

Structural Immunity: How Platform Dominance and Mandatory Arbitration Create Corporate Impunity

Jeremy McEntire
Cage & Mirror Press

Working Paper — Version 2.0
March 2026

Abstract

When platform market dominance combines with pre-dispute mandatory arbitration, the result is *structural immunity*: a collapse in conversion rates across the dispute pipeline such that systematic, low-dollar harms become non-justiciable regardless of merit. This paper develops a six-filter multiplicative pipeline model that formalizes structural immunity as a suppression cascade: awareness, rational apathy, information asymmetry, mandatory arbitration, class action waiver, and regulatory capture each independently suppress a fraction of surviving claims, and the compound survival rate is the product of all pass-through rates.

Calibrating against CFPB data—600 observed filings from 80,000 expected claims—yields a baseline survival rate of 0.75%. Under realistic parameters for consumer finance, rational apathy alone suppresses 99.1% of claims, producing a compound survival rate of 0.002% (2 claims from 80,000 harmed consumers). Intervention analysis reveals the model’s central policy implication: removing any single filter improves survival by at most $108\times$ (from 0.002% to 0.27%), but the remaining five filters still block 99.7% of claims. Achieving even 10% survival requires dismantling three or more filters simultaneously. Cross-sector extension to gig economy, social media, and healthcare reveals that the dominant filter varies by industry—rational apathy in consumer finance and social media, class action waivers in the gig economy, information asymmetry in healthcare—implying that sector-specific intervention strategies are necessary.

The multiplicative structure of the pipeline explains why decades of piecemeal reform have failed to restore access to adjudication and why structural immunity is robust to incremental policy change.

1 Introduction

Structural immunity is the practical inability to obtain legal remedy for documented harms—not through adjudication on merits but through architectural barriers preventing adjudication from occurring. The term is operational rather than moral: it describes system outputs, not actor intentions. A system can produce structural immunity without any individual intending that result.

Formally, structural immunity is a collapse in conversion rates across the dispute pipeline. Harms occur in the world. Some fraction become recognized claims. Some fraction of claims become active disputes. Some fraction of disputes become legal events—filings in forums capable of adjudication. The architecture described in this paper operates on these conversion rates, reducing each transition probability until the final output (legal events) approaches zero regardless of input volume (harms).

This framing implies falsification conditions: if arbitration filing rates approached litigation baselines for comparable covered populations, or if repeat-player effects vanished after controlling for case characteristics, the framework would require substantial revision. Section 13 elaborates these conditions.

This paper makes three contributions beyond the qualitative framework presented in earlier work (McEntire, 2026). First, Section 5 develops a formal six-filter multiplicative pipeline model that captures the compound suppression structure. Second, Section 6 calibrates the model against CFPB data and derives per-filter suppression rates consistent with observed filing patterns. Third, Sections 7 and 8 use the calibrated model to analyze single-filter interventions and cross-sector variation, establishing the paper’s central claim: structural immunity is robust to piecemeal reform because of the multiplicative architecture.

2 Definitions and Scope

Mandatory arbitration as used here means pre-dispute arbitration agreements imposed as conditions of platform access, combined with class action waivers and confidentiality provisions. The critical features are: (1) agreement occurs before any dispute arises, when consumers cannot evaluate whether they will need legal recourse; (2) agreement is required for platform access, not offered as one option among several; (3) class treatment is prohibited, requiring each harmed consumer to pursue individually; and (4) proceedings and outcomes remain confidential, preventing pattern recognition across cases. Post-dispute arbitration agreements, where parties choose arbitration after a specific dispute arises, present different dynamics and fall outside this analysis.

Platform indispensability describes conditions where non-participation in a platform carries substantial professional or economic costs, creating practical coercion despite formal voluntari-

ness. A professional who refuses LinkedIn’s terms can still technically have a career; a seller who refuses Amazon’s terms can still technically sell products; a worker who refuses an employer’s arbitration clause can still technically find employment. But “technically possible” is not the same as “meaningfully available.” Platform indispensability exists on a continuum and is distinct from antitrust “monopoly,” which requires specific market definition and market power analysis beyond this paper’s scope.

Selection mechanisms are structural features that filter outcomes through differential propagation rather than explicit prohibition. The concept is borrowed from evolutionary biology: selection does not prevent variation but determines which variants survive to influence future states. Applied here: selection mechanisms determine which disputes survive to influence legal outcomes. No one need decide to suppress claims. The structure creates conditions where claims systematically fail to become legal events.

Claim suppression refers to the elimination of disputes before they become legally cognizable—distinct from claim denial (adjudication against claimants) or claim diversion (resolution in alternative forums). A suppressed claim is never filed, never heard, never resolved. It ceases to exist as a legal phenomenon. The harm persists; the legal event does not. This distinction matters because defenders of mandatory arbitration often point to arbitration outcomes (win rates, award amounts, processing times) as evidence the system works. Those metrics describe the small fraction of disputes that survive to filing. They tell us nothing about the vast majority that never appear in any forum.

2.1 Scope Limitations

This analysis focuses primarily on US federal arbitration doctrine as applied to systematic low-dollar harms in consumer and employment contexts. Important sectoral differences exist: securities arbitration operates under FINRA rules with different procedural dynamics; international commercial arbitration involves sophisticated parties with meaningful bargaining power; labor arbitration under collective bargaining agreements reflects negotiated terms rather than adhesion contracts. The structural immunity framework is designed to explain claim suppression patterns where three conditions converge: (1) mandatory pre-dispute arbitration with class waivers, (2) individually small but collectively significant harms, and (3) substantial power asymmetries between parties.

2.2 Terminology: Harms, Claims, Disputes, Legal Events

A **harm** is an injury experienced by a person: money taken improperly, service not delivered as promised, wages unpaid. A **claim** is a harm that the injured party recognizes as potentially actionable. A **dispute** is a claim that the injured party has attempted to pursue through some

channel. A **legal event** is a dispute that enters a forum capable of adjudication: a court filing, an arbitration demand, a regulatory complaint that triggers investigation.

Structural immunity operates by filtering at each transition. The architecture's primary effect is ensuring that the vast majority of harms never survive to legal-event status. Claim suppression—the elimination of disputes before filing—dominates claim denial as a mechanism.

3 The Selection Framework

3.1 Organizational Precedent

Selection mechanisms that filter inconvenient information operate throughout organizational life. In previous work, I introduced the concept of dysmemic pressure—the selection dynamics causing organizations to systematically filter accurate information when that information conflicts with institutional interests (McEntire, 2025). The mechanism is structural, not conspiratorial. Organizational structures create conditions where accurate information fails to propagate: promotion criteria reward confidence over accuracy, communication channels compress nuance into simple signals, feedback loops reward telling leadership what it wants to hear. The cumulative result is institutional self-deception that no individual intends but the system reliably produces.

