



Legal Justification and Artificial Agents

David Opderbeck¹

Received: 24 April 2025 / Accepted: 24 April 2026 / Published online: 19 May 2026
© The Author(s) 2026

1 Introduction

The advent of artificial agents as legal decision-makers challenges virtue-based approaches to legal justification. Virtue jurisprudence, like theories of virtue ethics and virtue epistemology that reflect the “aretaic turn,” seeks a path between consequentialist and deontological theories.¹ As Colin Farelly and Lawrence Solum have noted, “the aretaic turn represents a renewed concern with human excellence as a unifying normative and explanatory concept.”² According to Farelly and Solum, “[f]or virtue jurisprudence, the final end of law is not to maximize preference satisfaction or to protect some set of rights and privileges: the final end of law is to promote human flourishing—to enable humans to live excellent lives.”³

Since Farelly and Solum’s groundbreaking work, much of the focus within scholarship on virtue jurisprudence has been on the role of the legal decisionmaker.⁴ This is a natural starting point for virtue jurisprudence, since the legislator or judge making legal decisions has, until very recently, always been a human being. Now that AI systems are beginning to make decisions and recommendations that affect people’s legal rights—for example, in commercial transactions, at an investigative stage in the criminal justice system, or in decisions about detention, bail, and sentencing—the need for human beings as to serve as legislators and judges has been questioned.⁵

¹ See Colin Farrelley and Lawrence B. Solum, “An Introduction to Aretaic Theories of Law,” in *Virtue Jurisprudence* (London: Palgrave-MacMillan 2007), 3.

² *Id.*

³ *Id.* at 2.

⁴ See, e.g., *id.*

⁵ See, e.g., Unesco, *AI and the Rule of Law: Capacity Building for Judicial Systems*, available at <https://www.unesco.org/en/artificial-intelligence/rule-law/mooc-judges>; *State v. Loomis*, 371 Wis.2d 235 (2016) (holding that use of AI risk assessment tool for sentencing purposes did not violate defendant’s due process rights).

✉ David Opderbeck
david.opderbeck@shu.edu

¹ Seton Hall University Law School, Newark, USA

As Gary Coglianese has suggested, “it is becoming increasingly possible to imagine governments shifting away from a predominant reliance on general rules and instead moving toward increased reliance on precise individual determinations—or on ‘personalized law,’” delivered by AI systems.⁶

From a virtue jurisprudence perspective, AI systems as legal decision makers seems like a nonstarter. If the end of the law is to promote human flourishing, and legal justification depends on the excellence of a human decisionmaker, law made by machine, whether general or personalized, seems unjustified. But under some aretaic theories of legal justification, this conclusion might not be so obvious. In particular, Amalia Amaya’s “counterfactual version” might apply regardless of whether the decisionmaker is in fact human. According to Amaya, in this counterfactual version, “[a] legal decision is justified if and only if it is a decision that a virtuous legal decision-maker would have taken in like circumstances.”⁷ Amaya eschews rule-based balancing models of legal reasoning and suggests that legal decision-makers should reflect a “connoisseurship” approach to case-specific reasons.⁸ If an artificial agent can be trained to evaluate legal options like a human legal “connoisseur,” and can make decision a virtuous human legal decision-maker would have made, then presumably the artificial agent passes Amaya’s counterfactual test even though it is not human.

This paper challenges that potential conclusion. In Part 2, it reviews the concept of legal justification. Part 3 examines Amaya’s aretaic twist on legal justification and her responses to the “publicity,” “authority,” and “disagreement” objections. This Part argues that some of those objections have more purchase the Amaya acknowledges and suggests that the notion of “connoisseurship” should be deepened in response to those objections. Part 4 argues that this deeper notion of legal connoisseurship means that legal decisions made solely by an AI system are not properly justified even if they satisfy Amaya’s counterfactual rule. The rise of AI in legal decision-making therefore demonstrates that the counterfactual rule alone cannot suffice as a norm of legal justification in an aretaic framework.

2 What Do We Mean by “Justification”?

Two different but related concepts of “justification” might relate to the role of aretaic theories. One is the internal justification for a particular legal decision. This is the kind of justification addressed by Amaya. Another is the justification or legitimacy of the law and the legal system within which specific legal decisions are taken. Both kinds of justification relate to the concept of the rule of law. They relate closely to each other because the criteria for justification of particular legal decisions must be rooted in the law and the legal system within which the decision is taken. A decision

⁶ Cary Coglianese, *Moving Toward Personalized Law*, University of Chicago Law Review Online, March 2022, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4051776.

