

Asset value and securitization under heterogeneous beliefs

Oliver Pardo¹

Departamento de Administración, Pontificia Universidad Javeriana, Colombia

ARTICLE INFO

Manuscript handled by Editor A. Toda

JEL classification:

G12
G21
G24
D84

Keywords:

Asset pricing
Securitization
Tranching
Financial innovation
Belief disagreement
Multiple equilibria

ABSTRACT

This paper studies how securitization affects asset prices and realized excess payoffs when investors hold heterogeneous beliefs. I extend the resale-option asset-pricing model of Harrison and Kreps (1978) by allowing the owner of a long-term asset to issue a finite set of asset-backed securities backed by its future cash flow. These securities are organized into tranches with different payment priorities. The no-tranching benchmark has a unique equilibrium price, but tranching can generate multiple equilibrium prices because separately priced tranches feed back into the asset's resale value. Securitization weakly raises the equilibrium prices of the underlying asset relative to the no-tranching benchmark. Furthermore, there exists a tranching that raises the lowest equilibrium price strictly above its no-tranching level if and only if, at some state, no agent perceives a distribution of the next-period asset price in the no-tranching benchmark that first-order stochastically dominates the distributions perceived by all other agents. Tranching has ambiguous effects on individual realized excess payoffs, but the aggregate realized excess payoff is nonpositive and weakly decreases as the tranching is refined. When one agent has the finest information partition, securitization may lower some agents' realized excess payoffs without raising anyone else's. The results characterize how belief disagreement shapes the effect of security design on asset prices and realized payoffs.

1. Introduction

Can securitization raise the market value of an asset without changing its cash flow? This paper studies this question in an infinite-horizon asset-pricing model with heterogeneous beliefs and short-selling constraints. I compare a benchmark economy in which investors trade a risk-free bond and a long-term asset with economies in which the owner can issue one-period asset-backed securities (ABS) collateralized by the asset and whose aggregate payoff equals its next-period resale value. The central mechanism is that an investor who is not the most optimistic buyer of the whole asset may be the most optimistic buyer of a particular tranche.

In the resale-option model of Harrison and Kreps (1978), the price of the whole asset reflects the valuation of the investor most optimistic about that asset. With securitization, the relevant object changes. Each tranche is priced by the investor most optimistic about the payoff range covered by that tranche. If perceived distributions over next-period asset prices cross, different tranches can be priced by different investors, and the sum of tranche prices can exceed the valuation of the whole asset by any single investor. Since ownership of the asset is necessary to issue the securities, the higher value of claims backed by the asset feeds back into the demand for the asset itself.

The paper formalizes this mechanism by allowing risk-neutral investors with a finite set of belief types to trade three types of assets in each period of a countably infinite sequence: i) an infinitely-supplied, risk-free, one-period bond, ii) a finitely-supplied, long-term asset that pays an uncertain dividend every period, and (iii) securities backed by the latter. By issuing ABS, the owner of the long-term asset commits to allocating the asset's next-period resale value among tranche holders. These securities have a waterfall payment structure: the most senior tranche promises a fixed amount unless the asset's resale value is below that amount; the next tranche is paid out of the resale value left after the senior tranche is paid; and so on. The process of issuing ABS with this priority structure is referred to as *tranching securitization*, or *tranching* for short.

The equilibrium price conditions rest on three assumptions. First, only owners of the long-term asset can issue ABS. Second, short positions in both the asset and the ABS are bounded. Third, one-period bonds are in infinite supply, so investors can borrow freely at the market interest rate. Because the long-term asset is in finite supply, the first two assumptions make the supply of ABS finite. Free borrowing, in turn, implies that each ABS is priced at the willingness to pay of the investor who is most optimistic about its next payout. The no-tranching benchmark has a unique equilibrium because its price-forming operator

E-mail address: pardoo@javeriana.edu.co.

¹ I thank Jason Donaldson, Erik Eyster, Leonardo Felli, Francesco Nava, Daniel Osorio, Petr Postl, Balazs Szentes, and participants at the LSE Summer Seminar for comments and suggestions. I also thank the anonymous referees for their careful and constructive reports. All errors and omissions are my own.

<https://doi.org/10.1016/j.jmateco.2026.103323>

Received 5 July 2026; Received in revised form 6 October 2026; Accepted 7 October 2026

Available online 15 October 2026

0304-4068/© 2026 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

is a contraction. With tranching, uniqueness can fail because changes in next-period asset prices can alter both which tranches pay and which theory has the highest willingness to pay for each tranche. Under any tranching, however, there is an equilibrium that assigns the lowest asset price in every state and another that assigns the highest.

Tranching weakly raises the price of the underlying asset relative to the no-tranching benchmark in every equilibrium. Furthermore, there exists a tranching that raises the lowest equilibrium price strictly above its no-tranching level if and only if, at some state, no agent perceives a distribution of the next-period asset price in the no-tranching benchmark that first-order stochastically dominates the distributions perceived by all other agents. When a dominant theory exists at every state, tranching cannot raise the lowest equilibrium price above the benchmark. If perceived distributions cross, tranching can create claims that are valued most highly by different investors. As a result, the gap between the asset price and any perceived present value of dividends can widen.

Tranching has ambiguous effects on individual realized excess payoffs. A realized excess payoff is the payoff from a purchased security minus the payoff that the same expenditure would have earned in the bond, averaged over time under the objective transition probabilities. Each position is perceived to break even at its equilibrium price, but its realized excess payoff need not be zero. For this analysis, I follow [Eyster and Piccione \(2013\)](#) and specialize beliefs to statistically consistent theories that may coarsen the state space. Securitization can transfer underperforming claims across investors: it may lower some investors' realized excess payoffs without raising anyone else's, but it can also make an investor's realized excess payoff less negative by shifting worse claims to someone else. At the aggregate level, the sum of realized excess payoffs from unit tranche purchases is nonpositive in every equilibrium and weakly decreases under refinement for both the lowest- and highest-price equilibria.

The mechanism is motivated by securitization before the 2007 financial crisis. Subprime mortgages were pooled into mortgage-backed securities (MBS); tranches of MBS were pooled into collateralized debt obligations (CDOs); and tranches of CDOs were repackaged into CDO-squared structures. Instruments such as MBS, CDOs, and CDO-squared structures are ABS, securities whose payout is derived from and collateralized by underlying assets. The possibility of creating highly rated securities backed by subprime mortgages eased credit to subprime borrowers, fueling an increase in house prices ([Brunnermeier, 2009](#)).² The purpose of the model is not to provide a full institutional account of that episode, but to isolate how belief disagreement and security design can interact to raise asset values.

This paper is organized as follows. The rest of this introduction presents a literature review. Section 2 presents the motivating example that illustrates the setup and gives a hint of the main results. Section 3 describes the general model and establishes existence and possible multiplicity of equilibrium. Section 4 studies the effects of tranching securitization on equilibrium prices of the underlying asset and on realized excess payoffs. Finally, Section 5 presents some conclusions and suggestions for further research.

1.1. Literature review

This paper builds on the asset-pricing literature under belief disagreement. [Milgrom and Stokey \(1982\)](#) show that, under common priors and an ex ante Pareto-optimal allocation, private information alone cannot generate gains from speculative trade when the trade's feasibility and acceptability to all traders are common knowledge. [Miller \(1977\)](#) emphasizes how disagreement and short-sale constraints lead asset prices to reflect the valuations of more optimistic investors, while [Harrison and Kreps \(1978\)](#) show that expected resale to a future

optimist can raise an asset's price above investors' valuations of its dividend stream. [Scheinkman and Xiong \(2004\)](#) survey this literature. The analysis of realized excess payoffs builds on [Eyster and Piccione \(2013\)](#), who study a long-lived asset under statistically consistent incomplete theories and evaluate investors' long-run realized returns, but do not allow the asset to be securitized. The present paper adds securitization to this line of analysis and studies how security design changes both the resale price of the underlying asset and the distribution of realized excess payoffs across agents.

The closest comparisons on securitization under diverse beliefs are [Guo et al. \(2014\)](#), [Broer \(2018\)](#), and [Ellis et al. \(2022\)](#). [Guo et al. \(2014\)](#) show that tranching a one-period payoff can raise an issuer's sale revenue and characterize optimal tranching. [Broer \(2018\)](#) shows how disagreement about the level and correlation of default risk can inflate the value of structured claims. [Ellis et al. \(2022\)](#) endogenize pooling and monotone security design in general equilibrium and characterize how investors sort across securities. Relative to these papers, the distinct object here is the feedback created by a long-lived underlying asset: securitization changes the asset's endogenous resale value, that resale value determines future tranche payoffs, and the loop can produce multiple equilibrium price functions. The long-lived structure also makes it possible to evaluate tranche buyers' long-run realized excess payoffs under the objective transition process.

The paper also relates to the literature on financial innovation, surveyed by [Duffie and Rahi \(1995\)](#) and [Allen and Barbalau \(2024\)](#). [Fostel and Geanakoplos \(2012\)](#) show that tranching and leverage can raise the price of an underlying asset, whereas credit default swaps can lower it. [Simsek \(2013a\)](#) shows how the form of disagreement affects investors' ability to borrow against an asset and how additional contracts redirect speculative positions. [Geanakoplos and Zame \(2014\)](#) develop a general-equilibrium framework with collateralized borrowing and default, while [Fostel and Geanakoplos \(2016\)](#) study how financial innovation affects asset prices and investment through the use of collateral. More generally, [Brock et al. \(2009\)](#), [Kubler and Schmedders \(2012\)](#), and [Simsek \(2013b\)](#) show that expanded asset menus can increase volatility or speculative risk. [Iachan et al. \(2021\)](#) study financial innovation as an expansion of portfolio choice in a two-period model, whereas the present paper examines how securitization feeds back into the endogenous resale value of a long-lived asset.

Finally, the paper relates to research on beliefs, financial innovation, and financial crises. [Gennaioli et al. \(2012\)](#) connect financial innovation with neglected risks and fragility. Models of diagnostic expectations, in which agents overreact to news when forming beliefs, generate credit and investment cycles ([Bordalo et al., 2018](#); [Maxted, 2024](#); [Bordalo et al., 2026](#)). Empirical and quantitative work links credit booms, asset-price growth, sentiment, and intermediation to financial crises ([Greenwood et al., 2022](#); [Krishnamurthy and Li, 2025](#)). These papers support the broader premise that distorted beliefs and financial intermediation can jointly generate price booms and fragility.

2. Motivating example

Consider an environment where trade takes place in a countably infinite sequence of periods. Each period there is an infinitely supplied bond used as a numeraire that entitles a payoff $R > 1$ next period. The bond can be traded for a finitely supplied asset that yields a random dividend d every period. There are three possible states of the world, collected in the state space $X := \{l, m, h\}$. The realizations of the dividend are such that $d(l) < d(m) < d(h)$. Securities backed by the long-term asset can be issued and traded.

There are two types of infinitely lived agents, $\mathcal{A}$ and $\mathcal{B}$. Both are risk-neutral. To keep the notation close to the general model, identify each type with the theory it uses. A theory is a forecasting rule for the next-period state. In the block interpretation developed in Section 3, theory $\mathcal{A}$ has a single block, $\{l, m, h\}$, so it does not condition its forecasts on the current state. Theory $\mathcal{B}$ has three singleton blocks,

² For a historical account of security innovation, see [Matthews \(1994\)](#).

$\{l\}$, $\{m\}$ and $\{h\}$, so it conditions its forecasts on the current state. For a theory $F \in \{A, B\}$, $Q_F(x, y)$ is the probability, according to F , of moving from current state x to next-period state y . In every state, theory A assigns probability $1/3$ to each state next period. Meanwhile, theory B assigns probability $1 - \epsilon$ to staying in the current state and is uniform otherwise:

$$Q_A(x, y) = \frac{1}{3} \quad \text{for all } x, y \in X, \quad (1)$$

$$Q_B(x, y) = \begin{cases} 1 - \epsilon, & \text{if } y = x, \\ \epsilon/2, & \text{if } y \neq x, \end{cases} \quad x, y \in X. \quad (2)$$

Let $P(x, y)$ denote the objective transition probability from state x to state y . Throughout the example, take $P = Q_B$. Thus, theory B coincides with the true transition process, whereas theory A is statistically consistent but does not condition on the current state: both theories imply the same long-run frequency for the three states, with each state occurring one third of the time.

The disagreement may emerge because agents use different forecasting tools or do not use all available information. The theories are maintained reduced-form forecasting rules, not Bayesian posteriors in a correctly specified common model. An agent may observe the current state yet use a theory that maps several states into the same forecast. Such a theory can be statistically consistent in the long run while remaining conditionally misspecified. The model deliberately abstracts from learning and switching among theories: the maintained rules persist as the state evolves.

2.1. Benchmark: A world without asset-backed securities

In the benchmark scenario, there are no asset-backed securities. Let $q : X \rightarrow [0, \infty)$ denote the underlying-asset price function. Because of limited short-selling, the market clears when the long-term asset's highest expected return matches the return of the short-term bond. Notation is simplified by assuming that the asset price is *cum-dividend* and that the declared dividend is to be paid next period.

Consider the state $x \in X$. When considering buying the long-term asset, both agents are willing to pay $d(x)/R$ plus the discounted asset price they expect for next period. For agent A , the expected price of the long-term asset is $\frac{1}{3}(q(l) + q(m) + q(h))$. For agent B , the same expectation is close to $q(x)$. If, as happens in equilibrium, $q(l) < q(m) < q(h)$, agent A buys the long-term asset in state l whereas agent B buys it in state h . Ex ante, it is not clear who will buy the asset in state m . Still, the solution to the price function q when ϵ goes to zero is given by

$$q(l) = \frac{d(l) + \frac{1}{3}(q(l) + q(m) + q(h))}{R} \quad (3)$$

$$q(m) = \frac{d(m) + \max \left\{ q(m), \frac{1}{3}(q(l) + q(m) + q(h)) \right\}}{R} \quad (4)$$

$$q(h) = \frac{d(h) + q(h)}{R} \quad (5)$$

For instance, if $R = 2$, $d(l) = 0$, $d(m) = 2$ and $d(h) = 3$, then $q(l) = 1$, $q(m) = 2$ and $q(h) = 3$.

To illustrate the conditions under which the introduction of ABS increases the long-term asset price, consider the perceived distributions regarding the asset price for the next period. Let $G_F(t | x)$ denote the cumulative distribution function (CDF) of the next-period benchmark price perceived by theory $F \in \{A, B\}$ conditional on current state $x \in X$, evaluated at $t \geq 0$. Fig. 1 plots these CDFs for each of the three current states. In states l and h , the distributions can be ranked according to first-order stochastic dominance, as one distribution always takes values no lower than the other. The agent whose beliefs first-order stochastically dominate those of the other will buy the asset. In state m , there is no stochastic dominance between $G_A(\cdot | m)$ and $G_B(\cdot | m)$. Still, whoever has highest expectation regarding the asset price will buy the asset. The nontrivial parts of the perceived distributions cross




Fig. 1. Perceived distributions on the asset price.

at $q(m)$: for $q(l) \leq t < q(m)$, $G_B(t | m) < G_A(t | m)$, whereas for $q(m) \leq t < q(h)$, $G_A(t | m) < G_B(t | m)$. The CDFs coincide below $q(l)$ and at or above $q(h)$. Thus, theory B is more optimistic about avoiding low price realizations, while theory A is more optimistic about the upper tail. This divergence can be exploited by dividing the asset's resale value between two securities: one that receives this value up to a fixed amount and another that receives any value above that amount. This is illustrated in the next subsection.

