WEIDU FANLING / 未读凡灵

THE CLAIMS MARKET

A Green Paper on Litigation Finance, Machine Intelligence and the Public Infrastructure of Remedy

PUBLIC DISCUSSION DRAFT • v0.3 • 17 AUGUST 2026

"When intelligence reduces the cost of reaching a remedy, who receives the dividend?"

TWO AUSTRALIAN LENSES: QANTAS FLIGHT CREDITS + ROBODEBT

Qantas makes the economics of collective remedy unusually visible. Robodebt makes the danger of machine-mediated legal consequence impossible to treat as hypothetical.

Status, scope and provenance

This is a green paper: an invitation to test a problem definition, evidence base and set of policy responses. It does not find misconduct by any person, law firm, litigation funder, airline, court, government agency or technology provider. It does not propose autonomous machine adjudication. Its purpose is narrower and more demanding: to ask how the cost of reaching legal remedy should change when information-processing costs fall, and how to prevent machine systems from turning uncertainty into coercive legal consequence.

ABOUT THIS PAPER

Prepared by Weidu Fanling / 未读凡灵 as an independent public-interest discussion project. No institutional affiliation or endorsement by a litigation funder, law firm, airline, court, government agency or AI provider is claimed. Load-bearing legal and quantitative claims were checked against primary or authoritative sources where reasonably available. The paper remains an exploratory policy document, not legal advice.

Method and evidence discipline

Version 0.3 follows two adversarial reviews. The first forced the paper to benchmark the Qantas figures, engage recent Australian class-action law and recognise that a partial claims market already exists. The second exposed a more consequential omission: the machine-intelligence half of the paper lacked a real Australian case study. Robodebt now provides that counterweight. It also required the paper to say explicitly when its own benchmark cuts against its initial intuition, identify the named legal and funding actors in the Qantas notice, qualify Victoria's Group Costs Order regime after Bogan, rank policy priorities, and apply its own transparency principles to its authorship and AI use.

Epistemic categories

Category	Meaning in this paper
Primary fact	Directly supported by a court document, statute, judgment, regulator, Royal Commission, government publication or user-supplied primary notice.
Comparison	A contextual benchmark that helps interpretation but may not be like-for-like across cases, periods or funding models.
Inference	A reasoned interpretation drawn from facts; it is labelled and should remain contestable.
Proposal	A policy or design option advanced for consultation, not a statement of current law.
Unknown / evidence gap	A material question for which this paper does not have a sufficiently current or comparable evidence base.

What this paper does not claim

- That litigation funding is gambling. The analogy is limited to risk pricing, contingent return and portfolio logic; funders also provide capital, adverse-cost protection and access to justice.
- That the Qantas proposed distribution is unusually poor. On the historical benchmark used here, its estimated 65% Group Member pool is above the 51% historical median reported for funded class actions, although the comparison is not like-for-like.
- That a 70% statutory floor is the correct answer. Australia consulted on that possibility in 2021; the proposal was contested and is treated here as policy history, not an endorsed rule.
- That Robodebt was an LLM or contemporary generative-AI system. It was an automated administrative scheme. Its relevance is institutional: machine-mediated legal consequence can scale error, opacity and burden-shifting long before modern LLMs enter the picture.
- That governance alone determines who receives AI savings. Competition, bargaining power, market concentration, court supervision, regulation and public infrastructure can all affect distribution.
- That legal judgment can be reduced to information processing. Evidence, credibility, advocacy, discretion, professional responsibility and rights-affecting decisions retain irreducibly human and institutional elements.

Executive Summary

The paper began with a simple observation. In the Qantas Flight Credits Class Action, a court-filed notice describes a proposed $105 million settlement and estimates that approximately $68 million, about 65%, will be available for distribution to Group Members after proposed deductions and the effect of interest, subject to Court approval.[1] That makes the economics of collective remedy unusually visible.

The first intuition was that roughly one dollar in three appearing between gross settlement and claimant pool revealed an unusually hungry legal-finance machine. The evidence does not support that framing. Treasury's 2021 consultation recorded historical analysis in which the median plaintiff return was 51% in funded class actions and 85% where no funder was involved.[2][3] On its face, Qantas's estimated 65% is fourteen percentage points above the historical funded median. Different cases, eras and cost structures prevent a clean ranking, but the direction matters: Qantas is not evidence of exceptional extraction. It is evidence that the distribution of collective remedies is a longstanding policy question.

The Qantas notice also makes clear who performs and finances the work. Echo Law and Piper Alderman are identified as the solicitors running the case; Omni Bridgeway and CASL are identified as joint litigation funders.[1] The proposed deductions include estimated pre-settlement legal costs and disbursements, settlement approval and administration costs, a litigation-funding commission, and representative payments. The funders are described as having financed legal costs and disbursements and indemnified the applicant against adverse costs.[1] Those functions are real. The public-interest question is therefore not whether intermediation exists, but whether its cost and return are proportionate, transparent and responsive when technology changes the amount of work required.

Robodebt supplies the missing second case study. It was not generative AI. It was an automated public-administration scheme in which income-averaging and data-matching processes were used to raise or pursue welfare debts at scale. The Royal Commission concluded that Robodebt was a costly failure of public administration and recommended that debt recovery be ethical, proportionate, consistent and transparent and that individual circumstances be taken into account.[15] The scheme shows the opposite danger from excessive private intermediation: public automation can collapse friction too far, too fast, turning uncertain or proxy data into coercive legal consequence without adequate authority, review, reasons or remedy.

The two cases therefore discipline each other. Qantas warns against pretending that access to collective remedy is costless. Robodebt warns against pretending that removing human and procedural friction is automatically progress. The correct objective is not maximum automation. It is proportionate friction: enough process to preserve legality, contestability and human dignity, but not so much informational cost that practical rights become economically inaccessible.

A partial claims market already exists. Litigation finance uses capital, risk assessment and portfolio logic; public legal-finance companies provide portfolio exposure; commercial claim-sourcing technologies have existed in forms that search court information. Future machine intelligence may not create this market so much as change its state by compressing research, intake, valuation, document review and administration costs. The distribution of those savings will depend on market competition and bargaining power as well as governance. The present paper does not have sufficiently current evidence to state whether Australia's litigation-funding market is competitive enough for those savings to be reliably bid through to claimants. Establishing that evidence is therefore a first-order research task, not a footnote.

THE TWO-YEAR POLICY STARTING POINT

Do the cheap, information-forcing things first. Priority 1: standardise class-action settlement transparency, including named legal/funding actors, calculation methods, adjustment triggers and timing. Priority 2: establish a Justice Productivity Dividend and market-evidence baseline so that automation savings, claimant outcomes, funding-model differences and competition can be observed before larger institutional reforms are attempted. High-consequence public automation should not be expanded merely because it is technically possible.