⁷ Amalia Amaya, “The Role of Virtue in Legal Justification,” in Amalia Amaya and Hock Lai Ho, eds., *Law, Virtue, and Justice* (London: Bloomsbury 2012), 62.

⁸ *Ibid.*, 64.

taken outside the law and the legal system simply is not a “legal” decision. They also relate closely because legal decisions that lack adequate internal justification undermine the broader legitimacy of the rule of law.

2.1 What Do We Mean by “Legal Decisions?”

The literature on internal justification sometimes seems to assume that “legal decisions” ordinarily are judicial in character. The chapter from Amaya under consideration, for example, focuses on judgments by judges or juries in contested cases, that is, in civil or criminal litigation. Litigation, however, comprises only a small subset of the kinds of legal judgments that require justification. Indeed, most people are never personally a party to litigation, and if they are, most litigation settles long before final judgment (at least in the U.S.). The legal decisions that govern most people’s lives are those encoded in constitutions, statutes, regulatory rules, administrative proceedings, and prior case law involving other parties (particularly in common law countries).⁹

These distinctions are important because the criteria of internal justification for particular judicial and other legal decisions may differ. These differences in internal justification, in turn, may relate to the justification for the basic law (e.g., a constitution) that establishes the governing legal institutions—in most modern democracies, the legislative, judicial, and executive branches of government.

2.2 Aretaic and Positivist Theories of Justification

Any effort to establish an aretaic basis for legal justification, whether internal to specific decisions or more broadly of the rule of law, must situate itself within modern debates between legal positivists and non-positivist theories. Positivist meta-theories about the nature of law and legal institutions will by definition exclude aretaic justifications, at least as aretaic theories traditionally have been conceived. Because of this it is difficult to see how a positivist theory of internal justification for particular decisions can admit aretaic warrants.

A core feature of legal positivism is the claim that law is defined by facts and not by its moral merits.¹⁰ The facts referred to could be the power of a sovereign, the existence of certain basic norms (such as a Constitution), or social facts about the existence of official customs and practices. Most legal positivists would agree that some laws need not or should not be obeyed on moral grounds even if they can properly be called “law.” They would, however, separate those moral grounds for civil disobedience from the question whether the rule or principle is properly called “law” or something else.¹¹

⁹ This is noted in MacCormack (2005), 7 (“Much of what is written about the Rule of Law and about legal reasoning takes for granted the forum of adjudication. But that of legislation matters scarcely less.”).

¹⁰ See Green, Leslie and Thomas Adams, “Legal Positivism”, *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), available at <https://plato.stanford.edu/archives/win2019/entries/legal-positivism/>.

¹¹ See *id.*

A non-positivist, in contrast, might argue that social facts alone cannot comprise “law” in any legitimate sense. For natural law theorists, legal legitimacy, meaning the definition of what comprises “law” that entails some obligations of respect and compliance from persons under the law, is at some level inseparable from more basic moral questions. The classic modern debate, first engaged between Lon Fuller and H.L.A. Hart in the aftermath of the Second World War, concerned the Nazi legal system.¹² Fuller argued from a natural law perspective that Nazi law was not “law” at all and therefore could be discarded or disobeyed. Hart argued that Nazi law was “law” based on the social facts of the Nazi legal system. Hart agreed with Fuller that Nazi law could be discarded or disobeyed on moral grounds, but he understood those moral questions to be distinct from the existence of “law” as such.

Since the Hart-Fuller debate, various positivist and non-positivist positions have emerged in the Anglo-American legal and philosophical academies. Joseph Raz, for example, argued for a strong positivist view that law can only be evaluated from a perspective internal to its practices.¹³ In contrast, Ronald Dworkin agreed with Hart that law and morality cannot be entirely separated, but disagreed with Hart’s belief that there can be an overarching rule that identifies which laws in a legal system are valid.¹⁴ Dworkin favored a more case-specific approach in which a right answer theoretically could be discerned even in the most difficult cases.