2.2. A world with asset-backed securities

Assume now that any owner of the long-term asset can issue two kinds of one-period ABS whose claims add up to the next-period value of the long-term asset. Denote the underlying-asset equilibrium price function in this securitized economy by q^{ABS} . The attachment point of a tranche is the asset-value threshold above which it begins to pay. In this example, the attachment point of the junior tranche is calibrated to coincide with the equilibrium asset price in state m , $q^{\text{ABS}}(m)$. The senior tranche pays $q^{\text{ABS}}(m)$ next period unless the state becomes l , in which case it pays $q^{\text{ABS}}(l)$. The junior tranche pays $q^{\text{ABS}}(h) - q^{\text{ABS}}(m)$ if the realized state is h and 0 otherwise. Let ϕ_s and ϕ_j denote the payoff functions of the senior and junior tranches, respectively. Then

$$\phi_s(t) = t - (t - q^{\text{ABS}}(m))^+, \quad (6)$$

$$\phi_j(t) = (t - q^{\text{ABS}}(m))^+, \quad (7)$$

Fig. 2 plots the payoff from each tranche as a function of the cum-dividend price of the underlying asset. Owning the junior tranche is equivalent to holding a one-period call option on the underlying asset with strike price $q^{\text{ABS}}(m)$. Owning the senior tranche is equivalent to holding a risk-free claim paying $q^{\text{ABS}}(m)$ and writing a one-period put option with the same strike. The threshold separating the two tranches is $q^{\text{ABS}}(m)$. It is the attachment point of the junior tranche and the detachment point of the senior tranche.




Fig. 2. Payoff functions from senior and junior tranches.

On an alternative interpretation, the buyer of the junior tranche becomes the owner of the asset, but has a debt liability to the buyer of the senior tranche and is using the underlying asset as collateral. In this example, the junior buyer promises to pay $q^{ABS}(m)$ to the senior buyer in the next period. However, if the value of the collateral is $q^{ABS}(l)$, the junior buyer defaults and the senior buyer seizes the collateral.

Because short sales are constrained, the price of each ABS equals the highest expected next-period payoff across agents, discounted at the interest rate. To start with, assume $q^{ABS}(l) < q^{ABS}(m) < q^{ABS}(h)$. At state m , agent B expects a payoff close to $q^{ABS}(m)$ from the senior tranche, while agent A expects $\frac{2}{3}q^{ABS}(m) + \frac{1}{3}q^{ABS}(l)$. Hence, the senior tranche will be bought by agent B . Meanwhile, agent B expects a payoff close to zero from the junior tranche, while agent A expects $q^{ABS}(h) - q^{ABS}(m)$ with one-third probability. Hence agent A will buy the junior tranche.

Let p_s and p_j denote the respective price functions of the senior and junior tranches. As ϵ tends to zero,

$$p_s(m) = \frac{q^{ABS}(m)}{R}$$

$$p_j(m) = \frac{\frac{1}{3}(q^{ABS}(h) - q^{ABS}(m))}{R}$$

Since different agents might have claims on the same asset, the total market value of the securities can exceed the value that any single agent assigns to those securities. To show this, note that at state m , the expectation on the asset price for next period is close to $q^{ABS}(m)$ for agent B and $\frac{1}{3}(q^{ABS}(l) + q^{ABS}(m) + q^{ABS}(h))$ for agent A . Then

$$p_s(m) + p_j(m) > \frac{q^{ABS}(m)}{R}$$

$$p_s(m) + p_j(m) > \frac{\frac{1}{3}(q^{ABS}(l) + q^{ABS}(m) + q^{ABS}(h))}{R}$$

which means that the payoff from selling the securities surpasses any expectation regarding the discounted price of the underlying asset.³

The equilibrium price for the long-term asset is given by a no-arbitrage condition. Any agent could buy the long-term asset and sell the asset-backed securities derived from it. In equilibrium, the long-term asset cum-dividend price has to equal the discounted declared dividend plus the market value of the asset-backed securities:

$$q^{ABS}(x) = \frac{d(x)}{R} + p_s(x) + p_j(x) \quad (8)$$

If the left-hand side of (8) were greater than the right-hand side, there would be no demand for the long-term asset. On the other hand, if the left-hand side were smaller than the right-hand side, demand would be infinite.

For states l and h , the equilibrium conditions can be derived by identifying who buys the tranches at these states. At state l , A 's

expected payoff from any tranche is greater than B 's. The opposite is true at state h . Hence both tranches are bought by A at state l and both tranches are bought by B at state h . Adding up the tranche prices yields the discounted long-term asset price expected by whoever buys both tranches. Hence the price equations for $q^{ABS}(l)$ and $q^{ABS}(h)$ resemble Eqs. (3) and (5). The same is not true for $q^{ABS}(m)$, however, since at state m each tranche is bought by different agents. In summary, the equilibrium price function q^{ABS} is given by

$$q^{ABS}(l) = \frac{d(l) + \frac{1}{3}(q^{ABS}(l) + q^{ABS}(m) + q^{ABS}(h))}{R} \quad (9)$$

$$q^{ABS}(m) = \frac{d(m) + q^{ABS}(m) + \frac{1}{3}(q^{ABS}(h) - q^{ABS}(m))}{R} \quad (10)$$

$$q^{ABS}(h) = \frac{d(h) + q^{ABS}(h)}{R} \quad (11)$$

By comparing Eq. (10) with Eq. (4), it can be seen that the introduction of ABS has increased the asset price at state m . This in turn increases the price at state l , since A 's expectation at state l is higher. Returning to the numerical example where $R = 2$, $d(l) = 0$, $d(m) = 2$ and $d(h) = 3$, the asset price increases by $\frac{1}{20}$ at state l and by $\frac{1}{4}$ at state m . Thanks to the reselling opportunity, the introduction of ABS increases the asset price in states in which ABS appear to be innocuous.

Because $P = Q_B$, theory B has rational expectations in this example and its realized excess payoff is zero both with and without ABS. This excess payoff is the one-period payoff from buying the asset or ABS, net of the payoff that the same expenditure would have obtained by buying the bond, averaged across states according to their long-run frequencies and evaluated under the objective transition P . Theory A , instead, earns a negative realized excess payoff. In the numerical example, its excess payoff falls from $-\frac{1}{3}$ without ABS to $-\frac{13}{30}$ with ABS. Appendix 6.1 reports the corresponding calculations. Thus, in this example, securitization increases the price of the underlying asset while making the less sophisticated agent's realized payoff worse. This payoff effect is not general: Section 4.3 later presents a case in which securitization makes an agent's realized excess payoff less negative, although it remains below zero.

3. The model

Consider an environment where trade takes place in each period of a countably infinite horizon. The assets traded in each period are of three kinds: 1) an infinitely supplied one-period bond used as a numeraire, 2) a single long-term asset that generates a random dividend every period, and 3) securities backed by the long-term asset.

3.1. States and transition probabilities

The background structure of the model is given by an exogenous Markov process defined over a finite state-space X . The transition probability from state $x \in X$ to state $y \in X$ is given by $P(x, y)$, which is assumed to be strictly positive for every pair of states. This assumption implies the existence of a unique stationary distribution μ , satisfying $\mu(x) > 0$ for every $x \in X$, $\sum_{x \in X} \mu(x) = 1$, and

$$\mu(y) = \sum_{x \in X} P(x, y)\mu(x) \quad (12)$$

for every $y \in X$. In the long-run, the world spends a fraction $\mu(x)$ of the time at state $x \in X$, regardless of the initial state.

3.2. Beliefs

All agents perceive that the state of the world evolves according to a Markov chain, but they may disagree on the transition probabilities. Each agent has a theory that represents her beliefs about the next-period state of the world. The finite set of theories in the market is represented by F . The transition probability from state $x \in X$ to state

³ See Fostel and Geanakoplos (2012) for a similar result.

$y \in X$, as perceived by theory $F \in \mathbf{F}$, is denoted by $Q_F(x, y)$. For the pricing analysis, $(Q_F)_{F \in \mathbf{F}}$ can be any finite collection of strictly positive stochastic matrices. In particular, the price-forming operator does not require the subjective transition matrices to be generated by partitions or to share the objective stationary distribution.

For the analysis of realized excess payoffs, agents' beliefs are restricted to the statistically consistent incomplete theories studied by Eyster and Piccione (2013). Under this restriction, every theory $F \in \mathbf{F}$ is a partition of X . An element in F is known as a block of theory F . If x and z are two different states in the same block of F , then $Q_F(x, y) = Q_F(z, y)$ for all $y \in X$. An agent using theory F therefore ignores the informational content of being at state x rather than at state z .

Let $F(x)$ be the block in F that contains $x \in X$. The statistically consistent transition matrix generated by the partition is

$$Q_F(x, y) = \frac{\sum_{z \in F(x)} P(z, y) \mu(z)}{\sum_{z \in F(x)} \mu(z)} \quad (13)$$

for all y in X . By averaging across states, it follows that $\sum_{x \in X} Q_F(x, y) \mu(x) = \mu(y)$, meaning that there is no disagreement in the long-run frequency of any state. Eq. (13) and the partition ordering are used only when explicitly invoked, primarily in the realized-payoff analysis. The pricing results require only the finite, full-support Markov beliefs described above.

This distinction also clarifies the interpretation of persistent belief errors. Statistical consistency is an unconditional calibration property: each partition-based theory predicts the correct long-run frequency of states. It does not imply that the theory uses all information in the current state or that repeated observations drive it toward P . Endogenous learning would require a model of how agents compare or revise forecasting rules and lies outside the equilibrium analyzed here.

If every block in a theory F is a singleton, it follows that $Q_F = P$. In this case, F implies rational expectations. At the other extreme, if $F = \{X\}$, then $Q_F(x, y) = \mu(y)$ for all x and y in X . This corresponds to the most naive theory, since its forecasts are not conditioned on the state.

The motivating example in Section 2 fits this notation by using the state space $X = \{l, m, h\}$, a constant interest rate R , and $\mathbf{F} = \{A, B\}$, with perceived transitions as in Eqs. (1) and (2). Consistent with the assumption that $P = Q_B$, theory A corresponds to the coarse partition $\{X\}$ and theory B corresponds to the singleton partition $\{\{l\}, \{m\}, \{h\}\}$.

One-period forward expectations play a pivotal role in the forthcoming analysis. To denote them, consider a random variable $s : X \rightarrow \mathbb{R}$. For any theory $F \in \mathbf{F}$, the expectation regarding s one period forward, conditional on being at state $x \in X$, is given by

$$E_F(s | x) := \sum_{y \in X} s(y) Q_F(x, y) \quad (14)$$

Denote by $E(s | x)$ the expectation for the case where $Q_F = P$. Eqs. (12), (13) and (14) imply that:

$$\sum_{x \in X} E_F(s | x) \mu(x) = \sum_{x \in X} E(s | x) \mu(x)$$

for all $F \in \mathbf{F}$. This means that the average one-period-ahead forecast of any theory matches the actual average of the random variable.

3.3. Asset-backed securities

The owner of the long-term asset can obtain liquidity (bonds) not only by reselling the asset but also by issuing and selling ABS. Each ABS has a maturity of one period and commits the owner of the long-term asset (the issuer) to transferring a nonnegative amount of bonds to its holder (the investor) next period. This amount depends on the realization of the asset price when payment is due. The collection of securities is backed by the long-term asset: the aggregate transfer from the issuer to investors equals the price of the underlying asset.




Fig. 3. Payoff function of a generic tranche τ .

Formally, write $\mathbb{R}_+ := [0, \infty)$ and let $\mathcal{T}$ be a finite collection of non-overlapping intervals that partition $\mathbb{R}_+$, the set of all conceivable realizations of the asset price. The partition $\mathcal{T}$ is known as a *tranching* and an interval $\tau \in \mathcal{T}$ is known as a *tranche*. Define $\underline{\tau} := \inf(\tau)$ and $\bar{\tau} := \sup(\tau)$. The payoff function $\phi_\tau : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ from a security associated with tranche τ is given by

$$\phi_\tau(t) := (t - \underline{\tau})^+ - (t - \bar{\tau})^+ \quad (15)$$

where $t \in \mathbb{R}_+$ is a realization for the underlying asset price. Under the convention $(t - \infty)^+ := 0$, Eq. (15) also applies when $\bar{\tau} = \infty$. A simple debt-equity structure is therefore a two-tranche waterfall: for a face value K , debt is the senior tranche $[0, K]$, with payoff $\min\{t, K\}$, and equity is the residual tranche $[K, \infty)$, with payoff $(t - K)^+$. For any price function $q : X \rightarrow \mathbb{R}_+$, $\phi_\tau(q)$ denotes the function $\phi_\tau \circ q : X \rightarrow \mathbb{R}_+$, so $\phi_\tau(q)(y) = \phi_\tau(q(y))$ for every $y \in X$.

The point $\underline{\tau}$ is known as the *attachment point* of tranche τ . This is because the tranche only starts to pay off when the realization for the asset price is above $\underline{\tau}$. Above $\bar{\tau}$, the payoff from tranche τ is capped at $\bar{\tau} - \underline{\tau}$ and the next more junior tranche (if any) begins to pay. Fig. 3 illustrates the payoff function for a generic tranche.

The securities are backed by the long-term asset in the sense that they are *budget balanced*: if t is the realization of the asset price at the expiration date, the tranche payoffs add up to t :

$$\sum_{\tau \in \mathcal{T}} \phi_\tau(t) = t$$

In addition, there is *limited liability* for the issuer: $\phi_\tau(t) \geq 0$ for all $\tau \in \mathcal{T}$ and $t \in \mathbb{R}_+$.

In the motivating example, the economy without ABS corresponds to $\mathcal{T}^0 := \{[0, \infty)\}$, while the securitized economy corresponds to $\mathcal{T}^{\text{ABS}} := \{\tau^s, \tau^j\}$, where $\tau^s = [0, q^{\text{ABS}}(m))$ and $\tau^j = [q^{\text{ABS}}(m), \infty)$. Eqs. (6) and (7) are the special cases of (15) associated with τ^s and τ^j , respectively.

3.4. The market

In each period, there is an infinitely-supplied bond used as a numeraire. Let $R : X \rightarrow (1, \infty)$ denote its payoff next period as a function of the current state, and let $d : X \rightarrow \mathbb{R}_+$ denote the nonnegative dividend declared by the long-term asset in the current state and paid next period. Assume that $d(x) > 0$ for at least one state $x \in X$. Since X is finite, set $\underline{R} := \min_{x \in X} R(x) > 1$. The bond can be traded for the long-term asset, which is in finite supply.