3.2 Extension to Legal Architecture

Mandatory arbitration operates through analogous selection dynamics, but at the interface between corporations and the legal system. Dysmemic pressure filters information before it reaches organizational decision-makers. Arbitration clauses filter disputes before they reach courts. The mechanisms differ but the function is identical: ensuring that inconvenient challenges never aggregate into actionable patterns.

The system's primary function is preventing disputes from existing as legal events in the first place, rather than resolving them unfairly once they exist. Critics who focus on arbitration outcomes (win rates, award amounts) are examining the residue that survives selection pressure, rather than the selection pressure itself.

3.3 Alternative Explanation

A natural counter-hypothesis: perhaps arbitration reduces filings because it resolves disputes more satisfactorily. If arbitration is faster, cheaper, and produces acceptable outcomes, the low filing rate could reflect successful informal resolution rather than suppression.

This hypothesis is difficult to reconcile with observed magnitudes. If consumer satisfaction explained low filing rates, we would expect filing rates to be low but non-trivial. Instead, filing rates approach zero in practical terms: approximately 600 annual filings across six consumer finance markets covering tens of millions of consumers. We would also expect no systematic outcome disparities once cases reach filing. Instead, repeat-player advantages persist and intensify with familiarity. And we would expect companies to welcome arbitration research demonstrating consumer satisfaction. Instead, arbitration confidentiality prevents such research. The satisfaction hypothesis cannot explain the combination of (1) near-zero filing rates, (2) persistent repeat-player advantages, and (3) resistance to transparency. Claim suppression can.

3.4 Predictions

This framing generates specific predictions. If structural immunity operates through selection rather than adjudication bias, we should observe:

1. Filing rates far below what harm prevalence would predict—because the architecture suppresses claims before filing.
2. Outcomes skewed by structural position (repeat-player effects) persisting after controls for observed case characteristics.
3. Information asymmetries compounding over time as corporations accumulate institutional knowledge.
4. Pattern recognition failures as confidentiality prevents cross-case analysis.

Observable refutations include: arbitration filing rates approaching litigation baselines; no repeat-player effect after controlling for case characteristics; arbitration producing public accountability despite confidentiality; or class action and arbitration producing comparable aggregate outcomes.

4 The Six Filters

Structural immunity emerges from multiple reinforcing barriers. Each filter operates independently; together they create cumulative selection pressure that the formal model in Section 5 quantifies.

4.1 Filter 1: Awareness Failure

Many consumers never recognize that they have experienced an actionable harm. Billing errors, hidden fees, data misuse, metric inflation, and algorithmic discrimination produce injuries that are individually small and technically opaque. A consumer charged \$3.50 in undisclosed fees may never notice; one whose advertising metrics are inflated may lack the expertise to detect the discrepancy. The awareness filter operates before any dispute process begins: harms that are never recognized as claims cannot enter any resolution pathway.

4.2 Filter 2: Rational Apathy

For harms that are recognized, rational actors compare expected recovery to pursuit costs. Justice Kagan stated it plainly in her *Italian Colors* dissent: “Only a lunatic or a fanatic sues for \$30” (ita, 2013). Filing for arbitration requires fees (typically \$200 or more), time, and attention. For claims involving small damages—a few dollars in improper fees, a few hours of unpaid wages—the cost of individual pursuit exceeds potential recovery. Rational actors do not pursue.

The economics are straightforward. If expected harm is h , filing cost is c , and win probability is p , a rational actor files only when $hp > c$. For the modal consumer finance harm ($h = \$35$, $c = \$200$, $p = 0.30$), the expected value of filing is \$10.50 against a cost of \$200. The claim is economically irrational to pursue regardless of merit.

4.3 Filter 3: Information Asymmetry

Arbitration proceedings are confidential by default. Each case resolves in isolation, invisible to everyone except the immediate parties. Corporations arbitrate repeatedly and accumulate institutional knowledge: which arbitrators favor which arguments, which procedural tactics succeed, which settlement ranges arbitrators find reasonable. Each new consumer starts from zero. The corporation has seen thousands of similar disputes; the consumer has seen none.

Confidentiality also prevents external pattern recognition. A platform might engage in systematic misconduct affecting millions of consumers. Each arbitration resolves confidentially. No public record accumulates. No journalist can FOIA arbitration filings. No regulator can identify systematic problems from court dockets. The misconduct remains invisible because the mechanism that would make it visible—public adjudication—has been eliminated.

4.4 Filter 4: Mandatory Arbitration Friction

Arbitration clauses are imposed through take-it-or-leave-it terms. The consumer either accepts the entire terms of service or forgoes platform access entirely. When the platform commands sufficient

market position, “leave it” carries substantial costs. Courts have historically policed contracts of adherence through unconscionability doctrine, but the Supreme Court has systematically preempted state unconscionability protections for arbitration clauses ([con, 2011](#))—though *Heckman v. Live Nation* demonstrates this preemption has limits when procedural engineering becomes sufficiently extreme.

Beyond the contractual mandate, arbitration itself introduces procedural friction. Truncated discovery prevents the extended investigation complex claims require. The abbreviated timeline—often 60–90 days from filing to hearing—prevents evidence development. These limitations systematically favor defendants with superior access to their own records.

4.5 Filter 5: Class Action Waiver

Class action waivers eliminate the only economically rational mechanism for addressing small-dollar harms. Consider a platform that improperly charges one million consumers \$30 each. Total harm: \$30 million. Individual harm: \$30. No individual consumer can rationally pursue a \$30 claim through arbitration. But if consumers could aggregate claims, the \$30 million harm becomes addressable. Class waivers eliminate this solution. They eliminate disputes rather than redirecting them to a different forum. The realistic alternative to class treatment is zero arbitrations rather than one million individual arbitrations.

The CFPB found approximately 600 consumer arbitration filings annually across six consumer finance markets despite tens of millions of covered consumers ([Consumer Financial Protection Bureau, 2015](#)). Over five years, class settlements made 160 million consumers eligible for approximately \$2.7 billion in relief; arbitration produced affirmative relief in 32 disputes totaling under \$400,000.

4.6 Filter 6: Regulatory Capture and Operational Edge Closure

The final filter combines two reinforcing mechanisms. First, *regulatory capture*: the legislative and regulatory apparatus that might dismantle the other filters is itself subject to selection effects. The FAIR Act has passed the House twice and stalled in the Senate both times. The CFPB’s 2017 rule prohibiting class action waivers in financial services was nullified via Congressional Review Act resolution before taking effect ([United States Senate, 2017](#)). State efforts face systematic federal preemption.

Second, *operational edge closure*: every customer-facing node may be structurally incapable of issuing the remedies that would moot the dispute. Support agents cannot issue refunds beyond nominal thresholds, cannot override system settings, cannot escalate to anyone with authority to provide substantive remedy. The system absorbs responsibility while distributing it to no one.

Unlike Filters 1–5, which rest on documented legal doctrine and peer-reviewed research, operational edge closure is proposed based on qualitative observation and warrants further empirical investigation.

