consiliumbots

Using Data & Technology to Help People Make Important Decisions

Working
Paper Series

Making Algorithmic Transparency Meaningful: Evidence from Chile's National School Admissions Platform

Christopher A. Neilson
Exequiel Medina
Leonardo Ortiz Mesías

N.°5 August 2026

About the ConsiliumBots Working Paper Series

The ConsiliumBots Working Paper Series disseminates high-quality research at the intersection of education, economics, and policy. Our mission is to spark dialogue among scholars, policymakers, and practitioners on critical topics such as school choice, digital tools for learning, and educational inequality. Each working paper undergoes independent peer review and is published with the aim of advancing both theoretical understanding and practical solutions. We welcome contributions that challenge conventional thinking, employ rigorous empirical methods, and offer actionable insights for real-world educational systems.

For the complete archive of titles, abstracts, and submission guidelines, please visit:
consiliumbots.github.io/working-papers-cb/

Making Algorithmic Transparency Meaningful: Evidence from Chile's National School Admissions Platform

Christopher A. Neilson[a,*], Exequiel Medina[b], Leonardo Ortiz Mesías[c]

[a] Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, United States
[b] Johns Hopkins School of Advanced International Studies (SAIS), Washington, DC, United States
[c] Facultad de Derecho, Universidad Diego Portales, Santiago, Chile

[*] Corresponding author. E-mail: christopher.neilson@yale.edu

Abstract

Governments worldwide increasingly deploy algorithmic systems to allocate scarce public goods. The dominant transparency paradigm focuses on expert audit rather than citizen comprehension. We propose a substantive standard, Meaningful Algorithmic Transparency (MAT): information delivered to affected users must be proactive (placed into awareness, not merely published), tailored (referring to the user's decision context), and actionable (delivered before the decision closes). We develop the argument through Chile's Sistema de Admisión Escolar (SAE), a national school admissions platform serving roughly 470,000 families per year. Two measurements document the residual comprehension gap. Linking the SAE Satisfaction Survey to administrative truth for four cycles (2020–2023, N≈140,000), we show that families systematically overstate their placement probability: among applicants the platform can pre-identify as high-risk, the median family believes the risk is around 20% when it is in fact around 80%. The bias is uneven: at the same objective risk, lower-SES applicants understate by more, with a conditional gradient that is statistically stable across all four cycles. In the SAE Parent Surveys (2023–25), self-reported familiarity with a key procedural rule nearly doubled across cycles while verifiable comprehension barely moved (2.3% to 4.9%). The comprehension gap is therefore persistent and unevenly distributed, falling more heavily on the families whose outcomes most depend on effective information delivery, even when the underlying mechanism gives those families priority. We propose six operational components and show how Chile's existing regulatory architecture (access-to-information law, algorithmic-transparency recommendations, data-protection safeguards, and sector-specific regulation) can be articulated to make the standard enforceable.

Keywords: algorithmic transparency; automated decision-making; meaningful transparency; citizen comprehension; algorithmic allocation; public sector algorithms; Chile.

1. Introduction

Governments at every level now rely on algorithmic systems to make or support decisions that directly affect citizens' rights and opportunities. Centralized platforms allocate school seats, scoring algorithms determine eligibility for social benefits, risk-assessment tools inform immigration decisions, and automated systems assign public housing. The scale of deployment has grown sharply over the last decade: in education alone, Neilson (2024) documents that roughly 60 percent of 149 countries surveyed now operate some form of coordinated choice and assignment system, with substantial heterogeneity in mechanism design, priority structures, and information provision. The transparency response has lagged this expansion. Algorithmic transparency has become a first-order governance concern, generating a rapidly expanding system of policy instruments: algorithm registries, transparency recording standards, impact assessments, and public audit frameworks. The GPAI (2025) report, the most comprehensive mapping to date, catalogues 83 public-algorithm repositories across four continents and organizes the instruments in use along a supply-driven to demand-driven spectrum.

Yet the same report diagnoses a structural limitation of the current paradigm. Most existing instruments focus on what we call Transparency 1.0 (Formal Transparency): making the algorithm and its documentation available so that experts, auditors, journalists, and motivated citizens can scrutinize the system. The implicit theory of change is that publication produces accountability. That theory is powerful, and insufficient. As Nieuwenhuizen (2024) asks of algorithm registers, the question is whether they are "a box-ticking exercise or a meaningful tool for transparency." Gutiérrez and Muñoz-Cadena (2025), studying Colombia's public-algorithm repositories, diagnose "a practice of performative transparency in which information about automated decision-making systems is used to project efficiency and innovation" without enabling citizens to

understand when a system is being used on them or what it is doing with their data. Cath and Jansen (2022) warn of "the dangers of romanticizing the registry as a governance solution." Grimmelikhuijsen (2023), in an experimental study published in Public Administration Review, shows that the explainability of algorithmic decisions has more pronounced effects on perceived trustworthiness than mere accessibility of the algorithm, a result that points to the centrality of user comprehension over document availability.

This paper argues that closing the gap between published information and citizen comprehension requires a substantively different standard: what we call Meaningful Algorithmic Transparency (MAT). MAT is defined by three requirements: information delivered to affected users must be (i) proactive, placed into the user's awareness through concrete notification acts, not merely published on a website; (ii) tailored, adapted to the characteristics, informational needs, and vulnerabilities of different groups of affected users, rather than assuming a homogeneous user of the system; and (iii) actionable, delivered before the process closes and in a format that allows the user to modify their decision[1]. The compliance criterion shifts from the existence of documents to measurable evidence that users can correctly anticipate how their actions map onto system outcomes.

We develop the argument through a case that is unusually well suited to testing it. Chile's Sistema de Admisión Escolar (SAE), introduced by the 2015 Inclusion Law and extended nationally since 2018, is a national school admissions platform that processes roughly 470,000 applications per year. Its design incorporates the strongest recommendations of the mechanism-

[1] The approach adopted here does not challenge Formal Transparency solely from a public-management perspective. It also advances a legal critique of formalistic understandings that place transparency obligations mainly at the stage of final decisions or legally enforceable administrative acts. In algorithmic systems, however, legally relevant effects may arise earlier through system design choices, interfaces, and informational environments that shape how individuals interact with public decision-making processes. These interactions can be understood as part of the Administration's material acts (Izquierdo, 2026). (Izquierdo, 2026).

design literature: a deferred-acceptance algorithm that eliminates strategic incentives for misreporting preferences, no cap on the length of preference lists, and full publication of the algorithm, the rules, the data, and a simulator for interested users (Abdulkadiroğlu & Sönmez, 2003; Haeringer & Klijn, 2009; Correa et al., 2022; Neilson, 2026). By any comparative metric, Chile sits, in this case, at the frontier of what publication-centered algorithmic transparency can achieve. If the comprehension gap persists here, it is unlikely to be smaller in systems that invest less in transparency.[2]

1.1 The comprehension gap

The accumulated evidence shows that it does persist. Arteaga, Kapor, et al. (2022), in a randomized experiment with Chilean parents, document that a significant share of families hold systematically biased beliefs about admission probabilities: they underestimate the risk of going unassigned, and that underestimation translates into preference lists that are too short. When the intervention provides real-time, tailored feedback on assignment risk, families add alternatives, risk falls, and the academic quality of the schools in which they ultimately enroll rises. Agte et al. (2024) show that these effects operate specifically through correction of false beliefs, not through general salience or exhortation, a finding that points directly to the informational architecture rather than to the quality of communication.

This paper contributes two further, independent measurements of the gap, at population scale. The first links the SAE Satisfaction Survey to administrative truth across four cycles (2020–2023, roughly 140,000 respondents): re-running the published algorithm on each family's submitted

[2]The argument does not imply that the SAE's design is exhausted in every dimension. Decisions regarding priority structures, tie-breaking rules, and sibling treatment admit further refinements. The claim in the text is narrower: on the specific dimension of making the system's operation transparent and minimizing the informational burden on families, Chile has implemented the standard recommendations of the mechanism-design and algorithmic-governance literatures.

application recovers their objective placement risk, which we compare against the probability they state. Families systematically overstate their chances: among applicants the platform can flag as high-risk, the median family believes the risk of going unassigned is around 20% when it is in fact around 80%. The bias is larger for lower-SES families at the same objective risk (Section4.2). The second measurement, using the 2023–2024 and 2024–2025 SAE Parent Surveys, shows that the share of parents who declare knowing how the family-application rule (Article 24, Decree 152/2016) works nearly doubled between cycles (from 36.8% to 60.4% in the group with the highest educational capital), while the share who correctly identify its operation barely rose, from 2.3% to 4.9%. Among those who declare knowing, conditional accuracy is only 6 to 8 percent. In plain terms: more citizens believe they understand the system without understanding it. This divergence, declared comprehension rising while verified comprehension flatlines, is the operational signature of what publication-centered transparency produces when it reaches its ceiling.

A complementary policy report by ConsiliumBots (2025), prepared while implementing the SAE algorithm for three consecutive cycles, documents an analogous pattern for joint family applications and describes the design of an informational intervention aimed at closing that gap.

The convergence of these results (randomized experiments on probabilistic beliefs, a population-scale calibration of those beliefs against administrative truth, operational reports from the system's implementer, and a national survey on rule comprehension) across distinct dimensions and moments of the SAE's lifecycle suggests that the gap is structural, not attributable to a poorly designed system or insufficiently published documents. It is the residual that remains once the dominant paradigm has done everything it can. That residual, we argue, cannot be closed with more publication: it requires a different institutional and legal approach.

1.2 The uneven distribution of the comprehension gap

A second feature of the comprehension gap, which we develop empirically in Section 4, is that it is not evenly distributed across the user population. At the same objective non-assignment risk, lower-SES applicants in the SAE understate that risk by more, and the conditional gradient is statistically stable across four annual cycles. The basic intuition is straightforward and does not require elaborate theoretical machinery: publication-centered transparency requires users to convert public documents, simulator interfaces, and FAQ pages into operative knowledge about their own case, and the capacity to perform that conversion is unequally distributed across families: by education, time available, prior familiarity with bureaucratic registers, and digital fluency. Related observations have been made in adjacent literatures, most prominently the administrative-burden framework on learning costs in public-service delivery (Herd & Moynihan, 2018), the second-order digital-divide literature on differential effectiveness of access (Hargittai, 2002; van Deursen & van Dijk, 2014), and the digital-government-inclusion framework recently developed in this journal (Peeters, Miller & Schuilenburg, 2025), which distinguishes practical access, procedural fairness, quality of participation, and equality of outcomes as four dimensions of inclusive government automation. The implication for the SAE, which the empirical analysis below quantifies, is that the dominant transparency paradigm is incomplete, and incomplete more for some families than for others. The comprehension shortfall falls more heavily on families whose placement outcomes depend most on the state delivering information effectively.