Legal positivists can now be identified as non-inclusivist, meaning moral judgments play no role in defining law, or normative / inclusivist (such as Jeremy Waldron), who argue that the descriptive project of legal positivism also can include normative dimensions in suggesting better or worse legal systems.¹⁵ Meanwhile, at opposite poles, Critical Legal Studies (CLS) tends to accept the positivist view that law is socially constructed all the way down, while New Natural Law (NNL) theory attempts to construct natural law arguments for rational first principles while leaving broader metaphysical and religious questions to the side.¹⁶

Neil MacCormick’s work on legal justification occupied an interesting place in this matrix. He came to describe his view as “post-positivistic.” In *Rhetoric and the Rule of Law*, MacCormick defined law as “institutional normative order”—a claim

¹² See H.L.A. Hart, *Positivism and the Separation of Law and Morals*, 71 Harv. L. Rev. 593 (1958); Lon L. Fuller, *Positivism and Fidelity to Law—A Reply to Professor Hart*, 71 Harv. L. Rev. 630 (1958).

¹³ See Joseph Raz, *The Concept of a Legal System* (Oxford: Oxford Univ. Press 1980); *The Authority of Law: Essays on Law and Morality* (Oxford: Oxford Univ. Press 2009).

¹⁴ See Ronald Dworkin, *Law’s Empire* (Cambridge, MA: Harvard Univ. Press 1986); *Justice for Hedgehogs* (Cambridge, MA: Harvard Univ. Press 2011).

¹⁵ An inclusive legal positivist might argue that moral norms provide social conventions that officials might be obligated to observe. An inclusive legal positivist might also argue that, although it is at least possible for there to exist at least one legal system in which morality and law are distinct, it is also possible that there are some legal systems in which law and morality are not separable. There are many varieties of inclusive legal positivism, a survey of which is beyond the scope of this paper. See Kenneth Einar Himma, “Inclusive Legal Positivism,” in *The Oxford Handbook of Jurisprudence and Philosophy of Law* (Oxford: Oxford Univ. Press 2012), 125–165.

¹⁶ See John Finnis, *Natural Law and Natural Rights* (Oxford: Oxford Univ. Press 1980); Kimberle Crenshaw, Neil Gotanda, Gary Peller, and Kendall Thomas, eds., *Critical Race Theory: The Key Writings that Formed the Movement* (New York: The New Press 1996).

with roots in legal positivism. However, citing Adam Smith, MacCormick stated that legal justification involves

showing that we have gone beyond an immediate reaction to a situation seen at first glance. We will have sought to be fully informed and to have achieved an ideal impartiality in order to move from mere anger, sympathetic or direct, to settled moral concern and perhaps condemnation (or, in the converse case, approval in various degrees of warmth). . . one literally rationalizes the passionate reaction, in seeking a view that can be common to all concerned persons.¹⁷

For a view to be potentially common to all concerned persons, for MacCormick, it must be universalizable: “To justify an act is to show that it was right. To show that it is right is to show that, upon any objective view of the matter, the act ought to have been done, or even had to be done, given the character of the act and the circumstances of the case.”¹⁸ In contrast, for MacCormick, “[m]otivating and explanatory reasons do not *justify* an act.”¹⁹

3 Amaya’s Aretaic Counterfactual Twist on MacCormick and Her Response to Objections

Amaya draws on MacCormick’s distinction between normative and explanatory reasons to argue against the “strong” version of aretaic justification. According to Amaya, a legal decision-maker can act for the wrong (i.e., non-virtuous) reasons and still reach the right decision.²⁰ But, she says, this does not make the normative and explanatory reasons “indistinguishable.”²¹ We can still ask the counter-factual question whether a virtuous decision-maker would have taken the same decision under like circumstances.

3.1 The Publicity Objection

After sketching out her proposal, Amaya responds to some potential objections. The first is what she calls the “publicity objection.” In the first version of the publicity objection, the decision maker cannot justify her decision merely on the grounds that she is virtuous. The decision maker might indeed be virtuous, but the reasons for the decision are otherwise in a black box. Here we approach an important connection to the question whether an AI as a legal decision maker could satisfy the requirements of a virtue-based jurisprudence. The black box problem is central to the emphasis in AI ethics and policy on “explainability.”

¹⁷ McCormack, *Rhetoric and the Rule of Law*, 87.

¹⁸ *Ibid.*, 88.

¹⁹ *Ibid.*, 99.

²⁰ Amaya, *The Role of Virtue in Legal Justification*, 59.

²¹ *Ibid.*

Amaya agrees that black box legal decisions, even from a virtuous decision-maker, would be unacceptable, although she does not say why. Instead, she suggests that the objection is simply misplaced. Although she agrees that backing legal decisions with reasons is essential to democratic governance, she states that justification for a legal decision is distinct from the promulgation of public reasons for the decision.

