AI Power and Discourse, vol. 2, núm. 2, 2026, pp. 1-10.

Delegation to a Probability Distribution: Rethinking Legal Authority in Stochastic Systems.

Agustin V. Startari.

Cita:

Agustin V. Startari (2026). *Delegation to a Probability Distribution: Rethinking Legal Authority in Stochastic Systems. AI Power and Discourse, 2 (2), 1-10.*

Dirección estable: https://www.aacademica.org/agustin.v.startari/260

ARK: https://n2t.net/ark:/13683/p0c2/2E7

Delegation to a Probability Distribution: Rethinking Legal Authority in Stochastic Systems

A Jurisprudential Theory of Stochastic Delegation

Author: Agustin V. Startari

Author Identifiers

- ResearcherID: K-5792-2016
- ORCID: https://orcid.org/0009-0001-4714-6539
- SSRN Author Page: https://papers.ssrn.com/sol3/cf_dev/AbsByAuth.cfm?per_id=7639915

Institutional Affiliations

- Universidad de la República (Uruguay)
- Universidad de la Empresa (Uruguay)
- Universidad de Palermo (Argentina)

Contact

- Email: astart@palermo.edu
- Alternate: agustin.startari@gmail.com

Date: August 31, 2026

DOI

- Primary archive: **10.5281/zenodo.22233599**
- SSRN: Pending assignment

Language: English

Series: *The Stochastic Legal Order*

Word count: 6,062 (main text)

Keywords: legal theory; jurisprudence; artificial intelligence; stochastic systems; legal delegation; algorithmic decision-making; legal authority; institutional authority; computational law; legal attribution; probabilistic decision-making; administrative law

Abstract

This article develops a jurisprudential account of delegation to stochastic decision architectures. Existing work on automated public decision-making has already established that algorithmic systems can disrupt familiar doctrines of delegation, discretion, accountability, and administrative responsibility. The unresolved problem addressed here is narrower. A generative or otherwise stochastic system may be institutionally authorized even though, under fixed legally relevant inputs, its operative procedure can instantiate more than one legally consequential output. The legal question is therefore not exhausted by asking whether authority was delegated to an algorithm, whether a human remains nominally in the loop, or whether the system is explainable. It is whether authorization of the decision-generating architecture entails authorization of each particular output that the architecture can produce.

The paper calls the discontinuity between these two levels the Distributional Authorization Gap. Let G denote an authorized generative procedure composed of a model M, a selection or decoding rule S, and institutional constraints C. For legally relevant input X, G induces a distribution P(O | X; G) over possible outputs. Institutional authorization of G does not, without a further rule of legal imputation, entail legal authorization of every o_i in the support of that distribution. This proposition is especially important where two technically valid executions under materially identical facts and law can produce legally non-equivalent outcomes, such as grant and denial, detention and release, eligibility and ineligibility, or sanction and non-sanction.

The argument is analytical rather than empirical. It reconstructs the presuppositions of delegation theory around identifiable authority, institutional continuity, imputability, and bounded competence; distinguishes the present problem from established debates on algorithmic delegation, randomness, discretion, and predictive uncertainty; and identifies the point at which stochastic generation changes the object that must be legally authorized. The paper does not claim that randomization is inherently unlawful, that all machine-learning systems are stochastic at inference, or that legal decision-making has historically required deterministic reasoning. Law has long used lotteries, probabilistic evidence, structured discretion, and decision under uncertainty. The distinctive problem arises when

an institution authorizes a generative procedure whose output space includes legally incompatible realizations without specifying why the realization actually sampled, rather than another equally procedurally available realization, counts as the institution's authorized act.

The article therefore reframes stochastic delegation as a problem of output-level validity. Its central claim is that architecture-level authorization and output-level authorization are analytically distinct and that legal systems require an explicit bridge between them. That bridge may be supplied by bounded output design, a legally specified selection rule, meaningful human ratification, reviewable reasons, or another institutional rule of imputation. Where no such bridge exists, the system may be technically authorized while the legal status of its instantiated decision remains underdetermined.

1. Introduction: The Decision Before the Decision

Imagine an immigration authority that has lawfully acquired and approved a generative system for first-instance eligibility decisions. The governing statute, the applicant's file, the model version, the institutional prompt, the retrieval material, and the operative decoding settings are held constant. One execution produces APPROVE. A second execution produces REJECT. Both executions satisfy the technical procedure the authority approved. Both are generated by the same architecture. Nothing legally relevant about the applicant has changed.