1.3 From the dominant paradigm to meaningful transparency

The core argument of the paper can be stated as a move from one generation of the algorithmic-transparency standard to the next, a generation designed to address both the residual comprehension gap and its unequal distribution across users. The GPAI (2025) report organizes

instruments in use along two dimensions: supply-driven versus demand-driven, and, within supply-driven instruments, four distinct tools (public algorithm registries, user-demand-driven proactive publications, automated responses triggered by interaction with a specific process, and publication of audits or impact assessments). Chile has advanced on the first and last of those tools (what we call Transparency 1.0), but has not yet consolidated the cell the GPAI report only sketches: proactive, tailored, automated responses delivered at the moment of decision, the only one of the four whose success criterion depends on what happens inside the head of the affected citizen. That cell requires a substantively different standard, which we call Meaningful Algorithmic Transparency (MAT), or Transparency 2.0. The GPAI report itself uses Chile's school admission system as an example of an automated decision-making process that directly affects families, lending external validation to the choice of the SAE as a test case for this argument.

We call Transparency 1.0 or Formal Transparency the paradigm that Chile (and, to varying degrees, most legal systems with algorithmic-transparency ambitions) has implemented: publish the algorithm and the data, provide generic audit tools, and make the documentation accessible so that any interested party can scrutinize the system. The function of Transparency 1.0, as the growing literature on algorithm registries documents (Nieuwenhuizen, 2024; Giest & Grimmelikhuijsen, 2020), is to enable supervision and external scrutiny of a deployed system. Its implicit addressee is the expert, the auditor, the researcher, the citizen with time and incentives to become informed. Its criterion of compliance is the existence and accessibility of the artifacts. This paradigm is necessarily complemented by data-protection law, which addresses a related but distinct question: what rights does an individual hold when their personal data are processed, profiled, or used in an automated decision? Together, the two regimes establish the system-level conditions of lawful algorithmic governance: transparency enables public scrutiny of the system,

while data protection safeguards the rights and agency of the data subject. Neither regime, however, necessarily requires the operator to transform the case-specific information that the system can generate throughout the process into a timely communication directed to the affected user. Neither requires that such information be delivered before the user's choice becomes final, in a form that enables operative understanding and meaningful action.

We call MAT (Transparency 2.0) the standard this paper proposes to consolidate as the successor to the previous one. Its function is to produce operative understanding in the affected user, at the moment and in the context in which that user must make a decision with foreseeable consequences. Its addressee is the applying family. Its criterion of compliance is measurable evidence that the user can correctly anticipate how their actions map onto system outcomes. Its substantive content is the proactive, tailored, and actionable delivery of relevant information before the process closes.

Table 1 summarizes the contrast.

Table 1. Formal Transparency ("Transparency 1.0") compared with Meaningful Algorithmic Transparency ("Transparency 2.0")

	Transparency 1.0 / Formal Transparency	**Transparency 2.0 / MAT**
Question answered	Can someone audit the algorithm?	Can this family anticipate what the algorithm will do with their application?
Normative function	External supervision and scrutiny	Decisional capacity of the affected user
Implicit addressee	Expert, auditor, researcher, general public	User at the moment of decision

	Transparency 1.0 / Formal Transparency	**Transparency 2.0 / MAT**
Typical content	Code, data, rules, administrative documents, generic simulator	Tailored risk estimate, actionable explanation, timely warnings
Moment of delivery	Permanent, on demand	Before the process closes, at the point of decision
Compliance criterion	Existence and accessibility of the artifacts	Operative comprehension verifiable in the user population
Chilean legal anchor	Law 20,285; CPLT Res. 372/2024; complemented by Law 21,719 where applicable	Not yet consolidated; proposed here through coordinated regulatory development
Status in SAE	Implemented with high ambition	Isolated pilot interventions; no regulatory frame

Note. Rows compare the two paradigms on seven dimensions. Conceptual table; no data are reported. Chilean legal anchors cited in the final rows are Law 20,285 on access to public information, CPLT Resolution 372/2024, and Law 21,719 on personal data protection, which enters into force in December 2026.

Two clarifications. First, MAT does not replace Formal Transparency. It presupposes it. A state that does not publish its algorithm cannot invoke MAT to be exempted from doing so. The relationship is cumulative complementarity, not substitution. Second, MAT is not satisfied by doing "more" of Formal Transparency: its content is qualitatively different. Publishing more documents, translating them into plain language, or improving the simulator's interface are improvements within the formal paradigm; none of them satisfies the central requirement of MAT, which is proactive, tailored delivery of relevant information about the user's concrete case at the moment of deciding.

Read in this light, the empirical evidence this paper presents plays a precise role: it is the demonstration that Chile has taken Formal Transparency close to its frontier (strategy-proof

mechanism, full publication, simulator) and that the comprehension gap nevertheless persists. The Chilean effort has reached the limit of what publication alone can achieve. That persistence is the empirical signal that the formal paradigm is exhausting what it can contribute, and that the next generation of the standard must be built on a different logic. Figure 1 synthesizes the three-layer architecture visually.

Figure 1. Three-layer architecture: from market design to meaningful algorithmic transparency

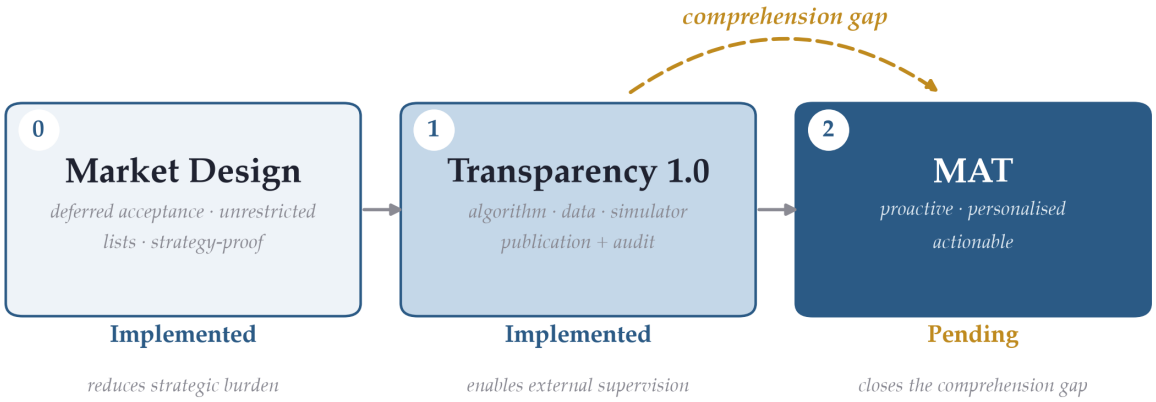

Note. Schematic; no data are plotted. Each layer is a necessary but not sufficient condition for the next. The dashed arc denotes the residual comprehension gap documented in Arteaga, Kapor, et al. (2022), Agte et al. (2024), ConsiliumBots (2025), and in Section 4 of this paper.

The rest of the paper develops this argument in four substantive steps before concluding. Section 2 presents the institutional context of the SAE as a case study of algorithmic allocation in the public sector. Section 3 formalizes the conceptual framework of MAT, distinguishes it from prior approaches, including the recent literature's diagnosis of performative transparency, and articulates the relationship between MAT and the dominant paradigm. Section 4 presents two measurements of the comprehension gap: a population-scale calibration of families' beliefs against administrative truth, and a survey of comprehension of a specific procedural rule. Section 5

proposes an operational standard and a legal pathway to make MAT enforceable, using Chile as a worked example with implications for other legal systems. Section 6 concludes with the argument's generalizability beyond Chile.

2. Institutional context of the SAE

Chile enrolled roughly 3.5 million primary and secondary students in 2025, of whom 89 percent attend schools receiving public funding, a mix of fully public municipal schools and privately operated schools that receive a per-student state subsidy. The remaining 11 percent attend private schools without state funding, which set their own admission criteria. The SAE operates on the publicly funded 89 percent, making it one of the largest centralized school admissions platforms in the developing world. In 2025 approximately 473,000 students applied through the SAE, distributed heterogeneously across the country, with a national average of 2.35 available seats per applicant (CEM 2025). In Santiago and other large urban centers, competition for places at the most sought-after schools is considerably more intense than the national average suggests.

From the family's perspective, the application process works as follows. During a window of roughly three weeks (typically in August for the following academic year), parents log into an online platform managed by the Ministry of Education. There they build a rank-ordered list of schools, placing their most preferred school first and continuing in descending order. There is no cap on the number of schools a family may list, and the Ministry explicitly recommends including at least six alternatives. Families may also activate special options: the family-application rule (Article 24, Decree 152/2016), which allows siblings to be considered jointly under a shared priority, and the joint-application option, which links the applications of children from different families who wish to attend the same school. Once the application window closes, the central

algorithm, a variant of deferred acceptance with random tie-breaking and legally defined priorities (siblings already enrolled, children of school staff, students classified as socioeconomically vulnerable), processes all applications simultaneously and produces a single assignment. Families then receive their results and have a brief period to accept, waitlist, or re-enter a complementary round.

The SAE operates in three stages. The principal round, described above, accounts for the majority of placements. A complementary round follows for students who were unassigned or who declined their initial placement. Finally, a regularization stage runs on a first-come-first-served basis through a platform known as "Anótate en la Lista" (Sign Up for the List), which assigns remaining vacancies according to order of inscription rather than by algorithm. This third stage introduces non-algorithmic criteria that complicate the user experience and can generate perceptions of arbitrariness: families who arrive late to the regularization platform may find no vacancies at nearby schools, regardless of their preferences or priority status.

The governing regulations contained in the General Education Law (Law 20,370) and its implementing regulation, Decree 152 of 2016 of the Ministry of Education, establish a fundamental principle: the past or potential academic merit of a student may not be considered in admissions to publicly funded schools. In place of merit-based selection, the regulation opted for a process designed to be transparent and equitable, safeguarding equal opportunity while leaving wide scope for the expression of family preferences. This normative choice is important context for the transparency argument: the SAE was designed as an institutional commitment to non-discrimination, which makes the question of whether families can effectively navigate it a matter of substantive equity as well as procedural efficiency.

Several features of the SAE's design are worth highlighting for readers unfamiliar with the system, because they bear directly on the transparency argument. First, the information environment during the application period is richer than in most comparable systems. The Ministry publishes a "school guide" (Guía de Establecimientos) with information on each school's location, academic performance, special programs, and available seats. A generic simulator allows users to explore how the algorithm works with hypothetical preference lists. Informational videos, help desks, and a dedicated website complement the documentation. By the standards of developing-country education systems, this is an unusually well-resourced informational effort, which makes the persistence of the comprehension gap all the more significant. If the gap persists under these conditions, it is unlikely to be smaller in systems that invest less in communication.

Second, the SAE's transition from the pre-2015 regime (in which schools selected their own students, often using interviews, entrance exams, or family background as criteria) to the current centralized system was politically contentious and remains so. The Inclusion Law's guarantee that no student would be selected on the basis of past or potential academic merit was a deliberate policy choice, and public debate about whether to reintroduce merit-based criteria resurfaces periodically. This political context matters because it means that misunderstandings about how the SAE works (for instance, the common belief that listing a school first increases one's chances mechanically, or that the algorithm "knows" which families are more deserving) can fuel political narratives about the system's fairness that are disconnected from its actual operation. Transparency that reaches the user is therefore both an informational good and a precondition for informed public debate.

