left[R\left(z\right)\right]$ whenever A is symmetric or dominated by upside risk. Moreover, $\lim_{\alpha\to 0}\hat{R}\left(z,\lambda^*\left(z\right)\right)<0$ for $z<\hat{z}$ and $\lim_{\alpha\to 0}\frac{\hat{R}(z,\lambda^*(z))}{R(z,\lambda^*(z))}=\infty$ for $z>\hat{z}$, so that

$$\lim_{\alpha\to 0}\frac{\mathbb{E}\left[\max\left\{0,\hat{R}\left(z,\lambda^*\left(z\right)\right)\right\}\right]}{\mathbb{E}[R(z,\lambda^*\left(z\right))]}=\lim_{\alpha\to 0}\frac{C'(k^{LF})}{C'(k^*)}=\infty,$$

i.e. as $\alpha\to 0$ the incentive to over-invest becomes unboundedly large.

Proof of Theorem 2:

For a given aggregate investment $\mathcal{C}=\int_{i\in\mathcal{I}}C_i\left(k_i\right)di$, an allocation is aggregate Pareto efficient only if it maximizes aggregate earnings, solving problem P1

$$\max_{k_i,\bar{Z}_i,\lambda_i}\int_{i\in\mathcal{I}}k_i\int_{\bar{Z}_i}R_i\left(z,\lambda_i\left(z\right)\right)\varphi(z)\,dzdi \text{ s.t. } \mathcal{C}\geq\int_{i\in\mathcal{I}}C_i\left(k_i\right)di.$$

It follows immediately that an allocation $\left(k_i,\bar{Z}_i,\lambda_i\left(\cdot\right)\right)$ is aggregate Pareto efficient only if $\lambda_i\left(\cdot\right)=\lambda_i^*\left(\cdot\right)$, $\bar{Z}_i=\mathbb{R}$, i.e. hiring is efficient and there is no liquidation, and $C_i'\left(k_i\right)/\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)=\psi^{-1}$ is equalized across all firms. Let $k_i^*\left(\mathcal{C}\right)$ solve $C_i'\left(k_i^*\right)/\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)=1/\psi\left(\mathcal{C}\right)$ and let $D\left(\mathcal{C}\right)=\int_{i\in\mathcal{I}}k_i^*\left(\mathcal{C}\right)C_i'\left(k_i^*\left(\mathcal{C}\right)\right)di$ denote the Pareto-efficient aggregate earnings, given investment spending $\mathcal{C}$. An allocation is then aggregate Pareto efficient if it solves problem P2

$$\max_{\mathcal{C}\geq 0}\left\{u'\left(D\left(\mathcal{C}\right)\right)D\left(\mathcal{C}\right)-\mathcal{C}+\Gamma\left(u\left(D\left(\mathcal{C}\right)\right)-u'\left(D\left(\mathcal{C}\right)\right)D\left(\mathcal{C}\right)\right)\right\}$$

for some $\Gamma\geq 0$, which yields

$$u'\left(D\left(\mathcal{C}\right)\right)\left(1+\left(\Gamma-1\right)\chi\left(D\left(\mathcal{C}\right)\right)\right)=\frac{1}{D'\left(\mathcal{C}\right)}=\frac{1}{\psi\left(\mathcal{C}\right)}=\frac{C_i'\left(k_i^*\right)}{\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)}$$

for all $i\in\mathcal{I}$. To recap, an allocation $\left(k_i,\bar{Z}_i,\lambda_i\left(\cdot\right)\right)$ is aggregate Pareto efficient, if and only if $\lambda_i\left(\cdot\right)=\lambda_i^*\left(\cdot\right)$, $\bar{Z}_i=\mathbb{R}$, and $C_i'\left(k_i^*\right)/\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)=1/\psi\left(\mathcal{C}\right)$ for all $i\in\mathcal{I}$, where $\mathcal{C}\equiv\int_{i\in\mathcal{I}}C_i\left(k_i^*\right)di$ satisfies $1/\psi\left(\mathcal{C}\right)\geq u'\left(D\left(\mathcal{C}\right)\right)\left(1-\chi\left(D\left(\mathcal{C}\right)\right)\right)$.