Short positions in ABS and in the long-term asset are assumed to be bounded. Under this assumption, the price of any ABS is determined by the willingness to pay of whoever is most optimistic about its future payoff. With unrestricted short selling and free borrowing, disagreement can generate unbounded positions, so equilibrium need not exist.

3.5. Equilibrium

In equilibrium, demand for the long-term asset and for each security must equal their respective supplies. The exogenous limits on short positions ensure that these supplies are finite. For $x \in X$, let $p_\tau(x)$ be the price of the security associated with tranche $\tau \in \mathcal{T}$ and $q(x)$ be the price of the long-term asset.

Definition 1. For a given tranching $\mathcal{T}$, a waterfall equilibrium is a function $q : X \rightarrow \mathbb{R}_+$ and a collection of functions $p_\tau : X \rightarrow \mathbb{R}_+$ for $\tau \in \mathcal{T}$ such that, for every $x \in X$,

$$q(x) = \frac{d(x)}{R(x)} + \sum_{\tau \in \mathcal{T}} p_\tau(x) \quad (16)$$

and, for every $\tau \in \mathcal{T}$,

$$p_\tau(x) = \frac{\max_{F \in \mathcal{F}} E_F(\phi_\tau(q) | x)}{R(x)} \quad (17)$$

Eq. (16) states that, in equilibrium, zero profits are made by buying the long-term asset and selling the securities derived from it. Eq. (17) states that the price of each security is given by the highest discounted expectation regarding the payoff from the corresponding tranche.

3.6. Existence and multiplicity

For any tranching $\mathcal{T}$, define a mapping $\Psi_{\mathcal{T}}$ from the set of price functions $\mathbb{R}_+^X$ to itself. For $q \in \mathbb{R}_+^X$ and $x \in X$, the mapping is given by

$$\Psi_{\mathcal{T}}(q)(x) = \frac{d(x) + \sum_{\tau \in \mathcal{T}} \max_{F \in \mathcal{F}} E_F(\phi_\tau(q) | x)}{R(x)} \quad (18)$$

By substituting Eq. (17) into Eq. (16), it follows that a long-term asset price function q is part of a waterfall equilibrium if and only if it is a fixed point of $\Psi_{\mathcal{T}}$. Denote the set of such price functions by

$$\mathcal{Q}_{\mathcal{T}} := \{q \in \mathbb{R}_+^X : q = \Psi_{\mathcal{T}}(q)\}.$$

The space $\mathbb{R}_+^X$ is endowed with the pointwise order: for $q, w \in \mathbb{R}_+^X$, write $q \leq w$ if $q(x) \leq w(x)$ for every $x \in X$. The symbols $\mathbf{0}$ and $\mathbf{1}$ denote the constant functions on X equal to zero and one, respectively. For any $q \in \mathbb{R}_+^X$, define $\|q\|_\infty := \max_{x \in X} |q(x)|$.

Proposition 1. For every finite tranching $\mathcal{T}$, the set $\mathcal{Q}_{\mathcal{T}}$ is a nonempty complete lattice under the pointwise order. In particular, there exist a least equilibrium price function $q_{\mathcal{T}}^{\min}$ and a greatest equilibrium price function $q_{\mathcal{T}}^{\max}$, satisfying

$$q_{\mathcal{T}}^{\min}(x) \leq q(x) \leq q_{\mathcal{T}}^{\max}(x)$$

for every $q \in \mathcal{Q}_{\mathcal{T}}$ and every $x \in X$.

Proof. Every ϕ_τ is nondecreasing, so $\Psi_{\mathcal{T}}$ is monotone. Because $\mathcal{T}$ is finite, it contains a unique tranche τ_∞ with $\bar{\tau}_\infty = \infty$. Define $a_{\mathcal{T}} := \frac{\tau_\infty}{R}$. The widths of the remaining tranches add up to $a_{\mathcal{T}}$. Hence, for every $q \in \mathbb{R}_+^X$ and $x \in X$,

$$\sum_{\substack{\tau \in \mathcal{T} \\ \bar{\tau} < \infty}} \max_{F \in \mathcal{F}} E_F(\phi_\tau(q) | x) \leq \sum_{\substack{\tau \in \mathcal{T} \\ \bar{\tau} < \infty}} (\bar{\tau} - \tau) = a_{\mathcal{T}}.$$

The payoff function of τ_∞ satisfies $\phi_{\tau_\infty}(q)(y) = (q(y) - a_{\mathcal{T}})^+ \leq q(y)$ for every $y \in X$, and therefore

$$\max_{F \in \mathcal{F}} E_F(\phi_{\tau_\infty}(q) | x) \leq \|q\|_\infty.$$

Let

$$\bar{d} := \max_{x \in X} d(x), \quad M_{\mathcal{T}} := \frac{\bar{d} + a_{\mathcal{T}}}{R - 1}.$$

It follows that

$$\|\Psi_{\mathcal{T}}(q)\|_\infty \leq \frac{\bar{d} + a_{\mathcal{T}} + \|q\|_\infty}{R}.$$

Consequently, $\Psi_{\mathcal{T}}$ maps $[0, M_{\mathcal{T}}]^X$ into itself. The same inequality implies that every fixed point of $\Psi_{\mathcal{T}}$ belongs to this set.

The interval $[0, M_{\mathcal{T}}]^X$ is a complete lattice under the pointwise order. Since $\Psi_{\mathcal{T}}$ is a monotone self-map of this interval, Tarski's fixed-point theorem implies that its fixed-point set is a nonempty complete lattice (Tarski, 1955). This fixed-point set is precisely $\mathcal{Q}_{\mathcal{T}}$. The following iterations identify its least and greatest elements constructively.

Define $q_{\mathcal{T}}^0 = \mathbf{0}$ and

$$q_{\mathcal{T}}^{n+1} = \Psi_{\mathcal{T}}(q_{\mathcal{T}}^n).$$

Since dividends are nonnegative, $q_{\mathcal{T}}^0 \leq q_{\mathcal{T}}^1$. Monotonicity then implies that the sequence is nondecreasing. It is bounded above by $M_{\mathcal{T}} \mathbf{1}$, so it converges pointwise to some $q_{\mathcal{T}}^{\min}$. The mapping $\Psi_{\mathcal{T}}$ is continuous because X , $\mathcal{F}$, and $\mathcal{T}$ are finite and every ϕ_τ is continuous. Thus $q_{\mathcal{T}}^{\min}$ is a fixed point. If $\hat{q}$ is any other fixed point, then $q_{\mathcal{T}}^0 = \mathbf{0} \leq \hat{q}$. Moreover, if $q_{\mathcal{T}}^n \leq \hat{q}$, monotonicity and the fixed-point property imply

$$q_{\mathcal{T}}^{n+1} = \Psi_{\mathcal{T}}(q_{\mathcal{T}}^n) \leq \Psi_{\mathcal{T}}(\hat{q}) = \hat{q}.$$

Therefore, $q_{\mathcal{T}}^n \leq \hat{q}$ for every n , and taking limits gives $q_{\mathcal{T}}^{\min} \leq \hat{q}$.

Similarly, define $\bar{q}_{\mathcal{T}}^0 = M_{\mathcal{T}} \mathbf{1}$ and

$$\bar{q}_{\mathcal{T}}^{n+1} = \Psi_{\mathcal{T}}(\bar{q}_{\mathcal{T}}^n).$$

This sequence is nonincreasing and bounded below by $\mathbf{0}$, so it converges to a fixed point $q_{\mathcal{T}}^{\max}$. Every fixed point is bounded above by $M_{\mathcal{T}} \mathbf{1}$; monotonicity therefore gives $\hat{q} \leq \bar{q}_{\mathcal{T}}^n$ for every fixed point $\hat{q}$ and every n . Taking limits shows that $\hat{q} \leq q_{\mathcal{T}}^{\max}$. $\square$

The proof is constructive: iterating $\Psi_{\mathcal{T}}$ from $\mathbf{0}$ converges to $q_{\mathcal{T}}^{\min}$, while iterating it from $M_{\mathcal{T}} \mathbf{1}$ converges to $q_{\mathcal{T}}^{\max}$.⁴

Waterfall equilibria need not be unique. To see why the contraction argument that guarantees uniqueness without tranching can fail, suppose that every next-period asset price rises by the same amount c . Without tranching, every theory's expected continuation value also rises by c . The current price at state x therefore rises by only $c/R(x)$. With tranching, however, the same increase changes how the asset's payoff is distributed across tranches. Although their combined payoff rises by only c in every future state, the sum of their highest expected payoffs can rise by more than c because each tranche is priced separately by the theory that values it most. The price-forming operator $\Psi_{\mathcal{T}}$ therefore need not be a contraction, and the resulting feedback can support multiple fixed points. The following example makes this possibility explicit.⁵

Example 1. Let $X = \{l, m, h\}$,

$$P = \frac{1}{8} \begin{pmatrix} 4 & 1 & 3 \\ 1 & 6 & 1 \\ 3 & 1 & 4 \end{pmatrix}, \quad R = \frac{13}{12}, \quad d = \frac{1}{12}(0, 3, 5).$$

The matrix P is strictly positive and doubly stochastic. Consider the coarse theory $\{X\}$ and the singleton theory. Their transition matrices are, respectively, the matrix with every row equal to $(1/3, 1/3, 1/3)$ and P . For the tranching

$$\mathcal{T} = \{[0, 3], [3, \infty)\},$$

the senior tranche $[0, 3]$ is debt with face value 3, and the junior tranche $[3, \infty)$ is residual equity. Direct substitution shows that $\Psi_{\mathcal{T}}$ has at least the following three fixed points, with coordinates ordered as (l, m, h) :

$$q^{(1)} = \frac{1}{327}(848, 930, 978) \approx (2.593, 2.844, 2.991),$$

⁴ Replication code for the numerical examples, including an executable notebook, is available in the [computational companion](#).

⁵ I thank an anonymous referee for suggesting the example on which this simplified version is based.

$$q^{(2)} = \frac{1}{97}(259, 285, 297) \approx (2.670, 2.938, 3.062),$$

and

$$q^{(3)} = \frac{1}{803}(2201, 2425, 2523) \approx (2.741, 3.020, 3.142).$$

Numerically implementing the monotone-iteration algorithm in [Proposition 1](#), starting from $\mathbf{0}$ and $M_{\mathcal{T}}\mathbf{1}$, yields $q^{(1)}$ and $q^{(3)}$, respectively. Hence, $q^{(1)} = q_{\mathcal{T}}^{\min}$ and $q^{(3)} = q_{\mathcal{T}}^{\max}$, while $q^{(2)}$ is an intermediate equilibrium. The first vector is also the unique equilibrium price without tranching. The approximate values make the payment patterns of the residual-equity tranche $[3, \infty)$ explicit. At $q^{(1)}$, all three prices are below 3, so this tranche never pays. At $q^{(2)}$, only $q^{(2)}(h) > 3$, so it pays only in state h . At $q^{(3)}$, both $q^{(3)}(m) > 3$ and $q^{(3)}(h) > 3$, so it pays in states m and h . Thus even a two-tranche debt–equity structure can support different patterns of active tranches and multiple equilibrium prices.

The multiplicity in this example is locally robust: sufficiently small changes in the transition probabilities, dividends, interest rate, and attachment point preserve three distinct equilibria. The equilibrium prices may change, but in each of the three equilibria the states in which residual equity pays remain as described above.⁶

Corollary 1. *The no-tranching benchmark always has a unique equilibrium price. In contrast, there exist economies in which a finite tranching has multiple equilibrium prices.*

Proof. Let $\mathcal{T}^0 = \{[0, \infty)\}$. Since $\phi_{[0, \infty)}(q) = q$, for any $q, w \in \mathbb{R}_+^X$ and $x \in X$,

$$\begin{aligned} & \left| \max_{F \in \mathcal{F}} E_F(q \mid x) - \max_{F \in \mathcal{F}} E_F(w \mid x) \right| \\ & \leq \max_{F \in \mathcal{F}} |E_F(q - w \mid x)| \leq \|q - w\|_{\infty}. \end{aligned}$$

Consequently,

$$\|\Psi_{\mathcal{T}^0}(q) - \Psi_{\mathcal{T}^0}(w)\|_{\infty} \leq \frac{1}{R} \|q - w\|_{\infty}.$$

Because $\frac{1}{R} > 1$, $\Psi_{\mathcal{T}^0}$ is a contraction and has a unique fixed point. [Example 1](#) establishes the second claim. $\square$

[Corollary 1](#) isolates multiplicity as a possible consequence of pricing tranches separately, rather than a feature of the no-tranching resale-option benchmark. It does not assert that every refinement increases the number of equilibria. [Proposition 1](#) nevertheless imposes order on the multiplicity: every collection of waterfall-equilibrium prices has a greatest lower bound and a least upper bound within the equilibrium set. [Example 1](#) shows that this complete-lattice structure may include equilibria strictly between its least and greatest elements.

3.7. Equilibrium selection

Comparisons based on a single equilibrium price function for each tranching require specifying which equilibrium is selected. To assess whether any alternative tranching can yield a higher asset price, the analysis compares tranchings using their least equilibrium prices.

The least equilibrium price has a finite-horizon interpretation. Consider an economy in which H dividend dates remain, the terminal continuation value is zero, and the same menu of one-period waterfall claims is available at every date. The following proposition shows that its price converges to the least waterfall-equilibrium price as the horizon increases.

⁶ Each equilibrium is determined locally by an invertible linear system. Its solution varies continuously with the parameters, and sufficiently small changes preserve the strict inequalities determining tranche payments and the theories pricing tranches with nonzero payoffs.

Proposition 2. *For every finite tranching $\mathcal{T}$ and every $H \geq 0$, the finite-horizon economy with H remaining dividend dates has a unique backward-induction price $q_{\mathcal{T}}^{[H]}$, determined by*

$$q_{\mathcal{T}}^{[0]} = \mathbf{0}, \quad q_{\mathcal{T}}^{[H+1]} = \Psi_{\mathcal{T}}(q_{\mathcal{T}}^{[H]}).$$

The sequence is nondecreasing and

$$q_{\mathcal{T}}^{[H]} \uparrow q_{\mathcal{T}}^{\min} \quad \text{as } H \rightarrow \infty.$$

Proof. When $H = 0$, the terminal condition uniquely determines $q_{\mathcal{T}}^{[0]} = \mathbf{0}$. Suppose that $q_{\mathcal{T}}^{[H]}$ has been determined. In the economy with $H + 1$ remaining dividend dates, the next-period payoff of tranche τ is $\phi_{\tau}(q_{\mathcal{T}}^{[H]})$. Eqs. (16) and (17) therefore uniquely determine

$$q_{\mathcal{T}}^{[H+1]} = \Psi_{\mathcal{T}}(q_{\mathcal{T}}^{[H]}).$$

Thus backward induction uniquely determines the finite-horizon price for every H .