Seven working propositions

1. Rights weaken when the cost of identifying, proving or enforcing them is disproportionate to the value of the remedy.
2. Litigation funding can simultaneously enable access to justice and operate as an investment activity; policy must analyse both functions without collapsing either into caricature.
3. A partial claims market already exists. Machine intelligence is more likely to alter its scale, speed and information advantage than to create it from nothing.
4. The Qantas 65% estimate is not evidence of anomalous extraction; it is a transparent trigger for asking how remedy costs are justified, compared and reported.
5. Robodebt demonstrates that reducing administrative friction can produce severe injustice when proxy data, weak legal authority, opacity and inadequate review are allowed to generate rights-affecting consequence at scale.
6. The justice productivity dividend is distributed by a combination of competition, bargaining power, market structure, court supervision, regulation and public design. Its destination should be measured rather than assumed.
7. The public-interest objective is to reduce unnecessary informational tolls between right and remedy while preserving the human, evidentiary and institutional friction required for legality, fairness, contestability and correction.

Contents

HOW TO READ THIS PAPER

Sections 1-3 establish the Australian facts and law. Sections 4-6 test the problem against international practice and machine-risk design. Section 7 ranks the policy programme. Sections 8-10 specify measurement, objections and sequencing. The appendices preserve the two worked case studies and the legal timeline.

1. Two case studies, two directions of risk

1.1 Qantas: the economics of collective remedy made visible

The Federal Court notice in VID650/2023 concerns the proposed settlement of the Flight Credits Class Action against Qantas Airways Limited. The proceeding alleges that customers whose flights were cancelled between 1 January 2020 and 1 November 2022 should have received refunds within a reasonable time rather than flight credits, or experienced delay or disadvantage associated with those credits. Qantas has agreed, without admission of liability, to pay $105 million subject to Court approval.[1]

Item in filed notice	Current estimate / description	Who / what is identified
Gross proposed settlement	$105m	Qantas Airways Limited; no admission of liability
Legal costs & disbursements to 13 Mar 2026	$9.23m incl GST	Applicant's legal costs; Echo Law and Piper Alderman are named as the solicitors running the case
Settlement approval & administration	$4.76m incl GST, partly offset by interest	Administrator had not yet been formally appointed when the notice was filed
Litigation-funding commission	$25m-$26.25m; described as 25% of settlement amount, adjusted based on claims registered	Omni Bridgeway and CASL are named as joint funders
Applicant / sample-member payments	$20,000 + four × $5,000 = $40,000	Representative recognition payments
Estimated Group Member pool	≈$68m, ≈65%, after deductions and interest effect	Eligible registered Group Members; minimum proposed return $50

The notice says all legal costs will be reviewed by an independent costs assessor and that no amount will be deducted unless the Court approves it.[1] It also explains the access-to-justice counterweight: the funders paid legal costs and disbursements, indemnified the applicant against adverse costs and provided security-related support on the basis that they would receive a return only if the proceeding succeeded.[1]

1.2 The benchmark correction: 65% is not the scandal

Treasury's 2021 consultation, drawing on the Parliamentary Joint Committee and ALRC material, recorded historical analysis in which the median return to plaintiffs was 51% in funded class actions compared with 85% where no litigation funder was involved.[2][3] On the face of those figures, the Qantas estimate of approximately 65% is fourteen percentage points above the historical funded median.

INTERPRETIVE CAUTION

That comparison is not apples-to-apples. The historical medians aggregate different matters, periods, risks, legal-cost structures and funding arrangements. They do not establish that Qantas is unusually favourable. But they do establish something equally important: a 65% claimant pool should not be used as evidence that this particular settlement is unusually extractive. The case study is valuable because it makes the architecture visible, not because it proves abuse.

Australia also formally consulted in 2021 on whether legislation should guarantee a minimum gross return to class members, including whether 70% should be a floor and whether a graduated approach should reflect risk, complexity, length and likely proceeds.[2] The proposal was contested. That policy history shows that claimant share is a recognised governance question, but not one with a simple universal percentage answer.

1.3 What the Qantas notice reveals - and what it does not

The notice is unusually informative, but it does not fully expose the calculation logic behind every figure. The funding commission is described as between $25 million and $26.25 million, being 25% of the settlement amount and adjusted based on claims registered.[1] The notice does not, on its face, set out a complete reproducible formula showing how registration numbers generate the final commission within that range. Nor can the approximately $37 million aggregate difference between $105 million and the estimated $68 million claimant pool be reconstructed simply by adding the listed estimates, because the notice says interest partly offsets administration costs and the estimates may change.

This is exactly the type of limitation a transparency standard should address: not merely a range, but the calculation method, variables, adjustment triggers, expected interest treatment and a reconciliation from gross fund to expected claimant distribution.

1.4 Timing is part of transparency

Eligible Group Members must register by 4pm AEST on 2 October 2026 to receive a settlement payment if the settlement is approved. The notice states that the Court will be asked to approve the proposed settlement and distribution on 13 October 2026.[1] Registration is not the same thing as judicial approval, and Group Members who do nothing may still be bound by an approved settlement while receiving no payment. The sequence is lawful only if the governing rules and Court orders permit it; this paper makes no finding about that. The policy point is narrower: participation materials should make the chronology of registration, objection, opt-out, approval and payment unmistakable.

1.5 Robodebt: the opposite failure mode

Qantas does little by itself to test the machine-intelligence half of this paper. Robodebt does. The scheme was not an LLM and should not be retrospectively relabelled as one. It was an automated public-administration system centred on data matching and income averaging for welfare-debt compliance. Its significance is that automation was connected to legal and financial consequence at national scale.[15]

The Royal Commission's final report concluded that Robodebt was a costly failure of public administration. Its recommendations include designing policies and processes around the people they serve; making debt recovery ethical, proportionate, consistent and transparent; taking individual circumstances into account; and refraining from recovery action while a debt is under review or dispute except where lawful authority supports otherwise.[15]

The Federal Court-approved class-action remediation itself continued long after the scheme ended. Services Australia states that the original settlement distributed approximately $112 million in 2022 and that a new settlement, approved on 23 June 2026, provides an additional $475 million in compensation, $13.5 million in reasonable legal costs and up to $60 million in administration costs.[16] Those figures are not presented here as a direct comparison with Qantas; the claims, legal bases, populations and remediation structures are different. They show the scale of institutional cost when a consequential administrative system fails.