The Laissez-faire equilibrium always satisfies $\lambda_i\left(\cdot\right)=\lambda_i^*\left(\cdot\right)$, so it is aggregate Pareto efficient if and only if it satisfies the other two conditions. Since k_i^{LF} satisfies $C_i'\left(k_i^{LF}\right)/\mathbb{E}\left(\hat{R}_i\left(z,\lambda_i^*\left(z\right)\right)\right)=1/Q^{LF}$ if there is no liquidation, the Laissez-faire equilibrium must be Pareto-inferior if $\bar{Z}_i\neq\mathbb{R}$ (inefficient liquidation) or $\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)/\mathbb{E}\left(\hat{R}_i\left(z,\lambda_i^*\left(z\right)\right)\right)\neq\mathbb{E}\left(R_j\left(z,\lambda_j^*\left(z\right)\right)\right)/\mathbb{E}\left(\hat{R}_j\left(z,\lambda_j^*\left(z\right)\right)\right)$ with positive probability for any two firms $i,j\in\mathcal{I}$ (heterogeneous capital market imperfections).

Conversely, if the Laissez-faire equilibrium satisfies $\bar{Z}_i=\mathbb{R}$ and $\mathbb{E}\left(R_i\left(z,\lambda_i^*\left(z\right)\right)\right)/\mathbb{E}\left(\hat{R}_i\left(z,\lambda_i^*\left(z\right)\right)\right)=$

$\Psi > 0$ for all $i \in \mathcal{I}$ (no liquidation and homogeneous capital market imperfections), then $\left(k_i^{LF}, \mathbb{R}, \lambda_i^*(\cdot)\right)$ solve problem P1 with $\psi\left(\mathcal{C}^{LF}\right) = Q^{LF}\Psi$. Moreover, from equation (16), $Q^{LF}\Psi = 1/u'\left(D^{LF}\right)$, and therefore the Laissez-faire equilibrium satisfies $u'\left(D\left(\mathcal{C}^{LF}\right)\right) = 1/\psi\left(\mathcal{C}^{LF}\right) = C_i'\left(k_i^{LF}\right)/\mathbb{E}\left(R_i\left(z, \lambda_i^*(z)\right)\right)$, which solves problem P1 when $\Gamma = 1$. Therefore, with no liquidation and homogeneous capital market imperfections, the Laissez-faire equilibrium maximiyes aggregate total surplus.

Lastly, $Q^{LF}\Psi = 1/u'\left(D^{LF}\right)$ implies $u'\left(D^{LF}\right)\mathbb{E}\left(R_i\left(z, \lambda_i^*(z)\right)\right) = 1/Q^{LF} \cdot \mathbb{E}\left(\hat{R}_i\left(z, \lambda_i^*(z)\right)\right)$ for all $i \in \mathcal{I}$ if and only if $\mathbb{E}\left(R_i\left(z, \lambda_i^*(z)\right)\right)/\mathbb{E}\left(\hat{R}_i\left(z, \lambda_i^*(z)\right)\right) = \Psi > 0$ for all $i \in \mathcal{I}$. Hence, the Laissez-faire equilibrium is intra-firm Pareto-efficient for all firms, iff it is aggregate Pareto efficient.

Proof of Proposition 5:

Whenever the laissez-faire equilibrium is Pareto-inferior aggregate dividends satisfy $D^{LF} = \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i\left(z, \lambda_i^*(z)\right)\varphi(z)\,dzdi < D^* = \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i\left(z, \lambda_i^*(z)\right)\varphi(z)\,dzdi$. Since $D^{LF}u'\left(D^{LF}\right) = P^{LF} = 1/Q^{LF}\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left(\max\left\{0, \hat{R}_i\left(z, \lambda_i^*(z)\right)\right\}\right)di = \int_{i\in\mathcal{I}} k_i^{LF} C_i'\left(k_i^{LF}\right)di$ at the Laissez-faire equilibrium and $D^*u'\left(D^*\right) = P^* = 1/Q^*\int_{i\in\mathcal{I}} k_i^*\mathbb{E}\left(R_i\left(z, \lambda_i^*(z)\right)\right)di = \int_{i\in\mathcal{I}} k_i^* C_i'\left(k_i^*\right)di$ at the first best, $D^* > D^{LF}$ implies $P^* > P^{LF}$ and $\int_{i\in\mathcal{I}} k_i^* C_i'\left(k_i^*\right)di > \int_{i\in\mathcal{I}} k_i^{LF} C_i'\left(k_i^{LF}\right)di$. The uniform constant elasticity assumption for $C_i(k)$ implies $\int_{i\in\mathcal{I}} k_i C_i'\left(k_i\right)di = (1+\kappa)\int_{i\in\mathcal{I}} C_i\left(k_i\right)di$ for all investment allocations $\{k_i\}$, and therefore $D^* > D^{LF}$ implies $\int_{i\in\mathcal{I}} C_i\left(k_i^*\right)di > \int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right)di$.