This initial condition and recursion coincide with those of the lower sequence $\underline{q}_{\mathcal{T}}^H$ constructed in the proof of [Proposition 1](#). Hence $q_{\mathcal{T}}^{[H]} = \underline{q}_{\mathcal{T}}^H$ for every H . The proof of [Proposition 1](#) shows that this sequence is nondecreasing and converges to $q_{\mathcal{T}}^{\min}$. $\square$

When $H = 0$, the asset has no remaining payoff. With one dividend date remaining, there is no resale value, so

$$q_{\mathcal{T}}^{[1]}(x) = \frac{d(x)}{R(x)}.$$

With two dividend dates remaining, next period's asset price is $q_{\mathcal{T}}^{[1]}$, so the current price also includes the value of the waterfall claims that pay $\phi_{\tau}(q_{\mathcal{T}}^{[1]})$. Each subsequent iteration extends the horizon by one dividend date and applies the same resale-and-tranching pricing rule.

4. The effects of tranching

This section studies the effect of tranching securitization, understood as the introduction of new ABS with a waterfall payment constraint, on the equilibrium price of the long-term asset. Formally, tranching securitization means a refinement of the existing tranching:

Definition 2. An economy engages in tranching securitization if it switches from a tranching $\mathcal{T}$ to a tranching S , where S is a refinement of $\mathcal{T}$.

This means that agents who buy the long-term asset are able to sell a greater variety of tranches. [Appendix 6.2](#) shows that every finite iterated securitization structure backed by existing ABS generates the same set of equilibrium prices for the underlying asset as a suitable one-shot waterfall tranching. In particular, their least and greatest equilibrium prices coincide.

4.1. Effect on asset prices

Further securitization cannot decrease either extreme equilibrium price of the underlying asset:

Lemma 1. *If S is a refinement of $\mathcal{T}$, then:*

- (i) *for every $q \in Q_{\mathcal{T}}$, there is a $w \in Q_S$ such that $q \leq w$; and*
- (ii) *for every $w \in Q_S$, there is a $q \in Q_{\mathcal{T}}$ such that $q \leq w$.*

Consequently,

$$q_{\mathcal{T}}^{\min} \leq q_S^{\min} \quad \text{and} \quad q_{\mathcal{T}}^{\max} \leq q_S^{\max}.$$

Proof. See [Appendix 6.3](#).

The intuition for Lemma 1 is that splitting a tranche into smaller claims cannot reduce the sum of their prices. The highest valuation of the original tranche cannot exceed the sum of the highest valuations of its components, because each component may be priced by a different investor. Refinement therefore cannot reduce either extreme equilibrium price of the underlying asset.

The operator comparison also implies $q_T^{[H]} \leq q_S^{[H]}$ at every finite horizon H . Thus, the weak price comparison does not depend on taking the stationary infinite-horizon limit.

It is worth comparing the benchmark scenario, where there is no tranching, with any scenario where there is some degree of tranching:

Corollary 2. Let $\mathcal{T}^0 := \{[0, \infty)\}$ denote the no-tranching benchmark and $q_{\mathcal{T}^0}$ its unique equilibrium price. Then, for any tranching $\mathcal{T}$ and any $q \in \mathcal{Q}_{\mathcal{T}}$,

$$q_{\mathcal{T}^0}(x) \leq q(x) \quad (19)$$

for every $x \in X$.

Proof. Every tranching $\mathcal{T}$ refines $\mathcal{T}^0$. Lemma 1(ii) therefore places every $q \in \mathcal{Q}_{\mathcal{T}}$ weakly above an equilibrium of $\mathcal{T}^0$, whose unique equilibrium price is $q_{\mathcal{T}^0}$. $\square$

The weak comparison raises the question of when further tranching can increase the asset price strictly. For the analysis of issuer optimality, tranchings are evaluated at their least equilibrium prices, motivated by the finite-horizon limit in Proposition 2. A tranching is *issuer-optimal* if no alternative tranching yields a higher least equilibrium price in any state:

Definition 3. A tranching $\mathcal{T}$ is issuer-optimal if

$$q_S^{\min}(x) \leq q_T^{\min}(x)$$

for any tranching S and all $x \in X$.

In order to characterize the necessary and sufficient conditions for a tranching to be issuer-optimal, the following object is introduced: Let I_A be the indicator function for event A . For $t \in \mathbb{R}_+$, define $G_F^T(t|x)$ as the probability at state x that the asset price next period is no higher than t , as perceived by theory F when the tranching is $\mathcal{T}$ and prices are given by $q_T^{\min}$:

$$G_F^T(t|x) := E_F \left(I_{\{q_T^{\min} \leq t\}} \mid x \right)$$

Proposition 3. A tranching $\mathcal{T}$ is issuer-optimal if and only if, for each $\tau \in \mathcal{T}$ and $x \in X$, there is a theory $G \in \mathbf{F}$ such that

$$G_G^T(t|x) \leq G_F^T(t|x)$$

for every t in the interior of τ and all $F \in \mathbf{F}$. Moreover, if this condition fails, there is a refinement S of $\mathcal{T}$ such that

$$q_T^{\min}(x) < q_S^{\min}(x)$$

for every $x \in X$.

Proof. See Appendix 6.3.

Proposition 3 states that a tranching is issuer-optimal if and only if, for every state and every tranche, there is a theory that is weakly more optimistic than every other theory over the entire range of asset prices covered by the tranche. If this is not the case, then there is a refinement of the tranching that increases the least equilibrium asset price strictly in every state. This is because, at some state, the current tranche can be split so that another theory is willing to outbid the theory that values the original tranche most for at least one of the resulting smaller tranches.

Applying Proposition 3 to the no-tranching benchmark gives the following corollary.

Corollary 3. The no-tranching benchmark $\mathcal{T}^0$ is issuer-optimal if and only if, for every $x \in X$, there is a theory $G \in \mathbf{F}$ whose perceived distribution $G_G^{\mathcal{T}^0}(\cdot|x)$ first-order stochastically dominates $G_F^{\mathcal{T}^0}(\cdot|x)$ for every $F \in \mathbf{F}$.

Proof. Immediate from Proposition 3. $\square$

In the motivating example, the perceived distributions cross at state m (Fig. 1), so $\mathcal{T}^0$ is not issuer-optimal.

Finally, it is worth remarking that whenever the introduction of ABS increases the asset price, the asset will be traded at a price higher than anyone thinks it is worth:

Proposition 4. Let $\mathcal{T}$ be any finite tranching and let $q \in \mathcal{Q}_{\mathcal{T}}$. Then

$$q(x) \geq \frac{d(x) + \max_{F \in \mathbf{F}} E_F(q|x)}{R(x)}$$

for every $x \in X$. Moreover, if $q \neq q_{\mathcal{T}^0}$, the inequality is strict in at least one state.

Proof. Let

$$\Psi_{\mathcal{T}^0}(w)(x) := \frac{d(x) + \max_{F \in \mathbf{F}} E_F(w|x)}{R(x)}$$

be the price-forming operator without tranching. Since every $\mathcal{T}$ is a refinement of $\mathcal{T}^0$,

$$\Psi_{\mathcal{T}^0}(w) \leq \Psi_{\mathcal{T}}(w)$$

for every $w \in \mathbb{R}_+^X$. Hence, if $q \in \mathcal{Q}_{\mathcal{T}}$,

$$\Psi_{\mathcal{T}^0}(q) \leq \Psi_{\mathcal{T}}(q) = q.$$

If equality holds in every state, then q is a fixed point of $\Psi_{\mathcal{T}^0}$. This operator is a contraction with modulus $R^{-1} < 1$, so its unique fixed point is $q_{\mathcal{T}^0}$. Thus equality in every state implies $q = q_{\mathcal{T}^0}$. Equivalently, if $q \neq q_{\mathcal{T}^0}$, the inequality is strict in at least one state. $\square$

Intuitively, in an equilibrium different from the no-securitization benchmark, there has to be at least one state in which the tranching allows for an investor who otherwise would not have bought the whole asset to buy a piece of it. Since investors with different theories will be buying different tranches of the same asset, each of them will think that the others are overpaying for theirs.

4.2. Relation to fundamentals

Consider the relationship between the asset price and the perceived fundamental values, namely the present value of dividends perceived by each theory in the market. Let $v_F(x)$ be the present value of dividends perceived by $F \in \mathbf{F}$ at state $x \in X$. This value function is given by the recursive solution to

$$v_F(x) = \frac{d(x) + E_F(v_F|x)}{R(x)} \quad (20)$$

for $x \in X$.

Proposition 5. For any finite tranching $\mathcal{T}$, any $q \in \mathcal{Q}_{\mathcal{T}}$, and any collection of theories $\mathbf{F}$,

$$v_F(x) \leq q(x)$$

for all $x \in X$ and all $F \in \mathbf{F}$. Furthermore, if $\{[0, \infty)\}$ is not issuer-optimal, then there is a tranching $\mathcal{T}$ such that

$$v_F(x) < q(x)$$

for every $q \in \mathcal{Q}_{\mathcal{T}}$, all $x \in X$, and all $F \in \mathbf{F}$.

Proof. For each theory F , define the operator

$$\Gamma_F(w)(x) := \frac{d(x) + E_F(w|x)}{R(x)}.$$

The perceived fundamental value v_F is the fixed point of Γ_F . The benchmark price q_{T0} is the fixed point of

$$\Psi_{T0}(w)(x) := \frac{d(x) + \max_{G \in \mathbf{F}} E_G(w | x)}{R(x)}.$$

For every nonnegative price function w and every F , $\Gamma_F(w)(x) \leq \Psi_{T0}(w)(x)$. Since both operators are monotone contractions, iterating from the zero function and taking limits gives

$$v_F(x) \leq q_{T0}(x) \quad (21)$$

for all $x \in X$ and $F \in \mathbf{F}$. The first claim follows from inequality (21) and Corollary 2. If the no-tranching benchmark is not issuer-optimal, Proposition 3 gives a tranching $\mathcal{T}$ such that $q_{T0}(x) < q_T^{\min}(x)$ in every state. Combining this strict inequality with (21) and $q_T^{\min} \leq q$ for every $q \in \mathcal{Q}_T$ proves the second claim. $\square$

Proposition 5 states that if the no-tranching benchmark is not issuer-optimal, then there exists a tranching under which every equilibrium asset price strictly exceeds every theory's perceived present value of dividends. The full-support assumption ensures that the resale opportunities created by tranching raise the asset price above these perceived fundamental values in every state.

4.3. Effect on realized excess payoffs

This subsection imposes the partition-based, statistically consistent belief structure in Eq. (13). It studies how the introduction of ABS affects the realized excess payoff obtained by each theory in the market. For each unit of a security purchased, the excess payoff is the security's realized payoff minus the payoff from investing its purchase price in the bond.

The effect of the introduction of new ABS on realized excess payoffs is subtle. Agents who are not optimistic about the overall expected payoff from an underlying asset could be the most optimistic about the payoff being above a certain threshold. Without ABS, they will not buy the asset. By introducing the appropriate tranching, they will buy a security with an attachment point at that threshold. Tranching thus allows the issuer to sell claims that other agents value more highly than the issuer does. On the other hand, by being outbid by less sophisticated agents, they may be prevented from buying tranches that do not have a promising payoff. Except for special cases, the effect of the introduction of ABS on individual realized excess payoffs is ambiguous.

Formally, fix a tranching $\mathcal{T}$ and any equilibrium price $q \in \mathcal{Q}_T$. Consider buying one unit of a tranche whenever theory $F \in \mathbf{F}$ is selected as its buyer and financing that purchase by forgoing $p_\tau(x)$ units of the bond. Denote by $\Delta_F^T(q)$ the long-run average realized payoff from these purchases in excess of the corresponding bond payoff. In the examples below, the argument q is omitted when the equilibrium price has been specified. Let $\psi_\tau := \phi_\tau(q)$ be the function that maps states into equilibrium payoffs for tranche $\tau \in \mathcal{T}$. Tranche prices and buyer choices below are evaluated at q , with this dependence left implicit. Denote by $\mathbf{B}(\tau, x)$ the set of theories with the highest willingness to pay for tranche $\tau \in \mathcal{T}$ at state $x \in X$:

$$\mathbf{B}(\tau, x) = \{F \in \mathbf{F} : E_F(\psi_\tau | x) \geq E_G(\psi_\tau | x) \text{ for all } G \in \mathbf{F}\}$$

Fix once and for all a strict total priority ordering $\succ$ over $\mathbf{F}$, independent of the current state and the tranche. Let

$$\beta_T(\tau, x) := \max_{\succ} \mathbf{B}(\tau, x)$$

denote the selected buyer. Thus, whenever several theories have the same highest willingness to pay, the one with the highest fixed priority buys the tranche. The equilibrium price is independent of the priority ordering, whereas realized excess payoffs can depend on it. The *actual* as opposed to the *perceived* payoff from buying the security in tranche $\tau \in \mathcal{T}$ at state $x \in X$ is given by:

$$T(\psi_\tau)(x) := \sum_{y \in X} \psi_\tau(y) P(x, y)$$

Finally, let p_τ be the equilibrium price function for tranche $\tau \in \mathcal{T}$. The realized excess payoff assigned to theory $F \in \mathbf{F}$ is:

$$\Delta_F^T(q) = \sum_{x \in X} \left(\sum_{\tau \in \mathcal{T}} (T(\psi_\tau)(x) - R(x)p_\tau(x)) I_{\{F = \beta_T(\tau, x)\}} \right) \mu(x) \quad (22)$$

At the equilibrium price $p_\tau(x)$, the selected buyer perceives a one-unit position in tranche τ to break even relative to investing $p_\tau(x)$ in the bond. The *actual* average excess payoff $\Delta_F^T(q)$, evaluated under the objective transition probabilities, is not necessarily zero.⁷

Eq. (22) could be rewritten as the average gap between actual and perceived payoffs from each tranche. Since the equilibrium price for tranche τ is given by

$$p_\tau(x) = \frac{\max_{F \in \mathbf{F}} E_F(\psi_\tau | x)}{R(x)}$$

it follows that

$$\Delta_F^T(q) = \sum_{x \in X} \left(\sum_{\tau \in \mathcal{T}} (T(\psi_\tau)(x) - E_F(\psi_\tau | x)) I_{\{F = \beta_T(\tau, x)\}} \right) \mu(x)$$

Thus, the realized excess payoff is the average gap between actual and perceived payoffs for the tranches the theory buys.

An immediate result obtained from the statistical consistency of beliefs is that if $\mathbf{F} = \{F\}$, then $\Delta_F^T(q) = 0$. Thus, without competing theories, the long-run average payoff from the securities purchased by the agent equals the payoff from investing the same amounts in the bond.

The aggregate per-unit excess payoff of the selected buyers has a definite sign even though individual excess payoffs need not. Define

$$\Delta_{\text{agg}}^T(q) := \sum_{F \in \mathbf{F}} \Delta_F^T(q). \quad (23)$$

Summing across theories eliminates the selected-buyer indicators because the priority rule assigns each tranche in each state to exactly one theory. The next proposition establishes the sign of the aggregate excess payoff.