1.6 The paired lesson

Lens	Primary problem made visible	Wrong lesson	Better lesson
Qantas	Collective remedy requires legal work, capital, risk bearing and administration; those costs reduce the gross fund available to claimants.	All intermediation is extraction.	Measure whether each layer of intermediation is necessary, proportionate, competitive and transparent.
Robodebt	Automation can reduce administrative friction while magnifying poor source data, weak legality, burden-shifting and inadequate review.	Automation is inherently efficient and therefore desirable.	Preserve proportionate human and institutional friction wherever consequence is coercive, rights-affecting or difficult to reverse.

2. The claims market already partly exists

2.1 A bounded definition

This paper uses "claims market" as a heuristic for arrangements through which legal claims or portfolios of claims are identified, valued, financed, risk-managed or monetised. The term is deliberately imperfect. A legal claim is not an ordinary security, and litigation is not a sports wager. Funding typically requires contractual relationships, legal-merits analysis, duration risk, adverse-cost exposure and court supervision. The market is partial and illiquid.

2.2 Why the market analogy still matters

The analogy becomes useful when the same economic grammar appears: uncertain future payoff, capital at risk, information advantage, selection among opportunities, portfolio construction and contingent return. Publicly listed legal-finance businesses provide investors with exposure to portfolios of disputes; commercial claim-sourcing systems have historically been promoted as methods for finding litigation opportunities.[27][28] These examples are used only to establish market form, not to endorse provider claims about performance.

2.3 Machine intelligence changes the bottleneck

The most plausible near-term discontinuity is not a live exchange quoting odds on individual lawsuits. It is the reduction of the slow, bespoke underwriting burden. Machine systems can assist with public-record retrieval, document classification, chronology building, legal research, settlement modelling, claimant intake and portfolio monitoring. If those functions become cheaper, a funder or law firm can evaluate more matters with the same professional workforce.

That does not mean selection becomes accurate or just. Legal outcomes are partly endogenous to evidence quality, advocacy, procedure, settlement incentives, judicial discretion and changing law. Historical success patterns can encode institutional inequality. A model that predicts likely outcome is not a machine for discovering legal merit.

2.4 Competition is a missing variable

Version 0.2 treated the destination of AI cost savings too much as a governance choice. That was incomplete. In a genuinely competitive funding and legal-services market, lower underwriting or administration costs may be competed through into lower commissions, lower fees or better claimant terms. In a concentrated, opaque or high-barrier market, intermediaries may retain more of the gain. Court supervision and regulation can alter both outcomes.

EVIDENCE GAP - DO NOT ASSUME THE MARKET STRUCTURE

This paper has not established a sufficiently current 2026 measure of Australian litigation-funder concentration, effective competition, barriers to entry or pricing dispersion. Historical sources show that competition has long been part of the policy debate, but they are not an adequate basis for a present-tense concentration claim. A current market-structure study is therefore part of Priority 2.

2.5 Courts are observable, but not frictionless datasets

Public judgments, class-action pages, listings and some case metadata make court activity observable. That does not make every filing freely available for bulk machine ingestion, nor does it erase suppression, privilege, privacy, inspection rules, data quality problems or unequal representation. Claim-scanning policy should therefore distinguish what is public to read from what is lawful, ethical and technically appropriate to aggregate, infer from or commercialise.

3. Australia's legal and regulatory architecture

Australia's class-action and litigation-funding architecture has repeatedly changed. That history matters because the central policy question - who may finance collective remedy, on what terms and under whose supervision - is not new.

Date	Development	Policy significance
2019	BMW Australia Ltd v Brewster [2019] HCA 45	Limits common fund orders made at an early stage under the general case-management powers considered by the Court.[4]
2020	Litigation funding moved into AFSL / managed investment scheme regulation	Illustrates financial-services regulation as one model for controlling funding activity.
2020	Victoria's Group Costs Order regime commenced	Creates a statutory exception allowing the Supreme Court of Victoria to order legal costs as a percentage of recovery in a group proceeding.[7]
2021	Treasury consultation on minimum gross return, including a 70% floor	Shows claimant-share policy was already a national regulatory question.[2]
2022	Litigation-funding arrangements largely exempted again from MIS / AFSL settings	Shows regulatory oscillation rather than settled consensus.[8]
2025	Bogan v Estate of Smedley [2025] HCA 7	A Victorian GCO is relevant to transfer and does not remain enforceable if the proceeding transfers to NSW; the GCO model has jurisdictional consequences.[5]
2025	Kain v R&B Investments Pty Ltd [2025] HCA 28	Confirms settlement/judgment common-fund-order power in the Federal Court while distinguishing solicitor remuneration constrained by professional fee rules.[6]
2025	ASIC extends specified litigation-funding relief to 31 Jan 2029	Current settings remain partly shaped by exemption instruments and ongoing policy choices.[9]
2026	Federal Court GPN-AI; Supreme Court of Victoria SC GEN 25 and Judicial Guidelines	AI is recognised as potentially useful but bounded by verification, duties, privacy, human authority and judicial responsibility.[11][12][13]
2026	Ba v Sterling Parts Australia Pty Ltd [2026] FedCFamC2G 1245	Concrete Australian warning: generative AI can fabricate or misdescribe evidentiary material; verification and responsibility remain with the person using it.[14]

3.1 Victoria's alternative is real, but it does not simply travel

Victoria's Group Costs Order regime is an important alternative to conventional third-party funding because the Supreme Court of Victoria may order that legal costs be calculated as a percentage of any award or settlement and shared among the plaintiff and group members.[7] But Bogan v The Estate of Peter John Smedley (Deceased) [2025] HCA 7 adds a critical qualification. The High Court held that the fact a Victorian GCO has been made is relevant to the statutory discretion whether to transfer a proceeding to another court; if transferred to the Supreme Court of New South Wales, the Victorian GCO would not remain in force and enforceable there.[5] Funding design can therefore affect forum and transfer questions, not merely price.

3.2 Kain: timing, recipient and professional regulation matter

Kain confirms that the Federal Court's powers at settlement or judgment can support a common-fund-style order for a litigation funder, while the Court did not permit the proposed equivalent for a law practice where that would conflict with the applicable prohibition on contingency fees.[6] Read with Brewster, the lesson is not "funding percentages are lawful" or "funding percentages are unlawful". The answer depends on statutory power, procedural stage, recipient and professional regulation.