From equation (16), $Q^{LF}u'(D^{LF}) = \frac{\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left(\max\left\{0, \hat{R}_i\left(z, \lambda_i^*(z)\right)\right\}\right)di}{\int_{i\in\mathcal{I}} k_i^{LF}\int_{\bar{Z}_i} R_i\left(z, \lambda_i^*(z)\right)\varphi(z)\,dzdi}$, so $Q^{LF}u'(D^{LF}) \gtreqless 1$ if and only if $\int_{i\in\mathcal{I}} k_i^{LF}\mathbb{E}\left(\max\left\{0, \hat{R}_i\left(z, \lambda_i^*(z)\right)\right\}\right)di \gtreqless \int_{i\in\mathcal{I}} k_i^{LF}\int_{\bar{Z}_i} R_i\left(z, \lambda_i^*(z)\right)\varphi(z)\,dzdi$.

Proof of Proposition 6:

1. Partial Equilibrium: Since $\hat{R}\left(z, \hat{\lambda}(z)\right) > 0$ for all z, initial shareholders never prefer liquidation. Moreover, $\hat{R}\left(z, \hat{\lambda}(z)\right) > \hat{R}\left(z, \lambda^*(z)\right)$, $R\left(z, \hat{\lambda}(z)\right) < R\left(z, \lambda^*(z)\right)$, and $C'\left(\hat{k}\right) = \mathbb{E}\left(\hat{R}\left(z, \hat{\lambda}(z)\right)\right) > \mathbb{E}\left(\max\left\{0, \hat{R}\left(z, \lambda^*(z)\right)\right\}\right) = C'\left(k^{LF}\right)$, and therefore $\hat{k} > k^{LF}$. Initial shareholders' expected payoffs equal $\hat{R}\left(z, \hat{\lambda}(z)\right)\hat{k} - C\left(\hat{k}\right) = C'\left(\hat{k}\right)\hat{k} - C\left(\hat{k}\right)$ under shareholder value maximization and $\mathbb{E}\left(\max\left\{0, R\left(z, \lambda^*(z)\right)\right\}\right)k^{LF} - C\left(k^{LF}\right) = C'\left(k^{LF}\right)k^{LF} - C\left(k^{LF}\right)$ at the original Laissez-faire equilibrium. Since $C'(k)k - C(k)$ is increasing in k and $\hat{k} > k^{LF}$, initial shareholders gain from shareholder value maximization. The change in Social surplus satisfies

$$\mathbb{E}\left(R\left(z, \hat{\lambda}(z)\right)\right)\cdot\hat{k} - C\left(\hat{k}\right) - \left\{\int_{\bar{Z}} R\left(z, \lambda^*(z)\right)\varphi(z)\,dz \cdot k^{LF} - C\left(k^{LF}\right)\right\}$$

$$= \left(k^{LF} - \hat{k}\right)\left(C'(k) - \int_{\bar{Z}} R\left(z, \lambda^*(z)\right)\varphi(z)\,dz\right) + \hat{k}\left(\mathbb{E}\left(R\left(z, \hat{\lambda}(z)\right)\right) - \int_{\bar{Z}} R\left(z, \lambda^*(z)\right)\varphi(z)\,dz\right)$$

for some $k \in \left[k^{LF}, \hat{k}\right]$. If $\mathbb{E}\left(\max\left\{0, R\left(z, \lambda^*\left(z\right)\right)\right\}\right) \geq \int_{\bar{Z}} R\left(z, \lambda^*\left(z\right)\right) \varphi(z)\, dz \geq \mathbb{E}\left(R\left(z, \hat{\lambda}\left(z\right)\right)\right)$, it follows that $C'\left(k\right) \geq \mathbb{E}\left(\max\left\{0, R\left(z, \lambda^*\left(z\right)\right)\right\}\right)$ and both terms in this expression are negative, hence Social surplus (and a fortiori final shareholders' expected payoffs) are lower under shareholder value maximization. $\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right) di < \int_{i\in\mathcal{I}} C_i\left(\hat{k}_i\right) di$