The immediate temptation is to treat this as a familiar problem under a new technological description. Perhaps the system is unreliable. Perhaps the decision is arbitrary. Perhaps a human official should review the result. Perhaps the algorithm has been delegated too much discretion. Each response identifies a real legal concern, but none states the prior conceptual problem. Before asking whether the result is reasonable, reviewable, explainable, discriminatory, or properly supervised, one must ask why either result counts as an exercise of the authority that the legal order conferred in the first place.

Traditional delegation problems normally have an identifiable structure. A legal authority A confers or sub-delegates a competence upon an organ or official B, and B produces a

decision D that the legal system imputes to the relevant institution. The doctrinal details vary across jurisdictions, but the basic relation can be represented as:

$$A \rightarrow B \rightarrow D$$

The validity problem concerns competence, scope, procedure, purpose, relevant considerations, reasons, review, and the legal conditions under which B's act is treated as the act of the state. Even when B enjoys broad discretion, the institutional identity of the decision-maker remains conceptually available. A court can ask who decided, under what power, within which limits, on what grounds, and through which legally recognized procedure.

A stochastic decision architecture introduces a different object between authority and act. The operational chain is better represented as:

$$A \rightarrow G \rightarrow P(O \mid X; G) \rightarrow o_i \rightarrow I$$

Here G is not merely a person-like substitute for B. It is a decision-generating procedure. For present purposes, G = (M, S, C): a model M, a selection or decoding mechanism S, and a set of institutional constraints C. Given legally relevant input X, the procedure induces a distribution $P(O \mid X; G)$ over an output space O. One realization o_i is then instantiated and, if the legal system accepts it, converted through an institutional rule of imputation I into a legally consequential act.

This decomposition matters because the legal act is not identical to the model, the model is not identical to the distribution, and the distribution is not identical to the particular output selected in one run. In modern generative systems, the same underlying language model can be paired with different decoding strategies, and stochastic sampling methods intentionally select among tokens or sequences assigned non-zero probability by the model. Research on neural text generation has long shown that decoding strategy can materially alter outputs generated from the same model (Holtzman et al., 2020). The point is not that every deployed system necessarily samples stochastically. Greedy or otherwise constrained decoding can make some deployments deterministic for practical purposes. The relevant class for this paper is narrower: systems in which the authorized operative procedure preserves a non-degenerate set of legally non-equivalent outputs under fixed legally relevant inputs.

That technical qualification prevents an overbroad thesis. A conventional calculator, a deterministic eligibility formula, and a generative model operated under a fully deterministic and legally bounded mapping do not raise the same problem simply because they are computational. Nor does every use of probability create stochastic delegation. Predictive models can generate probabilities that a human decision-maker then evaluates. A model may estimate a 0.71 risk score without itself selecting between detention and release. The present problem arises when stochasticity enters the procedure by which the legally consequential outcome is instantiated or when the institution treats a sampled generative output as dispositive in practice.

Existing legal scholarship has already moved well beyond the simplistic question whether algorithms can participate in public decision-making. Hofmann's work on automated decision-making and cyber-delegation in European Union public law examines software that supports or replaces elements of human decision-making and places delegation, procedure, accountability, and public-private development at the center of the analysis (Hofmann, 2022, 2024). Langer (2024) treats automated administration through the non-delegation principle and emphasizes that algorithmic decision-making can shift decision power both to private software developers and to algorithms. Butler (2025) shows that modern non-delegation doctrine can be poorly fitted to public algorithmic decision-making because it focuses heavily on the final decision while important value-laden choices are made earlier in system design. These works establish the legal importance of the decision architecture and make it impossible to claim that delegation to computational systems is itself an unexplored field.

McCann (2023) narrows the issue further. His account of the automated administrative state argues that an algorithm is not itself a site of discretion. Importantly for the present paper, he expressly recognizes that algorithmic systems can produce outputs that are unpredictable in advance or that vary even when the same input data are supplied. He rejects the inference that such variation is legal discretion. Randomness, on his account, does not become discretion merely because the dispersion of outputs resembles the dispersion associated with human decision-making. That argument occupies substantial territory that a theory of stochastic delegation must not redescribe as new.