Chile's system invites comparison with centralized school-choice mechanisms in other countries. New York City and Boston reformed their systems in the early 2000s following the

theoretical contributions of Abdulkadiroğlu and Sönmez (2003) and subsequent work (Abdulkadiroğlu et al., 2005, 2009; Pathak & Sönmez, 2008), but those are high-income urban contexts with very different institutional constraints. Among developing countries, Colombia's SISBEN-based targeting for educational subsidies, Brazil's SISU system for university placement, Kenya's centralized secondary-school placement, and India's state-level experiments with school-choice reforms all face analogous challenges of algorithmic allocation under limited information. What distinguishes the Chilean case is the combination of national scale, a decade of operational experience, and a sustained investment in Formal Transparency that provides an unusually rich evidentiary base for studying the limits of the dominant paradigm.

3. Conceptual framework: from algorithmic transparency to MAT

3.1 The SAE as an automated decision-making system

The SAE is more than a technical platform in the narrow sense. It is an automated decision-making system (ADMS) that structures a high-stakes environment in which families must identify alternatives, evaluate them, estimate their relative probabilities, and rank them under time pressure with incomplete information. As Correa et al. (2022) show in their formal description of the algorithm, the deferred-acceptance procedure with priorities and random tie-breaking that operates in Chile is computationally transparent: given a set of rankings and priorities, the outcome is deterministic and reproducible. But computational transparency, the property that the system can be technically audited and its outputs reproduced, does not automatically transfer to operative comprehension by the user. Knowing that the algorithm is deterministic is not the same as knowing what effect it will have on my particular assignment if I put this school in second rather than third place, or if I activate the family-application option. This distinction between system-level transparency and user-level comprehension is precisely the gap that the algorithmic governance

literature has begun to identify (Molina & Medina, 2025; Grimmelikhuijsen, 2023; Veale & Brass, 2019; Bullock, 2019) but has not yet operationalized into an enforceable standard. Rizk and Lindgren (2025) provide a useful conceptual vocabulary for this point: automated decision-making in public administration reconfigures the "decision space" between officials and citizens, and the operative question is what happens within that space. The MAT standard we develop below (Section 3.3) can be read as a substantive specification of what the state owes the citizen in that decision space.

The behavioral-science literature (Tversky & Kahneman, 1974; Kahneman, 2011) describes well why this translation is difficult. Citizens interacting with ADMS rely on heuristics (simple mental shortcuts) that produce systematic rather than random errors: prioritizing a better-known school over a higher-performing one, assuming that ranking an option first mechanically increases admission probability, underestimating non-assignment risk, trusting informal recommendations over official information. The SAE's mechanism design acknowledges this reality and responds by choosing an algorithm that minimizes what families need to understand in order to act optimally. What the evidence shows is that even under the mechanism that minimizes that burden, a substantive residual of operative comprehension remains that system design, on its own, cannot eliminate.

3.2 Anatomy of the current paradigm: two strata within Formal Transparency

Formal Transparency is not monolithic. Within the current paradigm two analytically separable strata coexist: a formal stratum, centered on the publicity of administrative acts and documents, and an expert stratum, centered on the ability to technically reproduce and audit the system's outputs. Both are necessary and both are compatible with MAT, which operates on a different dimension, the decisional capacity of the affected user. Table 2 synthesizes the three resulting

strata. The three strata are complementary: the formal is necessary for the expert, and both together are necessary for MAT. But neither individually nor jointly is either stratum sufficient for MAT. That insufficiency is an empirical fact: even when the two strata of formal are well implemented, the gap between declared and operative comprehension persists.

Table 2. Three strata of algorithmic transparency

	Formal (1.0)	**Expert (1.0)**	**MAT (2.0)**
Function	Legality control	Technical supervision	Decisional capacity of the user
Addressee	General public	Auditors and researchers	User affected by the system
Compliance	Existence of documents	Technical reproducibility	Verifiable operative comprehension
Conceptual referent	Publicity principle	Brauneis and Goodman (2018)	Definition in Section 3.3

Note. Conceptual table; no data are reported. The first two columns are the formal and expert strata internal to Formal Transparency; the third is MAT. The conceptual referent for the expert stratum is Brauneis and Goodman (2018); MAT is defined in Section 3.3.

Legal systems also vary widely in how completely they implement the two strata of Formal Transparency, and that variation matters for the comparative claim of this paper. Chile sits near the implementation frontier: the algorithm, the data, the rules, and a publicly available simulator are formally published in audit-grade form. Most other legal systems are at less complete points. Denmark's gymnasium admissions, reformed in 2022, allocate applicants among requested schools using transport time from the applicant's address to each school as a key criterion; after the assignment runs, the cutoff transport time at which each school stopped admitting students that cycle is made public. The publication is real and generates accountability of a kind, but it does not enable an external party to reproduce the assignment (the full priority structure and within-priority

tie-breaking are not released at audit grade), it does not enable a regulator to monitor the system's ongoing operation (the data arrive after the cycle concludes), and it cannot help a family decide where to apply (the cutoff for this year is unknowable until next year). Denmark's practice illustrates a partial position within Formal Transparency: public after-the-fact reporting that satisfies neither the formal-publicity test in its strong form nor the expert-reproducibility test. The relevant comparative claim is therefore stronger than it would be if Chile were taken as representative: Chile is what Formal Transparency looks like when it is implemented near its frontier, and even at that frontier the comprehension gap persists. Most legal systems are starting from positions closer to Denmark's partial implementation, which makes the distance to MAT, from a typical starting point, wider still.

The recent comparative literature on algorithmic governance has coined a precise term for this pathology: performative transparency. GPAI (2025), in its global mapping of 83 public-algorithm repositories, reports that most fail to reach what Gutiérrez and Castellanos-Sánchez (2023) call meaningful transparency, information that actually allows one to assess a system's performance. The report echoes the warning of Cath and Jansen (2022) about "the dangers of romanticizing the registry as a governance solution" and cites the diagnosis of Gutiérrez and Muñoz-Cadena (2025) on the Colombian case: the state "seems to favor a practice of performative transparency in which information about automated decision-making systems is used to project efficiency and innovation" without that information being sufficient for citizens to know when a system is being used on them or for interested actors to evaluate its performance. MAT is, in this light, the standard that makes the distinction between substantive and performative transparency institutionally enforceable: a verifiable property whose absence is measurable in the distance between declared familiarity and verifiable comprehension. The Chilean case documented in Section 4, where the

first nearly doubles while the second barely moves, is a direct quantitative measurement of that distance.

3.3 An operational definition of MAT

Definition (MAT). An automated decision-making system that affects rights or public interests satisfies the standard of *Meaningful Algorithmic Transparency* (MAT) when, at each point in the process where a user must make a decision with foreseeable consequences for the outcome, the system operator provides affected users, without requiring any request, with information that is simultaneously *(i) proactive*: delivered at the operator's initiative through concrete acts of notification (not merely made available through channels that users are expected to consult) and presented in a way that makes its relevance salient; *(ii) tailored:* adapted to the user's concrete decision context and, where relevant, to the informational needs and vulnerabilities of different groups of affected users, rather than assuming a homogeneous user of the system; and *(iii) actionable*: expressed in a format and with enough advance notice that the user can modify their choices before those choices produce significant practical or legal consequences.

Among the three requirements, proactivity carries the greatest conceptual weight. It is not enough for the operator merely to publish information or make it available; proactivity requires institutional action on three simultaneous planes. On the plane of initiative, delivery cannot depend on the user knowing what to look for: the asymmetry over what is relevant is precisely part of the gap to be closed, so delegating the first move to the user reproduces the gap. On the plane of salience, the information must arrive with a design that makes its importance for the current decision evident: treated as a critical notification, not as one option among many in a help menu. On the plane of record, the operator must be able to show that the information was actually placed

into the user's awareness, via a read receipt, a dedicated channel, or an analogous mechanism, not merely posted on a website "anyone could have consulted".

The practical difference between "information available" and "information placed into awareness" is precisely the empirical distance that the gap between declared familiarity and verifiable comprehension records: when a rule is published on the Ministry website but is not actively highlighted at the moment of application, 95% of parents, even in the group with the highest educational capital, fail to identify it correctly. This strong reading of proactivity is what distinguishes MAT from any maximalist version of Formal Transparency. A state can publish every line of code, every dataset, and every manual and still fail the MAT standard if it does not execute concrete acts that place the critical information in front of the user at the moment of decision. This failure is not neutral. It risks widening inequalities, especially for underserved groups, because those with fewer resources are less likely to bridge the gap between formal availability and operative comprehension on their own.

A concrete example may help clarify what MAT looks like in practice. Consider a family in Santiago that has listed five schools in its application. Under Formal Transparency, the family can visit the Ministry's website, download the algorithm's documentation, and consult a generic simulator that shows historical admission cutoffs. Under MAT, the system would additionally send the family a tailored notification, before the application deadline, informing them that, based on their specific list and the historical demand patterns for each school, their estimated probability of being assigned to at least one of their top three choices is, say, 62 percent, and that adding two more alternatives would raise that probability to 89 percent. The notification would arrive through a dedicated channel (not buried in a general information page), would include a clear explanation of what the family can do to improve its position, and would generate a record confirming that the

family received and opened the message. The risk-warning intervention tested by Arteaga, Kapor, et al. (2022) is a functional prototype of exactly this kind of communication: when families received tailored feedback on their assignment risk during the application process, a significant share of them added alternatives to their lists, and their eventual match quality improved.

This definition differs from the approach of Brauneis and Goodman (2018) on one important point. Their work conceives of algorithmic transparency primarily as a disclosure standard aimed at enabling expert supervision and external control. Our definition accepts that layer as necessary but introduces an additional, user-centered requirement. The reason is empirical: the evidence accumulated on the SAE shows that expert supervision and technical disclosure, even when well implemented, do not guarantee operative comprehension in the user population.

3.4 Converging evidence on the comprehension gap

Five lines of evidence support empirically the need to introduce MAT as a layer distinct from the previous ones. None is conclusive on its own; their force comes from convergence across distinct dimensions, methods, and sources. Table 3 summarizes them.

Table 3. Converging lines of evidence on the SAE comprehension gap

Source	Method	Central finding relevant to the argument
Arteaga, Kapor, et al. (2022)	RCT with Chilean parents	Families systematically underestimate their risk of non-assignment; a tailored proactive warning closes part of the gap and improves match quality.
Agte et al. (2024)	Causal-channel analysis	Behavioral effects operate specifically through correction of false beliefs, not through general salience or exhortation.
ConsiliumBots (2025)	Operator report (3 cycles)	Comprehension of the joint-application rule is lower than declared; the report describes an intervention designed to close that gap.

Source	Method	Central finding relevant to the argument
This paper, Section 4.2	Linked Satisfaction Survey + admin truth, N≈140k, 2020–23	At the same objective risk, families systematically understate non-assignment risk; median gap in the high-risk slice is 49–62 pp across all four cycles; monotone in SES across two proxies; conditional gradient statistically equivalent within ±2 pp across cycles.
This paper, Section 4.3	SAE Parent Surveys 2023–25, scenario-based verification	For the family-application rule, declared familiarity nearly doubles between cycles while verifiable comprehension barely moves.