4.7 Cumulative Architecture

The filters are not merely additive—they are *multiplicative*. A consumer who overcomes rational apathy still faces information asymmetry, mandatory arbitration friction, and class action waivers. A regulator who identifies a pattern still cannot aggregate affected consumers. A journalist who obtains information about one arbitration still cannot build a story across confidential proceedings.

In employment contexts, [Estlund \(2018\)](#) estimates that between 320,000 and 727,000 claims annually simply disappear—what she terms a “black hole.” The National Employment Law Project estimated that forced arbitration enabled employers to pocket \$9.27 billion in stolen wages from low-paid workers in 2019 alone ([National Employment Law Project, 2021](#)). These are not disputes resolved through arbitration. They are disputes that the arbitration architecture prevents from existing as legal events.

The next section formalizes this multiplicative structure.

5 The Pipeline Model

The qualitative account in Section 4 identifies six filters that suppress claims before adjudication. This section formalizes the suppression architecture as a multiplicative cascade, enabling quantitative analysis of intervention effects.

Definition 1 (Suppression Pipeline). *Let N denote the number of consumers experiencing actionable harm. The suppression pipeline consists of k sequential filters F_1, F_2, \dots, F_k , each characterized by a suppression rate $s_i \in [0, 1]$. Filter F_i eliminates a fraction s_i of the claims that reach it, passing through the remaining fraction $1 - s_i$. The six filters are:*

1. **Awareness failure** (s_1): *fraction who never recognize they have a claim.*
2. **Rational apathy** (s_2): *fraction for whom pursuit costs exceed expected recovery.*
3. **Information asymmetry** (s_3): *fraction lost to confidentiality and repeat-player knowledge gaps.*
4. **Mandatory arbitration** (s_4): *fraction deterred by arbitration procedure and friction.*
5. **Class action waiver** (s_5): *fraction whose claims are only viable when aggregated.*

6. **Regulatory capture / edge closure** (s_6): fraction suppressed by legislative failure and operational design.

Proposition 1 (Compound Survival Rate). *The number of claims surviving all k filters is:*

$$C = N \prod_{i=1}^k (1 - s_i) \quad (1)$$

and the compound survival rate is:

$$P(\text{survival}) = \prod_{i=1}^k (1 - s_i) \quad (2)$$

Proof. Each filter passes through $(1 - s_i)$ of the claims reaching it. Under the independence assumption (see Section 14), the filters compose multiplicatively. Starting from N harmed consumers, after Filter 1, $N(1 - s_1)$ survive; after Filter 2, $N(1 - s_1)(1 - s_2)$; continuing through all k filters yields Equation (1). \square

The multiplicative structure has a critical implication: even moderate suppression at each stage compounds to near-total suppression overall.

Corollary 1 (Marginal Effect of a Single Filter). *Removing filter F_j (setting $s_j = 0$) increases the survival rate by a factor of $\frac{1}{1 - s_j}$. The new survival rate becomes:*

$$P'(\text{survival}) = \frac{P(\text{survival})}{1 - s_j} \quad (3)$$

However, the absolute survival rate after removal is still $\prod_{i \neq j} (1 - s_i)$, which may remain very small if the remaining filters have high suppression rates.

This corollary captures the paper's central policy finding: removing a single filter yields a multiplicative improvement (potentially large in relative terms) but leaves the remaining filters' compound suppression intact (potentially devastating in absolute terms).

Proposition 2 (Minimum Intervention Threshold). *To achieve a target survival rate τ , the number of filters that must be removed depends on which filters are chosen. Removing the set $\mathcal{R} \subseteq \{1, \dots, k\}$ achieves:*

$$P_{\mathcal{R}}(\text{survival}) = \prod_{i \notin \mathcal{R}} (1 - s_i) \geq \tau \quad (4)$$

When filters have heterogeneous suppression rates, the optimal removal strategy is to eliminate the highest-suppression filters first.

6 Calibration

The pipeline model requires empirical grounding. This section calibrates per-filter suppression rates against CFPB data on consumer finance arbitration.

6.1 Empirical Anchor

The CFPB’s 2015 study provides the most comprehensive federal data on consumer arbitration ([Consumer Financial Protection Bureau, 2015](#)). Across six consumer finance markets—credit cards, checking accounts, prepaid cards, payday loans, private student loans, and mobile wireless—consumers filed approximately 600 arbitration cases annually, despite tens of millions of consumers bound by arbitration clauses.

The calibration proceeds from conservative assumptions:

- **Covered population:** approximately 80 million consumers bound by arbitration clauses across the six markets.
- **Harm rate:** 0.1% of covered consumers experience actionable harm annually—a conservative estimate given the prevalence of billing disputes, service failures, and deceptive practices.
- **Expected claims:** $80,000,000 \times 0.001 = 80,000$.
- **Observed filings:** 600.
- **Observed survival rate:** $600/80,000 = 0.75\%$.

6.2 Equal-Suppression Baseline

As a baseline, assume all six filters have identical suppression rates. If each filter has pass-through rate p , then $p^6 = 0.0075$, yielding:

$$p = 0.0075^{1/6} \approx 0.4424 \quad (5)$$

Each filter would suppress approximately 55.8% of claims reaching it, and the product $0.4424^6 = 0.0075$ reproduces the observed survival rate. This establishes a symmetric benchmark: if suppression were distributed equally, each filter would need to block roughly half of surviving claims.

6.3 Realistic Heterogeneous Rates: Consumer Finance

In practice, filters do not suppress equally. For consumer finance, the parameter-derived rates from the model are:

Table 1: Calibrated suppression rates for consumer finance

Filter	Suppression	Pass-through	Basis
Awareness failure	85.0%	15.0%	CFPB awareness data
Rational apathy	99.1%	0.87%	$h = \$35, c = \$200, p = 0.30$
Information asymmetry	40.0%	60.0%	Confidentiality losses
Mandatory arbitration	60.0%	40.0%	Procedural friction
Class action waiver	85.0%	15.0%	Aggregation suppression
Regulatory capture / edge closure	50.0%	50.0%	Operational barriers

The product of pass-through rates is:

$$0.15 \times 0.0087 \times 0.60 \times 0.40 \times 0.15 \times 0.50 = 0.0000235 \quad (6)$$

yielding a compound survival rate of 0.002%—approximately 2 claims from 80,000 harmed consumers. This is substantially below the CFPB’s observed 600 filings, which is expected: the CFPB figure includes filings that may not represent the “typical” suppression pathway (e.g., attorney-coordinated mass filings, high-dollar outliers), and our conservative harm rate assumption may understate actual harm prevalence.

The key insight from calibration is not the precise survival rate but the *dominance of rational apathy* in consumer finance. At 99.1% suppression, rational apathy alone eliminates nearly all claims before any other filter operates. The remaining five filters then suppress 99.7% of the residual.

6.4 Sensitivity Analysis

The calibration depends on the assumed harm rate. Table 2 shows how the equal-suppression rate varies:

Across the full range of plausible harm rates, the compound survival rate remains far below 5%, and each filter must suppress at least half of surviving claims to reproduce observed filings. The qualitative conclusion—massive claim suppression through compound filtering—is robust to the harm rate assumption.