The law's relation to randomness is older still. Duxbury's study of lotteries and legal decision-making demonstrates that deliberate randomization has long been a jurisprudential possibility and, in some contexts, an institutional practice (Duxbury, 1999). Vermeule (2015) argues that administrative agencies can sometimes be rationally arbitrary under genuine uncertainty and that even random selection may be defensible when first-order reasons cannot distinguish between alternatives but second-order reasons justify making a choice. Contemporary machine-learning scholarship also increasingly distinguishes sources of predictive uncertainty. Sargeant (2026), for example, analyzes aleatoric and epistemic uncertainty in relation to anti-discrimination law and shows why probabilistic systems cannot be evaluated solely through generic claims about opacity or bias.

These literatures eliminate several weak novelty claims. This paper cannot plausibly claim that law has never confronted randomized decision procedures, that algorithms have never been analyzed as delegates, that probabilistic outputs are absent from legal decision-making, or that computational stochasticity is equivalent to a new form of human discretion. The contribution lies elsewhere.

The central proposition is a non-entailment claim:

$$\mathrm{Auth}(G) \not\Rightarrow \mathrm{Auth}_l(o_i), \text{ for every } o_i \in \mathrm{supp}(P(O \mid X; G)).$$

Authorization of the generative procedure does not by itself entail legal authorization of every output that the procedure can instantiate. The legal significance of this proposition becomes visible when the support of the distribution crosses a normative boundary. If all possible outputs are legally equivalent - for example, different phrasings of the same approved notice - stochasticity may be irrelevant to delegation. If the output space contains decisions that change rights, duties, status, liberty, benefits, liability, or access to public resources in incompatible ways, the authorization problem is different. The institution has not merely authorized a method of expression. It has authorized a process capable of selecting among distinct legal worlds.

The article calls the resulting discontinuity the Distributional Authorization Gap. The gap is the normative distance between institutional authorization of a stochastic decision-generating architecture and legal authorization of the particular decision instantiated by

one realization of that architecture. The word distributional identifies the structure of the problem; it does not imply that a legal institution must literally inspect the full mathematical probability distribution of every modern model. What matters is that the authorized procedure leaves more than one legally non-equivalent outcome live and that the selection among them is not fully determined by legally relevant reasons specified ex ante or supplied through a subsequent valid act of judgment.

This framing also clarifies why the familiar instruction to keep a human in the loop is insufficient as a conceptual answer. A genuinely independent human decision can close the gap by supplying a new legally cognizable act of judgment. A human official who evaluates the record, treats the machine output as non-binding, gives legally sufficient reasons, and can choose otherwise may remain the relevant delegate. But a nominal human click that merely ratifies the sampled output does not explain why that sampled output, rather than another output available under the same authorized procedure, deserved ratification. Human presence and human authorization are not synonyms. The legal question is whether a human act performs the normative work required for imputation.

The same point applies to explainability. An explanation may reconstruct why a particular output was generated, or it may provide post hoc reasons that make the result intelligible. Neither function necessarily establishes authority. A perfectly traceable output can still exceed the scope of what the institution lawfully authorized. Conversely, an institution may have lawful authority to make a decision even when the decision-maker's internal cognitive path is not fully reconstructible. The problem here is not epistemic access to the machine's reasoning as such. It is the legal relation between an authorized procedure and the instantiated act.

Nor is the argument that stochastic decisions are necessarily arbitrary. Law can expressly authorize random selection. A statute can require a lottery among equally eligible applicants; a regulation can prescribe random inspection; a court can use random assignment rules; a public body can employ tie-breaking mechanisms. In those cases, the legal order can define in advance both the bounded set of permissible outcomes and the rule by which chance selects among them. The eventual outcome is not unauthorized merely because chance contributed to its selection. The authorization relation is supplied

at the procedural level: the law authorizes this lottery, among these eligible alternatives, for this purpose, under these conditions, and treats the selected result as legally effective.

Generative systems can differ in a legally significant way. Their output spaces can be semantically open, context-sensitive, and difficult to enumerate ex ante. The same architecture can produce not merely different members of an explicitly authorized finite set but different interpretations, classifications, rationales, or decisions. If an institution authorizes the architecture without specifying the legal boundary of the output space and without a rule explaining why the sampled realization counts as the institution's act, the analogy to a legally authorized lottery breaks down. Randomness is not the defect. Unspecified normative reach is.

The research question is therefore: Can legal authority be validly delegated through a stochastic decision architecture when the particular legally consequential output is not fully determined ex ante? The answer developed in this paper is conditional. Legal authority can be exercised through stochastic procedures, but architecture-level authorization is insufficient where legally non-equivalent outputs remain possible. Valid stochastic delegation requires a bridge from procedure to output - a legally cognizable rule that bounds, selects, ratifies, or otherwise imputes the instantiated decision.