Amaya makes this distinction in part because, she suggests, certain kinds of legal decisions are taken without disclosure of public reasons, including jury verdicts and decisions that might be correct but that are supported by bad reasons. The example of jury verdicts raises bigger questions about whether yes / no jury verdicts without more really *are* justifiable under rule of law principles. But even in jurisdictions, such as the U.S. federal courts, in which yes / no jury verdicts are possible, (1) the law of the case is public (in the form of statutes, regulations, and case law); (2) the procedural and evidentiary rules are public; and (3) the trial (absent exceptional circumstances) occurs in public, including any legal rulings made by the court and the jury instructions given by the court. Further, although most U.S. courts have adopted a broad non-impeachment rule for jurors (that is, jurors cannot be forced to testify about their secret deliberations), this rule includes important exceptions.²² The U.S. Supreme Court has ruled, for example, that racial prejudice can be a reason to impeach a jury.²³

The example of correct decisions that are badly reasoned begs the question whether a badly reasoned decision really is “correct.” The relationship between legal justification, publicity, and the soundness of reasons is at the core of debates between natural law theorists and legal positivists. The fundamental question is whether there is a distinction between “law” and the mere exercise of power. Under a strong version of legal positivism, something might be called “law” so long as it is recognized as such by the legal decisionmakers who have the power of enforcement. If there is a requirement of public accessibility and soundness of reasoning, it is a minimal one through which subjects can at least broadly discern how to act or not act through the pattern of action of their rulers.²⁴

For a normative positivist such as Jeremy Waldron, however, public promulgation supplies the most basic separation between the rule of law and the mere exercise of private power.²⁵ As Waldron has noted,

The public existence of such norms is to be contrasted with a situation where there are no abiding points of reference in one’s relations with state authorities, where significant portions of the population are kept in the dark as to what is expected of them, and where the tribunals-whose job it is to resolve particular disputes-apply secret directives that are neither known to nor knowable by the

²² Federal Rules of Evidence 606(b).

²³ *Pena-Rodriguez v. Colorado*, 580 U.S. 206 (2017).

²⁴ See David Dyzenhaus, *Law and Public Reason*, MacGill L.J. 38:2 (1993), available at <https://lawjournal.mcgill.ca/article/law-and-public-reason/>.

²⁵ Waldron, Jeremy, “The Rule of Law”, *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), available at <https://plato.stanford.edu/archives/fall2023/entries/rule-of-law/>.

non-state parties affected. Such a system of secret norms is not a legal system, for it lacks the dimension of publicity that most people associate with law.²⁶

The requirement of publicity is even more urgent for non-positivists and it connects to the soundness of reasoning. Consider Lon Fuller's famous eight requirements for something to be called "law": (1) the creation of rules; (2) publication or at least availability to the affected party; (3) prospective rather than retroactive application; (4) clarity (understandability); (5) consistency (avoiding contradictions); (6) possibility of performance; (7) stability (avoiding frequent changes); and (8) congruence between rule and enforcement.²⁷ Fuller said that "[a] total failure in any one of these eight directions does not simply result in a bad system of law; it results in something that is not properly called a legal system at all. . .".²⁸

Under either Waldron's or Fuller's views, a legal decision that lacks opacity or that is badly reasoned is an anomaly. A legal system, applying what is properly called law, will over time correct such anomalies. If the anomaly becomes the norm, the rule of law is in trouble. The publicity objection to Amaya's counterfactual version of aretaic legal justification therefore has more bite than she suggests.

In fact, Amaya implicitly recognizes both this problem and its solution in her second version of the publicity objection. The second version of this objection, as Amaya summarizes it, is that an aretaic approach would leave legal justification to some opaque, mystical perceptual sensitivity of the decision maker.²⁹ Amaya's response to this objection is that, even if the virtuous decision maker possesses a perceptual sensitivity that is not immediately visible to others, this does not preclude the decision maker from, in fact, offering reasons. She cites McCormick's observation that "[p]ractically wise people do not only issue Delphic assurances as to what is right, they explain why it is."³⁰ In other words, the virtuous legal decision-maker *will* disclose reasons for her decision *precisely because virtue requires she do so*.