3.3 Australian courts are building AI guardrails now

The Federal Court's April 2026 GPN-AI recognises potential benefits for efficiency, cost reduction and access to justice while emphasising responsible use consistent with existing legal and professional obligations.[11] The Supreme Court of Victoria's SC GEN 25 similarly recognises access-to-justice potential while addressing verification, privacy and sanctions; separate Judicial Guidelines state that generative AI must not be used for judicial decision-making, although Court-approved tools may support defined judicial and staff functions.[12][13]

Ba v Sterling Parts provides the practical caution that abstract guidance needs. Judge Symons described the matter as a cautionary tale about AI-powered chatbots masquerading as legal advisers and inferred that an evidence list identified documents that did not exist; the claim was dismissed after failures to comply with production orders.[14] The lesson is not that self-represented litigants should be denied AI. It is that access-to-justice assistance must preserve verification, provenance, responsibility and the distinction between organising evidence and inventing it.

4. International and public-infrastructure comparators

4.1 United Kingdom: PACCAR shows that funding legitimacy remains contested

In PACCAR [2023] UKSC 28, the UK Supreme Court treated certain percentage-of-damages litigation funding agreements as damages-based agreements, creating enforceability consequences.[20] In December 2025 the UK Government announced reform intended to restore enforceability while developing a wider regulatory framework.[21] That response is not a neutral international endorsement of third-party funding. It is a policy choice favouring legal certainty and access to finance while regulation remains contested. The comparator is valuable precisely because similar tensions - return, enforceability, access and oversight - produce different institutional answers.

4.2 British Columbia: digital justice can be public without being autonomous

British Columbia's Civil Resolution Tribunal is a public online tribunal designed to resolve specified disputes through digital pathways under statutory authority.[22][23] It is not a blueprint for automated class-action adjudication. Its relevance is more modest: public digital justice can combine online service, legal authority, human adjudication and accessible process instead of treating commercial AI as the only path to lower transaction cost.

4.3 European and multilateral safeguards

The EU AI Act treats specified justice-related AI uses as high-risk and embeds risk-management, data-governance, documentation, oversight and accountability requirements.[24] The Council of Europe's CEPEJ charter emphasises fundamental rights, non-discrimination, quality/security, transparency and user control in judicial AI.[25] OECD people-centred justice work similarly emphasises accessibility, evidence and service design rather than technology for its own sake.[26]

4.4 Australia's legal-assistance data strategy is an institutional opening

The National Legal Assistance Data Strategy 2026-2030, endorsed by the Standing Council of Attorneys-General and published in August 2026, commits governments and the legal-assistance sector to more effective and responsible use of data for evidence-based decisions and service outcomes.[17] That is a plausible home for evidence-building about access, unmet need and service outcomes. It is not authority for automated entitlement decisions. The distinction between data for service improvement and data for coercive legal consequence must remain explicit.

5. Machine intelligence: three futures

5.1 The relevant stack is broader than LLMs

Future legal intelligence is likely to combine authenticated document ingestion, retrieval, legal databases, rules engines, deterministic calculators, statistical models, language models, workflow systems and human review. Treating all of this as "an LLM" obscures the important question: which component is authorised to do what, using which source, at what level of consequence?

5.2 Future A - specialist claims finance becomes more efficient

Existing firms and funders use machine assistance to reduce research, intake, document review, drafting, administration and monitoring costs. Professional underwriting remains bespoke. If competition is effective, some savings may pass through to lower prices or better claimant terms. If competition is weak, they may become margin. The policy task is measurement before intervention.

5.3 Future B - high-throughput claims finance

Scanning, clustering and valuation become sufficiently cheap that thousands of candidate matters can be continuously screened. This does not create a conventional securities exchange, but it can make legal grievances behave more like financial inventory: selected, scored and allocated capital according to expected risk-adjusted return. The danger is not merely more litigation. It is that socially important but low-monetary-value harms remain invisible while commercially attractive harms receive extraordinary machine attention.

5.4 Future C - public rights-and-remedy infrastructure

The same capabilities can be directed toward rights discovery rather than claim acquisition. A public or court-adjacent system could help a person identify a potential entitlement, gather authoritative records, calculate deterministic components, explain uncertainty and route genuine disputes to qualified humans. Its purpose would be to reduce informational tolls, not to automate contested legal judgment.

5.5 The distribution equation

REVISED CLAIM

Governance matters, but it does not determine the dividend alone. The distribution of AI-driven cost savings is a function of at least: competition among funders and firms; bargaining power; market concentration and entry barriers; court supervision; professional duties; funding and fee rules; public alternatives; data access; and the residual risk that cannot be automated. Policy should observe these variables before deciding which lever is actually failing.

This yields a more disciplined research question than “Will AI replace lawyers?”: when a defined legal function becomes materially cheaper, does the saving lower the total cost of remedy, expand access, improve claimant share, shorten delay, improve respondent efficiency, increase intermediary profit, or some combination? The answer should be measured case by case and over time.

6. A safe rights-and-remedy architecture

The goal of this section is not to design an AI judge. It is to identify the minimum architecture required before machine systems materially affect access to money, legal status or remedy. Robodebt converts what was previously a counterfactual exercise into a real institutional stress test.

6.1 Layer 1 - authoritative source truth

Consequential systems should begin with authoritative records and explicit provenance. Statistical proxies, inferred values or aggregated data should not silently become legal facts. If a proxy is legally authorised, the legal basis and evidentiary limitation should be visible. If data conflicts, the conflict should be surfaced rather than averaged away.

6.2 Layer 2 - legal authority, rule versioning and burden allocation

Every material rule should point to its legal or scheme authority and effective version. The system should identify who bears the burden of establishing a fact. It should not reverse that burden merely because machine-scale administration makes it convenient to demand that an individual disprove an inferred liability.

6.3 Layer 3 - uncertainty routing

Deterministic calculations can be automated. Ambiguous facts and law require a different route. A system should be able to say “not enough information”, identify what is missing and route the issue to an authorised human. Confidence should never substitute for legal authority.

6.4 Layer 4 - notice, reasons and meaningful participation

People should receive a plain-language explanation of what data was used, what rule was applied, what the result means, which choices and deadlines exist, and what happens if they do nothing. The explanation must be usable by people with disability, low literacy, limited English, unstable housing or limited digital access. Digital-first must not become digital-only.

6.5 Layer 5 - human review and independent escalation

A material adverse consequence should have an accessible human review path. The reviewer must be able to reconstruct the machine-assisted process from source records, rules and calculation history. Review that merely repeats the original machine output is not meaningful review. High-consequence systems require an escalation path sufficiently independent of the original decision channel.

6.6 Layer 6 - calculation, payment and change provenance

Every entitlement or payment calculation should be reproducible. Formula versions, source changes, overrides, corrections and payment events should be logged. If a rule or dataset changes, the system should be able to identify who was affected and whether retrospective correction is required.