The analysis proceeds in stages. Section 2 reconstructs what existing delegation theory and contemporary algorithmic-delegation scholarship already explain. Section 3 identifies the missing unit of authorization and distinguishes architecture-level from output-level authorization. Section 4 separates model, distribution, selection, and instantiated output as legally distinct objects of analysis. Section 5 develops the Distributional Authorization Gap as a formal jurisprudential problem. Section 6 distinguishes stochasticity from legal discretion and from legally authorized randomization. Section 7 identifies conditions of output-level imputability. Section 8 develops a Stochastic Delegation Test. Section 9 considers consequences for competence, institutional responsibility, review, and reason-giving. The conclusion returns to the initial immigration case and asks what, exactly, must be true before the sampled answer can count as the state's answer.

2. Delegation to Computational Systems: What Existing Theory Already Explains

The law of delegation is jurisdiction-specific, but the jurisprudential structure beneath it is comparatively stable. Legal systems allocate powers to institutions and officeholders, permit some powers to be exercised through subordinates or agents, restrict other powers to designated decision-makers, and provide doctrines for determining when an act is attributable to the relevant public authority. The vocabulary differs - delegation, sub-delegation, devolution, authorization, agency, competence, discretion - yet the central problem is recognizably institutional: who is entitled to change the legal position of another, within what limits, and under whose name?

Analytical jurisprudence does not supply a single doctrine of administrative delegation, but it provides the conceptual materials from which delegation becomes intelligible. Kelsen's dynamic account of legal validity emphasizes the production of lower-level norms through normatively authorized acts within a hierarchy of competence (Kelsen, 1967). Hart's account of power-conferring rules explains how legal systems create facilities for altering legal positions and how secondary rules identify the conditions under which authoritative acts count as legally valid (Hart, 2012). Raz's account of authority centers the normative significance of authoritative directives and the institutional relation between an authority and those subject to it (Raz, 1979). Dworkin's treatment of discretion, although developed in a different jurisprudential contest, is useful for rejecting the idea that discretion means the absence of standards; discretion is intelligible only relative to a surrounding structure of standards that defines the decision-maker's space (Dworkin, 1977).

Taken together, these approaches do not prove a single theory of stochastic delegation. They do, however, reveal that legal authority is not simply causal efficacy. A mechanism can cause a legal system to behave without thereby possessing legal competence. A clerk can press a button, a contractor can design a scoring system, a database can supply a record, and a model can generate a recommendation; the fact that each element causally contributes to an outcome does not answer which element exercises legally attributed authority. Delegation doctrine exists precisely because causal participation and normative authorization are different relations.

Four presuppositions are especially important for the present argument: identifiability, continuity, imputability, and boundedness.

First, delegation presupposes identifiability. The relevant delegate need not always be a named natural person. Modern administration frequently acts through offices, departments, boards, teams, statutory bodies, and hierarchical bureaucracies. Yet the legal system must be able to identify the institutional location of the competence with sufficient precision to ask whether the act was made by or on behalf of an authorized decision-maker. Butler's reconstruction of the British non-delegation principle makes this accountability concern explicit: delegation doctrine is connected to Parliament's allocation of functions, retention of responsibility, and transparency about where decision-making authority resides (Butler, 2025). Identifiability is therefore institutional rather than merely biological. The law can recognize an office as the relevant bearer of authority even when the individual official changes.

Second, delegation presupposes continuity. The delegate must persist as a legally recognizable decision locus across the decision-making process. A minister remains the minister for purposes of the statutory power even though civil servants perform preparatory tasks. A tribunal remains the tribunal even though it receives expert evidence, uses staff, or changes members according to lawful procedures. Continuity does not require psychological sameness. It requires that the legal system can connect the decision to an institutional identity whose competence, duties, and accountability remain applicable from authorization through act.

Third, delegation presupposes imputability. The ultimate legal question is not simply whether an authorized official participated, but whether the resulting act is treated by law as the act of the relevant institution. Public law repeatedly separates preparation from decision, advice from judgment, and assistance from abdication. The rules by which an act is attributed to a minister, agency, board, or court are therefore constitutive. An output can exist as information without yet existing as a legal act.