Note. Rows report source, method, and the finding relevant to this paper. Rows 1 to 3 summarize prior work; rows 4 and 5 are the two measurements reported in Section 4. Row 4 uses the SAE Satisfaction Survey for cycles 2020 to 2023 linked to simulated objective risk (N = 140,218 paired observations). Row 5 uses the SAE Parent Surveys 2023–24 and 2024–25.

Three observations make this body of evidence especially robust. First, the Arteaga, Kapor, et al. (2022) experiment was run on a population of families that already had access to the SAE simulator and public documentation: the difference between treated and control households was produced by the proactive delivery of that information at the moment of decision. That isolates the specific effect of MAT relative to Formal Transparency. Second, the calibration analysis we report in Section 4.2 uses a separate survey instrument (Satisfaction Survey) from the rule-comprehension analysis (Section 4.3, Parent Survey), with a four-cycle window and population scale that lets us identify the SES gradient and test its stability over time. Third, the rule-comprehension result is obtained on a procedural rather than probabilistic dimension and in a more recent cycle window. The same pattern, residual comprehension gap that publication does not close, appears across these distinct methodologies, dimensions, and moments.

Taken together, the five lines allow us to assert, with greater confidence than any one of them would support individually, that the comprehension gap that publication-centered transparency leaves open is persistent in time and unevenly distributed across the user population. That property

is what makes the combination of market design and Formal Transparency insufficient, and what justifies consolidating MAT as a third, enforceable layer.

4. Empirical evidence: two measurements of the comprehension gap

This section presents two independent measurements of the comprehension gap, on different dimensions of the system's complexity and using different survey instruments. Section 4.1 introduces the data sources. Section 4.2 uses the SAE Satisfaction Survey 2020–2023 linked to administrative truth to measure miscalibration of placement probability, including a population-scale socioeconomic gradient. Section 4.3 uses the SAE Parent Survey 2023–25 to measure declared versus verified comprehension of a specific procedural rule. Section 4.4 reads the two strands together.

4.1 The Parent Survey and the family-application rule

This subsection sets up the rule-comprehension measurement of Section 4.3; the calibration measurement of Section 4.2 draws on a separate instrument, introduced there. We draw on the 2023–2024 and 2024–2025 waves of the SAE Parent Survey (Encuesta a Apoderados del SAE), administered by the Ministry of Education in coordination with the implementer of the allocation algorithm. The survey is fielded nationally each year to a sample of parents who participated in that cycle's application process, covering families across all regions and school types in the publicly funded system. The instrument is designed to capture both self-reported perceptions and verifiable knowledge about how the SAE works. Across both waves the survey includes two types of questions relevant to this paper. The first type asks parents whether they are familiar with specific procedural rules of the system, a self-assessment of knowledge. The second type presents verification items that test whether a parent who reports familiarity can in fact identify how the

rule operates when presented with a concrete scenario. The combination of the two is what allows us to measure the gap between declared familiarity and operative comprehension.

Our headline measurement concerns the family-application rule (Article 24, Decree 152/2016). This rule allows parents who have more than one child applying to schools through the SAE to link their applications so that siblings are considered jointly, with a shared priority. The rule's logic is straightforward in principle (siblings applying together receive an additional priority at schools where one of them is already enrolled or where both are applying), but its interaction with the rest of the algorithm (tie-breaking, priority stacking, the distinction between "family application" and "joint application") makes its practical implications difficult for most families to anticipate. The family-application rule is an especially useful test case for the MAT argument precisely because it is procedural rather than probabilistic: unlike the overall assignment risk studied by Arteaga, Kapor, et al. (2022), it does not require families to reason about probabilities. If the comprehension gap appears even for a procedural rule, the case for a structural problem with the current transparency paradigm is stronger.

As disclosed in the front matter, these surveys were collected in the operational context of the SAE's implementation; the findings rest on the survey microdata, and the analysis and interpretations are those of the authors.

The verification instrument is structured as follows. The self-assessment item asks respondents whether they know how the family-application rule works (yes/no). The verification item presents a scenario in which a family with two children applies to the SAE, one of whom already has a sibling enrolled at a school on the list, and asks the respondent to identify which of four statements correctly describes what happens when the family-application option is activated. The correct answer requires understanding that the siblings receive shared priority at the school where one is

already enrolled and that this priority applies to the joint application. The distractors capture common misconceptions: that activating the option guarantees admission, that it assigns all siblings to the same school regardless of ranking, or that it merely groups the applications administratively without conferring any priority. The gap between declared and verified comprehension is therefore a gap between claiming to know how the rule works and being able to identify its basic operation in a realistic scenario.

The surveys were administered online in both cycles to the universe of families who completed an application, with response rates of approximately 18 percent in 2023–2024 and 22 percent in 2024–2025. Respondents are stratified by educational capital using the parent's highest attained level of education. The analysis is descriptive: we do not model why declared familiarity grew faster than verified comprehension or attribute the change to any specific campaign. The contribution is measurement.

4.2 Miscalibration of placement chances, 2020–2023

Data and simulation. Our measurement of objective placement risk takes advantage of exactly the kind of Formal Transparency the SAE provides. The allocation algorithm and its implementation code are publicly available, as are the anonymized application records (preference lists, applicant priorities, school capacities) and a published simulator. We use these public artifacts to compute, for each application in cycles 2020 through 2023, the deadline-locked probability that the family is placed in any school on its submitted list, fixing the application as the family submitted it and re-running the deferred-acceptance algorithm under the system's published rules. This procedure is what an external analyst with sufficient technical capacity can

do using only the SAE's publicly released materials; in that precise sense, it is an existence proof of expert-level Formal Transparency in this setting.

We then merge the simulated objective probabilities to the administrative application records and, in turn, to the SAE Satisfaction Survey, fielded annually to all applicants after the regular round closes and containing a subjective probability item asking each family their stated probability of being placed in any of the schools they listed. The merge yields, for each respondent in cycles 2020–2023, a paired observation of subjective and objective placement risk computed under the same submitted application: 140,218 respondents with both items observed (between 20,617 and 68,719 per cycle). To examine the socioeconomic gradient, we additionally merge two SES indicators on individual identifiers: the prioritario flag (Chile's SEP eligibility marker for lower-income households, from administrative records) and mother's education from the SIMCE NSE panel. Define gap = subjective risk − objective risk; a negative gap indicates that the respondent understates the probability of going unassigned. The merged panel lets us measure the gap, decompose it across socioeconomic groups, and track both across four annual cycles. This measurement uses a different survey instrument from the analysis in Section 4.3 (the Satisfaction Survey rather than the Parent Survey) and a different cycle window (2020–2023 rather than 2023–25), so the two analyses are independent.

The methodological asymmetry is itself worth flagging. The objective risk we compute uses precisely the materials that publication-centered transparency makes available to interested experts. The families participating in the system have nominal access to the same materials. What the rest of this section asks is whether the families themselves act on beliefs consistent with what the published materials imply and, when they do not, how that gap is distributed across the user population. Figure 2 summarizes the pipeline.

Figure 2. From public artifacts to a measured comprehension gap

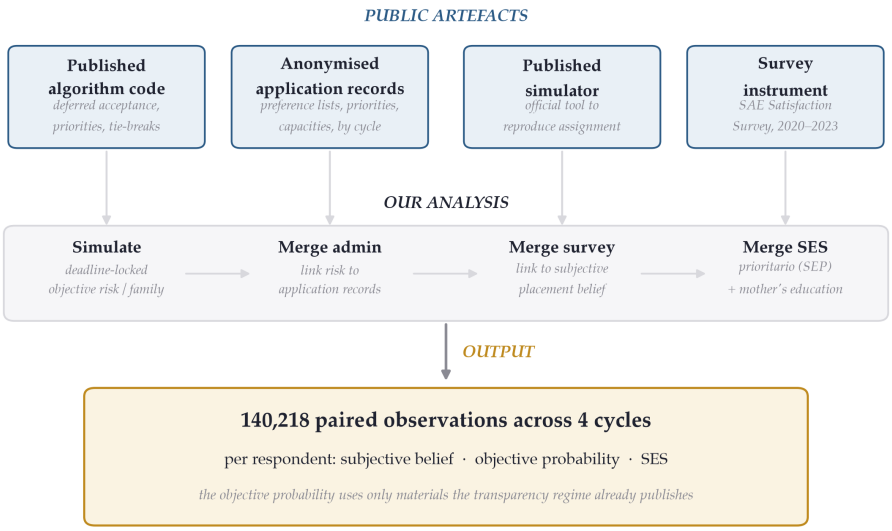

Note. Schematic of the measurement pipeline; no data are plotted. Objective risk is the simulated deadline-locked probability of assignment to no listed preference, from re-running the published deferred-acceptance algorithm on each family's final application. It is merged to the SAE Satisfaction Survey subjective item (pb_asignacionenpref) and to two socioeconomic indicators, the prioritario flag and mother's education. Pooled N = 140,218, cycles 2020 to 2023.

Aggregate miscalibration. In every cycle, mean subjective non-assignment risk lies 5 to 9 percentage points below the mean objective risk, with the bias concentrated in the upper tail of the objective-risk distribution. Figure 3 plots the calibration curve by cycle: within each objective-risk bin, median subjective risk lies visibly below the 45-degree calibration line in every cycle, and the gap grows monotonically as objective risk rises. Among applicants the platform can pre-identify as high-risk (objective non-assignment risk ≥ 30%, the operational threshold used by the SAE warnings layer), median true risk lies between 78% and 82% in every cycle while median stated belief lies between 19% and 30%, a median belief gap of 49 to 62 percentage points across the four-cycle panel (Figure 4). The level of aggregate miscalibration is modestly lower in 2023 than

in 2020–2022; we do not have a clean attribution for this level shift, and the substantive claim below rests on the conditional SES gradient rather than on the level. A striking aggregate signature appears in the modal subjective belief: across the four cycles, between 21% and 45% of respondents report a subjective placement probability of exactly 100%, regardless of the objective risk implied by their submitted portfolio.