6.7 Layer 7 - aggregate audit, contestability and remedy

Public reporting should measure exclusion, reversal, delay, participation, error, demographic disparity, human overrides and correction. The system must support remedy when it fails: reopening, recalculation, compensation where legally available, notification of affected people and institutional learning. Audit without a correction pathway is observation, not accountability.

6.8 Robodebt safeguard map

Observed institutional failure / risk	Design safeguard it implies	Why it matters
Income averaging / proxy data used in debt process	Authoritative source truth; explicit legal authority for any proxy	Convenient inference must not silently become a legally operative fact.
Large-scale debt consequence	Consequence classification and proportional human oversight	Scale multiplies small rule or data errors into systemic harm.
Burden placed on individuals to resolve inferred debt	Explicit burden allocation; no convenience-based reversal	Administrative efficiency cannot redefine who must prove liability.
Unclear or contested legality persisted	Rule-source trace; legal sign-off; stop / escalation trigger	A machine workflow must not outrun the legal authority on which it depends.
Weak capacity for individual circumstances to alter process	Meaningful human review; contextual evidence; hardship routing	A lawful general rule can still operate unjustly if individual facts cannot enter.
Opaque and distressing debt recovery	Plain reasons, review status, pause mechanisms and humane communications	Procedure itself can impose harm even before final liability is resolved.
Systemic failure required years of remediation	Audit logs, cohort impact analysis, rollback and remedy architecture	Reversibility and cohort tracing lower the cost of correcting system-wide error.

6.9 Eight public-interest design constraints

1. Source and authority: every consequential output must be traceable to authoritative data and an identified legal or scheme rule.
2. Human accountability and proportionality: the level of human authority, verification and pause must rise with consequence.
3. Reasons, review and remedy: people need understandable reasons, correction rights, meaningful human review and a route to repair.
4. Privacy, purpose limitation and security: collect the minimum data necessary; prohibit unrelated commercial reuse without lawful basis; secure and log access.[18][19]
5. Accessibility, equity and Indigenous Data Governance: digital systems must serve people experiencing disadvantage and respect appropriate governance of Aboriginal and Torres Strait Islander data.[17]
6. Respondent fairness: machine assistance must not erode evidentiary standards, duplicate claims, obscure who makes allegations or impair meaningful defence.
7. Transparency and conflicts: operators, funders, legal representatives, data providers and decision authorities should be identifiable, with material conflicts disclosed.
8. Measure the dividend: lower technical cost is not itself a justice outcome; track who benefits, who is excluded and whether delay, accuracy and practical access improve.

7. Priority policy programme

The earlier version presented ten co-equal options. That was analytically comprehensive and politically weak. Version 0.3 ranks six moves by leverage, reversibility and institutional burden. The first two can begin without building a new tribunal or AI agency.

Priority 1 - National settlement transparency standard

Time horizon: 0-12 months. Develop a standard public template for class-action settlement approval materials, adaptable by courts to particular proceedings. It should report gross settlement, legal costs and disbursements, funding return, administration, interest, representative payments, claimant pool, participation assumptions, named legal and funding actors, and whether figures are estimate or actual.

Crucially, the standard should disclose calculation methods rather than only outputs: the funding formula, adjustment variables, sensitivity to participation or registration, interest treatment, administration assumptions and a reconciliation from gross fund to expected claimant distribution. It should also show the chronology of notice, objection, opt-out, registration, approval and expected payment.

Priority 2 - Justice Productivity Dividend + market evidence baseline

Time horizon: 0-18 months. Require or encourage settlement and justice-sector pilots to state, where reasonably measurable, which functions were automated or materially accelerated, what cost or time changed, what residual risk remained and where the benefit appeared: claimant share, fee, commission, administrative cost, delay, respondent cost or intermediary margin.

At the same time, commission an empirical Australian market study covering funder count, concentration, entry barriers, co-funding, pricing structures, duration, capital deployed, adverse-cost exposure, funded / unfunded / GCO outcomes, and post-2020 claimant distributions. Without this evidence, claims that competition will solve the dividend problem - or that regulation must solve it - remain underdetermined.

Priority 3 - Purpose-bound claim origination and data governance

Time horizon: 1-2 years. Establish rules for machine-assisted claim discovery and clustering. Publicly accessible information is not automatically free of privacy, suppression, solicitation, profiling or secondary-use constraints. Data collected for court-directed settlement administration or public entitlement delivery should not become a commercial claim-sourcing feed without lawful authority and appropriate consent.

Priority 4 - Human review, reasons and correction for material machine-assisted determinations

Time horizon: 1-2 years. Where a machine system materially contributes to excluding a person from a remedy, lowering a payment, asserting a debt, generating evidence, missing a deadline path or otherwise affecting a legal interest, require an accessible human review route, identified source data, rule provenance and a record of correction or override.

Priority 5 - Narrow public or court-adjacent entitlement pilots

Time horizon: 2-4 years. Pilot only high-volume, rule-like, low-discretion domains with good source data and clear legal authority. Systems may help discover potential eligibility, collect evidence and calculate deterministic components, but adverse or contested outcomes should route to humans. Every pilot should publish exclusion, reversal, accessibility and privacy results and include a stop mechanism.

Priority 6 - Plural access-to-justice finance and later reassessment

Time horizon: 3-5 years. Test public-interest, mutual, philanthropic or levy-supported finance for meritorious matters that are socially significant but commercially unattractive. Reassess commercial funding architecture only after evidence shows whether technology has reduced routine underwriting or administration costs and whether competition has transmitted those savings.

> START HERE
>
> Priorities 1 and 2 are the immediate programme. They demand better information before demanding a new institution. Priorities 3 and 4 create the minimum governance layer for machine-assisted consequence. Priorities 5 and 6 are conditional, longer-horizon reforms that should proceed only if the evidence and pilots justify them.

8. Justice Dividend Scorecard

A single claimant percentage is too thin, and a large composite index would conceal trade-offs. The proposed scorecard therefore uses six outcome families. Each should be reported separately; no single aggregate “justice score” is proposed.

Outcome family	Example measures	Failure question
1. Claimant recovery	Net share of gross recovery; distribution spread; deductions by category; unclaimed / residual funds	Did a higher percentage conceal low participation or long delay?
2. Time to remedy	Time from harm / claim / settlement to payment; time saved by automation	Did lower cost create faster practical remedy?

Outcome family	Example measures	Failure question
3. Accuracy & correction	Reversal rate; recalculation rate; human override; source-data corrections; complaint outcomes	Can errors be found and repaired, or merely processed faster?
4. Participation & exclusion	Registration / take-up; missed deadlines; accessibility support; digitally excluded pathways; subgroup disparities	Who did not reach the remedy at all?
5. Cost & residual risk	Legal, administration and funding cost; capital deployed; adverse-cost exposure; expert / evidence burden	Which costs are genuine risk and complexity rather than information-processing friction?
6. Productivity distribution	Observed automated functions; time/cost change; claimant, respondent, professional, funder and public-sector incidence	Who captured the efficiency gain?