Table 2: Sensitivity of calibration to assumed harm rate

Harm rate	Expected claims	Survival rate	Equal suppression
0.05%	40,000	1.50%	49.6%
0.10%	80,000	0.75%	55.8%
0.50%	400,000	0.15%	65.9%
1.00%	800,000	0.075%	71.7%
5.00%	4,000,000	0.015%	79.2%

7 Intervention Analysis

The pipeline model’s most policy-relevant application is analyzing the effect of removing individual filters—simulating targeted reforms.

7.1 Single-Filter Removal

Using the calibrated consumer finance rates from Table 1, Table 3 and Figure 1 show the effect of removing each filter individually (setting its suppression rate to zero while leaving all others unchanged).

Table 3: Effect of removing a single suppression filter (consumer finance)

Filter removed	Suppression	New survival	Claims	vs. baseline
<i>Baseline (none)</i>	—	0.0023%	2	1.0×
Awareness failure	85.0%	0.0156%	13	6.5×
Rational apathy	99.1%	0.2700%	216	108.0×
Information asymmetry	40.0%	0.0039%	3	1.5×
Mandatory arbitration	60.0%	0.0059%	5	2.5×
Class action waiver	85.0%	0.0156%	13	6.5×
Regulatory capture	50.0%	0.0047%	4	2.0×

The results reveal a striking asymmetry. Removing rational apathy—the dominant filter—produces a 108× improvement, from 2 to 216 claims. But 216 out of 80,000 is still only 0.27%. The remaining five filters block 99.7% of claims even after the most impactful single intervention. No other single filter removal exceeds a 6.5× improvement. Removing information asymmetry or regulatory capture produces negligible absolute gains (1–2 additional claims).

This is the paper’s strongest quantitative finding: *structural immunity is robust to piecemeal reform*. The multiplicative architecture means that removing one filter—however suppressive—cannot overcome the compound effect of the remaining five.

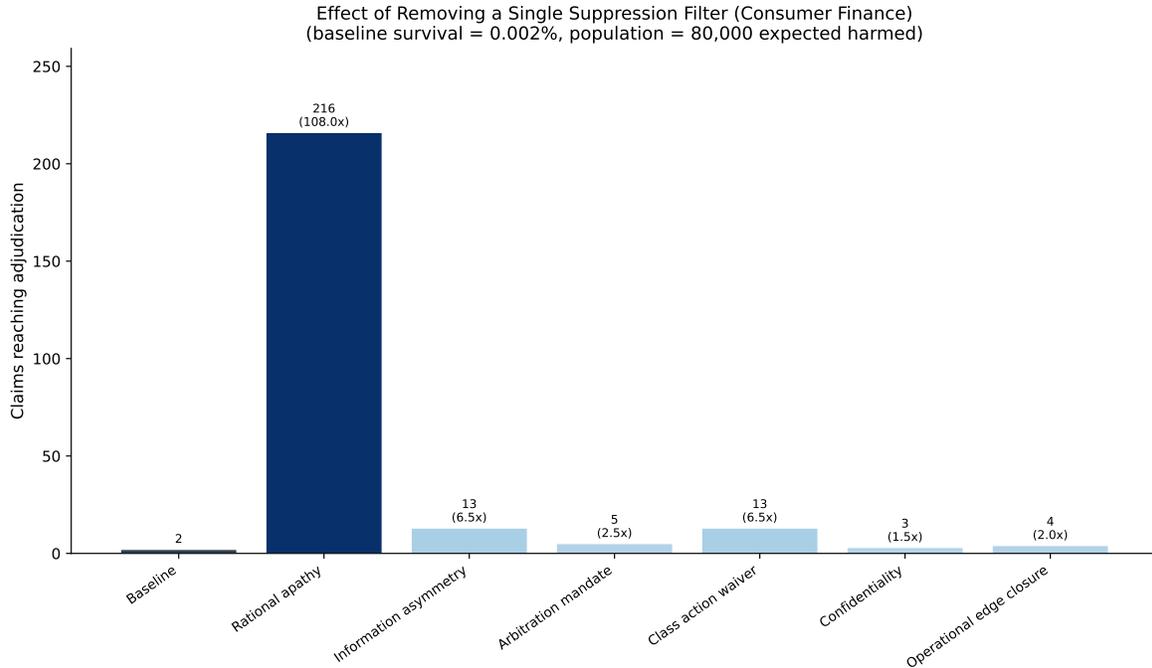


Figure 1: Effect of removing a single suppression filter on claims reaching adjudication (consumer finance sector). Removing rational apathy—the dominant filter at 99.1% suppression—produces a 108× improvement, but surviving claims increase only from 2 to 216 out of 80,000 harmed consumers. The remaining five filters still block 99.7% of claims.

7.2 Pairwise Filter Removal

Figure 2 presents a heatmap of all single and pairwise filter removals. The five most effective pairwise interventions are:

Table 4: Top pairwise filter removal effects (consumer finance)

Filters removed	Survival rate	Claims
Rational apathy + Awareness failure	1.80%	1,440
Rational apathy + Class action waiver	1.80%	1,440
Rational apathy + Mandatory arbitration	0.68%	540
Rational apathy + Regulatory capture	0.54%	432
Rational apathy + Information asymmetry	0.45%	360

Every effective pair includes rational apathy. The best pair not involving rational apathy—awareness failure plus class action waiver—achieves only 0.10% survival (83 claims). This confirms that the dominant filter anchors the system: reforms that do not address the dominant filter produce marginal improvements regardless of which other filters they target.