Proposition 6. For every finite tranching $\mathcal{T}$, every $q \in \mathcal{Q}_T$, and every fixed priority ordering,

$$\Delta_{\text{agg}}^T(q) \leq 0. \quad (24)$$

Proof. Summing Eq. (22) across theories and using definition (23) eliminates the selected-buyer indicators because each tranche in each state is assigned to exactly one theory. Hence

$$\Delta_{\text{agg}}^T(q) = \sum_{x \in X} \mu(x) \sum_{\tau \in \mathcal{T}} [T(\psi_\tau)(x) - R(x)p_\tau(x)]. \quad (25)$$

Budget balance implies

$$\sum_{\tau \in \mathcal{T}} T(\psi_\tau)(x) = \sum_{y \in X} P(x, y)q(y), \quad (26)$$

while Eq. (16) gives

$$R(x) \sum_{\tau \in \mathcal{T}} p_\tau(x) = R(x)q(x) - d(x). \quad (27)$$

Substituting Eqs. (26) and (27) into Eq. (25) yields

$$\Delta_{\text{agg}}^T(q) = \sum_{x \in X} \mu(x) \left[d(x) + \sum_{y \in X} P(x, y)q(y) - R(x)q(x) \right]. \quad (28)$$

Because μ is stationary under P ,

$$\sum_{x \in X} \mu(x) \sum_{y \in X} P(x, y)q(y) = \sum_{x \in X} \mu(x)q(x), \quad (29)$$

⁷ An individual theory's actual average excess payoff can be positive: see Example 5 in Eyster and Piccione (2013).

Substituting Eq. (29) into Eq. (28) gives the equivalent representation

$$\Delta_{\text{agg}}^T(q) = \sum_{x \in X} \mu(x) [d(x) - (R(x) - 1)q(x)]. \quad (30)$$

It remains to establish the sign. Tranche pricing implies

$$R(x)p_\tau(x) = \max_{F \in \mathcal{F}} E_F(\psi_\tau | x). \quad (31)$$

Substituting Eq. (31) into Eq. (25) gives

$$\Delta_{\text{agg}}^T(q) = \sum_{\tau \in \mathcal{T}} \left[\sum_{x \in X} \mu(x) T(\psi_\tau)(x) - \sum_{x \in X} \mu(x) \max_{F \in \mathcal{F}} E_F(\psi_\tau | x) \right] \quad (32)$$

$$\leq \sum_{\tau \in \mathcal{T}} \left[\sum_{x \in X} \mu(x) T(\psi_\tau)(x) - \max_{F \in \mathcal{F}} \sum_{x \in X} \mu(x) E_F(\psi_\tau | x) \right]. \quad (33)$$

The inequality uses that the average of pointwise maxima is at least the maximum of the averages. For every theory F and tranche τ , statistical consistency implies

$$\sum_{x \in X} \mu(x) E_F(\psi_\tau | x) = \sum_{x \in X} \mu(x) T(\psi_\tau)(x). \quad (34)$$

Eq. (34) implies that the right-hand side of Eq. (33) is zero, which proves Eq. (24). $\square$

Corollary 4. For every finite tranching $\mathcal{T}$ and every $q \in \mathcal{Q}_{\mathcal{T}}$,

$$\Delta_{\text{agg}}^T(q_{\mathcal{T}}^{\max}) \leq \Delta_{\text{agg}}^T(q) \leq \Delta_{\text{agg}}^T(q_{\mathcal{T}}^{\min}) \leq 0.$$

The aggregate payoff is strictly lower at $q_{\mathcal{T}}^{\max}$ than at $q_{\mathcal{T}}^{\min}$ whenever these price functions differ.

Proof. For any $q, w \in \mathcal{Q}_{\mathcal{T}}$ with $q \leq w$, Eq. (30) gives

$$\Delta_{\text{agg}}^T(w) - \Delta_{\text{agg}}^T(q) = - \sum_{x \in X} \mu(x) (R(x) - 1) [w(x) - q(x)] \leq 0.$$

Since $\mu(x) > 0$ and $R(x) > 1$ for every x , the inequality is strict if $q \neq w$. Proposition 1 gives $q_{\mathcal{T}}^{\min} \leq q \leq q_{\mathcal{T}}^{\max}$, and Proposition 6 gives the final inequality. $\square$

Corollary 5. If S refines $\mathcal{T}$, then

$$\Delta_{\text{agg}}^S(q_S^{\min}) \leq \Delta_{\text{agg}}^T(q_{\mathcal{T}}^{\min}),$$

$$\Delta_{\text{agg}}^S(q_S^{\max}) \leq \Delta_{\text{agg}}^T(q_{\mathcal{T}}^{\max}).$$

Each inequality is strict whenever the corresponding equilibrium price functions differ.

Proof. For any $q \in \mathcal{Q}_{\mathcal{T}}$ and $w \in \mathcal{Q}_S$ with $q \leq w$, Eq. (30) gives

$$\Delta_{\text{agg}}^S(w) - \Delta_{\text{agg}}^T(q) = - \sum_{x \in X} \mu(x) (R(x) - 1) [w(x) - q(x)] \leq 0.$$

Lemma 1 orders both pairs of extreme equilibrium prices. Applying this identity to each pair proves the comparisons. Because $\mu(x) > 0$ and $R(x) > 1$ for every x , each comparison is strict whenever the corresponding price functions differ. $\square$

Proposition 6 establishes a nonpositive aggregate realized excess payoff at every equilibrium. For a fixed tranching, Corollary 4 shows that the least equilibrium maximizes this payoff and the greatest equilibrium minimizes it, with a strict difference when there is multiplicity. The least equilibrium therefore yields the smallest aggregate loss. Corollary 5 shows that refinement weakly reduces the aggregate payoff under either extreme-equilibrium selection, strictly whenever the corresponding prices change. This measure sums the realized excess payoffs from buying one unit of each tranche, not from agents' actual portfolio holdings. The ranking across equilibria need not hold for each theory and is not a ranking of overall welfare. A nonpositive aggregate payoff does not rule out a positive excess payoff for a particular theory.

When one theory refines every other theory, the realized excess payoff is nonpositive for every theory.

Proposition 7. Suppose that $\mathcal{G} \in \mathcal{F}$ refines every $F \in \mathcal{F}$. Then, for every finite tranching $\mathcal{T}$, every $q \in \mathcal{Q}_{\mathcal{T}}$, and every fixed priority ordering,

$$\Delta_{\mathcal{G}}^T(q) = 0 \quad \text{and} \quad \Delta_F^T(q) \leq 0 \quad \text{for every } F \in \mathcal{F} \setminus \{\mathcal{G}\}.$$

Proof. See Appendix 6.4.

Proposition 7 extends the no-securitization result in Proposition 5 of Eyster and Piccione (2013): when one theory refines every other theory, no theory outperforms the bond on average at any waterfall equilibrium under any finite tranching. The proposition does not determine how a change in tranching affects each theory's payoff. In the motivating example, however, securitization lowers $\mathcal{A}$'s realized excess payoff while $\mathcal{B}$'s remains unchanged at zero. In states where, without securitization, the most sophisticated agent would have bought the whole asset, the tranches not bought by this agent underperform the same expenditure in the bond on average. Thus, securitization can lower one theory's realized excess payoff without increasing the payoff of any other theory.

Nevertheless, the introduction of ABS can increase a theory's realized excess payoff. Appendix 6.5 presents a full-support, three-theory example in which securitization induces a coarse theory to buy a mezzanine tranche. As a consequence, theory $\mathcal{B}$ has a higher realized excess payoff than it would have without securitization, even though its payoff remains negative:

$$\Delta_{\mathcal{B}}^{\tau^1} = -\frac{1}{16} > -\frac{373}{4464} = \Delta_{\mathcal{B}}^{\tau^0}.$$

Theory $\mathcal{A}$ has zero realized excess payoff under both tranchings, while theory $\mathcal{C}$'s payoff falls from zero to $-1/36$. Thus, securitization can transfer losses across agents: it may lower one theory's realized excess payoff while improving another theory's payoff relative to the benchmark, even if the improved payoff remains negative.

4.4. Other types of contracts

Apart from having their real-life counterparts, the only appealing features of securities with a waterfall structure are: (i) their payoff is monotone in the cash flow from the underlying asset, (ii) as shown in Appendix 6.2, they implicitly represent any case where only debt and equity can be issued and debt itself can be used as collateral. However, no argument has been given in this paper for agents to choose or be constrained by this class of contracts.

Without being constrained to issue ABS with a waterfall payment structure, agents could exploit the divergence of beliefs even further by creating contracts that do not satisfy this constraint. For instance, consider the equilibrium of Section 2.2. The senior tranche pays $q^{\text{ABS}}(m)$ if the state in the next period is m or h and $q^{\text{ABS}}(l)$ otherwise. At state m , this tranche is bought by agent $\mathcal{B}$. This agent could create a new security, backed by the senior tranche, that pays $q^{\text{ABS}}(m)$ at state h and 0 otherwise. This security cannot be written as a waterfall claim on the senior tranche because its payoff is not determined solely by the senior tranche's payoff: the senior tranche pays $q^{\text{ABS}}(m)$ in both states m and h , whereas the new security pays zero in m and $q^{\text{ABS}}(m)$ in h . This new security is almost worthless for $\mathcal{B}$, given that a realization h next period looks very unlikely. However, agent $\mathcal{A}$ would be willing to pay $\frac{1}{3} \frac{q^{\text{ABS}}(m)}{R}$ for it.

As an alternative to the waterfall equilibrium, consider the following asset-backed securities: for $y \in X$, the y -security pays $q(y)$ in the next period if the realized state is y and 0 otherwise. Define the vector b and the nonnegative matrix A by

$$b(x) := \frac{d(x)}{R(x)}, \quad A(x, y) := \frac{\max_{F \in \mathcal{F}} Q_F(x, y)}{R(x)}.$$

An equilibrium price $\bar{q}$ under this state-contingent securitization structure must satisfy

$$\bar{q} = b + A\bar{q}. \quad (35)$$

The matrix A is not a transition matrix: because the maximization is performed separately across next-period states, different entries in the same row may be determined by different theories, and the row sum can exceed one. Therefore, $R(x) > 1$ alone does not guarantee the existence of a nonnegative solution to Eq. (35).

Proposition 8. *Eq. (35) admits a nonnegative solution if and only if all eigenvalues of A lie strictly inside the unit circle. In that case, the solution is unique and*

$$\bar{q} = (I - A)^{-1}b = \sum_{n=0}^{\infty} A^n b.$$

Proof. If all eigenvalues of A lie strictly inside the unit circle, then the matrix series converges and

$$(I - A)^{-1} = \sum_{n=0}^{\infty} A^n$$

exists and is nonnegative. This proves existence, uniqueness, and nonnegativity of $\bar{q} = (I - A)^{-1}b$.

Conversely, suppose that Eq. (35) has a nonnegative solution $\bar{q}$. Since every subjective transition matrix is strictly positive, A is strictly positive. Hence A has a positive eigenvalue λ^* and a positive left eigenvector v such that $v'A = \lambda^*v'$, with $|\lambda| \leq \lambda^*$ for every eigenvalue λ of A . Multiplying Eq. (35) by v' gives

$$(1 - \lambda^*)v'\bar{q} = v'b.$$

The right-hand side is strictly positive because $b \geq 0$ and at least one dividend is positive. Since $v'\bar{q} \geq 0$, the equality requires $\lambda^* < 1$. Therefore, all eigenvalues of A lie strictly inside the unit circle. $\square$

The term $A^n b$ captures the contribution of dividends reached after n rounds of resale-price feedback. Convergence of the matrix series $\sum_{n=0}^{\infty} A^n$ means that discounting dominates the amplification created by selling each state-contingent security to its most optimistic buyer, so the contribution of increasingly distant resale opportunities vanishes.

Proposition 9. *Suppose that Eq. (35) admits a nonnegative solution $\bar{q}$. Then every finite tranching $\mathcal{T}$ has a unique waterfall-equilibrium price $q_{\mathcal{T}}$. Moreover,*

$$q_{\mathcal{T}} \leq \bar{q}.$$

Proof. For any $r, s \in \mathbb{R}_+$, the monotonicity and budget balance of the tranche payoff functions imply

$$\sum_{\tau \in \mathcal{T}} |\phi_{\tau}(r) - \phi_{\tau}(s)| = |r - s|.$$

Using the triangle inequality and $|\max_i a_i - \max_i b_i| \leq \max_i |a_i - b_i|$, it follows that, for any nonnegative price functions q, w ,

$$\begin{aligned} |\Psi_{\mathcal{T}}(q)(x) - \Psi_{\mathcal{T}}(w)(x)| &\leq \frac{1}{R(x)} \sum_{\tau \in \mathcal{T}} \max_{F \in \mathcal{F}} \sum_{y \in X} Q_F(x, y) |\phi_{\tau}(q(y)) - \phi_{\tau}(w(y))| \\ &\leq \frac{1}{R(x)} \sum_{\tau \in \mathcal{T}} \sum_{y \in X} \max_{F \in \mathcal{F}} Q_F(x, y) |\phi_{\tau}(q(y)) - \phi_{\tau}(w(y))| \\ &= \sum_{y \in X} A(x, y) |q(y) - w(y)|. \end{aligned}$$

If q and w are both fixed points, this inequality gives $|q - w| \leq A|q - w|$. Since A is nonnegative, repeated multiplication gives

$$|q - w| \leq A^n |q - w|$$

for every $n \geq 1$. Proposition 8 implies that A^n converges to zero. Hence $q = w$. Existence follows from Proposition 1, so $Q_{\mathcal{T}}$ is a singleton; denote its element by $q_{\mathcal{T}}$.

It remains to prove the upper bound. For every nonnegative price function w and every tranche τ ,

$$\max_{F \in \mathcal{F}} E_F(\phi_{\tau}(w) | x) = \max_{F \in \mathcal{F}} \sum_{y \in X} Q_F(x, y) \phi_{\tau}(w(y))$$

$$\leq \sum_{y \in X} \max_{F \in \mathcal{F}} Q_F(x, y) \phi_{\tau}(w(y)).$$

Summing over $\tau \in \mathcal{T}$ and using budget balance gives

$$\Psi_{\mathcal{T}}(w) \leq b + Aw.$$

Applying this inequality to the unique fixed point gives

$$q_{\mathcal{T}} \leq b + Aq_{\mathcal{T}},$$

or $(I - A)q_{\mathcal{T}} \leq b$. By Proposition 8, $(I - A)^{-1}$ is nonnegative. Multiplication by this matrix yields $q_{\mathcal{T}} \leq \bar{q}$. $\square$

The matrix A depends on beliefs and interest rates, not on the tranching. When all its eigenvalues lie strictly inside the unit circle, every finite waterfall tranching has a unique equilibrium. Tranching may still raise the asset price, but it cannot generate multiplicity. Failure of the eigenvalue condition does not by itself imply multiplicity.