8.1 Auditability

Where machine systems supply scorecard data, the definitions, source lineage, exclusions and calculation methods should be independently auditable. A metric that cannot identify its source data and rule should not be treated as a performance fact. Unknown values should remain unknown rather than being imputed for the sake of a complete dashboard.

8.2 Distribution without moral arithmetic

A higher claimant percentage is not automatically better if it makes meritorious high-risk cases impossible to finance. A lower percentage is not automatically fair merely because a funder assumed risk. Faster resolution is not automatically better if accuracy or review deteriorates. The scorecard exists to force these trade-offs into view rather than collapse them into one emotionally convenient number.

9. Objections and counter-reasons

Objection 1 - Commercial funding may be the only realistic route for many group claims

Strong. Removing or over-capping return can make meritorious cases unfinanceable. The paper therefore does not recommend abolishing commercial funding or imposing a universal percentage floor. The immediate proposal is transparency plus evidence.

Objection 2 - Competition may pass AI savings through without regulation

Plausible and important. If funders and law firms compete effectively, lower underwriting and administration costs can improve terms. That is why the paper now treats current market structure as an empirical question. Priority 2 is designed to find out before prescribing a capture rule.

Objection 3 - Cheaper claim discovery can industrialise weak or strategic claims

Real. Lower search and drafting costs remove friction from both strong and weak cases. Claim discovery must remain distinct from legal merit; accountable lawyers and authorised decision-makers remain responsible for assertions, evidence and procedure. Respondents need clear provenance and challenge rights.

Objection 4 - Public automation can become more coercive than private finance

Robodebt makes this objection decisive. Public ownership does not guarantee fairness. A state system can combine data access, legal authority and scale in ways no private funder can. The answer is not “public AI”; it is bounded public infrastructure with legality, proportionality, source trace, human review, reasons, pause and remedy designed in from the beginning.

Objection 5 - Rights infrastructure can become surveillance infrastructure

Serious. Combining legal, financial, identity and behavioural data creates power as well as convenience. Data minimisation, purpose limitation, access logging, retention controls, privacy impact assessment, security, Indigenous Data Governance and prohibitions on unrelated secondary commercialisation are foundational, not optional.[17][18][19]

Objection 6 - Legal judgment and human care are not reducible to machine throughput

Correct. Advocacy, credibility, strategic judgment, ethical duty, negotiation, trauma-informed client care, discretion and legitimacy do not disappear because language models become competent. The paper’s target is unnecessary informational scarcity. Where friction protects rights, dignity or evidentiary quality, removing it is not efficiency.

Adversarial conclusion

After the final red-team pass, the strongest surviving thesis is narrower than the first draft and more useful. The problem is not that Qantas proves excessive legal extraction, and not that AI will inevitably turn courts into a casino. The problem is that legal remedy already sits inside systems of capital, professional labour, public authority and information scarcity. Machine intelligence changes the cost of information work. Society needs evidence and institutions capable of distinguishing productive savings from dangerous shortcuts, and competitive pass-through from private capture.

10. Staged roadmap and research agenda

Stage 0 - 0 to 12 months: establish the evidence

- Build a longitudinal Australian dataset of class-action settlements using consistent categories for gross recovery, legal costs, funding return, administration, interest, duration, participation and claimant distribution.
- Separate funded, unfunded and Victorian Group Costs Order matters rather than treating them as one funding architecture.
- Commission a current study of Australian litigation-funding market structure, competition, barriers to entry, co-funding and pricing dispersion.
- Map actual AI and litigation-analytics use by courts, legal-assistance providers, law firms, funders and settlement administrators, distinguishing verified practice from vendor claims.
- Publish missing-data fields and methodological limitations rather than filling gaps with assumptions.

Stage 1 - 0 to 18 months: standardise transparency

Develop the settlement transparency template and Justice Productivity Dividend statement. Trial them first as voluntary or practice-guidance instruments where formal rule change is unnecessary, then evaluate whether stronger requirements are justified.

Stage 2 - 1 to 2 years: govern machine-assisted consequence

Create model rules for purpose-bound claim origination, provenance, privacy, human review and correction. Align them with existing court AI practice directions rather than inventing a disconnected AI-law regime.

Stage 3 - 2 to 4 years: narrow public pilots

Test source-grounded assistance in domains with clear statutory or scheme criteria. Publish false-exclusion, reversal, delay, accessibility and user-experience results. No pilot should expand merely because throughput improved.

Stage 4 - 3 to 5 years: reassess finance after the costs actually change

Only once measured cost and competition evidence exists should policymakers decide whether changes to funding, fee or public-finance architecture are warranted. The question is not whether funders made money; it is whether the observed price of capital, residual risk and professional work remains proportionate after technological productivity changes.

Research questions that must be answered before stronger intervention

1. What are the post-2020 claimant distributions, total deductions and durations across funded, unfunded and GCO class actions?
2. How concentrated and contestable is the Australian litigation-funding market in 2026, and what prevents entry or switching?
3. Which underwriting, legal and administration functions are actually being automated, and what measurable cost has changed?
4. Do lower internal costs pass through into funding terms or legal fees in competitive matters?
5. Which entitlement domains have sufficiently authoritative source data and deterministic law to justify public digital pilots?
6. What error, reversal and exclusion thresholds should stop a consequential justice-AI pilot?
7. How should public court data be structured for legitimate research and access without becoming an unrestricted commercial surveillance feed?
8. Which independent institution is capable of auditing high-consequence justice systems across legal accuracy, privacy, accessibility and distributional impact?

11. Consultation and conclusion

Eight consultation questions

1. Should Australian class-action settlement materials use a standard public disclosure template, and which calculation formulas and actors must it identify?
2. What evidence is needed to determine whether competition already transmits lower legal and underwriting costs to group members?
3. When significant automation is used, which productivity gains should be disclosed to courts and group members, and how should uncertainty in those estimates be expressed?
4. What limits should apply to machine-assisted claim scouting using public court data, consumer data or public-sector entitlement data?
5. Which machine-assisted legal consequences should always carry a human-review right, a reason and an auditable source trail?
6. What domains, if any, are suitable for narrow public or court-adjacent entitlement pilots after the lessons of Robodebt?
7. How should claimant access and respondent procedural fairness be protected simultaneously as claim discovery and document production become cheaper?
8. Which outcome measures should determine whether justice technology is actually reducing the distance between having a right and exercising it?