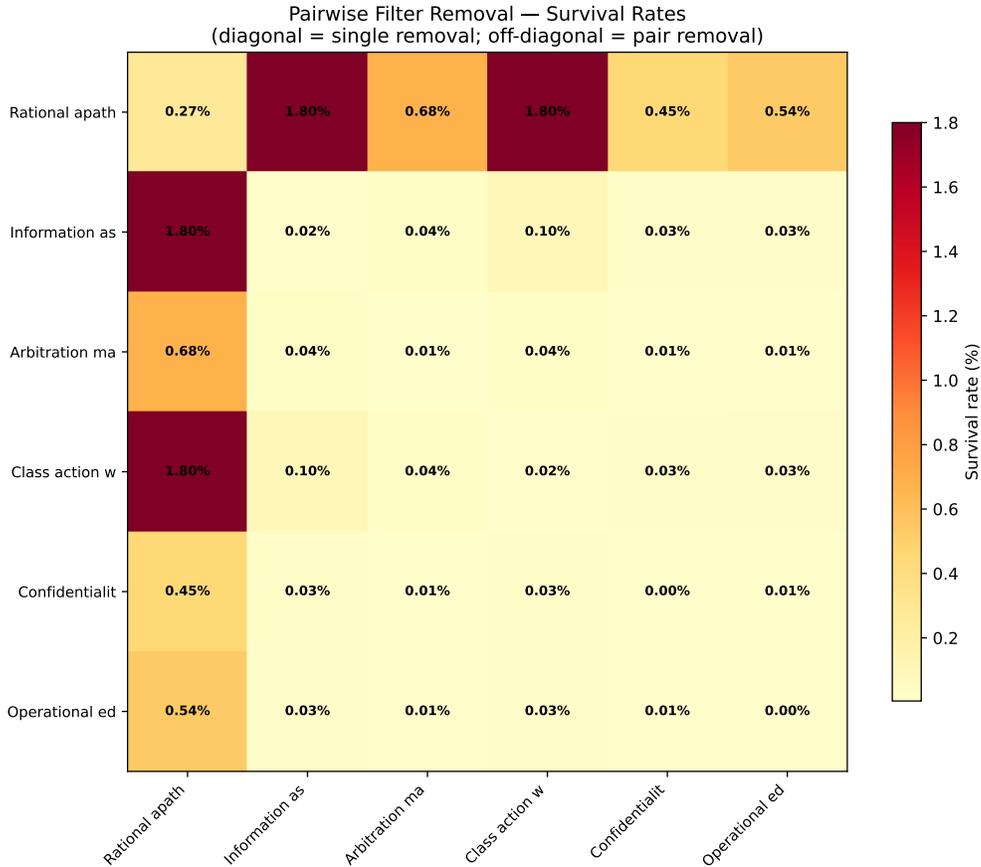


Figure 2: Pairwise filter removal heatmap. Diagonal cells show single-filter removal survival rates; off-diagonal cells show the survival rate when both filters are removed simultaneously. Every effective pair includes rational apathy—the dominant filter. Pairs not involving rational apathy produce survival rates below 0.11%.

7.3 Minimum Intervention Thresholds

Table 5 shows the minimum number of filters that must be removed to achieve given survival targets, using the optimal combination at each level.

Achieving 10% survival—still leaving 90% of meritorious claims unsatisfied—requires dismantling at least three filters simultaneously. Achieving 50% survival requires dismantling five of six. This explains why decades of incremental reform—a narrow carve-out here, a regulatory initiative there—have failed to restore meaningful access to adjudication.

8 Cross-Sector Extension

The pipeline model is not specific to consumer finance. This section parameterizes the model for three additional sectors—gig economy, social media, and healthcare—to test whether the multi-

Table 5: Minimum filters to remove for target survival rates (consumer finance)

Target	Filters	Achieved	Filters to remove
1%	2	1.80%	Rational apathy, Awareness failure
5%	3	12.00%	Rational apathy, Awareness, Class waiver
10%	3	12.00%	Rational apathy, Awareness, Class waiver
25%	4	30.00%	Rational apathy, Awareness, Arbitration, Class waiver
50%	5	60.00%	Rational apathy, Awareness, Arbitration, Class waiver, Reg. capture

plicative structure and dominant-filter dynamics generalize.

8.1 Sector Parameterization

Table 6 summarizes the sector-specific parameters and model outputs.

Table 6: Cross-sector comparison of pipeline model outputs

	Consumer Finance	Gig Economy	Social Media	Health- care
Population	80M	7M	250M	30M
Harm rate	0.1%	10.0%	2.0%	5.0%
Expected harmed	80,000	700,000	5,000,000	1,500,000
Model surviving claims	2	10,584	4	23,901
Observed filings	600	400	50	800
Compound survival	0.002%	1.51%	0.00007%	1.59%
Dominant filter	Rational apathy	Class action waiver	Rational apathy	Information asymmetry
Dominant suppression	99.1%	80.0%	99.3%	75.0%

8.2 Sector-Specific Findings

Gig economy. Gig workers experience high-dollar harms (misclassification, lost benefits, wage theft averaging \$3,000 per worker), so rational apathy is effectively zero—workers know they are being harmed and the amounts are large enough to justify pursuit. Instead, *class action waivers* dominate at 80% suppression: workers’ claims are individually viable but far more efficiently pursued collectively, and platform contracts prohibit aggregation. The model predicts 10,584 surviving claims from 700,000 harmed workers (1.51% survival), substantially above the approximately 400 observed filings—suggesting additional suppression mechanisms (platform retaliation risk, immigration status concerns) not captured in the six-filter model.

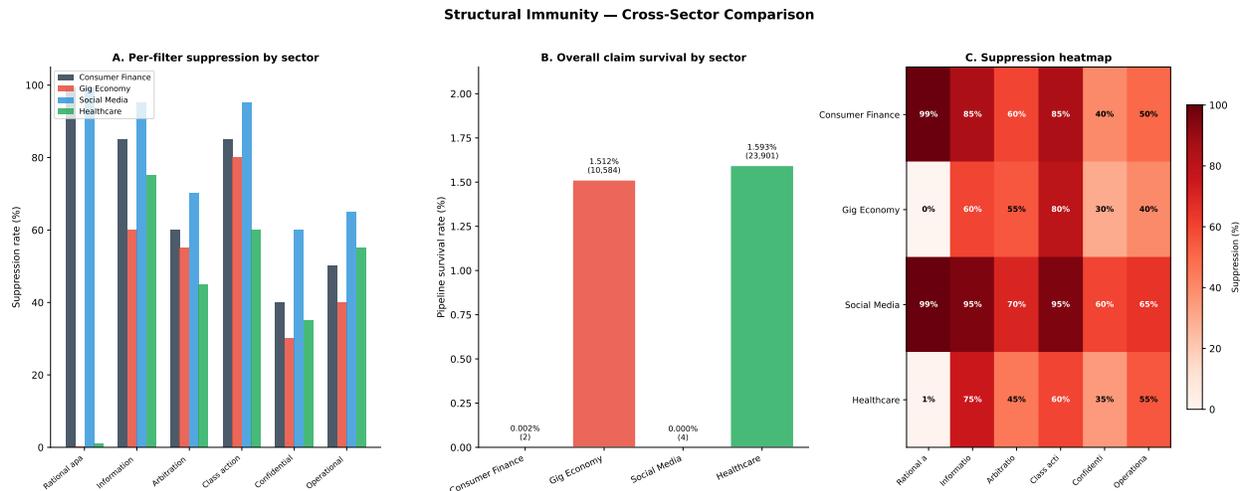


Figure 3: Cross-sector comparison of structural immunity. Panel A: per-filter suppression rates by sector, showing that different filters dominate in different industries. Panel B: overall claim survival rates, with social media showing near-total immunity and gig economy/healthcare showing higher (but still low) survival. Panel C: suppression rate heatmap across sectors and filters.

Social media. Structural immunity is near-total. Per-user dollar harm is minimal (\$5), making rational apathy extreme (99.3%). But *every* filter operates at high suppression: only 5% of users are even aware they have a claim (95% awareness suppression), class action waivers suppress 95%, and information asymmetry and operational edge closure each exceed 60%. The compound survival rate is $7.2 \times 10^{-5}\%$ —approximately 4 claims from 5 million harmed users. With 250 million users and diffuse harms, social media platforms enjoy the most complete structural immunity of any sector analyzed.