Figure 3. Calibration of subjective against objective non-assignment risk, by cycle

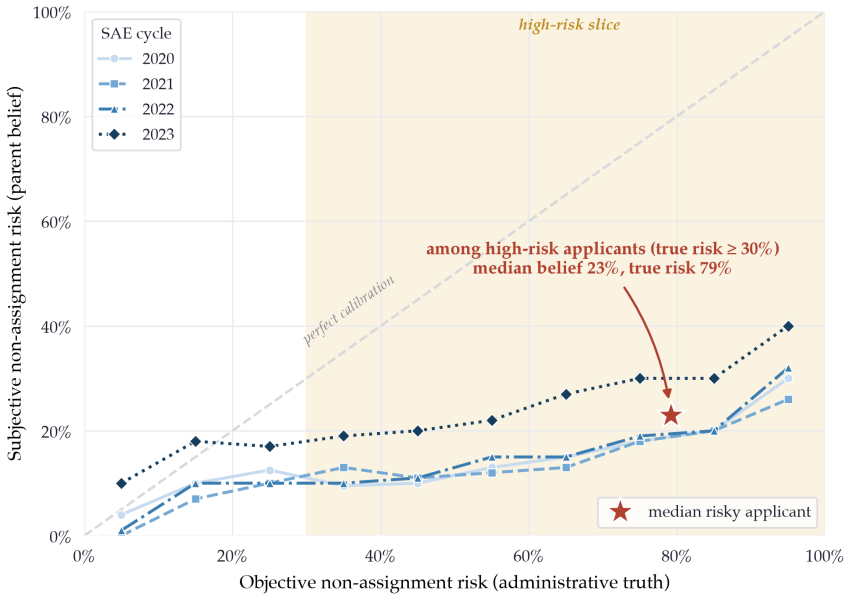

Note. Lines plot within-bin median subjective non-assignment risk against objective risk by cycle, with objective risk grouped in deciles. Subjective risk is one minus the stated placement probability (item pb_asignacionenpref); objective risk is the simulated deadline-locked probability of assignment to no listed preference. The dashed line is perfect calibration; the shaded band marks the 30 percent SAE warnings threshold. N = 140,218, cycles 2020 to 2023.

Figure 4. Median believed and actual non-assignment risk among high-risk applicants, by cycle

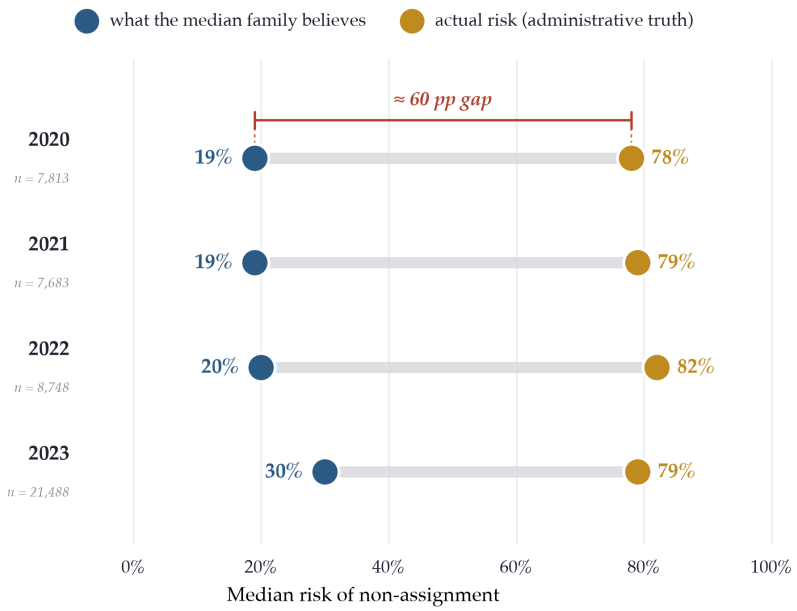

Note. Restricted to applicants whose objective non-assignment risk is 30 percent or above, the operational threshold used by the SAE warnings layer. Points are within-cycle medians of believed risk, from the SAE Satisfaction Survey, and of simulated objective risk; the connecting bar is their difference in percentage points. Cycles 2020 to 2023.

The SES gradient. A central finding is that the miscalibration is not uniformly distributed across the user population. We use two independent socioeconomic proxies. The first is the prioritario indicator, the administrative marker for students from lower-income households via Chile's SEP program; of the 140,218 respondents, 68,484 are prioritario and 71,734 are not. A naive aggregate comparison suggests that prioritario applicants are better calibrated: their mean gap is smaller in absolute value in every cycle. This is a composition effect. The SAE algorithm gives prioritario applicants priority through a reserved-seat quota and a tiebreaker boost, so their mean objective non-assignment risk is 6 to 12 percentage points lower than that of non-prioritarios in every cycle. Because the bias concentrates in the high-risk tail, a group facing less true risk on average mechanically has a smaller raw gap. Partitioning all respondents into five fixed objective-risk bins and recomputing mean belief gaps within each bin reverses the picture. At the same true

risk, prioritarios understate by more in every bin above 10%. In the highest-risk bin (75–100%, where objective risks are essentially identical between groups at 93.4% and 93.1%), the mean belief gap is −55.3 percentage points for non-prioritarios and −60.4 for prioritarios (Figure 5).

Figure 5. Socioeconomic gradient in belief miscalibration and its stability across cycles

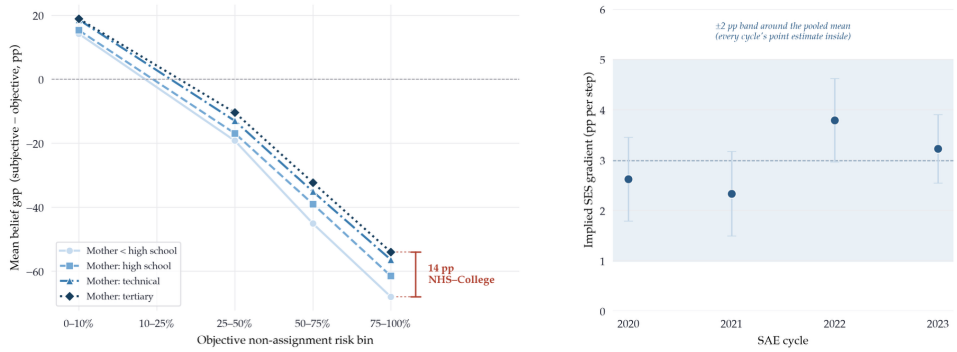

(a) (b)

Note. The belief gap is subjective minus objective non-assignment risk in percentage points; negative values indicate understatement. Panel (a): mean gap within five objective-risk bins, pooled 2020 to 2023, by mother's education (SIMCE panel, four ordinal levels), N = 59,185 of 140,218. Panel (b): implied gradient per one-step move on that scale at full objective risk, by cycle, OLS with HC1 standard errors, ±2 point equivalence band.

Pooled OLS confirms the conditional differential. Estimating gap = $\alpha + \beta \cdot$obj_risk + $\gamma \cdot$prio + $\delta \cdot$(prio × obj_risk) + ε across the four cycles (N = 140,218, heteroskedasticity-consistent standard errors), the interaction coefficient is $\delta = -0.041$ (SE 0.004): the marginal prioritario penalty in the gap is −1.4 percentage points at zero true risk and −5.5 at full true risk. The aggregate "smaller gap" for prioritarios is a composition effect, not better calibration.

The second proxy is mother's education from the SIMCE NSE panel, available for 59,185 of the 140,218 respondents. Pooled across cycles, the within-bin gradient is monotone in education: in the highest-risk bin, applicants whose mothers reported less than completed high school understate their risk by a mean of 68 percentage points, compared with 54 percentage points for

applicants whose mothers reported tertiary education, a 14-percentage-point conditional difference at the high end of the risk distribution. An analogous OLS specification yields a between-extreme conditional differential of roughly 17 percentage points at full true risk. The two proxies come from independent administrative sources and capture different facets of socioeconomic position; their consistency rules out the most obvious instrument-specific explanations.

Stability across cycles. The implied SES gradient (per one-step move on an ordinal mother-education scale, evaluated at full true risk) is 2.6 percentage points in 2020, 2.3 in 2021, 3.8 in 2022, and 3.2 in 2023. A joint test of equality across cycles is rejected ($F(3, 59{,}175) = 4.59$, $p = 0.003$), but the magnitudes are small: a two one-sided test of equivalence against a ±2 percentage-point tolerance accepts equivalence for every individual cycle's deviation from the 2020 baseline. A linear-trend specification yields a year-on-year drift of +0.24 percentage points per cycle, in the direction of mild attenuation. The conditional gradient is statistically stable across the panel; the level moves modestly.

Reading the evidence. At Chile's frontier implementation of Formal Transparency, families systematically overstate their placement probability; the bias concentrates in the high-risk tail; it is monotone in socioeconomic status, consistent across two independent SES proxies, and statistically stable in its conditional form across four cycles. Arteaga, Kapor, et al. (2022)'s randomized experiment demonstrated that delivering tailored risk information closes part of the gap and improves match quality; the analysis above provides the population-scale descriptive complement to that result and identifies the segments where the policy return is largest.

4.3 Declared versus verifiable comprehension

The miscalibration of placement probability documented above is one operational signature of the comprehension gap that publication-centered transparency leaves unaddressed. A second, distinct signature appears on a different dimension of the system's complexity: comprehension of specific procedural rules. We turn now to this second measurement, which uses a different survey instrument (the Parent Survey rather than the Satisfaction Survey), a different cycle window (2023–24 and 2024–25), and a different verification method (a scenario-based item rather than a subjective-probability item).

Table 4 presents the headline results for the highest-educational-capital group, parents with tertiary education or above. We focus on this group deliberately, as the most conservative test: these are the families most likely to read official communications, use the simulator, and extract information from published materials, so a comprehension gap here is a lower bound on the gap in the population. Between the 2023–2024 and 2024–2025 cycles, declared familiarity in this group rose by 23.6 percentage points in a single year (Table 4).

Table 4. Declared familiarity and verifiable comprehension of the family-application rule

	2023–24	**2024–25**	**Change (pp)**	**Conditional accuracy**
Declared familiarity (%)	36.8	60.4	+23.6	—
Verifiable comprehension (%)	2.3	4.9	+2.6	6–8%

Note. Respondents with parental education tertiary or above, the highest-educational-capital group (N = 19,562 in 2023–24 and 2,396 in 2024–25). Declared familiarity is the share reporting they know how the family-application rule (Article 24, Decree 152/2016) works; verifiable comprehension is the share correctly identifying its operation in a scenario-based item; conditional accuracy is the latter among the former. Change is in percentage points. Source: SAE Parent Surveys, administered online with response rates of approximately 18 and 22 percent, Ministry of Education of Chile.

Over the same period, verifiable comprehension, as captured by the survey's scenario-based verification items, barely moved. Among parents in the highest educational-capital group, the

share who correctly identified how the family-application rule operates rose from 2.3 percent to 4.9 percent, a gain of just 2.6 percentage points. The contrast is striking: declared familiarity increased nearly ten times faster than verified understanding (Figure 6). Among parents who declared familiarity, conditional accuracy, the probability of actually getting the verification question right given that the respondent claims to know the answer, was only 6 to 8 percent across both cycles. In plain language: for every 100 parents who say they understand the rule, between 92 and 94 do not.

The pattern is robust to disaggregation. The gap between declared and verifiable comprehension appears in every educational-capital group, and the fastest growth in declared familiarity occurs precisely in the groups where verified comprehension remains lowest. This rules out the interpretation that the gap is a product of a particular subgroup's unfamiliarity with the survey instrument. It also rules out the interpretation that "more information" is gradually closing the gap: information campaigns between the two cycles clearly succeeded in making more families aware of the rule's existence, but awareness and understanding are moving on different tracks.