Consultation mechanics

This is an independent discussion draft, not a statutory consultation. There is no formal closing date. The intended publication home is Weidu Fanling / 未读凡灵 (weidufanling.au). Any public release should retain the version number, date, provenance note, source list and revision history so that corrections can be distinguished from silent edits.

Conclusion: the goal is proportionate friction

The paper began with discomfort at a $105 million proposed settlement becoming an estimated $68 million pool for Group Members. The benchmark correction matters: Qantas's approximate 65% is above the historical 51% median Treasury reported for funded matters, subject to major comparability limits.[1][2] The case therefore cannot responsibly carry the accusation that the system is exceptionally extractive.

But the underlying question survives. Collective remedy has a cost. Some of that cost is unavoidable legal work, evidence, professional duty, capital risk and administration. Some of it may be information-processing scarcity that machine systems can reduce. If that scarcity falls, policy should be able to see whether the benefit reaches claimants, respondents, courts, legal professionals, funders, the public purse, or several at once.

Robodebt supplies the essential boundary condition. Friction is not always waste. Human review, source verification, legal authority, reasons, contestability and pause can be the machinery that prevents a cheap administrative process from becoming an unlawful or cruel one. The objective is therefore neither maximum intermediation nor maximum automation.

> CENTRAL PROPOSITION
>
> The public-interest goal is to remove informational tollbooths while preserving the friction that protects legality, evidence, dignity and correction. When machine intelligence makes a legal function cheaper, the justice system should be able to show what cost disappeared, what risk remained, who received the benefit and what new risk was created.

That is the claims market question after the reality check: not whether courts should become investment markets, and not whether AI should replace lawyers, but whether society can ensure that cheaper intelligence improves practical access to rights without making either private capital or public automation an unaccountable gatekeeper.

Appendix A. Qantas evidence table

Primary-source fact	What the paper takes from it	What the paper does NOT infer
Proposed settlement: $105m, subject to Court approval; no admission of liability.[1]	Gross reference point for distribution analysis.	That Qantas admits wrongdoing.
Approx. $68m / 65% estimated for Group Members after deductions and interest effect.[1]	Visible claimant-pool estimate.	That 65% is inherently unfair or unusually low.
Historical funded median reported by Treasury: 51%; unfunded: 85%.[2][3]	Context: 65% is above the historical funded median on its face.	That Qantas is directly comparable to the historical dataset.
Solicitors named: Echo Law and Piper Alderman.[1]	Legal operators should be identified in the worked example.	That either firm's charges are improper.
Funders named: Omni Bridgeway and CASL.[1]	Capital providers should be identified in the worked example.	That the funding return is excessive without risk evidence and Court assessment.
Funding commission: $25m-$26.25m; 25% adjusted based on claims registered.[1]	Need for formula and adjustment transparency.	A calculation formula not actually disclosed in the notice.
Register by 2 Oct 2026; approval hearing scheduled 13 Oct 2026.[1]	Participation chronology should be visually explicit.	That the sequence itself is unlawful.
Separate ACCC proceeding resulted in $100m penalties for misleading conduct concerning cancelled flights.[29]	There is distinct regulatory enforcement concerning Qantas cancellation conduct.	That the ACCC proceeding and the Flight Credits Class Action are the same claim or remedy.

Appendix B. Robodebt safeguard map

Robodebt is included as an institutional precedent for automated legal consequence, not as an example of generative AI. The mapping below translates documented failure modes into prospective design tests rather than pretending the historical scheme used today's technology.

Test for any future system	Minimum evidence before deployment	Stop / escalation condition
Source truth	Authoritative data sources, known limitations, provenance and reconciliation process.	Proxy or inferred data becomes decisive without explicit lawful authority.
Legal authority	Rule-to-statute / instrument trace; version control; qualified legal sign-off.	Material uncertainty about legal basis or unresolved adverse authority.
Burden allocation	Documented account of who must prove what and why.	Automation shifts evidentiary burden to the affected person for administrative convenience.
Human review	Named reviewer role, access channel, review standard, capacity to consider individual circumstances.	Reviewer cannot depart from or reconstruct machine output.
Reasons & notice	Plain-language explanation, accessible formats, deadlines and consequences.	Person cannot understand the basis or practical path to challenge.
Pause & hardship	Mechanism to suspend adverse action while review / dispute is active where law permits.	Recovery or coercive action continues despite unresolved material contest.
Cohort audit & remedy	Logging, impact analysis, reversal / recalculation capability, notification plan.	System cannot identify who was affected by a discovered rule or data defect.

Appendix C. Australian timeline

Year	Instrument / case	Relevance to this paper
2018-19	ALRC Report 134; Brewster [2019] HCA 45	Funding economics and limits of early common fund orders.
2020	Victorian Group Costs Orders; Commonwealth MIS / AFSL settings	Competing models for financing and regulating representative proceedings.
2021	Treasury minimum-return consultation	Explicit national debate about claimant share and proportionality.
2022	Litigation funding exemptions	Regulatory reversal / recalibration.
2023	Robodebt Royal Commission report	National postmortem on automated administration, legality, transparency, review and human consequence.
2025	Bogan [2025] HCA 7; Kain [2025] HCA 28; ASIC relief extension	Forum consequences of GCOs; settlement-stage funding powers; current exemption horizon.
2026	Federal Court GPN-AI; Supreme Court Victoria SC GEN 25 + Judicial Guidelines; Ba v Sterling Parts	Current judicial guardrails and a concrete generative-AI evidence failure.
2026	National Legal Assistance Data Strategy 2026-2030	Current public-sector foundation for responsible justice-data improvement.
2026	Qantas proposed class-action settlement; new Robodebt settlement approval	Two contemporary remedies expose different cost, governance and consequence questions.

Sources and notes

[1] Federal Court of Australia, Flight Credits Class Action Against Qantas - Notice of Proposed Settlement, Registration and Opt Out, Proceeding VID650/2023, filed 7 July 2026. User-supplied sealed notice used as the primary source for settlement figures, named legal representatives and funders, deadlines and proposed deductions.