Fourth, delegation presupposes boundedness. Every competence has limits, even when broad. Statutory language, purpose, jurisdiction, procedural duties, legally relevant considerations, constitutional constraints, reasonableness, proportionality, equality norms,

and review standards delimit what an authorized decision-maker may do. The precise doctrinal configuration varies, but the conceptual point is stable: the fact that B is authorized to decide does not mean that every conceivable act by B is legally authorized. Authority over a domain is not a blank cheque over all outputs that the decision-maker can physically produce.

This fourth presupposition is the immediate bridge to stochastic systems. Classical delegation already distinguishes authorization of a decision-maker from validity of each decision. A minister may be authorized to decide immigration cases yet still make an unlawful decision in an individual case. A tribunal may possess jurisdiction yet exceed statutory limits. The thesis of this paper is not that stochastic systems invent the distinction between institutional competence and act-level validity. They intensify and relocate it.

In ordinary human delegation, the invalid act is usually analyzed as an exercise - or purported exercise - of an identified delegate's power. The legal system can say: this official had competence, but this decision was unreasonable; this tribunal had jurisdiction, but it considered an irrelevant factor; this minister could act, but not for this purpose. The institutional subject remains stable while the legality of the act is evaluated.

In stochastic delegation, the institution may authorize a generative architecture before any particular legally consequential output exists. That authorization can define model version, data access, prompt structure, retrieval sources, decoding settings, confidence thresholds, escalation rules, and downstream workflows. Yet if the architecture's operative distribution includes multiple legally non-equivalent outputs under the same legally relevant input, there is an additional question between competence and act validity: what makes this realized output the institutionally selected exercise of competence?

Contemporary scholarship on algorithmic delegation approaches this question from adjacent directions. Hofmann (2022) emphasizes that software can support or replace human decision-making and that delegation analysis must account for public-private cooperation, data infrastructures, and the relation between software and law. His later work on cyber-delegation places automated decision-making within EU public-law principles governing delegated powers and accountability (Hofmann, 2024). Langer (2024) argues that automated administration can involve a transfer of decision-making power both to

software developers and to algorithms, creating non-delegation concerns linked to reviewability, contestability, and governmental responsibility. Butler (2025) focuses on the limits of a doctrine that asks too much about the final decision and too little about design choices that shape it. His analysis is particularly important because it shows that an official can nominally retain the final act while substantive discretionary choices have migrated upstream into system design.

These accounts make a major contribution: they disaggregate the administrative decision process. They show that authority can be functionally displaced even when formal legal doctrine continues to locate the final decision in a human officeholder. That insight is a necessary premise of stochastic delegation theory. The present paper adds a different disaggregation. Once design choices have been identified, the generated output itself must be separated into at least four objects: the model, the induced distribution, the selection rule, and the instantiated output. Delegation scholarship can identify who authorized the system and who designed it while still leaving unclear what authorizes the particular realization selected from the system's output space.

McCann's analysis supplies another boundary. He argues that the algorithm is not a site of discretion and that even output variation under identical inputs should not be mistaken for legal discretion (McCann, 2023). That position is compatible with the present thesis and, in one respect, makes the problem sharper. If stochastic variation is not an exercise of discretion, then it cannot automatically inherit the jurisprudential legitimacy that legal systems sometimes attach to discretionary judgment. The fact that a human decision-maker could lawfully choose between two permissible outcomes for reasons that emerge through deliberation does not imply that a sampling process may select between the same outcomes without an additional legal rule.

At the same time, McCann's argument should not be converted into a general claim that randomization is legally defective. Duxbury (1999) documents legal and political uses of lotteries and treats randomization as a serious decision procedure rather than an absurdity. Vermeule (2015) goes further in administrative law, arguing that under genuine uncertainty there are circumstances in which an agency may have sound second-order reasons to choose arbitrarily when no first-order reason distinguishes the options. The jurisprudential

lesson is not that law always demands reason-determined selection at the final step. It is that chance must itself occupy a legally authorized place in the decision structure.

A statutory lottery illustrates the point. Suppose a legislature creates one hundred identical licenses for one thousand equally qualified applicants and directs the agency to select recipients by lottery. The law has not delegated authority to an unknowable result. It has specified the eligible class, the number of awards, the random-selection procedure, the purpose of the procedure, and the legal effect of selection. The eventual identity of the winners is unknown ex ante, but the normative structure is not. The rule answers why winner W rather than applicant L receives the license: because the legally authorized lottery selected W from the legally defined class under the legally prescribed procedure.