Figure 6. Declared familiarity and verifiable comprehension of the family-application rule, 2023–24 to 2024–25

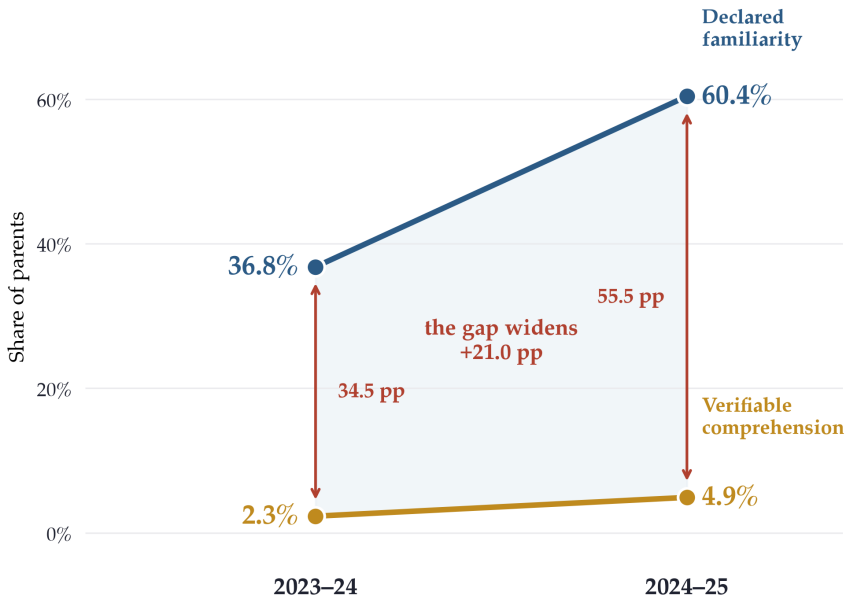

Note. Respondents with parental education tertiary or above. The upper series is the share reporting they know how the family-application rule (Article 24, Decree 152/2016) works; the lower series is the share correctly identifying its operation in a scenario-based item. Brackets give the gap in each cycle and the change between them. Source: SAE Parent Surveys 2023–24 and 2024–25, Ministry of Education of Chile.

4.4 Interpretation

Two independent measurements, on different dimensions of the system's complexity, in distinct cycle windows and using distinct instruments, reach a consistent conclusion: at Chile's frontier implementation of Formal Transparency, a substantial residual comprehension gap remains; it is concentrated in the high-risk tail of the user population; it is uneven across socioeconomic status, with lower-SES applicants showing larger belief errors at the same objective risk; and the conditional SES gradient is stable across the four cycles for which we observe it. The convergence across the two analyses strengthens the diagnosis: the most plausible alternative explanations of either finding in isolation (instrument artefact, item-interpretation effects, cycle-

specific events, sample composition shifts) would not produce a consistent pattern across two instruments, two cycle windows, and two dimensions of complexity.

The empirical pattern is inconsistent with three commonly held theories of how publication-centered transparency closes citizen comprehension. It is inconsistent with a slow-convergence theory in which accumulated publication closes the gap over time, because the conditional SES gradient is statistically equivalent across four annual cycles. It is inconsistent with a self-selection theory in which the most engaged citizens reach comprehension through their own initiative, because verifiable comprehension is near the floor even in the highest-educational-capital segment of the Parent Survey, and the high-risk tail of the calibration analysis includes the families with the strongest incentive to seek information. And it is inconsistent with a communication-quality theory in which a better FAQ or a clearer simulator would close the gap, because Chile already maintains an unusually well-resourced informational effort and the gap persists nonetheless.

What the pattern is consistent with is a structural feature of publication-centered transparency itself: the paradigm requires users to convert available information into operative knowledge through their own initiative, and that conversion task imposes a cost that is unequally distributed across the user population. Equal nominal access produces unequal effective uptake. The comprehension shortfall is at once persistent (the underlying mechanism does not change with publication volume) and unevenly distributed (the unequal distribution of conversion capacity does not change with publication volume either).

Three precisions should be made explicit. First, the distributional claim is about the information layer of the system, not its assignment outcomes. The SAE's underlying allocation mechanism is equity-oriented, and prioritario applicants face lower objective non-assignment risk than non-prioritarios by mechanical design; the unevenness we document is in the informational

architecture that surrounds the equity-oriented mechanism, not in the mechanism itself. Second, the causal link from miscalibration to assignment outcomes is established by Arteaga, Kapor, et al. (2022), whose intervention closed the gap in subjective beliefs and improved both portfolio choice and match quality; this paper's contribution is the population-scale descriptive evidence that the mechanism Arteaga identified is real, structural, SES-graded, and persistent. Third, the SES gradient in subjective beliefs may include a measurement component (rounding to focal probabilities, item-interpretation differences, numeracy effects); the policy implication is robust to this caveat because the published information is not producing well-formed beliefs in the affected population whatever the precise mix of mechanisms, and because Arteaga's experimental evidence demonstrates that delivering information changes behavior, not merely stated beliefs.

One feature of the broader evidence base reinforces this reading while remaining secondary to it. The systematic overestimation of admission chances we document for Chile is not setting-specific: Arteaga, Elacqua, et al. (2022), implementing the Chilean "smart matching platform" design in Peru and Ecuador, report that the same underlying frictions (limited awareness of nearby alternatives and systematic optimism about admission chances) are "remarkably consistent across Chile, Peru, and Ecuador," and that delivering tailored placement-risk warnings reduces non-placement risk where the warning reaches families (closely matching the Chilean benchmark in Peru, with near-null effects in Ecuador that they attribute to limited delivery reach rather than to a different underlying gap). Two implications follow, both secondary to our contribution: the comprehension gap itself travels across systems, and proactive delivery of the missing information has demonstrated benefits that are additive to Formal Transparency, not a substitute for it. Establishing that efficacy is not this paper's contribution: the cited experiments make that case.

Ours is prior and complementary: we show that the need for such delivery persists even after a legal system has implemented the full Formal Transparency's checklist.

A final observation concerns the research program behind the evidence. Four of the five lines of evidence summarized in Table 3 involve one or more of the present authors, reflecting a connected research program on the SAE that has had unusual depth of operational access for nearly a decade. The convergence across distinct methods, instruments, dimensions of system complexity, and moments in the SAE's lifecycle is what distinguishes a structural diagnosis from a one-off finding; it is also what the depth of access made possible.

5. Making MAT enforceable: an operational standard and a legal pathway

The conceptual and empirical argument converges on an institutional question: how can MAT be translated into concrete and enforceable obligation for public algorithmic systems? This section first proposes six operational components of MAT, designed to be adaptable across different legal systems, and then examines how these requirements can be grounded within existing legal frameworks. The Chilean case is presented as a worked example: not as the only possible institutional pathway, but as an illustration of how a legal system can transform MAT from a normative standard into a set of enforceable obligations.

The regulatory challenge addressed by MAT is one that any legal system deploying algorithmic allocation systems will recognize. Formal Transparency (making information about administrative decision-making processes, rules and criteria publicly available) and expert transparency (enabling specialists to reproduce and audit algorithmic outputs) are primarily concerned with making the system visible, reviewable and auditable. However, neither layer, even in legal systems where both are well developed, creates an obligation for system operators to ensure that affected users receive information in a way that enables meaningful understanding and action. That additional obligation

is what MAT seeks to introduce. The European Union's Artificial Intelligence Act (European Parliament & Council of the European Union, 2024) illustrates the point at the current regulatory frontier: Annex III classifies systems that determine access or admission to educational institutions as high-risk, placing a mechanism of the SAE's kind squarely within its scope. Its transparency obligations nonetheless run principally to deployers (Article 13), while the affected person's own entitlement is a right to obtain an explanation from the deployer once an adverse decision has been taken (Article 86). That duty is reactive and ex post, leaving the applicant's comprehension during the application itself outside the standard.

5.1 Starting point: Chile as a worked example

Chile's legal framework illustrates both the potential and the remaining gap. First, Law 20,285 on access to public information, read together with the publicity principle in Article 8 of the Chilean Constitution and its expert enforcement authority, the Council for Transparency (CPLT), covers Formal Transparency, the obligation to make government information available upon request and proactively. More recently, the CPLT issued Recommendations on Algorithmic Transparency (CPLT, 2024, Resolution 372/2024), highlighted by the GPAI (2025) report as one of the leading international developments in the area. These recommendations move beyond traditional disclosure by extending transparency duties to the design, operation, data inputs, and impact assessment of automated decision-making systems. However, their main limitation is that they remain non-enforceable.

Second, Law 21,719 on personal data protection, which enters into force in December 2026 together with the creation of a new supervisory authority, the new Agency for the Protection of Personal Data (APDP), introduces specific safeguards regarding decisions based on automated processing of personal data, including profiling. Article 8 bis recognizes the data subject's right to

object and not to be subject to such decisions, as well as the right to information and transparency, the right to obtain an explanation, the right to human intervention, the opportunity to express their point of view, and the right to request a review of the decision. Nevertheless, specialized literature warns that these safeguards are significantly limited in the context of public automated decision-making due to the potential limitation resulting from Article 21 (Ponce de León, 2026, 81-85). Until the new APDP adopts a different interpretation through its regulatory and enforcement powers, these transparency safeguards may need to be reconstructed through fragmented and sector-specific competences, depending on the regulatory powers and priorities of each public authority (Contreras et al., 2022), including those derived from Law 20,285.

The diagnosis is therefore twofold. On the one hand, Chile already has the legal scaffolding for Formal Transparency, and is an explicitly recognized international reference in that dimension. On the other hand, MAT is not yet articulated as an enforceable legal standard: the legal order does not require the operator of an automated decision-making system to provide information under conditions that enable effective user comprehension and action, does not define procedures to evaluate whether transparency measures achieve that purpose, and does not establish corrective mechanisms when they fail to do so. Filling that gap does not necessarily require new legislation: it requires articulating existing legal materials into a coherent normative framework. This insight, that the normative foundations for MAT may already exist within legal systems and need to be interpreted and systematized rather than created from scratch, can apply across different legal systems.

5.2 Six operational components of the standard

Building on Section 3.3 and the converging evidence, this section operationalizes MAT as six institutional requirements. The first two capture proactivity: identifying moments of

consequential choice and placing relevant information into users' awareness. The third captures tailoring: adapting information to the user's concrete decision context and foreseeable informational needs. The fourth captures actionability: explaining what the user can still do while action remains possible. The final two make the standard enforceable through periodic measurement of operative comprehension and corrective intervention when persistent or unequal gaps remain. Table 5 summarizes these components.