State-contingent securities allow the issuer to exploit disagreement more flexibly than waterfall tranches. Proposition 9 also shows that, under this eigenvalue condition, the state-contingent price bounds the waterfall-equilibrium price from above.⁸ The lack of a belief that stochastically dominates all others in at least one state is a necessary condition for waterfall tranching to increase the asset price only because of the waterfall constraint. State-contingent securitization, in contrast, can exploit any disagreement about transition probabilities, provided that the resulting price system satisfies the existence condition above.

5. Conclusion

This paper has shown the effect of the introduction of asset-backed securities under heterogeneous beliefs and short-selling constraints. The heterogeneity is motivated by assuming that agents are not capable of processing all information available. Hence, different agents may ignore different information, leading to divergent expectations about the assets' payouts. This heterogeneity can be exploited by financial institutions through the issuing of asset-backed securities to be traded among those agents with divergent beliefs.

The no-tranching benchmark has a unique equilibrium, but waterfall tranching can create multiplicity because tranche activation and the identity of the marginal buyer feed back into the resale price of the underlying asset. Consequently, the same dividends, interest rates, beliefs, and tranching can support different equilibrium price functions. Example 1 shows that this multiplicity can arise even with a two-tranche debt-equity structure. Securitization can therefore affect both the level of asset prices and whether the equilibrium price is unique.

Refining a tranching weakly increases both the least and the greatest equilibrium prices of the underlying asset in every state. For issuer optimality, tranchings are compared by their least equilibrium prices. A tranching is issuer-optimal precisely when, at its least equilibrium and within each tranche and state, one perceived distribution first-order stochastically dominates the others on the relevant payoff range. If this dominance condition fails, there exists a refinement that raises the least equilibrium price strictly in every state.

The paper has also studied the effect of securitization on agents' realized excess payoffs. Whoever buys an asset may be subject to a winner's curse: the payoff they expect on the asset may exceed its objective average payoff, even though the position is perceived to break even relative to the bond. Under statistically consistent beliefs, the aggregate realized excess payoff from unit tranche purchases is non-positive in every equilibrium and weakly decreases under refinement

⁸ A related result appears in Allen and Gale (1988): in a different security-design environment with short-sale constraints, market value is maximized by designing securities so that, in each state, the asset's entire payoff goes to the investor group that values it most.

for both the least and greatest equilibria. Individual effects remain ambiguous: securitization may allow one agent to transfer part of an underperforming claim to another, but it may also decrease some agents' realized excess payoffs without increasing anyone else's.

Because agents are risk-neutral, ABS do not generate risk-sharing gains in the present model. With risk-averse agents, however, securitization could generate such gains. Determining whether these gains outweigh the potential welfare costs of belief-driven speculation would require a richer model. The belief-neutral efficiency criterion of Brunnermeier et al. (2014) provides one possible approach.

In a full theoretical account of financial crises, both the boom and the bust of asset prices should be explained. This paper isolates how securitization allows belief disagreement to raise asset prices, but it does not generate a collapse. One extension would endogenize learning or switching among theories after persistently disappointing realized payoffs; such behavior is not part of the maintained-belief equilibrium studied here. Another would introduce liquidity constraints, so investors holding underperforming ABS eventually lose the capacity to roll over their debts. A formalization of these narratives is left for future research.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the revision of this work, the author used OpenAI Codex to help edit the text, review proofs, check notation, and write and test the code in the computational companion. The author reviewed and edited the resulting text and code and takes full responsibility for the manuscript and computational companion.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix 6

6.1. Price equations for the motivating example

This appendix collects the systems solved in Section 2. Consider the limiting case in which ϵ tends to zero and both price functions are increasing across states: $q(l) < q(m) < q(h)$ and $q^{\text{ABS}}(l) < q^{\text{ABS}}(m) < q^{\text{ABS}}(h)$.

In the benchmark without ABS, the price function q solves

$$\begin{aligned} q(l) &= \frac{d(l) + \frac{1}{3}(q(l) + q(m) + q(h))}{R}, \\ q(m) &= \frac{d(m) + \max \left\{ q(m), \frac{1}{3}(q(l) + q(m) + q(h)) \right\}}{R}, \\ q(h) &= \frac{d(h) + q(h)}{R}. \end{aligned}$$

For $R = 2$, $d(l) = 0$, $d(m) = 2$ and $d(h) = 3$, the solution is $q(l) = 1$, $q(m) = 2$ and $q(h) = 3$.

With senior and junior ABS, the price function q^{ABS} solves

$$\begin{aligned} q^{\text{ABS}}(l) &= \frac{d(l) + \frac{1}{3}(q^{\text{ABS}}(l) + q^{\text{ABS}}(m) + q^{\text{ABS}}(h))}{R}, \\ q^{\text{ABS}}(m) &= \frac{d(m) + q^{\text{ABS}}(m) + \frac{1}{3}(q^{\text{ABS}}(h) - q^{\text{ABS}}(m))}{R}, \\ q^{\text{ABS}}(h) &= \frac{d(h) + q^{\text{ABS}}(h)}{R}. \end{aligned}$$

For the same parameter values, the solution is $q^{\text{ABS}}(l) = \frac{21}{20}$, $q^{\text{ABS}}(m) = \frac{9}{4}$ and $q^{\text{ABS}}(h) = 3$.

The same numerical example illustrates the realized-payoff effect discussed in Section 4.3. To keep notation consistent with that subsection, let $T(\psi)(x) = \sum_{y \in X} \psi(y)P(x, y)$ denote the objective expectation of a payoff vector $\psi : X \rightarrow \mathbb{R}_+$ from state x . If theory F buys a security with payoff ψ at price $p_\psi(x)$, its contribution to the realized excess payoff is $T(\psi)(x) - Rp_\psi(x)$. Since the buyer's equilibrium price satisfies $Rp_\psi(x) = E_F(\psi | x)$, this contribution can be written as

$$T(\psi)(x) - E_F(\psi | x).$$

The long-run average gives weight $\frac{1}{3}$ to each state.

In the benchmark without ABS, take $\psi = q$; the current dividend is known at state x , so it cancels from the actual-versus-perceived payoff gap. Use a priority ordering in which $A > B$. Theory A buys the asset in state l and wins the tie with B in state m , where the contribution is zero. Hence

$$\Delta_A = \frac{1}{3} [T(q)(l) - E_A(q | l)] + \frac{1}{3} [T(q)(m) - E_A(q | m)] = \frac{1}{3}(1-2) + \frac{1}{3}(2-2) = -\frac{1}{3}.$$

With ABS, let ψ_s and ψ_j denote the payoff vectors of the senior and junior tranches. At the numerical solution,

$$\psi_s = \left(\frac{21}{20}, \frac{9}{4}, \frac{9}{4} \right) \quad \text{and} \quad \psi_j = \left(0, 0, \frac{3}{4} \right).$$

Theory A buys both tranches in state l and the junior tranche in state m . Moreover, $E_A(\psi_s | l) = \frac{37}{20}$ and $E_A(\psi_j | l) = E_A(\psi_j | m) = \frac{1}{4}$. Therefore

$$\begin{aligned} \Delta_A &= \frac{1}{3} [T(\psi_s)(l) - E_A(\psi_s | l) + T(\psi_j)(l) - E_A(\psi_j | l)] + \frac{1}{3} [T(\psi_j)(m) - E_A(\psi_j | m)] \\ &= \frac{1}{3} \left(\frac{21}{20} - \frac{37}{20} + 0 - \frac{1}{4} \right) + \frac{1}{3} \left(0 - \frac{1}{4} \right) = -\frac{13}{30}. \end{aligned}$$

Since $P = Q_B$, theory B 's objective and perceived expectations coincide for every payoff vector, so $\Delta_B = \Delta_B = 0$. Thus, in this example, securitization makes theory A 's realized excess payoff lower by $\frac{1}{10}$ without increasing theory B 's realized excess payoff.

6.2. Iterated securitization

In Section 4, increased securitization has been modeled as a refinement of the existing tranching, which is equivalent to creating more tranches for the cash flow coming from the long-term asset. Securitization, however, is also seen as issuing securities backed by asset-backed securities, as when CDO are created based on existing MBS or squared CDO are created based on existing CDO. This appendix shows the equivalence between these two interpretations. Remarkably, any tranching is equivalent to issuing debt using debt securities as collateral and iterating this process a finite number of times.

Fig. 4 sketches the kind of repeated securitization structure that motivates the formal equivalence result below. Subprime mortgages may be pooled into MBS, the MBS cash flow may back CDO, and the CDO cash flow may in turn back squared CDO.

6.2.1. Example

Consider the limit of a symmetric transition process in which the probability of returning to the current state tends to one and the remaining probability is divided uniformly among the other states. The stationary distribution is uniform for every positive transition probability, so the system spends a quarter of the time in each state. The calculations below report the limit as the probability of leaving the current state tends to zero. The interest rate is assumed to be constant. The collection of theories in the market is given by $F = \{A, B, C\}$, where

$$A = \{\{a, d\}, \{b, c\}\}$$

$$B = \{\{a, b, d\}, \{c\}\}$$

$$C = \{\{a, b, c, d\}\}$$

All theories in the market achieve statistical consistency: the perceived transition probability for agent A from states a or d to states a or d tends to one half, the perceived transition probability for agent B from




Fig. 4. Multiple rounds of ABS.

Table 1
Security payoffs for the issuer-optimal tranching.

States	ϕ_e	ϕ_j	ϕ_m	ϕ_s
a	$q(a) - q(b)$	$q(b) - q(c)$	$q(c) - q(d)$	$q(d)$
b	0	$q(b) - q(c)$	$q(c) - q(d)$	$q(d)$
c	0	0	$q(c) - q(d)$	$q(d)$
d	0	0	0	$q(d)$

states a , b or d to states a , b or d tends to one third and the perceived transition probability for agent C from any state to any other tends to one fourth.

For this example, it can be shown that a tranching that is issuer-optimal is one where there are four tranches and the nonzero attachment points coincide with the realizations of the asset price. Let q be the least waterfall-equilibrium price function for the issuer-optimal tranching. Assume that the dividend function is such that $q(a) > q(b) > q(c) > q(d)$. In order of seniority, the tranches are labeled s (senior), m (mezzanine), j (junior) and e (equity). Table 1 presents the payoff function for each tranche.

In state d , theory A buys the tranche e , theory B buys the tranche j and theory C buys the tranche m . All agents have the same willingness to pay regarding tranche s . In state d , three different types of agents have claims on the underlying asset.

Instead of a unique round of securitization with four tranches, consider three rounds of securitization. In the first round, the underlying asset is securitized into debt and equity tranches. In the second round, the debt tranche of the first round is securitized into other debt and equity tranches. In the third round, the debt tranche of the second round is securitized into more debt and equity tranches (Fig. 5). Let ϕ_e^i and ϕ_f^i be the payoffs from the equity and the debt tranches in the i th round of securitization, where $i \in \{1, 2, 3\}$. Table 2 shows the assumed payoff for each tranche. By construction, the payoff from the debt and the equity tranches add up to the payoff from the debt tranche from the previous round.

Consider the equilibrium in state d . The equity tranche of the 1st round will be bought by theory A , the equity tranche of the 2nd round will be bought by theory B , and the equity tranche of the 3rd round will be bought by theory C . The value of the debt tranche on the 3rd round will be $q(d)$, discounted by the interest rate. Given this and the willingness to pay by agent C on the equity tranche of the 3rd round, the value of the debt tranche in the 2nd round will be $q(d)$ plus three




Fig. 5. Iterated securitization.

Table 2
Security payoffs for the iterated case.

States	ϕ_e^1	ϕ_f^1	ϕ_e^2	ϕ_f^2	ϕ_e^3	ϕ_f^3
a	$q(a) - q(b)$	$q(b)$	$q(b) - q(c)$	$q(c)$	$q(c) - q(d)$	$q(d)$
b	0	$q(b)$	$q(b) - q(c)$	$q(c)$	$q(c) - q(d)$	$q(d)$
c	0	$q(c)$	0	$q(c)$	$q(c) - q(d)$	$q(d)$
d	0	$q(d)$	0	$q(d)$	0	$q(d)$

quarters of $q(c) - q(d)$, discounted by the interest rate. Given this and the willingness to pay by agent B on the equity tranche of the 2nd round, the value of the debt tranche in the 1st round will be $q(d)$ plus three quarters of $q(c) - q(d)$ plus two thirds of $q(b) - q(c)$, discounted by the interest rate. Given this and the willingness to pay by agent A on the equity tranche of the 1st round, the value of the asset at state d will be the dividend at d plus $q(d)$ plus three quarters of $q(c) - q(d)$ plus two thirds of $q(b) - q(c)$ plus half of $q(a) - q(b)$, discounted by the interest rate. This is the same asset price equation for the aforementioned case of one round of securitization with four tranches.

A way of reading the equilibrium transactions at state d is that agent A is borrowing from agent B using the underlying asset as collateral. In turn, agent B is borrowing from agent C using the debt issued by A as collateral. The equilibrium features a *pyramiding arrangement*, as defined by Geanakoplos (1996).

6.2.2. General case

In order to define what will be labeled an *iterated securitization equilibrium*, consider a tree with finite node set $\mathbb{N}$ and terminal-node set $\mathbb{Z} \subset \mathbb{N}$. Let $r \in \mathbb{N} \setminus \mathbb{Z}$ denote the node representing the long-term asset. Let $a : \mathbb{N} \setminus \{r\} \rightarrow \mathbb{N} \setminus \mathbb{Z}$ be a function where $a(n)$ denotes the predecessor

of n . Each node $n \in \mathbb{N} \setminus \{r\}$ represents a security backed by asset $a(n)$. Nodes in $\mathbb{Z}$ represent securities that are not further securitized.

For $m \in \mathbb{N} \setminus \mathbb{Z}$, denote by $S(m)$ the set of m 's successors:

$$S(m) = \{n \in \mathbb{N} \setminus \{r\} : a(n) = m\}$$

Let $(\tau_n)_{n \in \mathbb{N}}$ be a collection of intervals such that

1. $\tau_r = [0, \infty)$ and
2. for $m \in \mathbb{N} \setminus \mathbb{Z}$, the collection $(\tau_n)_{n \in S(m)}$ partitions the interval $[0, \sup(\tau_m) - \inf(\tau_m)]$

Each security in $\mathbb{N} \setminus \{r\}$ is a one-period waterfall tranche. Its payoff is a function of the cash flow from its predecessor. Specifically, if t is the cash flow from asset $a(n)$, then the cash flow for security $n \in \mathbb{N} \setminus \{r\}$ would be given by:

$$\zeta_n(t) = (t - \inf(\tau_n))^+ - (t - \sup(\tau_n))^+ \quad (36)$$

Hence, for $m \in \mathbb{N} \setminus \mathbb{Z}$ and every feasible cash flow t from asset m ,

$$\sum_{n \in S(m)} \zeta_n(t) = t.$$

This confirms that the cash flows from securities in $S(m)$ add up to the cash flow from asset m .