The generative case can be structurally thinner. Suppose the immigration authority merely approves a language model and a prompt stating, "Apply the statute to this file and return APPROVE or REJECT with reasons." The model and prompt are authorized. The same legally relevant input can nevertheless produce opposite outcomes across valid runs. If the institution has not specified a legal selection rule among those outputs, has not required a determinative human act, and has not bounded the generative procedure so that legally incompatible outputs cannot both be procedurally admissible, then the lottery analogy does not solve the case. There is no equivalent legal rule saying why the first sampled answer rather than the second is the authorized exercise of the statutory power.

This exposes an important distinction between ex ante indeterminacy of outcome and ex ante indeterminacy of authority. Law routinely permits the former. A trial outcome is unknown before judgment. A jury verdict is unknown before deliberation. A lottery winner is unknown before the draw. A discretionary administrative decision is unknown before the officer evaluates the case. None of this means the source of legal authority is unknown. The institutional procedure identifies who or what is authorized to turn uncertainty into a legally effective result.

Stochastic architectures can create a different problem when the procedure specifies how to generate candidate outcomes but not the normative basis on which one candidate becomes the legally imputable act. The architecture may be technically complete and legally incomplete. It can answer the engineering question "how is an output produced?"

without answering the jurisprudential question "why does this output count as the authorized decision?"

The distinction can be stated formally. Let Comp(A,G,d) mean that authority A is legally competent to authorize generative procedure G for decision domain d. Let ProcValid(G,X,o_i) mean that o_i was generated according to the institutionally approved technical procedure on input X. Let ActValid(o_i,d) mean that o_i is a legally valid exercise of authority in domain d. Then:

$$\mathrm{Comp}(A,G,d) \wedge \mathrm{ProcValid}(G,X,o_i) \not\Rightarrow \mathrm{ActValid}(o_i,d).$$

This is not controversial as a general proposition; legal competence never guarantees substantive validity of every act. The distinctive issue is what the "without more" must contain when the source of variation is the authorized stochastic procedure itself rather than an identifiable official's exercise of judgment. In human delegation, law can evaluate reasons, purpose, process, and the decision-maker's legally attributed act. In stochastic delegation, the legal order must decide whether the sampling or selection step itself is normatively authorized, whether a human or institutional process ratifies the result, and whether the output space has been bounded so that technical validity does not extend beyond legal permissibility.

The problem therefore sits between delegation and validity. It is not reducible to either. Delegation asks whether the architecture may participate in or exercise the competence. Validity asks whether the resulting act satisfies applicable law. The Distributional Authorization Gap asks what connects the two when the authorized architecture can instantiate more than one legally non-equivalent act from the same legally relevant input.

That gap is easy to miss if the system is anthropomorphized. Phrases such as "the AI decided" or "the model exercised discretion" compress model, decoding procedure, developer choices, runtime configuration, operator action, and institutional adoption into a single artificial agent. The compression may be convenient in ordinary speech, but it is analytically destructive. A legal theory of stochastic delegation must resist it. The next section therefore asks what, exactly, is the unit that an institution authorizes when it deploys a stochastic decision architecture.

3. The Missing Unit of Authorization

The simplest account of computational delegation treats the system as the delegate. Authority moves from A to M, and M produces D. That representation is useful when the computational process is effectively deterministic and legally bounded. It becomes inadequate once the operative procedure can instantiate multiple legally non-equivalent outputs from the same legally relevant input.

The missing unit is not the artificial "agent" but the authorization relation itself. A legal institution can authorize at least four different things that are too often collapsed into one: a model, a generation procedure, a class of permissible outputs, and an instantiated act. These objects can overlap in practice, but they are not normatively identical.

Model authorization concerns the approval of M: a specified computational model or model family. It can include procurement, version control, training provenance, security requirements, validation standards, and permitted domains of use. Model authorization answers whether this model may be used. It does not by itself specify how outputs are selected from the model's conditional probabilities, which external tools or retrieval sources may be called, or which downstream decisions may be made dispositively.

Procedure authorization concerns G = (M, S, C): the model together with the decoding or selection rule S and institutional constraints C. This level is closer to the actual decision-generating process. A language model paired with greedy decoding is not operationally identical to the same model paired with top-p sampling. A model allowed to retrieve from a controlled statutory database is not operationally identical to the same model with unrestricted retrieval. A system required to abstain below a confidence threshold is not equivalent to one required to produce a binary decision in every case. Procedure authorization therefore identifies the computational process the institution has actually chosen.