McCormick's statement against "Delphic assurances" correctly notes that a virtuous legal decisionmaker will give reasons, but he mistakenly locates this imperative to explain under the virtue of *phronesis* or practical wisdom. Of course, practical wisdom would recommend the promulgation of reasons, since the craft of the legal decisionmaker involves precisely that kind of activity. To employ Aristotle's distinction between *sophia* and *phronesis*, a legal decisionmaker cannot only reflect on the abstract principles of law (*sophia*), she must make concrete decisions in circumstances of contingent experience (*phronesis*)—that is the job.

But the giving of sound reasons also, and more fundamentally, entails the classical cardinal virtue of justice (*dikaïosyne*). Justice, giving everyone their due, entails respect for the subject of a legal decision as a fellow human being, a fellow subject of the law, a fellow citizen or sojourner, with a rational capacity equal to that of the deci-

²⁶ Jeremy Waldron, *The Concept and the Rule of Law*, 43 Ga. L. Rev. 1, 26 (2008).

²⁷ Lon Fuller, *The Morality of Law: Revised Edition* (New Haven: Yale Univ. Press 1969), 39.

²⁸ *Id.*

²⁹ Amaya, *The Role of Virtues in Legal Justification*, 59.

³⁰ McCormick, *Practical Reason in Law and Morality*, 17.

sionmaker's (in principle, if not always in absolute quantity).³¹ A legal decisionmaker who fails to offer sound reasons for her decision fails to exercise the virtue of justice because she treats the subject as unworthy or incapable of hearing and understanding the reasons for the decision, that is, as something less than human. Likewise, a legal decisionmaker who fails to make reasons accessible to the *polis* fails to exercise the virtue of justice because she treats her fellow citizens as unworthy or less able and because she undermines the basis for future decisions including refinements and corrections of mistakes she may have made or responses to changed circumstances.

Justice is both an intellectual and an affective virtue. "Affective" here means more than emotions, although emotions are involved. It means at least that justice involves a habitual disposition to regard the worth of others. Justice (contra Rawls) is not merely "fairness" of a type that could be secured behind a hypothetical veil of ignorance. And contra MacCormick, motivating reasons *do* matter to justice. Justice, in fact, requires intimate knowledge of the other that motivates action and therefore is convertible with love.³²

3.2 The Authority Objection

The second objection Amaya addresses is what she calls the "authority" objection. A legal decision is only justified, Amaya suggests, if it is "grounded in law."³³ A virtuous legal decisionmaker might take a decision contrary to what the law requires if the law would produce an unjust result. Amaya responds that a virtue of "integrity" would require a virtuous legal decisionmaker to assess the case from an "internal point of view," which would preclude the possibility of virtue trumping law.

Amaya cites H.L.A. Hart's *The Concept of Law* for this idea.³⁴ Hart's notion of the "internal point of view," however, is a hermeneutic, and not a moral, concept. A person might take the internal point of view for numerous reasons, some of which would contravene the expectations of virtue. According to Hart,

Not only may vast numbers be coerced by laws which they do not regard as morally binding, but it is not even true that those who do accept the system voluntarily, must conceive of themselves as morally bound to do so, though the system will be most stable when they do so. In fact, their allegiance to the system may be based on many different considerations: calculations of long-

³¹ This formulation raises questions about non-human legal subjects, whether of property *in rem*, corporations and other artificial legal entities, or animals, and about human subjects with diminished capacities, such as infants and mentally disabled people. In the case of artificial legal entities or *in rem* property claims, the law still ultimately operates because human agents, such as shareholders or title holders, ultimately are affected. In cases involving human beings with diminished capacities, the law often recognizes that special considerations are due to such persons, typically resulting in the recognition of some fiduciary to whom legal reasons are due on behalf of the incapacitated person.

³² Nicholas Wolterstorff, *Justice in Love* (Grand Rapids: Eerdmans 2011).

³³ Amaya, "The Role of Virtue in Legal Justification," 60.

³⁴ See *id.* at 61, n. 23.

term self-interest; disinterested interest in others; an unreflecting inherited or traditional attitude; or the mere wish to do as others do.³⁵

Hart does use language analogous to virtue ethics when he speaks of the internal point of view as a “standing disposition” to act in accordance with the rules.³⁶ Indeed, this “standing disposition” is what makes the perspective internal, and thus a standard for evaluating conduct, rather than external, and merely descriptive. But Hart emphasized that a person could possess such a standing disposition for many reasons:

Some rules may be accepted simply out of deference to tradition or the wish to identify with others or in the belief that society knows best what is to the advantage of individuals. These attitudes may coexist with a more or less vivid realization that the rules are morally objectionable.³⁷

Further, Hart’s notion of the standing disposition to act differs from the notion of “habits” as a source of norms, which Hart explicitly rejected in response to the sanction-centered theory of J.L. Austin.³⁸ Austin’s notion of the role of habit, of course, differed from the concept of habit in virtue ethics, but in any event it is clear that Hart’s theory of the internal point of view is incommensurable with the role of habits and dispositions in virtue ethics.