Let $\gamma_n : X \rightarrow \mathbb{R}_+$ be the cash flow for security $n \in \mathbb{N}$ as a function of the state of the world. If $q : X \rightarrow \mathbb{R}_+$ is the price function for the long-term asset, then $\gamma_r = q$ and

$$\gamma_n(x) = \zeta_n(\gamma_{a(n)}(x)) \quad (37)$$

for $n \in \mathbb{N} \setminus \{r\}$ and $x \in X$.

A tuple $\Sigma = (\mathbb{N}, r, \mathbb{Z}, a, (\tau_n)_{n \in \mathbb{N}})$ is labeled an *iterated securitization structure*.

Definition 4. For a given iterated securitization structure Σ , an iterated securitization equilibrium is a function $p_\Sigma : \mathbb{N} \times X \rightarrow \mathbb{R}_+$ such that, for $x \in X$,

$$1. \quad p_\Sigma(r, x) = \frac{d(x)}{R(x)} + \sum_{n \in S(r)} p_\Sigma(n, x) \quad (38)$$

$$2. \text{ For } m \in \mathbb{N} \setminus (\mathbb{Z} \cup \{r\}), \quad p_\Sigma(m, x) = \sum_{n \in S(m)} p_\Sigma(n, x) \quad (39)$$

$$3. \text{ For } z \in \mathbb{Z}, \quad p_\Sigma(z, x) = \frac{\max_{F \in \mathcal{F}} E_F(\gamma_z | x)}{R(x)} \quad (40)$$

where $\gamma_r = p_\Sigma(r, \cdot)$ and, for $n \in \mathbb{N} \setminus \{r\}$, γ_n is given by Eq. (37).

The function $p_\Sigma(r, \cdot)$ is the price function for the long-term asset. Eq. (38) states that it is given by the discounted declared dividend plus the added prices of the securities derived from it. Eq. (39) states that for the securities that are not terminal nodes, their price is given by the added prices of the securities derived from them. Finally, Eq. (40) gives the equilibrium for the terminal nodes, i.e. the securities that are not further securitized. Their price is given by the highest discounted expectation regarding the cash flow coming from them. Ultimately, the cash flow from the securities at the terminal nodes will depend on the price of the long-term asset in the next period.

Proposition 10. For every iterated securitization structure Σ , there is a tranching $\mathcal{T}_\Sigma$ such that

$$\{p_\Sigma(r, \cdot) : p_\Sigma \text{ is an iterated securitization equilibrium}\} = \mathcal{Q}_{\mathcal{T}_\Sigma}.$$

Consequently, under the pointwise order, the least and greatest underlying-asset equilibrium price functions generated by Σ are $q_{\mathcal{T}_\Sigma}^{\min}$ and $q_{\mathcal{T}_\Sigma}^{\max}$, respectively.

Proof. For each terminal node $z \in \mathbb{Z}$, let $b(z)$ be the set of non-terminal nodes on the branch from r to the predecessor of z , and define

$$c_z := \sum_{n \in b(z)} \tau_n, \quad \hat{\tau}_z := [c_z + \underline{\tau}_z, c_z + \bar{\tau}_z].$$

Repeated application of Eqs. (36) and (37) gives

$$\gamma_z(x) = (p_\Sigma(r, x) - c_z - \underline{\tau}_z)^+ - (p_\Sigma(r, x) - c_z - \bar{\tau}_z)^+ = \phi_{\hat{\tau}_z}(p_\Sigma(r, x)). \quad (41)$$

Because the successors of every non-terminal node partition the feasible payoff of that node, the collection

$$\mathcal{T}_\Sigma := \{\hat{\tau}_z : z \in \mathbb{Z}\}$$

is a finite partition of $[0, \infty)$.

Let p_Σ be an iterated securitization equilibrium and set $q = p_\Sigma(r, \cdot)$. Recursively substituting Eq. (39) into Eq. (38) and then using Eq. (40) gives

$$q(x) = \frac{d(x)}{R(x)} + \sum_{z \in \mathbb{Z}} \frac{\max_{F \in \mathcal{F}} E_F(\gamma_z | x)}{R(x)}.$$

Eq. (41) therefore implies

$$q = \Psi_{\mathcal{T}_\Sigma}(q),$$

so $q \in \mathcal{Q}_{\mathcal{T}_\Sigma}$.

Conversely, let $q \in \mathcal{Q}_{\mathcal{T}_\Sigma}$. Set $p_\Sigma(r, \cdot) = q$ and, for every terminal node z , define

$$p_\Sigma(z, x) := \frac{\max_{F \in \mathcal{F}} E_F(\phi_{\hat{\tau}_z}(q) | x)}{R(x)}.$$

For each non-terminal node other than r , define its price recursively as the sum of the prices of its successors. These prices satisfy Eqs. (39) and (40) by construction. Since q is a fixed point of $\Psi_{\mathcal{T}_\Sigma}$, recursive addition of the terminal-node prices also gives Eq. (38). Hence the constructed p_Σ is an iterated securitization equilibrium whose underlying-asset price is q . The two sets of prices are therefore equal, and their least and greatest elements coincide. $\square$

Proposition 10 states the equivalence between the equilibrium concepts in Definitions 1 and 4. The full set of underlying-asset price functions generated by an iterated structure is the set of waterfall-equilibrium prices under a suitable one-shot tranching. Intuitively, multiple rounds of securitization are redundant since the same terminal payoffs can be obtained by tranching the long-term asset into $|\mathbb{Z}|$ securities in a single round.

6.3. Proofs of the asset-price results

Proof of Lemma 1. For each $\tau \in \mathcal{T}$, let

$$S(\tau) := \{v \in S : v \subseteq \tau\}.$$

Because S is a refinement of $\mathcal{T}$,

$$\phi_\tau(g) = \sum_{v \in S(\tau)} \phi_v(g).$$

Consequently, for any $g : X \rightarrow \mathbb{R}_+$ and $x \in X$,

$$\begin{aligned} \max_{F \in \mathcal{F}} E_F(\phi_\tau(g) | x) &= \max_{F \in \mathcal{F}} \sum_{v \in S(\tau)} E_F(\phi_v(g) | x) \\ &\leq \sum_{v \in S(\tau)} \max_{F \in \mathcal{F}} E_F(\phi_v(g) | x). \end{aligned}$$

Summing over $\tau \in \mathcal{T}$ gives

$$\Psi_{\mathcal{T}}(g) \leq \Psi_S(g)$$

for every $g \in \mathbb{R}_+^X$.

Fix $q \in Q_{\mathcal{T}}$. The operator comparison gives

$$q = \Psi_{\mathcal{T}}(q) \leq \Psi_S(q).$$

Iterating Ψ_S from q therefore produces a nondecreasing sequence. In the notation of the proof of [Proposition 1](#), refinement implies $a_{\mathcal{T}} \leq a_S$ and hence $M_{\mathcal{T}} \leq M_S$. Since $q \leq M_{\mathcal{T}} \mathbf{1}$, the sequence is bounded above by the invariant bound $M_S \mathbf{1}$. Continuity therefore implies that it converges to some $w \in Q_S$ with $q \leq w$.

Conversely, fix $w \in Q_S$. Then

$$\Psi_{\mathcal{T}}(w) \leq \Psi_S(w) = w.$$

Iterating $\Psi_{\mathcal{T}}$ from w produces a nonincreasing sequence bounded below by $\mathbf{0}$. Its limit is a fixed point $q \in Q_{\mathcal{T}}$ satisfying $q \leq w$.

Applying the second comparison to $q_S^{\min}$ and the first comparison to $q_{\mathcal{T}}^{\max}$ yields

$$q_{\mathcal{T}}^{\min} \leq q_S^{\min} \quad \text{and} \quad q_{\mathcal{T}}^{\max} \leq q_S^{\max},$$

respectively. $\square$

[Definition 5](#) and [Lemmas 2, 3](#) and [4](#) will be used to prove [Proposition 3](#).

Definition 5. A tranching $\mathcal{T}$ satisfies the dominance condition if, for each $\tau \in \mathcal{T}$ and $x \in X$, there is a theory $G \in \mathbf{F}$ such that

$$G_G^T(t | x) \leq G_F^T(t | x)$$

for every t in the interior of τ and all $F \in \mathbf{F}$.

Lemma 2. If $\mathcal{T}$ satisfies the dominance condition and S is a refinement of $\mathcal{T}$, then $q_S^{\min} = q_{\mathcal{T}}^{\min}$.

Proof. For $q = q_{\mathcal{T}}^{\min}$, the payoff from a bounded tranche τ expected for next period by theory F at state x can be written as

$$E_F(\phi_{\tau}(q) | x) = \int_{\inf(\tau)}^{\sup(\tau)} [1 - G_F^T(t | x)] dt. \quad (42)$$

If τ is unbounded, the upper limit is infinity. The integral remains finite because its integrand is zero above $\max_{x \in X} q_{\mathcal{T}}^{\min}(x)$.

For each $\tau \in \mathcal{T}$ and $x \in X$, let $G_{\tau,x}$ be a theory satisfying the dominance condition. [Eq. \(42\)](#) implies that $G_{\tau,x}$ maximizes the expected payoff of every subtranche $v \in S$ contained in τ . The payoff of each such subtranche integrates $1 - G_F^T(t | x)$ over a subinterval of τ , on which the same dominance inequality applies. Hence one theory attains the maximum for every subtranche contained in τ , and

$$\begin{aligned} \sum_{v \in S: v \subseteq \tau} \max_{F \in \mathbf{F}} E_F(\phi_v(q_{\mathcal{T}}^{\min}) | x) &= \max_{F \in \mathbf{F}} E_F \left(\sum_{v \in S: v \subseteq \tau} \phi_v(q_{\mathcal{T}}^{\min}) | x \right) \\ &= \max_{F \in \mathbf{F}} E_F(\phi_{\tau}(q_{\mathcal{T}}^{\min}) | x) \end{aligned}$$

for any $\tau \in \mathcal{T}$ and $x \in X$. Summing this equality over $\tau \in \mathcal{T}$ and using [Eq. \(18\)](#) gives

$$\begin{aligned} \Psi_S(q_{\mathcal{T}}^{\min}) &= \Psi_{\mathcal{T}}(q_{\mathcal{T}}^{\min}) \\ &= q_{\mathcal{T}}^{\min} \end{aligned}$$

Thus $q_{\mathcal{T}}^{\min}$ is a fixed point of Ψ_S , and the minimality of $q_S^{\min}$ gives $q_S^{\min} \leq q_{\mathcal{T}}^{\min}$. Since S refines $\mathcal{T}$, [Lemma 1](#) gives the reverse inequality. Hence $q_S^{\min} = q_{\mathcal{T}}^{\min}$. $\square$

Lemma 3. If $\mathcal{T}$ satisfies the dominance condition, then

$$q_S^{\min}(x) \leq q_{\mathcal{T}}^{\min}(x)$$

for all $x \in X$ and any tranching S .

Proof. Let $S \vee \mathcal{T}$ be the coarsest common refinement between $\mathcal{T}$ and S . By [Lemma 1](#),

$$q_S^{\min}(x) \leq q_{S \vee \mathcal{T}}^{\min}(x)$$

for $x \in X$. By [Lemma 2](#),

$$q_{S \vee \mathcal{T}}^{\min}(x) = q_{\mathcal{T}}^{\min}(x)$$

for $x \in X$. $\square$

Lemma 4. If $\mathcal{T}$ does not satisfy the dominance condition, then there is a refinement S of $\mathcal{T}$ such that

$$q_{\mathcal{T}}^{\min}(x) < q_S^{\min}(x)$$

for all $x \in X$.

Proof. Let $q = q_{\mathcal{T}}^{\min}$. Since $\mathcal{T}$ does not satisfy the dominance condition, there are a state $x^* \in X$ and a tranche $\tau \in \mathcal{T}$ such that no theory minimizes $G_F^T(t | x^*)$ at every point in the interior of τ . Because X is finite, these perceived distribution functions are step functions. Choose $B > \max_y q(y)$, truncate the active part of an unbounded tranche at B , and partition it at every value $q(y)$ in its interior. Complete this collection with the inactive tail $[B, \infty)$ when τ is unbounded, and denote the resulting finite partition of τ by $\mathcal{P}(\tau)$.

Write $\tau^B := \tau \cap [0, B)$ when τ is unbounded and $\tau^B := \tau$ otherwise. For every cell $I = [a, b) \in \mathcal{P}(\tau)$ contained in τ^B ,

$$E_F(\phi_I(q) | x^*) = \int_a^b [1 - G_F^T(t | x^*)] dt.$$

It follows that

$$\sum_{\substack{I \in \mathcal{P}(\tau) \\ I \subseteq \tau^B}} \min_{F \in \mathbf{F}} \int_I G_F^T(t | x^*) dt < \min_{F \in \mathbf{F}} \int_{\tau^B} G_F^T(t | x^*) dt.$$

Indeed, let F^* minimize the integral over τ^B . Each integral of $G_{F^*}^T$ over a cell is weakly above the minimum for that cell. Equality in the display would therefore require equality cell by cell. Since the CDFs are right-continuous step functions, F^* would then minimize the CDF at every point in the interior of τ , contradicting the failure of the dominance condition. Consequently,

$$\sum_{I \in \mathcal{P}(\tau)} \max_{F \in \mathbf{F}} E_F(\phi_I(q) | x^*) > \max_{F \in \mathbf{F}} E_F(\phi_{\tau}(q) | x^*).$$

Let S be the refinement of $\mathcal{T}$ that replaces τ by $\mathcal{P}(\tau)$. Then

$$q(x^*) = \Psi_{\mathcal{T}}(q)(x^*) < \Psi_S(q)(x^*).$$

[Lemma 1](#) gives $q \leq q_S^{\min}$. Monotonicity therefore implies

$$\Psi_S(q) \leq \Psi_S(q_S^{\min}) = q_S^{\min},$$

and hence $q(x^*) < q_S^{\min}(x^*)$.