[2] Australian Government, Treasury and Attorney-General's Department, Guaranteeing a minimum return of class action proceeds to class members, Consultation Paper, June 2021. https://treasury.gov.au/sites/default/files/2021-06/176658_consultation_paper.pdf

[3] Australian Law Reform Commission, Integrity, Fairness and Efficiency - An Inquiry into Class Action Proceedings and Third-Party Litigation Funders, ALRC Report 134, 2018. https://www.alrc.gov.au/publication/integrity-fairness-and-efficiency-an-inquiry-into-class-action-proceedings-and-third-party-litigation-funders-alrc-report-134/

[4] High Court of Australia, BMW Australia Ltd v Brewster; Westpac Banking Corporation v Lenthall [2019] HCA 45. https://www.hcourt.gov.au/cases-and-judgments/judgments/judgments-1998-current/bmw-australia-ltd-v-brewster

[5] High Court of Australia, Bogan v The Estate of Peter John Smedley (Deceased) [2025] HCA 7. https://www.hcourt.gov.au/cases-and-judgments/judgments/judgments-1998-current/bogan-v-estate-peter-john-smedley-deceased

[6] High Court of Australia, Kain v R&B Investments Pty Ltd; Ernst & Young v R&B Investments Pty Ltd; Shand v R&B Investments Pty Ltd [2025] HCA 28. https://www.hcourt.gov.au/cases-and-judgments/judgments/judgments-1998-current/kain-v-rb-investments-pty-ltd

[7] Supreme Court of Victoria, Practice Note SC GEN 10, Conduct of Group Proceedings (Class Actions). https://www.supremecourt.vic.gov.au/areas/legal-resources/practice-notes/sc-gen-10-conduct-of-group-proceedings-class-actions-second

[8] Australian Securities and Investments Commission, ASIC amends relief for litigation funding arrangements, 19 December 2022. https://asic.gov.au/about-asic/news-centre/find-a-media-release/2022-releases/22-368mr-asic-amends-relief-for-litigation-funding-arrangements/

[9] Australian Securities and Investments Commission, ASIC extends relief for litigation funding arrangements and conditional costs schemes, 17 December 2025. https://asic.gov.au/about-asic/news-centre/news-items/asic-extends-relief-for-litigation-funding-arrangements-and-conditional-costs-schemes/

[10] Federal Court of Australia, Class Actions Practice Note (GPN-CA). https://www.fedcourt.gov.au/law-and-practice/practice-documents/practice-notes/gpn-ca

[11] Federal Court of Australia, Use of Generative Artificial Intelligence Practice Note (GPN-AI), 16 April 2026. https://www.fedcourt.gov.au/law-and-practice/practice-documents/practice-notes/gpn-ai

[12] Supreme Court of Victoria, SC GEN 25 - The use of Artificial Intelligence by Court users, commenced 14 May 2026. https://www.supremecourt.vic.gov.au/areas/legal-resources/practice-notes/sc-gen-25-the-use-of-artificial-intelligence-by-court-users

[13] Supreme Court of Victoria, Guidelines: The use of Artificial Intelligence by Judicial Officers, 13 May 2026. https://www.supremecourt.vic.gov.au/forms-fees-and-services/forms-templates-and-guidelines/guidelines-the-use-of-artificial-intelligence-by-judicial-officers

[14] Federal Circuit and Family Court of Australia (Division 2), Ba v Sterling Parts Australia Pty Ltd [2026] FedCFamC2G 1245, 1 July 2026; judgment accessed via AustLII mirror. https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/FedCFamC2G/2026/1245.html

[15] Royal Commission into the Robodebt Scheme, Report, 7 July 2023 (updated 11 July 2023). https://robodebt.royalcommission.gov.au/publications/report

[16] Services Australia, New robodebt class action settlement, updated 26 June 2026. https://www.servicesaustralia.gov.au/new-robodebt-class-action-settlement

[17] Australian Attorney-General's Department, National Legal Assistance Data Strategy 2026-2030, 3 August 2026. https://www.ag.gov.au/legal-system/publications/national-legal-assistance-data-strategy-2026-2030

[18] Privacy Act 1988 (Cth), current compilation. https://www.legislation.gov.au/C2004A03712/latest

[19] Office of the Australian Information Commissioner, Privacy by design and AI/privacy guidance. https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies

[20] UK Supreme Court, R (PACCAR Inc and others) v Competition Appeal Tribunal and others [2023] UKSC 28. https://www.supremecourt.uk/cases/uksc-2021-0078

[21] UK Government, Increased access to justice for claimants to take on powerful organisations in court, 19 December 2025. https://www.gov.uk/government/news/increased-access-to-justice-for-claimants-to-take-on-powerful-organisations-in-court

[22] British Columbia Civil Resolution Tribunal, official site and process description. https://civilresolutionbc.ca/

[23] British Columbia, Civil Resolution Tribunal Act, SBC 2012 c 25. https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/12025_01

[24] European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act). https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[25] Council of Europe, CEPEJ European Ethical Charter on the use of artificial intelligence in judicial systems and their environment. https://www.coe.int/en/web/cepej/cepej-european-ethical-charter-on-the-use-of-artificial-intelligence-in-judicial-systems-and-their-environment

[26] OECD, Toolkit for Access to Justice and People-Centred Justice Systems, 12 November 2025. https://www.oecd.org/en/publications/toolkit-for-access-to-justice-and-people-centred-justice-systems_aecf7f78-en.html

[27] Burford Capital, corporate and investor materials. Used only as self-description evidence of a publicly listed legal-finance / portfolio model, not as independent performance evidence. https://www.burfordcapital.com/

[28] LexShares, historical company / press materials. Used only as historical evidence of claim-sourcing and marketplace form, not as independent performance evidence. https://www.lexshares.com/

[29] Australian Competition and Consumer Commission, Federal Court orders Qantas to pay $100m in penalties for misleading consumers, 8 October 2024. https://www.accc.gov.au/media-release/federal-court-orders-qantas-to-pay-100m-in-penalties-for-misleading-consumers

Source hierarchy and limitations

Primary and official sources are preferred for legal status, regulation, institutional practice and settlement facts. Commercial-provider material [27]-[28] is used only to establish that particular market forms or claim-sourcing propositions have existed; it is not treated as independent evidence of return, efficacy or prevalence. Robodebt is described from the Royal Commission and official settlement information rather than contemporary commentary. The paper deliberately leaves current litigation-funding market concentration as an evidence gap because a sufficiently current, comparable primary dataset was not established for this revision.

Revision note: v0.2 to v0.3

Version 0.3 is the final adversarial rebuild of this discussion draft. It adds Robodebt as the machine-consequence case study; states explicitly that Qantas's estimated 65% claimant pool is above the historical 51% funded median cited by Treasury; names Echo Law, Piper Alderman, Omni Bridgeway and CASL from the filed notice; adds Bogan as a jurisdictional qualification to Victoria's Group Costs Order model; expands current Australian court-AI material with the Supreme Court of Victoria and Ba v Sterling Parts; treats competition and market structure as unresolved variables rather than assuming governance determines the productivity dividend; ranks the policy programme; reduces the metric and consultation sprawl; and adds an explicit provenance and AI-use disclosure.