Output-space authorization concerns the legally acceptable set O_l within the broader technically reachable output space O^G. This distinction is elementary but decisive. A generative model may be technically capable of producing outputs outside the institution's lawful competence: unlawful reasons, irrelevant criteria, discriminatory classifications, remedies the institution cannot grant, factual propositions unsupported by the record, or

decisions inconsistent with binding legal rules. No sensible account of delegation treats those outputs as authorized merely because the approved model can generate them. The legal output space must therefore be narrower than, or at most coextensive with, the technically reachable space:

$$O_l \subseteq O^G.$$

The existence of this subset relation shows why authorization of the model or procedure cannot be equivalent to authorization of every technically possible output. The institution may authorize use of G while simultaneously prohibiting entire classes of outputs that G can physically generate. This is familiar in software governance: systems have capabilities that operators are not legally entitled to use. The same distinction must be carried into jurisprudence.

Instantiated-act authorization concerns the particular output o_i that is selected in a concrete case and converted into a legal act. Even if o_i belongs to O_l, another question remains: what makes this permissible output the institution's actual decision? Legal permissibility and legal selection are separate. If the law permits both A and B, it does not follow that a machine may choose between them by any process whatever. A power to choose must be exercised through a legally cognizable decision rule.

This structure can be represented as four nested questions:

1. May the institution use M?
2. May it use procedure G = (M, S, C)?
3. Which outputs in O^G are legally admissible, forming O_l?
4. What rule makes the realized o_i in O_l the legally imputable act in this case?

The Distributional Authorization Gap exists when the legal system answers the first questions but leaves the fourth under-specified while the procedure can produce multiple legally non-equivalent realizations. The gap is therefore not identical to model risk or output error. A particular output can be substantively lawful and still expose an authorization defect if there is no legally recognized reason why that realization, rather than another equally procedure-valid realization, became dispositive.

Return to the immigration example. Assume that both APPROVE and REJECT are, in an abstract sense, outcomes the agency has power to issue. Both therefore lie within the agency's remedial competence. Assume further that the model has been validly procured,

the prompt lawfully approved, and the generation procedure technically validated. None of these facts explains why the first sampled output should bind the applicant. If the second run could have produced the opposite legal status without any change in legally relevant facts, the institution must identify the norm that converts one realization into the authoritative act.

Several candidate bridges are possible. The institution could require deterministic decoding after a legally specified decision boundary, thereby eliminating legally relevant distributional variation. It could authorize a lottery-like selection rule for a bounded class of equally permissible outcomes. It could require repeated sampling followed by a legally specified aggregation rule. It could treat every generated result as non-binding advice and require an official to exercise independent judgment. It could establish a rule that only outputs satisfying a set of reviewable reasons and consistency constraints may be ratified. The merits of these devices differ, but they share one feature: each supplies normative structure at the point where a distribution becomes an act.

This is the precise sense in which the paper describes stochastic delegation as delegation to a probability distribution. The phrase is shorthand, not metaphysics. Legal authority is not literally transferred to a mathematical object capable of holding office. Rather, the institution authorizes a procedure whose legally relevant behavior is distributional. The legal theory must therefore evaluate not only the identity of the software or developer but the mapping from an authorized distribution of possibilities to an instantiated legal consequence.

The distinction also disciplines the use of "agency" in AI jurisprudence. Nothing in the argument requires treating the model as a legal or moral agent. Searle's account of institutional facts is useful at this level because legal status depends on constitutive rules that make some event count as an institutional act under specified conditions (Searle, 1995). Whether a model possesses intentions is secondary to the legal question. The state can decide that an electronically generated record counts as a notice, that a digital signature counts as authentication, or that an automatically issued assessment counts as an agency act, provided the constitutive legal conditions are satisfied. The stochastic problem asks

what those conditions must contain when the electronic event is one realization among several legally non-equivalent possibilities.

The argument therefore avoids two symmetrical errors. The first is anthropomorphism: assuming that because the system outputs a decision, it must be treated as the decision-maker. The second is formalism: assuming that because a statute still names a human official, every computationally generated outcome automatically remains the human official's act. Both positions skip the imputation rule. In practice, institutional authority can be distributed across designers, procurement officials, model operators, front-line officers, supervisors, and automated components. The law must decide which of those causal contributions are legally constitutive of the final act.