3.3 The Disagreement Objection

The final objection to aretaic theories of legal justification noted by Amaya is what she calls the “disagreement” objection. The question here is how to evaluate the justification for legal decisions when there exists a plurality of different views, based in differing moral, religious, political, and social frameworks, about what would comprise the right legal decision in a contested case.

Amaya suggests first that the objection is overstated because, in many or most cases, it should be possible to discern a right answer. She cites Dworkin for this position.³⁹ But Dworkin’s view about why there is always a right answer relies on justifications other than virtue, so this citation is inapposite.

More promisingly, Amaya then moves on to note that “a virtue theory of adjudication is not committed to the claim that there is only one right answer.”⁴⁰ It is possible that two equally virtuous judges could reach different conclusions in a difficult case. This is, in fact, a feature of virtue jurisprudence, not a bug.

As Amaya notes in her concluding section on virtue and “reason” in law, a virtue approach differs from deontological approaches because virtue ethics are not rule-

³⁵ Hart, *The Concept of Law*, 203.

³⁶ See *id.* at 255.

³⁷ *Id.* at 257.

³⁸ See *id.* at 56; Scott J. Shapiro, *What is the Internal Point of View*, 75 Fordham L. Rev. 1157 (2006).

³⁹ Amaya, “The Role of Virtue in Legal Justification,” 61.

⁴⁰ *Id.*

based and subject to codification.⁴¹ The key virtue of *phronesis*, or practical reason, can be understood as a form of “connoisseurship,” which entails an “ability to appreciate heterogeneous, case-specific reasons for choice or preference by means of informed judgment or perception rather than the application of general principles or procedures.”⁴² The virtuous legal decisionmaker, then, does not parse the rules ever more closely until the “right” decision is reached. Rather, the virtuous legal decisionmaker offers an “informed judgment.”

To illustrate the notion of connoisseurship, Amaya, drawing on Wallace, offers the example of a wine contest.⁴³ Wine connoisseurs make judgments based on what they recognize as characteristics of a good wine. According to Wallace, “there is no independent set of rules or procedures which could be applied by anyone to determine the quality of an arbitrary bottle of wine; rather, a good wine simply is a wine that connoisseurs agree in finding case-specific reasons to approve of.”⁴⁴

The problem is that Wallace is wrong about what wine connoisseurs do (or think they are doing). Wine connoisseurs judge quality based on characteristics of mouth feel, scent, flavor components, tannins, color, and balance.⁴⁵ Organizations such as the Court of Master Sommeliers (CMS) and the Wine & Spirit Education Trust (WSET) publish “tasting grids” that enable scoring based on such characteristics. The tasting grid “provides a guide for an objective description of a wine’s character and quality. . . which leads to an unbiased conclusion of quality, age and development.”⁴⁶ WSET offers structured courses leading to credentials as a wine taster and the CMS offers similar courses to become certified as a sommelier.⁴⁷

As Amaya deploys this example, she seems to recognize that wine judging entails a recognized set of characteristics applied by connoisseurs that could apply across different bottles of wine. According to Amaya, if a non-connoisseur made contrary judgments based on very different characteristics, we would not likely take the non-connoisseur’s judgment seriously—but why not?

It could be only a matter of social signaling rather than virtue. The general public associates wine snobs with luxury and status and therefore credits their otherwise arbitrary or merely subjective judgments. From a classical virtue ethics perspective, in fact, this kind of dynamic would entail vices such as envy and dishonesty, not

⁴¹ *Id.* at 64.

⁴² *Id.* (quoting J. Wallace, *Normativity and the Will: Selected Essays in Moral Psychology and Practical Reason* (Oxford: Clarendon Press 2006), 257; and citing John Henry McDowell, *Mond, Value, and Reality* (Cambridge: Harvard Univ. Press 1998)).

⁴³ Amaya, “The Role of Virtue in Legal Justification,” 62 n. 29.

⁴⁴ Wallace, *Normativity and the Will*, 258.