Table 5. Operational components of the MAT standard

Component	**Concrete obligation for the operator**
(i) Identification of moments of consequential choice	Document ex ante the points in the process at which users make choices with foreseeable consequences for the outcome. For the SAE, these include initial list construction, activation of family or joint-application options, and modifications before the deadline.
(ii) Concrete act of placing information into awareness	At each identified moment, execute a concrete notification act initiated by the operator. The information must not merely be available on a website or in a help menu; it must be brought to the affected user's attention through a dedicated channel.
(iii) Tailored, case-specific information	Deliver information adapted to the user's concrete decision context and, where relevant, to the informational needs and vulnerabilities of different groups of affected users. For the SAE, this may include an estimate of the risk associated with the user's current list and relevant warnings linked to the choices actually entered.
(iv) Actionable explanation and adequate time window	Accompany the information with a plain-language explanation of which available actions can change the outcome and how, delivered with enough time for the user to revise their choices before the process closes.
(v) Periodic evaluation of operative comprehension	Field instruments that measure whether users understand the relevant rule or risk, including verification questions, pilots, and user-testing. Results should be published in aggregate form, and the distance between declared familiarity and verifiable comprehension should become an institutional indicator.
(vi) External audit and corrective intervention	When comprehension indicators show persistent or unequal gaps, the operator must document corrective measures and their effects. A competent oversight body (such as the CPLT, the APDP where

Component	Concrete obligation for the operator
	legally applicable, or a sector-specific authority) should be able to audit the measures and require adjustments.

Note. Conceptual table; no data are reported. Components (i) and (ii) operationalize proactivity, (iii) tailoring, and (iv) actionability, as defined in Section 3.3; (v) and (vi) supply measurement and enforcement. The oversight bodies named are the Council for Transparency (CPLT), the Agency for the Protection of Personal Data (APDP), and sector-specific authorities.

5.3 From existing legal foundations to enforceable obligations

The implementation of MAT within the Chilean legal system does not require creating a new algorithmic transparency statute. The SAE case illustrates, in fact, how meaningful transparency obligations can emerge from adapting existing legal practices and frameworks. This requires reconstructing existing algorithmic transparency obligations from a systematic perspective that allows the incorporation of MAT components identified above.

Due to the strong influence of traditional administrative law and public-management practices, transparency rules have often operated, in practice, mainly through access to existing information concerning final administrative decisions. This occurs despite the fact that such an approach is not an explicit normative requirement (Camacho, 2018, 87-89). In this sense, algorithmic decision-making systems challenge this approach because relevant effects on individuals may arise before a final administrative decision is adopted, through system design choices, interfaces, and informational environments. In other words, automated public decisions involve a sequence of "intermediate administrative acts" that, in practice, produce effects on individuals comparable to those traditionally associated with "final administrative acts" (Bermúdez, 2014, 142-143). In all these instances, individuals must be able to understand how automated systems affect their legal position and available choices. Transparency in algorithmic environments challenges traditional doctrinal categories: as has been said, these interactions can be understood as part of the Administration's material acts (Izquierdo, 2026).

In this context, MAT requires interpreting transparency as extending beyond access to information already held by the Administration to a duty to create the informational conditions that allow affected individuals to understand and exercise their rights when interacting with automated systems. Therefore, it should be noted that this departure from the traditional administrative paradigm is, in fact, a practice that has already been implemented by the CPLT itself by instructing State administrative bodies to adopt a standard of "proactive transparency", understood as the obligation to "publish all information relevant to citizens […] whose publicity is determined by the principles of relevance, freedom of information, openness, maximum disclosure, facilitation, and timeliness" (Articles 2 and 114, Resolution 80/2026). Within this framework, the CPLT has issued specific recommendations on proactive transparency for State administrative bodies when adopting automated decision-making systems (Resolution 372/2024). The fact that the GPAI (2025) report explicitly cites Chile's CPLT recommendations as a global reference gives the initiative international legitimacy that can support domestic advocacy.

Second, personal data protection regulation contributes to a specific set of safeguards addressing the risks created by automated processing. Law 21,719 represents an important step in this direction by recognizing rights related to automated decision-making, including information, explanation, human intervention, and review mechanisms (Article 8 bis). Although its application to the public sector could be contested due to the interpretation of Article 21, its underlying rationale reinforces a broader legal transformation. In any case, these safeguards apply to the private sector; therefore, they may become indirectly applicable to public-sector contexts when private actors provide automated decision-making services to State administrative bodies.

The future regulatory and enforcement activity of the APDP will therefore play a central role in specifying how these guarantees operate in practice. However, as previously noted, these

transparency safeguards may be reconstructed through fragmented and sector-specific competences, including, of course, those of the CPLT (Article 33(m), Law 20,285), which already has oversight authority on access to public information when data protection is involved. If this occurs, as required by Law 21,719, the APDP should coordinate with the CPLT when adopting substantive doctrinal criteria on this matter.

Finally, sector-specific regulation provides an institutional framework through which MAT can become operational. The SAE itself provides the institutional infrastructure necessary for implementation. Within this framework, general transparency obligations can be translated into concrete procedural requirements adapted to the characteristics of the allocation system. The system already has a digital platform through which families interact, a notification system for deadlines and results, and a data infrastructure that records application behavior in real time. Delivering tailored information through the existing platform therefore requires adapting existing technical capacities, rather than creating an entirely new institutional infrastructure. The risk-warning intervention tested by Arteaga, Kapor, et al. (2022) was implemented precisely within this infrastructure, demonstrating that the technological and operational capacities already exist. In other words, proactive algorithmic transparency in this public policy would require identifying the relevant moments in which users make consequential decisions, defining proactive communication duties, establishing adequate timing requirements, and incorporating mechanisms to evaluate whether transparency measures effectively enable informed decision-making.

This architecture is generalizable. Many legal systems with mature access-to-information regimes, emerging data-protection frameworks, and sector-specific digital platforms already possess functional equivalents of these legal materials. What they often lack is their deliberate articulation around the specific objective of user comprehension. A design articulated in this way

would turn operative comprehension into an institutionally enforceable requirement of how the SAE functions, with clear measurement and correction mechanisms when the gap persists. That is the normative-legal contribution of this paper: to consolidate a decade of SAE learning into a standard that can be adapted to other automated decision systems.

5.4 A note on the merit debate

This paper focuses on how families understand the algorithm and leaves the normative debate over whether the SAE should incorporate academic-merit criteria outside its scope. The MAT argument is independent of that debate's resolution: whatever configuration of criteria the legislator chooses, the state owes affected families an informational architecture that lets them understand how their decisions map to the outcome.

6. Conclusion: the SAE as a laboratory for meaningful algorithmic transparency

The SAE's first decade offers a lesson that exceeds the Chilean case. A system can choose the best mechanism available from a market-design perspective, publish its algorithm, its data, and its simulator in full (that is, it can implement Formal Transparency in its most ambitious form) and still leave unresolved a systematic gap between what families believe they understand and what they actually understand about how their decisions map to outcomes. That gap has measurable consequences for match quality and therefore for the distributive objectives of the system. Its persistence after roughly a decade of implementation and refinement marks the point at which the Formal Transparency paradigm reaches its natural performance ceiling, and the next generation of algorithmic transparency must move from availability of information to legally meaningful conditions for comprehension and action: proactive, tailored and actionable.

Chile is unusually well positioned to do it because its legal system already contains the normative foundations required for this transition: algorithmic transparency regulation, emerging safeguards for automated decision-making and sector-specific regulation. Meaningful transparency does not depend exclusively on legislative reform. What it requires is that existing legal categories be adapted to the specific way in which algorithmic systems exercise public authority.

This approach also complements the recommendations of the 2025 Technical Committee on the School Admission System, which called for tailored notifications, interactive guidance tools, and timely support for families during the application process (Mesa Técnica del Sistema de Admisión Escolar, 2025, pp. 33–35). What they require is the institutional decision to treat the operative comprehension of families as a constitutive part of the right to participate, not as a desirable by-product of formal compliance. That change of frame is what we label the move to Meaningful Algorithmic Transparency.

The paper can also be read as a reply to the empirical vacuum the GPAI (2025) report itself declares in its conclusion. GPAI maps 83 public-algorithm repositories across four continents, diagnoses that most fail to reach meaningful transparency, and notes that "the literature has not sufficiently explored the actual contributions of repositories" to citizen comprehension and trust. The SAE is one of the examples the report itself uses to illustrate an automated decision-making system that a Chilean family interacts with. This paper offers the first longitudinal quantitative look from inside that example: the algorithm is published, and the question is whether the families who must use it understand it. The answer, for now, is: not enough. The question is which institutional standard should be adopted in order to change that.

A further observation is worth stating at the close. The SAE was created by the 2015 Inclusion Law as an equity-oriented intervention: a centralized system designed to displace school-level selection on academic merit and family background, with explicit priority structures for lower-SES applicants. Its substantive architecture delivers more equal placement outcomes than the regime it replaced. The empirical analysis in Section4 suggests that its informational architecture has a different distributional property: the comprehension gap that publication-centered transparency leaves open is larger for lower-SES applicants at the same objective risk, and four cycles of operating experience under the dominant paradigm do not narrow that conditional gap. MAT, in this reading, can be thought of as extending the equity logic of the underlying mechanism into the layer through which families actually interact with the system. The point generalizes: in any algorithmic allocation system with substantive equity objectives, the equity logic is not exhausted by the mechanism itself; whether the families the design intends to serve are equipped to navigate the system depends on the informational architecture as well.

These findings change the normative status of the choice between paradigms. Were the comprehension gap uniform, leaving it unaddressed would be a question of administrative efficiency. But the gap is unevenly distributed in the direction the system's equity design tries to correct, and the evidence reviewed in Section 4 suggests that proactive, tailored delivery can close part of it, at low cost and without weakening the publication regime it supplements. When a remediable shortfall in comprehension is concentrated on the disadvantaged, and the remedy is known, demonstrated in randomized evidence (Arteaga, Kapor, et al., 2022), and operationally feasible within existing infrastructure, remaining at Formal Transparency becomes a distributive choice rather than a neutral technical default. That is the sense in which consolidating MAT is an obligation of algorithmic transparency, not a refinement of it.

The conditions under which this argument applies extend well beyond Chile. The cross-country picture sketched in Section1 is one of substantial heterogeneity in the design dimensions that bear most directly on the comprehension gap: whether preference lists are restricted in length, whether tie-breaking is randomized or quasi-deterministic, and what information families receive before the application window closes. That heterogeneity matters for MAT in two ways. First, the empirical magnitudes documented here for the SAE will not transfer one-to-one across legal systems; the conceptual standard is neutral, but the operational components in Section5 must be tuned to the specific design and informational architecture of each system. Second, and more important, the same heterogeneity guarantees that comprehension gaps will exist in most of these systems and will be larger, in expectation, in legal systems that publish less than Chile does, which, as the Denmark contrast in Section 3 illustrates, is most of them.

Several comparative experiences illustrate the policy question concretely. Brazil's SISU, which centrally allocates university seats to roughly two million applicants per year using a score-based mechanism, has already generated evidence of strategic errors by applicants who do not understand how the algorithm processes their preferences. Kenya's centralized secondary-school placement system assigns students to national, extra-county, and county schools using an algorithm whose criteria are published but whose practical implications for individual students are opaque to most families. India's state-level experiments with school-choice reforms, notably the Right to Education Act's 25 percent reservation for economically weaker sections, involve algorithmic allocation at scale with minimal informational support for the affected families. Colombia's SISBEN scoring system, which determines eligibility for a range of social programs including educational subsidies, was flagged by the GPAI (2025) report itself as an example of an

automated decision system where algorithmic transparency has not translated into citizen comprehension.