Healthcare. Average harm is high (\$1,500 for surprise bills), so rational apathy is near zero (1.0%). The dominant filter is *information asymmetry* at 75%: patients often know they received an unexpected bill but cannot navigate the complex insurer-provider dispute system. The No Surprises Act has partially reduced arbitration friction (45% vs. 60% in consumer finance), and class action waivers are less relevant because individual claims can be viable. The model predicts 23,901 surviving claims from 1.5 million harmed patients (1.59% survival)—well above the 800 observed IDR filings, suggesting the healthcare system’s complexity introduces additional barriers beyond the six-filter model.

8.3 The Dominant-Filter Insight

The cross-sector analysis reveals a result with direct policy implications: *the dominant filter varies by industry* (Figure 3). Rational apathy dominates where per-unit harm is low (consumer finance, social media). Class action waivers dominate where harm is high but claims benefit from ag-

gregation (gig economy). Information asymmetry dominates where harm is high and claims are individually viable but procedurally opaque (healthcare).

This variation means that sector-specific intervention strategies are necessary. A reform targeting rational apathy (e.g., fee-shifting, filing subsidies) would be highly effective in consumer finance but irrelevant in the gig economy where rational apathy is already zero. Conversely, restoring class action rights would be transformative for gig workers but less impactful in healthcare where individual claims are often viable. *One-size-fits-all reform is mathematically suboptimal* under the pipeline model: the optimal intervention always targets the dominant filter first.

9 Legal Architecture: How We Got Here

The current regime emerged from a series of Supreme Court decisions transforming the Federal Arbitration Act from a statute governing commercial disputes between sophisticated businesses into a mechanism immunizing consumer-facing corporations from collective accountability.

9.1 The Doctrinal Trajectory

Congress enacted the FAA in 1925 to address courts' refusal to enforce arbitration agreements between commercial parties of relatively equal bargaining power. For decades, it operated as intended. The transformation came through three decisions.

AT&T Mobility v. Concepcion (2011) held that the FAA preempts state unconscionability doctrines that “disfavor” arbitration ([con, 2011](#)). California's *Discover Bank* rule—which found class waivers unconscionable when three conditions converged: adhesion, small damages, and alleged deliberate schemes—was eliminated. Corporate adoption accelerated immediately: mandatory arbitration rose from 22.3% to 41% among 100 consumer-facing businesses after *Concepcion*, and to 66% after *Italian Colors* ([Miller, 2019](#)). [Horton \(2025\)](#) reports that approximately 80% of Fortune 500 companies now mandate arbitration for customers or workers.

American Express v. Italian Colors (2013) closed the “effective vindication” escape hatch ([ita, 2013](#)). Even when expert costs to prove antitrust claims exceeded several hundred thousand dollars against maximum recovery of \$38,549, the Court enforced the class waiver. Justice Kagan's dissent warned of a “get out of jail free card.”

Epic Systems v. Lewis (2018) extended these principles to employment ([epi, 2018](#)). A 2017 EPI study found 60.1 million American workers—56.2% of private-sector, non-union employees—bound by mandatory arbitration ([Colvin, 2017](#)).

9.2 Legislative Failure as Filter 6

Legislative responses illustrate the regulatory capture filter in operation. The FAIR Act has passed the House twice and stalled in the Senate both times. The CFPB’s 2017 rule was nullified via Congressional Review Act resolution—Vice President Pence cast the tie-breaking vote—with a poison pill barring substantially similar rules ([United States Senate, 2017](#)). California’s AB 51 was enjoined by the Ninth Circuit as preempted, and the state was ordered to pay \$822,496 in attorneys’ fees. The pattern is systematic: concentrated corporate lobbying interests prevail over diffuse consumer interests regardless of majority preference.

10 The Limits of Immunity: *Heckman v. Live Nation*

Structural immunity is not absolute. The Ninth Circuit’s decision in *Heckman v. Live Nation Entertainment, Inc.*, No. 23-55770 (9th Cir. Oct. 28, 2024), cert. denied Oct. 6, 2025, demonstrates that immunity has limits when corporate procedural engineering becomes sufficiently extreme ([hec, 2024](#)).

Ticketmaster had adopted “New Era ADR,” a boutique arbitration provider whose rules included batching claims, binding absent claimants to bellwether outcomes, 10-page complaint limits, and 10-document evidence caps. The Ninth Circuit found both procedural unconscionability (adhesion plus extreme power imbalance, unilateral retroactive modification) and substantive unconscionability (bellwether system violating due process, asymmetrical appeal rights). Critically, the court held that *Concepcion*’s FAA preemption does not protect procedures this extreme—distinguishing rules that “disfavor arbitration” from unconscionability doctrines applied to procedures that abandon genuine dispute resolution.

In the pipeline model’s terms, *Heckman* constrains the maximum suppression rate achievable through any single filter. Corporations cannot push arbitration friction to 100% without triggering judicial intervention. But the ordinary operation of mandatory arbitration—producing the compound suppression documented throughout this paper—remains largely untouched. Standard AAA procedures that produce 99.7% compound suppression are not unconscionable under current doctrine. Structural immunity exists in fragile equilibrium: robust against ordinary challenges, vulnerable when corporate procedural engineering triggers judicial intervention.

11 Reform Analysis

The pipeline model explains both why piecemeal reform fails and what successful reform requires.

11.1 Why Single-Filter Interventions Fail

The mathematical structure is clear from Equation (3): removing filter j multiplies survival by $1/(1-s_j)$. Even for the dominant filter (rational apathy at $s = 0.991$), this yields a $115\times$ multiplier—applied to a base of 0.002%, producing 0.27%. The remaining five filters’ compound suppression ($1 - 0.0027 = 99.73\%$) is barely affected.

This is not a contingent empirical finding. It is a mathematical consequence of the multiplicative structure. For k independent filters each with suppression rate s , removing one filter changes survival from $(1-s)^k$ to $(1-s)^{k-1}$ —an improvement factor of $1/(1-s)$. But $k-1$ filters at high suppression still produce near-total blockage. With six filters each at the equal-calibrated rate of 55.8%, removing one changes survival from 0.75% to 1.69%—better, but still catastrophically low.

11.2 Successful Reform Requires Multi-Filter Intervention

Historical examples confirm the model’s prediction. The No Surprises Act (2022) in healthcare simultaneously addressed multiple filters: it reduced information asymmetry (standardized billing disclosures), lowered arbitration friction (created an accessible IDR process), and partially circumvented class action waivers (created a regulatory pathway that functions as de facto aggregation for common billing patterns). The result was a measurable increase in dispute resolution—from near-zero to hundreds of IDR filings per month. While still insufficient (our model predicts 23,901 potentially viable claims versus 800 observed filings), it represents the most significant single-legislation improvement in any sector analyzed.