⁴⁵ See Marshal Tilden III, “How to Taste Wine Like a Pro,” *Wine Enthusiast*, available at <https://www.wineenthusiast.com/culture/wine/wine-tasting-grids/>.

⁴⁶ *Id.* The WSET tasting grids are available at https://www.wsetglobal.com/media/6916/wset_12wines_sat_en_may2019.pdf; https://www.wsetglobal.com/media/3119/wset_13_wines_sat_en_jun-2016.pdf; https://www.wsetglobal.com/media/13271/wset_14wines_sat_en_aug2023.pdf.

⁴⁷ See <https://www.wsetglobal.com/>; https://www.mastersommeliers.org/resources?gad_source=1&gclid=Cj0KCQjw-r-vBhC-ARIsAGbUO2AqWz0rpooLnEMeeX9YW3wpNYuTTSRFaFCLxUtyAhdeCxxPqP-y6XcaAmMpEALw_wcB.

virtue. No doubt, as applied to status goods such as fine wines, this sort of dynamic is at least partly in play.

Or, it could be that the qualities listed on WSET and CMS's tasting grids reflect something about common human aesthetic experience and about the effort required from vintners to produce wines that embody these qualities in high degrees. Human beings have been producing wine for at least 8,000 years.⁴⁸ The Epic of Gilgamesh refers to "wine fit for a king."⁴⁹ If there is such a thing as a virtuous wine judge, she stands in one of the deepest known streams of human cultural tradition, and her judgments have something to do with the *arete* of the wine and the winemaker. The virtuous wine judge's *phronesis* is cultivated through deep exposure to the fruits of this long tradition (drinking plenty of different wines!); training of the senses required to discern subtle variations in smell, texture, and taste; participation in professional societies that provide structure, education, and accountability; and adherence to protocols that protect procedural integrity (blind tastings, tasting card metrics, averages of multiple judges' scores, and so-on). In short, the *virtuous* wine judge, contra Wallace, *refuses* ad hoc judgments based only on the circumstances of the particular case, and instead brings the subjective experience into conversation with the objective markers of excellence embedded in the tradition, based on a long process of common human experience.

The wine judge example remains helpful in relation to the disagreement objection to virtue jurisprudence, however, because it remains possible that different judges with roughly equivalent training and experience using the same tasting cards might give different scores to the same bottle—indeed, this is common. Do the potential variations among judges mean there are no objective virtues in wine judging? Not necessarily.

First, outside a winner-takes-the-gold competition, the practice of wine judging is not meant to convey absolutely quantitative information to the relevant public. Some variation among tasters is expected, within a reasonable range. If a wine receives a score of 91 from Wine Spectator magazine, an 89 from Wine Advocate, and a 93 from Bob the Sommelier's Wine Blog, all based on equivalent WSET training and tasting cards and sampling the same production batch, a reasonable consumer can see a quality profile within a certain range. If there were enough of a sample size, a statistical analysis probably would show confidence bands within five points or so are common among similarly trained and experienced tasters. Of course, if Bob had awarded this wine a 65 on contrast to other scores in the 90's, a reasonable consumer would suspect Bob is an outlier for some reason unrelated to the wine itself—maybe Bob just had a particularly bad day.

In a winner-takes-the-gold competition the professional standard is evaluation by a panel of judges, whose scores are averaged. This accounts for typical reasonable variations in individual scores, although it might not account for a single egregious outlier. Note that, in either case, disagreements among the experts are expected to be relatively small, absent a major anomaly. But if there must be a single gold medal

⁴⁸ See Patrick McGovern, et al., *Early Neolithic Wine of Georgia in the South Caucasus*, PNAS, Nov. 13, 2017, <https://www.pnas.org/doi/https://doi.org/10.1073/pnas.1714728114>.

⁴⁹ <https://www.ancienttexts.org/library/mesopotamian/gilgamesh/tab6.htm>.

winner, the process to which virtuous professionals are bound accounts for inevitable individual variations.

The same is true for legal decisions. Remember that legal decisions entail far more than the adjudication of contested court cases by a single jury or single judge. In fact, *most* legal decisions involve many decisionmakers in legislative and executive domains creating legislation and administrative rules, and even contested litigated cases, if they are appealed, are finally reviewed by a panel of judges. In this sense, disagreement is accounted for by the virtues of virtuous legal decisionmakers who function within legal systems—also, ideally, crafted by virtuous legal decisionmakers—that account for inevitable disagreements within a reasonable range of experience.