2. General equilibrium: Aggregate dividends satisfy $\hat{D} = \int_{i\in\mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i\left(z\right)\right)\right) \varphi(z)\, dzdi$ and $\hat{D}u'\left(\hat{D}\right) = \hat{P} = (1+\kappa) \int_{i\in\mathcal{I}} C_i\left(\hat{k}_i\right) di$, under SHV, and $D^{LF} = \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i\left(z, \lambda_i^*\left(z\right)\right) \varphi(z)\, dzdi$ and $D^{LF}u'\left(D^{LF}\right) = P^{LF} = (1+\kappa) \int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right) di$ at the Laissez-faire equilibrium, where the last equality uses the uniform constant elasticity assumption for $C_i\left(k\right)$. Therefore $\hat{P} \gtreqless P^{LF}$ iff $\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right) di \gtreqless \int_{i\in\mathcal{I}} C_i\left(\hat{k}_i\right) di$ iff $\hat{D} \gtreqless D^{LF}$, which holds iff $\int_{i\in\mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i\left(z\right)\right)\right) \varphi(z)\, dzdi \gtreqless \int_{i\in\mathcal{I}} k_i^{LF} \int_{\bar{Z}_i} R_i\left(z, \lambda_i^*\left(z\right)\right) \varphi(z)\, dzdi$.

Since final shareholders must have $D > 0$ to justify $P > 0$, there is no SHV equilibrium with positive investment and equity trade if $\int_{i\in\mathcal{I}} \hat{k}_i \mathbb{E}\left(R_i\left(z, \hat{\lambda}_i\left(z\right)\right)\right) \varphi(z)\, dzdi \leq 0$. This result holds whenever $R_i\left(z, \hat{\lambda}\left(z\right)\right) < 0$ for high z and returns to scale $\alpha_i$are sufficiently low for a positive measure of firms.

Proof of Proposition 7:

Since $\lambda^B\left(z\right) = \arg\max_{\lambda\in[0,\lambda^*(z)]} \hat{R}\left(z, \lambda\right) = \min\left\{\hat{\lambda}\left(z\right), \lambda^*\left(z\right)\right\}$, $\hat{R}\left(z, \lambda^B\left(z\right)\right) > 0$ for all z, so initial shareholders never prefer liquidation. Moreover, $\lambda^B\left(z\right) \neq \hat{\lambda}\left(z\right)$ whenever $\hat{\lambda}\left(z\right) > \lambda^*\left(z\right)$, and $\lambda^B\left(z\right) \neq \lambda^*\left(z\right)$ whenever $\hat{\lambda}\left(z\right) < \lambda^*\left(z\right)$. Consider now the Paretian Social planner's problem given by equation (9) and suppose that $\lambda^B\left(z\right)$ maximized the social planner's return $R^{SP}\left(z, \lambda\right)$. For z such that $\lambda^B\left(z\right) = \hat{\lambda}\left(z\right) < \lambda^*\left(z\right)$, $\lambda^B\left(z\right)$ satisfies the Social Planner's FOC only if $\Gamma R_\lambda\left(z, \hat{\lambda}\left(z\right)\right) = 0$, which implies $\Gamma = 0$ since $R_\lambda\left(z, \hat{\lambda}\left(z\right)\right) = 0$ is ruled out by $\hat{\lambda}\left(z\right) < \lambda^*\left(z\right)$. At the same time, For z such that $\lambda^B\left(z\right) = \lambda^*\left(z\right) < \hat{\lambda}\left(z\right)$, $\lambda^B\left(z\right)$ satisfies the Social Planner's FOC only if $\left(1-\Gamma\right) \hat{R}_\lambda\left(z, \lambda^*\left(z\right)\right) = 0$, or $\Gamma = 1$. It follows that there does not exist any $\Gamma \geq 0$, for which $\lambda^B\left(z\right) = \min\left\{\hat{\lambda}\left(z\right), \lambda^*\left(z\right)\right\}$ maximizes the social planner's return for all z. Hence the Laissez-faire equilibrium with debt remains intra-firm Pareto-inferior. In general equilibrium, since $\lambda_i^B\left(\cdot\right) \neq \lambda_i^*\left(\cdot\right)$ for a positive measure of i it follows that $D^B < D^*$ and therefore that the Laissez-faire equilibrium with debt is also aggregate Pareto inferior.