Butler's critique of final-decision formalism becomes especially important here. If a human official simply accepts a generated output under severe automation bias, time pressure, or institutional expectation, the nominal retention of the final click may conceal functional delegation upstream (Butler, 2025). The Distributional Authorization Gap adds that the institution must also explain the relationship between the sampled output and the human act. A meaningful human decision can provide the missing bridge; a rubber stamp may merely inherit the architecture's unresolved selection problem.

The same analysis applies beyond public administration. Courts may use generative systems for triage or drafting; prosecutors may use them for charging recommendations; insurers may use them for claim decisions; universities may use them for admissions or discipline; platforms may use them for content moderation; private employers may use them for hiring. The public-law doctrine of delegation will not apply identically across these domains. Yet the jurisprudential distinction between authorizing a stochastic procedure and authorizing the instantiated legally consequential output remains available wherever institutional rules transform generated content into a decision that changes another person's legal or practical position.

The concept should therefore be stated at a sufficiently general level:

Distributional Authorization Gap: the normative discontinuity between authorization of a stochastic decision-generating procedure and authorization of the particular legally consequential output instantiated from that procedure's set or distribution of possible outputs.

Three conditions are necessary for the gap to matter. First, the procedure must preserve more than one outcome under fixed legally relevant inputs. Second, at least two of those outcomes must be legally non-equivalent. Third, no independent legal rule may already determine how the institution validly selects, ratifies, or imputes one realization as the act. If any condition fails, the strong version of the problem disappears.

This definition also supplies falsification conditions. If repeated executions differ only stylistically while legal effect remains identical, there is no distributional authorization problem. If a deterministic post-processing rule maps every generative variation to the same legal result, the relevant distribution collapses at the legal level. If the law expressly authorizes a bounded random-selection procedure, output uncertainty is normatively structured. If a human official genuinely exercises independent judgment and the machine output is only evidence or advice, the relevant act may remain the human decision. The theory does not require finding stochastic delegation everywhere generative AI is present.

The next analytical step is therefore technical and legal at once. Model, distribution, sampling mechanism, and output must be separated with enough precision to determine where the institution's authorization attaches. Only then can the paper ask whether legal indeterminacy and computational stochasticity are genuinely different, and whether the authorization gap can be closed by boundedness, reproducibility, traceability, reason-giving, or institutional imputation.

References

Butler, O. (2025). Algorithmic decision-making, delegation and the modern machinery of government. Oxford Journal of Legal Studies, 45(3), 727-752. https://doi.org/10.1093/ojls/gqaf018

Duxbury, N. (1999). Random justice: On lotteries and legal decision-making. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198268253.001.0001

Dworkin, R. (1977). Taking rights seriously. Harvard University Press.

Hart, H. L. A. (2012). The concept of law (3rd ed., L. Green, Ed.). Oxford University Press.

Hofmann, H. C. H. (2022). Automated decision-making and delegation: Discussing implications for EU public law. In H. Hofmann & R. Weaver (Eds.), Digitalisation of administrative law and the pandemic-reaction (pp. 91-115). Cambridge Scholars Publishing.

Hofmann, H. C. H. (2024). Assessing cyber delegation in European Union public law. In H. C. H. Hofmann & F. Pflücke (Eds.), Governance of automated decision-making and EU law (pp. 32-51). Oxford University Press. https://doi.org/10.1093/9780198919575.003.0002

Holtzman, A., Buys, J., Du, L., Forbes, M., & Choi, Y. (2020). The curious case of neural text degeneration. International Conference on Learning Representations. https://arxiv.org/abs/1904.09751

Kelsen, H. (1967). Pure theory of law (M. Knight, Trans.). University of California Press.

Langer, C. (2024). Decision-making power and responsibility in an automated administration. Discover Artificial Intelligence, 4, Article 59. https://doi.org/10.1007/s44163-024-00152-1

McCann, S. (2023). Discretion in the automated administrative state. Canadian Journal of Law & Jurisprudence, 36(1), 171-194. https://doi.org/10.1017/cjlj.2022.25

Raz, J. (1979). The authority of law: Essays on law and morality. Oxford University Press.

Sargeant, H. (2026). From estimation to discrimination: Algorithmic bias, predictive uncertainty, and anti-discrimination law. The Modern Law Review, 89, 593-627. https://doi.org/10.1111/1468-2230.70045

Searle, J. R. (1995). The construction of social reality. Free Press.

Vermeule, A. (2015). Rationally arbitrary decisions in administrative law. The Journal of Legal Studies, 44(S2), S475-S507. https://doi.org/10.1086/676332