In each of these cases, the policy question is the same one Chile faces: at what point does the state's obligation shift from making the algorithm auditable to making it understandable to the people whose lives it affects? MAT is the generic form of the institutional response to that question. The SAE, by virtue of its maturity and the depth of evidence it has generated, is a privileged laboratory for formulating the standard. This paper offers a first version of what that standard should require. The next step, for Chile and for any legal system that takes algorithmic transparency seriously, is to test whether the standard, once implemented, actually closes the gap.

Several directions for future research follow. The SAE survey data can be extended to track the gap longitudinally and, if MAT-style interventions are adopted, to measure their effect on operative comprehension over time, with attention to within-family learning dynamics across cycles. The relationship between comprehension and match quality deserves direct investigation: Arteaga, Kapor, et al. (2022) documented the link in an experimental setting, and the surveys' coverage of the full applicant population could test whether the relationship holds observationally at national scale. Finally, the six-component MAT standard proposed in Section 5 needs to be tested operationally. A randomized pilot in which a subset of SAE applicants receives tailored, proactive, actionable information (following the Arteaga et al. template but at scale and with institutional backing) would provide the first direct evidence on whether MAT closes the gap when implemented as a regulatory obligation rather than as a research intervention.

Beyond these research directions, the paper opens a broader conceptual question about the relationship between algorithmic transparency and democratic legitimacy in the context of the increasingly complex relationship between algorithms and citizenship. The SAE was adopted as

part of a legislative commitment to non-discrimination in school admissions, a commitment that required displacing the previous regime of school-level selection. The political sustainability of that commitment depends, in part, on whether families perceive the new system as fair. That perception, in turn, depends on comprehension: a family that does not understand how the algorithm processed its application cannot meaningfully evaluate whether the outcome was fair. In this sense, MAT is both an informational standard and a precondition for the democratic legitimacy of algorithmic governance. The SAE's experience suggests that Formal Transparency, however ambitious, does not automatically produce that legitimacy.

The challenge for the next decade is therefore to make algorithms visible to external observers and to ensure that the exercise of algorithmic public power remains understandable and contestable for the individuals subject to it.

Declaration of competing interest

Christopher A. Neilson advised MINEDUC in 2014–2015 on the design of the SAE and, from 2021 to 2024, led the external technical supervision of the execution of the SAE algorithm through a publicly procured contract held by the technology NGO ConsiliumBots. After the initial SAE evidence on platform information tools was produced and published, he founded TetherEducation, which provides software and services for application, assignment, and coordination platforms to local and national governments.

Exequiel Medina and Leonardo Ortiz Mesías have no competing interests to declare.

The surveys analyzed in Section 4 were collected in the operational context of the SAE's implementation. The findings rest on the survey microdata; the analysis and interpretations are

those of the authors and do not represent the views of the Ministry of Education of Chile or its affiliated institutions.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

CRediT authorship contribution statement

Christopher A. Neilson: Conceptualization, Data curation, Methodology, Supervision, Visualization, Writing – review & editing. Exequiel Medina: Formal analysis, Writing – original draft. Leonardo Ortiz Mesías: Investigation, Writing – review & editing.

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

During the preparation of this work the authors used Claude (Anthropic) in order to improve the wording of previously drafted text and to assist with editorial revision and reference formatting. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Data and code availability

The empirical analysis in Section 4 combines public and restricted-access inputs. The published SAE assignment algorithm and the anonymized application records (preference lists, applicant priorities, school capacities) for cycles 2020–2023 are available through MINEDUC's open-data portal. The SAE Satisfaction Survey microdata (2020–2023) and the SAE Parent Survey microdata (2023–2024 and 2024–2025) are held by Chile's Ministry of Education and can be

requested for research purposes through MINEDUC's data-access procedure. AI assistance (Claude, Anthropic) was used in developing and testing the replication code; the authors reviewed all code and take full responsibility for the analysis.

References

Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. E. (2009). Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. American Economic Review, 99(5), 1954–1978. https://doi.org/10.1257/aer.99.5.1954

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., & Sönmez, T. (2005). The Boston public school match. American Economic Review, 95(2), 368–371. https://doi.org/10.1257/000282805774669637

Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. American Economic Review, 93(3), 729–747. https://doi.org/10.1257/000282803322157061

Agte, P., Allende, C., Kapor, A., Neilson, C., & Ochoa, F. (2024). *Search and biased beliefs in education markets* (NBER Working Paper No. 32670). National Bureau of Economic Research.

Arteaga, F., Elacqua, G., Krussig, T., Méndez, C., & Neilson, C. A. (2022). *Can information on school attributes and placement probabilities direct search and choice? Evidence from choice platforms in Ecuador and Peru*. Inter-American Development Bank. https://doi.org/10.18235/0004672

Arteaga, F., Kapor, A. J., Neilson, C. A., & Zimmerman, S. D. (2022). Smart matching platforms and heterogeneous beliefs in centralized school choice. The Quarterly Journal of Economics, 137(3), 1791–1848. https://doi.org/10.1093/qje/qjac013

Bermúdez, J. (2014). Derecho Administrativo General. 3rd ed. Santiago: Thomson Reuters.

Brauneis, R., & Goodman, E. P. (2018). Algorithmic transparency for the smart city. Yale Journal of Law and Technology, 20, 103–176. https://doi.org/10.7282/00000058

Bullock, J. B. (2019). Artificial intelligence, discretion, and bureaucracy. American Review of Public Administration, 49(7), 751–761. https://doi.org/10.1177/0275074019856123

Camacho, G. (2018). Sobre la transparencia, el derecho de acceso a la información pública y los deberes de publicidad: Su diferencia conceptual y práctica. Revista de Derecho Público, Special Issue, 75–95. https://doi.org/10.5354/0719-5249.2018.51252

Cath, C., & Jansen, F. (2022). Dutch comfort: The limits of AI governance through municipal registers. Techné: Research in Philosophy and Technology, 26(3), 395–412. https://doi.org/10.5840/techne202323172

CEM (Centro de Estudios del Ministerio de Educación). (2025). Elección y admisión escolar en el contexto de implementación del SAE en Chile (Documento de Trabajo 36). Ministerio de Educación.

ConsiliumBots. (2025). Policy report on comprehension of the school admission system: The joint-application rule and the design of informational interventions. Report prepared for the Ministry of Education of Chile.

Contreras, P., Trigo, P., & Ortiz, L. (2022). Un Sistema fragmentado: la protección sectorial de los datos personales en Chile. Revista de Derecho Administrativo Económico, 35. https://doi.org/10.7764/redae.35.2

Correa, J., Epstein, R., Escobar, J., Ríos, I., Bahamondes, B., Bonet, C., Epstein, N., Aramayo, N., Castillo, M., Cristi, A., & Epstein, B. (2022). School choice in Chile. Operations Research, 70(2), 1066–1087. https://doi.org/10.1287/opre.2021.2184

CPLT (Consejo para la Transparencia). (2024). Resolution 372: Recommendations on algorithmic transparency. Santiago.

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

Giest, S., & Grimmelikhuijsen, S. (2020). Introduction to special issue: Algorithmic transparency in government: Towards a multi-level perspective. Information Polity, 25(3), 267–274. https://doi.org/10.3233/IP-200010

GPAI (Global Partnership on Artificial Intelligence). (2025). Algorithmic transparency in the public sector: A state of the art report on algorithmic transparency instruments (Co-leads: J. D. Gutiérrez & A. Gillwald). GPAI/OECD. https://wp.oecd.ai/app/uploads/2025/05/algorithmic-transparency-in-the-public-sector.pdf

Grimmelikhuijsen, S. (2023). Explaining why the computer says no: Algorithmic transparency affects the perceived trustworthiness of automated decision-making. Public Administration Review, 83(2), 241–262. https://doi.org/10.1111/puar.13483

Gutiérrez, J. D., & Castellanos-Sánchez, M. (2023). Transparencia algorítmica y Estado Abierto en Colombia. Revista Reflexión Política, 25(52). https://doi.org/10.29375/01240781.4789

Gutiérrez, J. D., & Muñoz-Cadena, S. (2025). Proactive algorithmic transparency in government: The case of the Colombian repositories of public algorithms. In Handbook on governance and data science (pp. 222–239). Edward Elgar. https://doi.org/10.4337/9781035301348.00019

Haeringer, G., & Klijn, F. (2009). Constrained school choice. Journal of Economic Theory, 144(5), 1921–1947. https://doi.org/10.1016/j.jet.2009.05.002

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. First Monday, 7(4). https://doi.org/10.5210/fm.v7i4.942

Herd, P., & Moynihan, D. P. (2018). Administrative burden: Policymaking by other means. Russell Sage Foundation.

Izquierdo, T. (2026). Las actuaciones materiales de la administración. Una aproximación. Revista de Derecho (Valdivia), 39(1). https://doi.org/10.4067/S0718-09502026000100155

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Mesa Técnica del Sistema de Admisión Escolar. (2025). Informe de recomendaciones. Ministerio de Educación de Chile.

Molina, E., & Medina, E. (2025). AI Revolution in Higher Education. What you need to know. In Digital Innovations in Education. World Bank.

Neilson, C. A. (2024). The rise of coordinated choice and assignment systems in education markets around the world (Background paper for the World Development Report). World Bank.

Neilson, C. A. (2026). A decade of the Chilean school admission system: Lessons from implementation and directions for reform (Working paper).

Nieuwenhuizen, E. (2024). Algorithm registers: A box-ticking exercise or meaningful tool for transparency? Information Polity, 29(4), 415–433. https://doi.org/10.1177/15701255241297107

Pathak, P. A., & Sönmez, T. (2008). Leveling the playing field: Sincere and sophisticated players in the Boston mechanism. American Economic Review, 98(4), 1636–1652. https://doi.org/10.1257/aer.98.4.1636

Peeters, R., Miller, S. M., & Schuilenburg, M. (2025). Digital government inclusion: Exploring strategies for inclusive government automation. Government Information Quarterly, 42(2), https://doi.org/10.1016/j.giq.2025.102028

Ponce de León, V. (2026). Protección de datos y vulnerabilidad socioeconómica: Los desafíos de la automatización en la asistencia social. Ius et Praxis, 32(1), 70–89.

Rizk, A., & Lindgren, I. (2025). Automated decision-making in public administration: Changing the decision space between public officials and citizens. Government Information Quarterly, 42(3), Article 102061. https://doi.org/10.1016/j.giq.2025.102061

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

van Deursen, A. J. A. M., & van Dijk, J. A. G. M. (2014). The digital divide shifts to differences in usage. New Media & Society, 16(3), 507–526. https://doi.org/10.1177/1461444813487959

Veale, M., & Brass, I. (2019). Administration by algorithm? Public management meets public sector machine learning. In K. Yeung & M. Lodge (Eds.), Algorithmic regulation (pp. 121–149). Oxford University Press. https://doi.org/10.1093/oso/9780198838494.003.0006