It remains to propagate strictness. Budget balance and monotonicity of tranche payoffs imply that there is some $\sigma^* \in S$ such that

$$\phi_{\sigma^*}(q_S^{\min}(x^*)) > \phi_{\sigma^*}(q(x^*)).$$

Every subjective transition matrix has full support, so for every current state y and every theory F ,

$$E_F(\phi_{\sigma^*}(q_S^{\min}) | y) > E_F(\phi_{\sigma^*}(q) | y).$$

The remaining tranche valuations increase weakly. Therefore

$$q_S^{\min}(y) = \Psi_S(q_S^{\min})(y) > \Psi_S(q)(y) \geq \Psi_{\mathcal{T}}(q)(y) = q(y)$$

for every $y \in X$. $\square$

Proof of Proposition 3. The condition in [Proposition 3](#) is the dominance condition in [Definition 5](#). If it holds, [Lemma 3](#) implies that $\mathcal{T}$ is issuer-optimal. If it fails, [Lemma 4](#) provides a refinement S of $\mathcal{T}$ such that $q_{\mathcal{T}}^{\min}(x) < q_S^{\min}(x)$ for every $x \in X$. Hence $\mathcal{T}$ is not issuer-optimal, and the final claim follows. $\square$

6.4. Proof of Proposition 7

Fix any $q \in Q_T$ and evaluate tranche payoffs and buyer choices at this equilibrium. For each tranche $\tau \in \mathcal{T}$, define

$$\delta_F^T(\tau) := \sum_{x \in X} (T(\psi_\tau)(x) - E_F(\psi_\tau | x)) I_{\{F=\beta_T(\tau, x)\}} \mu(x)$$

as the average realized excess payoff from tranche τ assigned to theory F . For any block $C \in \mathcal{G}$ and $x \in C$, Eqs. (13) and (14), together with the definition of T , imply

$$E_G(\psi_\tau | x) = \sum_{y \in X} \psi_\tau(y) Q_G(x, y) = \frac{\sum_{z \in C} \mu(z) T(\psi_\tau)(z)}{\sum_{z \in C} \mu(z)}.$$

Since $\mathcal{G}$ refines every $F \in \mathbf{F}$, each block $C \in \mathcal{G}$ is contained in a block of F . Thus, $E_F(\psi_\tau | x)$ is constant over $x \in C$. Let $e_F^\tau(C)$ denote this common value. Hence $\mathbf{B}(\tau, x)$ and, under the fixed priority ordering, $\beta_T(\tau, x)$ are also constant on C . If theory F is selected on C , its willingness to pay is at least that of $\mathcal{G}$, and therefore

$$\begin{aligned} \sum_{x \in C} \mu(x) [T(\psi_\tau)(x) - E_F(\psi_\tau | x)] \\ = \left(\sum_{x \in C} \mu(x) \right) [e_G^\tau(C) - e_F^\tau(C)] \leq 0. \end{aligned}$$

If $F = \mathcal{G}$, the expression is zero. On blocks where $\mathcal{G}$ is selected, its contribution is therefore zero; on every other block, its selected-buyer indicator is zero. Hence $\delta_G^T(\tau) = 0$. For every $F \neq \mathcal{G}$, summing the nonpositive contributions over the blocks where F is selected gives $\delta_F^T(\tau) \leq 0$. Finally, since $\Delta_F^T(q) = \sum_{\tau \in \mathcal{T}} \delta_F^T(\tau)$, summing over tranches proves the result. $\square$

6.5. Realized-payoff example with three theories

This appendix provides the example summarized in Section 4.3, in which securitization increases one theory's realized excess payoff even though that payoff remains negative. Let the objective transition matrix be

$$P = \begin{pmatrix} 1/2 & 1/4 & 1/4 \\ 1/4 & 1/2 & 1/4 \\ 1/4 & 1/4 & 1/2 \end{pmatrix}.$$

This matrix has full support and stationary distribution $\mu = (1/3, 1/3, 1/3)$. The collection of theories is $\mathbf{F} = \{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$, where

$$\mathcal{A} = \{\{l\}, \{m, h\}\},$$

$$\mathcal{B} = \{\{l, h\}, \{m\}\},$$

$$\mathcal{C} = \{\{l, m, h\}\}.$$

The corresponding statistically consistent transition matrices are obtained from Eq. (13). Set $R = 4$ and

$$d = \left(\frac{245}{48}, \frac{113}{16}, \frac{209}{16} \right),$$

and fix the priority ordering $\mathcal{A} > \mathcal{B} > \mathcal{C}$.

Under the no-tranching benchmark $\mathcal{T}^0$, the unique equilibrium price is

$$q_{\mathcal{T}^0} = \left(\frac{4451}{2232}, \frac{5579}{2232}, \frac{8927}{2232} \right).$$

Theory $\mathcal{B}$ buys the whole claim at state l , while theory $\mathcal{A}$ buys it at states m and h . Because

$$T(q_{\mathcal{T}^0})(l) - E_B(q_{\mathcal{T}^0} | l) = \frac{q_{\mathcal{T}^0}(l) - q_{\mathcal{T}^0}(h)}{8},$$

the realized excess payoffs are

$$\Delta_{\mathcal{A}}^{\mathcal{T}^0} = 0, \quad \Delta_{\mathcal{B}}^{\mathcal{T}^0} = \frac{q_{\mathcal{T}^0}(l) - q_{\mathcal{T}^0}(h)}{24} = -\frac{373}{4464}, \quad \Delta_{\mathcal{C}}^{\mathcal{T}^0} = 0.$$

Now consider the three-tranche waterfall

$$\mathcal{T}^1 = \{[0, 2], [2, 5/2], [5/2, \infty)\}.$$

Table 3

Selected buyers in the realized-payoff example.

Tranche	l	m	h
Senior	$\mathcal{A}$	$\mathcal{A}$	$\mathcal{A}$
Mezzanine	$\mathcal{C}$	$\mathcal{A}$	$\mathcal{A}$
Junior	$\mathcal{B}$	$\mathcal{A}$	$\mathcal{A}$

The equilibrium price used below is

$$q_{\mathcal{T}^1} = (2, 5/2, 4).$$

Direct substitution gives $\Psi_{\mathcal{T}^1}(q_{\mathcal{T}^1}) = q_{\mathcal{T}^1}$. Because each tranche payoff changes by no more than the change in the underlying asset price, expectations and maximization over theories do not amplify this change. Summing across the three tranches and dividing by $R = 4$ therefore gives, for any price functions q and $\hat{q}$,

$$\|\Psi_{\mathcal{T}^1}(q) - \Psi_{\mathcal{T}^1}(\hat{q})\|_\infty \leq \frac{3}{4} \|q - \hat{q}\|_\infty.$$

If q and $\hat{q}$ were both fixed points, this inequality would imply $\|q - \hat{q}\|_\infty \leq (3/4)\|q - \hat{q}\|_\infty$. Thus they must coincide, and the displayed price is the unique equilibrium price. The resulting tranche payoff vectors are

$$\psi_s = (2, 2, 2), \quad \psi_m = (0, 1/2, 1/2), \quad \psi_j = (0, 0, 3/2).$$

The selected buyers are shown in Table 3.

For the mezzanine tranche at state l , $T(\psi_m)(l) = 1/4$ and $E_C(\psi_m | l) = 1/3$. For the junior tranche at that state, $T(\psi_j)(l) = 3/8$ and $E_B(\psi_j | l) = 9/16$. The realized excess payoffs are therefore

$$\Delta_{\mathcal{A}}^{\mathcal{T}^1} = 0, \quad \Delta_{\mathcal{B}}^{\mathcal{T}^1} = \frac{1}{3} \left(\frac{3}{8} - \frac{9}{16} \right) = -\frac{1}{16}, \quad \Delta_{\mathcal{C}}^{\mathcal{T}^1} = \frac{1}{3} \left(\frac{1}{4} - \frac{1}{3} \right) = -\frac{1}{36}.$$

Consequently,

$$\Delta_{\mathcal{B}}^{\mathcal{T}^1} = -\frac{1}{16} > -\frac{373}{4464} = \Delta_{\mathcal{B}}^{\mathcal{T}^0}.$$

Securitization induces theory $\mathcal{C}$ to buy the mezzanine tranche at state l , allowing theory $\mathcal{B}$ to avoid part of its benchmark loss. At the same time, the aggregate realized excess payoff falls from $-373/4464$ to

$$\Delta_{\text{agg}}^{\mathcal{T}^1} = -\frac{1}{16} - \frac{1}{36} = -\frac{13}{144} < -\frac{373}{4464},$$

as required by Corollary 5.

Data availability

No data was used for the research described in the article.

References

- Allen, Franklin, Barbalau, Adelina, 2024. Security design: A review. *J. Financ. Intermediation* 60, 101113. <http://dx.doi.org/10.1016/j.jfi.2024.101113>.
- Allen, Franklin, Gale, Douglas, 1988. Optimal security design. *Rev. Financ. Stud.* 1 (3), 229–263.
- Bordalo, Pedro, Gennaioli, Nicola, Shleifer, Andrei, 2018. Diagnostic expectations and credit cycles. *J. Financ.* 73 (1), 199–227. <http://dx.doi.org/10.1111/jofi.12586>.
- Bordalo, Pedro, Gennaioli, Nicola, Shleifer, Andrei, Terry, Stephen J., 2026. Real credit cycles. *Am. Econ. Rev.* 116 (4), 1274–1308. <http://dx.doi.org/10.1257/aer.2021.1820>.
- Brock, W.A., Hommes, C.H., Wagener, F.O.O., 2009. More hedging instruments may destabilize markets. *J. Econom. Dynam. Control* (ISSN: 0165-1889) 33 (11), 1912–1928. <http://dx.doi.org/10.1016/j.jedc.2009.05.004>, URL <http://www.sciencedirect.com/science/article/pii/S0165188909001158>.
- Broer, Tobias, 2018. Securitization bubbles: Structured finance with disagreement about default risk. *J. Financ. Econ.* 127 (3), 505–518. <http://dx.doi.org/10.1016/j.jfineco.2017.12.001>.
- Brunnermeier, Markus K., 2009. Deciphering the liquidity and credit crunch 2007–2008. *J. Econ. Perspect.* 23 (1), 77–100. <http://dx.doi.org/10.1257/jep.23.1.77>, URL <http://www.aeaweb.org/articles.php?doi=10.1257/jep.23.1.77>.
- Brunnermeier, Markus K., Simsek, Alp, Xiong, Wei, 2014. A welfare criterion for models with distorted beliefs. *Q. J. Econ.* 129 (4), 1753–1797. <http://dx.doi.org/10.1093/qje/qju025>.

- Duffie, Darrell, Rahi, Rohit, 1995. Financial market innovation and security design: An introduction. *J. Econ. Theory* 65 (1), 1–42. URL <http://ideas.repec.org/a/eee/jetheo/v65y1995i1p1-42.html>.
- Ellis, Andrew, Piccione, Michele, Zhang, Shengxing, 2022. Equilibrium securitization with diverse beliefs. *Theor. Econ.* 17 (1), 121–152. <http://dx.doi.org/10.3982/TE4157>.
- Eyster, Erik, Piccione, Michele, 2013. An approach to asset pricing under incomplete and diverse perceptions. *Econometrica* (ISSN: 1468-0262) 81 (4), 1483–1506. <http://dx.doi.org/10.3982/ECTA10499>.
- Fostel, Ana, Geanakoplos, John, 2012. Tranching, CDS, and asset prices: how financial innovation can cause bubbles and crashes. *Am. Econ. J.: Macroecon.* 4 (1), 190–225. <http://dx.doi.org/10.1257/mac.4.1.190>, URL <http://www.aeaweb.org/articles.php?doi=10.1257/mac.4.1.190>.
- Fostel, Ana, Geanakoplos, John, 2016. Financial innovation, collateral and investment. *Am. Econ. J.: Macroecon.* 8 (1), 242–284. <http://dx.doi.org/10.1257/mac.20130183>.
- Geanakoplos, John, 1996. Promises Promises. Cowles Foundation Discussion Paper No. 1143, Cowles Foundation for Research in Economics, Yale University, URL <http://ideas.repec.org/p/cwl/cwldpp/1143.html>.
- Geanakoplos, John, Zame, William R., 2014. Collateral equilibrium: I: A basic framework. *Econom. Theory* 56 (3), 443–492. <http://dx.doi.org/10.1007/s00199-013-0797-4>.
- Gennaioli, Nicola, Shleifer, Andrei, Vishny, Robert, 2012. Neglected risks, financial innovation and financial fragility. *J. Financ. Econ.* 104 (3), 452–468. <http://dx.doi.org/10.1016/j.jfineco.2011.05.005>.
- Greenwood, Robin, Hanson, Samuel G., Shleifer, Andrei, Sorensen, Jakob Ahm, 2022. Predictable financial crises. *J. Financ.* 77 (2), 863–921. <http://dx.doi.org/10.1111/jofi.13105>.
- Guo, Guixia, Wang, Frank Yong, Wei, Xu, 2014. Optimal tranching with diverse beliefs. *Econom. Lett.* 124 (2), 222–226. <http://dx.doi.org/10.1016/j.econlet.2014.05.030>.
- Harrison, J. Michael, Kreps, David M., 1978. Speculative investor behavior in a stock market with heterogeneous expectations. *Q. J. Econ.* 92 (2), 323–336. URL <http://ideas.repec.org/a/tpr/qjecon/v92y1978i2p323-36.html>.
- Iachan, Felipe S., Nenov, Plamen T., Simsek, Alp, 2021. The choice channel of financial innovation. *Am. Econ. J.: Macroecon.* 13 (2), 333–372. <http://dx.doi.org/10.1257/mac.20180429>.
- Krishnamurthy, Arvind, Li, Wenhao, 2025. Dissecting mechanisms of financial crises: Intermediation and sentiment. *J. Political Econ.* 133 (3), 935–985. <http://dx.doi.org/10.1086/733423>.
- Kubler, Felix, Schmedders, Karl, 2012. Financial innovation and asset price volatility. *Am. Econ. Rev.* 102 (3), 147–151. <http://dx.doi.org/10.1257/aer.102.3.147>, URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.102.3.147>.
- Matthews, John O., 1994. *Struggle and Survival on Wall Street: the Economics of Competition Among Securities Firms*. Oxford University Press, New York.
- Maxted, Peter, 2024. A macro-finance model with sentiment. *Rev. Econ. Stud.* 91 (1), 438–475. <http://dx.doi.org/10.1093/restud/rdad023>.
- Milgrom, Paul, Stokey, Nancy, 1982. Information, trade and common knowledge. *J. Econom. Theory* (ISSN: 0022-0531) 26 (1), 17–27. [http://dx.doi.org/10.1016/0022-0531\(82\)90046-1](http://dx.doi.org/10.1016/0022-0531(82)90046-1), URL <http://www.sciencedirect.com/science/article/pii/0022053182900461>.
- Miller, Edward M., 1977. Risk, uncertainty, and divergence of opinion. *J. Financ.* (ISSN: 00221082) 32 (4), 1151–1168, URL <http://www.jstor.org/stable/2326520>.
- Scheinkman, José, Xiong, Wei, 2004. Heterogeneous beliefs, speculation and trading in financial markets. In: *Paris-Princeton Lectures on Mathematical Finance 2003*. Springer, pp. 217–250.
- Simsek, Alp, 2013a. Belief disagreements and collateral constraints. *Econometrica* (ISSN: 1468-0262) 81 (1), 1–53. <http://dx.doi.org/10.3982/ECTA9956>.
- Simsek, Alp, 2013b. Speculation and risk sharing with new financial assets. *Q. J. Econ.* 128 (3), 1365–1396. <http://dx.doi.org/10.1093/qje/qjt007>, URL <http://qje.oxfordjournals.org/content/128/3/1365.abstract>.
- Tarski, Alfred, 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* 5 (2), 285–309. <http://dx.doi.org/10.2140/pjm.1955.5.285>.