By contrast, the Ending Forced Arbitration of Sexual Assault and Sexual Harassment Act (2022) addressed only one filter (mandatory arbitration) for only one claim type. Its impact, while important for affected claimants, does not address the structural immunity architecture for the vast majority of consumer and employment disputes.

11.3 Designing Effective Intervention

The pipeline model suggests a reform design principle: *effective intervention must target the dominant filter and at least two others simultaneously*. For consumer finance, this means:

1. **Address rational apathy** (99.1%): filing subsidies, mandatory fee-shifting, statutory minimum damages making small claims economically rational.
2. **Restore aggregation** (85.0%): eliminate or limit class action waivers for claims below a dollar threshold, enabling collective pursuit.

3. **Increase awareness** (85.0%): mandatory disclosure of arbitration outcomes, public databases of claim patterns, regulatory reporting requirements.

Removing these three filters would achieve 12% survival (Table 5)—still imperfect but an improvement of over 5,000× from the current 0.002%. The remaining three filters (information asymmetry, arbitration friction, regulatory capture) would continue suppressing 88% of claims, but 12% survival creates enough volume for market feedback, regulatory detection, and precedent development to begin operating.

12 Empirical Record

12.1 A Note on Identification

Arbitration confidentiality and private settlement practices make causal inference difficult. We cannot directly observe suppressed claims—by definition, they leave no trace. What we can observe are filing rates, outcome disparities, and recovery comparisons. Multiple independent research programs using different methodologies converge on the same findings: low filing rates, repeat-player advantages, outcome gaps. This convergence is suggestive but not conclusive.

12.2 Filing Rates (Prediction 1)

The CFPB found approximately 600 consumer arbitration filings annually across six markets despite tens of millions of covered consumers ([Consumer Financial Protection Bureau, 2015](#)). If 80 million consumers are bound by arbitration clauses and 600 file annually, the filing rate is approximately 0.001%. [Estlund \(2018\)](#) estimates 320,000 to 727,000 employment claims annually never reach any forum.

The pipeline model provides a structural explanation: the compound survival rate under calibrated parameters ranges from 0.002% (consumer finance) to 0.00007% (social media), consistent with observed near-zero filing rates across sectors.

12.3 Outcome Disparities (Predictions 2–3)

[Colvin \(2011\)](#) found employee win rates of 21.4% in arbitration versus 36.4% in federal court and 59% in state court trials. Median awards were approximately 21% of federal court medians. [Bingham \(1997\)](#) found employees recovered 48% of claims against first-time employers but only 11% against repeat players. [Colvin and Gough \(2015\)](#) found awards dropped to \$7,451 when the same arbitrator-employer pairing recurred. [Egan et al. \(2025\)](#) found random arbitrator selection would increase consumer awards by approximately \$60,000 in securities arbitration.

12.4 Pattern Recognition Failure (Prediction 4)

Over five years, class settlements made 160 million consumers eligible for approximately \$2.7 billion in relief. Arbitration produced affirmative consumer relief in 32 disputes totaling under \$400,000 ([Consumer Financial Protection Bureau, 2015](#)). The ratio illustrates what aggregation accomplishes and what its elimination costs.

13 Falsification Conditions

This analysis makes empirical claims that could be wrong.

If arbitration filing rates were high relative to covered populations, the claim-suppression thesis would weaken substantially. Current data shows approximately 600 consumer filings annually against tens of millions of covered consumers. If data emerged showing filing rates approaching even 1% of covered populations, the suppression claim would require substantial revision.

If arbitration outcomes showed no repeat-player effect after controlling for case characteristics, the structural advantage thesis would weaken. Multiple studies using different methodologies find the effect persists. Methodologically superior research finding no effect would require revision.

If the pipeline model's multiplicative structure were contradicted by evidence of filter dependence, the quantitative analysis would require reformulation. If removing one filter substantially weakened others (positive correlation) or strengthened others (negative correlation), the independence assumption would fail and the intervention analysis would need revision.

If single-filter reforms produced survival rate improvements substantially exceeding the model's predictions, the multiplicative structure would be falsified. If removing class action waivers alone (without addressing rational apathy) produced survival rates above 5% in consumer finance, the model's parameterization or structure would require revision.

If opt-out rates were substantial among ordinary users, the coercion thesis would weaken. Available evidence suggests opt-out rates are negligible, but platforms do not disclose this data.

The analysis stands or falls on these empirical foundations, not on normative assertions about fairness.

14 Limitations

The pipeline model introduces formal structure at the cost of simplifying assumptions that warrant explicit acknowledgment.

Estimated, not directly measured, suppression rates. The per-filter suppression rates in Table 1 are not directly observed. They are back-computed from aggregate data (total filings vs. expected claims) and informed by qualitative evidence about each filter’s operation. Different parameterizations consistent with the same aggregate survival rate would produce different intervention analysis results. The qualitative conclusions—multiplicative compounding, dominant-filter dynamics—are robust to reasonable re-parameterization, but the specific numerical results (e.g., “108× improvement”) depend on the chosen rates.

Independence assumption. The model assumes filters operate independently: a consumer’s probability of being suppressed by Filter 3 is unaffected by whether they survived Filter 2. In practice, filters may be correlated. Consumers who overcome rational apathy (high motivation, high harm) may also be more likely to overcome information asymmetry (they invest in learning the system). If filters are positively correlated in pass-through, the model overestimates compound suppression; if negatively correlated, it underestimates. The direction and magnitude of correlation are unknown.

Static model. The pipeline captures a snapshot of suppression rates but does not model dynamic effects. Reform in one filter may change others: restoring class action rights might increase awareness (through media coverage of class actions), reduce rational apathy (through attorney involvement), and undermine information asymmetry (through public discovery). These feedback effects are not captured by the static multiplicative model and could make multi-filter intervention more effective than the model predicts.

Ordinal insensitivity. The multiplicative model is insensitive to filter ordering—the product is the same regardless of which filter operates first. In practice, ordering may matter: a consumer who never becomes aware of a harm (Filter 1) never faces rational apathy (Filter 2). The model treats all filters as operating on the same population, which is appropriate for computing compound survival but obscures the sequential decision process individual consumers face.

Sector parameterization. The cross-sector extension relies on parameter estimates derived from a combination of published data, industry reports, and informed judgment. The gig economy and healthcare parameters are less empirically grounded than the consumer finance calibration. The sector comparison should be understood as illustrative of the dominant-filter phenomenon rather than as precise quantitative predictions for each industry.

15 Implications

15.1 For Policy Design

The pipeline model’s central lesson is negative: incremental reform within the current architecture cannot meaningfully restore access to adjudication. The multiplicative structure ensures that any single-filter intervention—however well-designed—will be overwhelmed by the remaining filters. Effective reform requires simultaneous, multi-filter intervention calibrated to the dominant filters in each sector.