1. Partial equilibrium: $\hat{R}\left(z, \lambda^B\left(z\right)\right) \geq \hat{R}\left(z, \lambda^*(z)\right)$ and $R\left(z, \lambda^B\left(z\right)\right) \leq R\left(z, \lambda^*(z)\right)$ with strict inequalities if $\lambda^*(z) > \hat{\lambda}\left(z\right)$, and $C'\left(k^B\right) = \mathbb{E}\left(\hat{R}\left(z, \lambda^B\left(z\right)\right)\right) > \mathbb{E}\left(\max\left\{0, \hat{R}\left(z, \lambda^*\left(z\right)\right)\right\}\right) = C'\left(k^{LF}\right)$, and therefore $k^B > k^{LF}$ and share prices are strictly higher with debt than at the equity-only equilibrium. Initial shareholders' expected payoffs equal $\mathbb{E}\left(\hat{R}\left(z, \lambda^B\left(z\right)\right)\right) k^B - C\left(k^B\right) = C'\left(k^B\right) k^B - C\left(k^B\right)$ with debt and $C'\left(k^{LF}\right) k^{LF} - C\left(k^{LF}\right)$ at the equity-only Laissez-faire

equilibrium. Since $C'(k)k - C(k)$ is increasing in k and $\hat{k} > k^{LF}$, issuing debt benefits initial shareholders. The change in Social surplus satisfies

$$\mathbb{E}\left(R\left(z,\lambda^{B}(z)\right)\right)\cdot k^{B} - C\left(k^{B}\right) - \left\{\int_{\bar{Z}} R\left(z,\lambda^{*}(z)\right)\varphi(z)\,dz\cdot k^{LF} - C\left(k^{LF}\right)\right\}$$

$$= \left(k^{LF}-k^{B}\right)\left(C'(k) - \int_{\bar{Z}} R\left(z,\lambda^{*}(z)\right)\varphi(z)\,dz\right) + k^{B}\left(\mathbb{E}\left(R\left(z,\lambda^{B}(z)\right)\right) - \int_{\bar{Z}} R\left(z,\lambda^{*}(z)\right)\varphi(z)\,dz\right)$$

for some $k\in\left[k^{LF},k^{B}\right]$. If $\mathbb{E}\left(\max\left\{0,R\left(z,\lambda^{*}(z)\right)\right\}\right) \geq \int_{\bar{Z}} R\left(z,\lambda^{*}(z)\right)\varphi(z)\,dz \geq \mathbb{E}\left(R\left(z,\lambda^{B}(z)\right)\right)$, where $\lambda^{B}(z) = \min\left\{\hat{\lambda}(z),\lambda^{*}(z)\right\}$, it follows that $C'(k) \geq \mathbb{E}\left(\max\left\{0,R\left(z,\lambda^{*}(z)\right)\right\}\right)$ and both terms in this expression are negative, hence debt reduces total surplus (and a fortiori final shareholders' expected payoffs).

2. General equilibrium: $D^{B}u'\left(D^{B}\right) = \hat{P} = (1+\kappa)\int_{i\in\mathcal{I}} C_i\left(k_i^{B}\right)di$ at the equilibrium with debt and $D^{LF}u'\left(D^{LF}\right) = P^{LF} = (1+\kappa)\int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right)di$ at the equity-only Laissez-faire equilibrium. Moreover, $D^{B} = \int_{i\in\mathcal{I}} k_i^{B}\mathbb{E}\left(R_i\left(z,\lambda_i^{B}(z)\right)\right)\varphi(z)\,dzdi$ and $D^{LF} = \int_{i\in\mathcal{I}} k_i^{LF}\int_{\bar{Z}_i} R_i\left(z,\lambda_i^{*}(z)\right)\varphi(z)\,dz\cdot di$. Therefore $P^{B} \gtreqless P^{LF}$, $D^{B} \gtreqless D^{LF}$ and $\int_{i\in\mathcal{I}} C_i\left(k_i^{B}\right)di \gtreqless \int_{i\in\mathcal{I}} C_i\left(k_i^{LF}\right)di$ if and only if $\int_{i\in\mathcal{I}} k_i^{B}\mathbb{E}\left(R_i\left(z,\lambda_i^{B}(z)\right)\right)\varphi(z)\,dzdi \gtreqless \int_{i\in\mathcal{I}} k_i^{LF}\int_{\bar{Z}_i} R_i\left(z,\lambda_i^{*}(z)\right)\varphi(z)\,dzdi$.