This implies that the traditional approach to arbitration reform—targeting one mechanism at a time through legislation, regulation, or judicial doctrine—is structurally inadequate. The FAIR Act, if enacted, would address mandatory arbitration (one filter) but leave rational apathy, information asymmetry, class action waivers, and regulatory capture intact. Under the pipeline model, even complete elimination of mandatory arbitration friction would increase consumer finance survival from 0.002% to 0.006%—a $2.5\times$ improvement that leaves 99.99% of claims suppressed.

15.2 For Legal Scholars

The pipeline model provides a formal framework for what legal scholars have described qualitatively. [Estlund \(2018\)](#)’s “black hole” is the compound survival rate approaching zero. [Resnik \(2015\)](#)’s account of the “erasure of rights” is the multiplicative elimination of claims across filters. The model makes these metaphors precise and enables quantitative comparison across sectors, interventions, and parameter assumptions.

The model also generates testable predictions about cross-sector variation that existing scholarship has not systematically explored. If the dominant-filter hypothesis is correct, the most effective intervention should vary predictably by sector—a claim that can be evaluated as reform efforts proceed in different industries.

15.3 For Consumers and Workers

Small-dollar disputes with dominant platforms are effectively unremedied under current architecture regardless of merit. The pipeline model quantifies this: even in the best case (healthcare, with 1.59% survival), more than 98% of meritorious claims are suppressed. Practical responses include exercising arbitration opt-out rights where available, maintaining documentation, and supporting organizations pursuing multi-filter structural reform.

16 Conclusion

Structural immunity is the predictable output of a multiplicative suppression architecture. This paper has formalized that architecture as a six-filter pipeline in which awareness failure, rational apathy, information asymmetry, mandatory arbitration, class action waivers, and regulatory capture each independently suppress a fraction of surviving claims. The compound survival rate—the product of all pass-through rates—approaches zero across every sector analyzed.

The pipeline model’s calibration against CFPB data yields concrete numbers: 0.75% aggregate survival in consumer finance, with the realistic heterogeneous parameterization producing 0.002% survival (2 claims from 80,000 harmed consumers). Cross-sector extension reveals that structural immunity takes different forms in different industries—dominated by rational apathy in consumer finance and social media, by class action waivers in the gig economy, by information asymmetry in healthcare—but the multiplicative architecture produces near-total claim suppression in all cases.

The model’s most important contribution is the intervention analysis. Removing any single filter improves survival by at most two orders of magnitude but cannot overcome the remaining filters’ compound suppression. Achieving meaningful access to adjudication (10% survival or above) requires dismantling three or more filters simultaneously. This mathematical finding explains why decades of incremental reform—a narrow statutory carve-out here, a regulatory initiative there, a judicial doctrine there—have failed to restore access to adjudication. The architecture is designed, whether intentionally or emergently, to be robust to piecemeal intervention.

Yet immunity is not absolute. *Heckman v. Live Nation* demonstrates that courts retain capacity to intervene when corporate procedural engineering becomes sufficiently extreme. And the pipeline model itself suggests the path forward: multi-filter intervention, calibrated to each sector’s dominant suppression mechanism, can break the compound suppression cycle. The model does not prescribe specific policies, but it establishes a design constraint that any effective reform must satisfy: *target the dominant filter and at least two others simultaneously*.

The empirical record is consistent with what the framework predicts: filing rates orders of magnitude below harm prevalence, persistent repeat-player advantages, and class settlements making hundreds of millions eligible for billions in relief while arbitration produces outcomes in dozens of cases totaling hundreds of thousands of dollars. Whether the precise suppression rates estimated here will prove accurate as better data becomes available is an empirical question. That the multiplicative architecture produces structural immunity is a mathematical certainty.

References

- (2011). *AT&T Mobility LLC v. Concepcion*. 563 U.S. 333. Supreme Court of the United States.
- (2013). *American Express Co. v. Italian Colors Restaurant*. 570 U.S. 228. Supreme Court of the United States.
- (2018). *Epic Systems Corp. v. Lewis*. 584 U.S. 497. Supreme Court of the United States.
- (2024). *Heckman v. Live Nation Entertainment, Inc.* No. 23-55770 (9th Cir. Oct. 28, 2024), cert. denied sub nom. *Live Nation Entertainment, Inc. v. Heckman*, No. 24-1145 (U.S. Oct. 6, 2025). United States Court of Appeals for the Ninth Circuit.
- Bingham, L. B. (1997). Employment arbitration: The repeat player effect. *Employee Rights and Employment Policy Journal*, 1(1):189–220.
- Colvin, A. J. S. (2011). An empirical study of employment arbitration: Case outcomes and processes. *Journal of Empirical Legal Studies*, 8(1):1–23.
- Colvin, A. J. S. (2017). The growing use of mandatory arbitration: Access to the courts is now barred for more than 60 million American workers. Technical report, Economic Policy Institute.
- Colvin, A. J. S. and Gough, M. D. (2015). Individual employment rights arbitration in the United States: Actors and outcomes. *ILR Review*, 68(5):1019–1042.
- Consumer Financial Protection Bureau (2015). Arbitration study: Report to Congress, pursuant to Dodd-Frank Wall Street Reform and Consumer Protection Act § 1028(a). Technical report.
- Egan, M., Matvos, G., and Seru, A. (2025). Arbitration with uninformed consumers. *Review of Economic Studies*, 92(6):3888–3923.
- Estlund, C. (2018). The black hole of mandatory arbitration. *North Carolina Law Review*, 96(3):679–710.
- Horton, D. (2025). Forced arbitration in the Fortune 500. *Minnesota Law Review*, 109(5):2165.
- McEntire, J. (2025). Dysmemic pressure: Selection dynamics in organizational information environments. *arXiv preprint arXiv:2512.14716*. Also available at SSRN: <https://ssrn.com/abstract=6015814>.
- McEntire, J. (2026). Structural immunity: How platform dominance and mandatory arbitration create corporate impunity. Working Paper, Version 1.0, January 2026. Available at SSRN.

Miller, R. (2019). Next-gen arbitration: An empirical study of how the Fortune 500 adapted to class action waivers. *Georgetown Journal of Legal Ethics*, 32:793.

National Employment Law Project (2021). Forced arbitration helped employers who committed wage theft pocket \$9.27 billion in 2019 from workers in low-paid jobs. Technical report.

Resnik, J. (2015). Diffusing disputes: The public in the private of arbitration, the private in courts, and the erasure of rights. *Yale Law Journal*, 124(8):2804–2939.

United States Senate (2017). Congressional review act resolution nullifying CFPB arbitration rule. 51–50 vote, Vice President Pence casting tie-breaking vote, October 24, 2017.

Jeremy McEntire is the author of “The Cage and the Mirror” and papers on organizational dysfunction and information selection dynamics. Contact: jmc@kindex.tools

Disclaimer: This paper is academic research intended to contribute to scholarly and policy discourse on dispute resolution architecture. It does not constitute legal advice, does not represent any party in litigation, and should not be relied upon for legal decisions. Readers with specific legal concerns should consult qualified counsel.