Even cases ultimately decided by a single judge or jury (at least at the trial level) entail the inputs of multiple legal decisionmakers as well as multiple legal decisions along the way to final judgment. Prosecutors must decide whether and what to charge (or in civil cases plaintiffs' lawyers must decide whether and what claims to file); in many cases a Grand Jury must pass the case along for trial; the parties' lawyers must decide what claims and defenses to press, what arguments to feature, and whether to plead out or settle all or some of the claims; and motion judges, who might not be the trial judge, will decide numerous dispositive and non-dispositive pre-trial motions, some of which might go up on interlocutory appeal. In the U.S., the very rare cases tried to a final judgment usually have run a long gauntlet of subsidiary legal decisions before reaching that final step of trial.

4 Applied to AI and Legal Justification

The concept of virtue and legal justification sketched above deepens Amalia's "connoisseurship" model of legal justification. Just as a professional wine critic stands in a long human tradition, a legal decisionmaker inherits a tradition stretching back to time immemorial. Both practices—making wine and making legal decisions—can be more or less oriented toward human flourishing, which is the focus of virtue ethics.⁵⁰ The virtuous legal "connoisseur" will embody practical reasonableness that reflect the additional criteria of reasonableness: "(1) 'reasonable' entails a set of epistemic and affective virtues; (2) the reasonable legal decision-maker engages in a practice of reflecting on the law's proper ends; and (3) the reasonable legal decision-maker cultivates habits of excellence (*phronesis*) in the process of deliberation."⁵¹

These criteria suggest difficulties for aretaic justification of legal decisions made by the kind of AI that now predominates—that is, by a machine learning system. First, a machine learning system can possess neither epistemic nor affective virtues because such systems neither know, think, nor sense. Rather, powerful machine learning systems such as Chat GPT are comprised of "transformers"—sets of algorithmic func-

⁵⁰ See Farrelly and Solum, "An Introduction to Aretaic Theories of Law," in *Virtue Jurisprudence* (London: Palgrave Macmillan 2007), 1–23, at 1 (noting that "For virtue jurisprudence, the final end of the law is to promote human flourishing—to enable humans to lead excellent lives.").

⁵¹ David W. Opderbeck, *Artificial Intelligence, Rights, and the Virtues*, 60 Washburn L.J. 445, 468 (2021).

tions—which encode initial mathematical conditions that produce comparisons of the statistical relationships between “tokens” (small pieces of training data).⁵² When queried, the machine will produce strings of tokens, that, based on the machine’s training data set and any fine tuning, would most probably appear in response to the query. None of this reflects a human-like process of knowledge acquisition, much less affective virtues such as commitment or love. Indeed, although machine learning systems are highly sophisticated prediction engines in a mathematical sense, they do not even *predict* anything in a conscious sense. They do nothing more than return a string of tokens based on statistical weights.

Second, machine learning systems are incapable of reflection on the proper ends of law (or on anything at all). The moniker machine “learning” suggests that these systems can learn, which would also seem to suggest that they can reflect on their experiences and the experiences of others, but this is misleading. As noted, these systems are trained on sets of existing data, which in the case of a large language model (LLM) such as ChatGPT might encompass the entire corpus of the public Internet.⁵³ The training data is tokenized and the associations among tokens are stored as statistical weights through the operation of transformers. This process might, for example, encode a probability weight that the word “data” should appear immediately after the word “training.”

After an initial foundational model training process, the machine can be fine-tuned in various ways. Human observers might test the machine and manually adjust model weights based on desired outputs.⁵⁴ User queries and user interactions with system outputs might feed back into the model weights. And new information produced after the foundational training may later be tokenized and used for model fine-tuning. This iterative process can produce impressive gains in a model’s predictive capabilities over time, but the output remains nothing more than a statistical prediction. The machine cannot ask “why,” make connections between raw information and sentiments, or weigh competing values to make judgments about the future. Its “learning” consists only in greater precision in associating tokens with each other.

Finally, a machine learning system cannot cultivate habits of excellence because it cannot refer outside itself to any standard of excellence. Nothing internal to the machine is drawn towards anything external to it that serves as a measure of excellence. The machine possesses no *telos* inherent to its own nature—and indeed, it possesses no nature of its own. It is merely a tool (in Aristotelian terms, *organon*)

⁵² See IBM, “What is a Transformer Model?,” available at <https://www.ibm.com/think/topics/transformer-model>; Jakob Uszkoreit, “Transformer: A Novel Neural Network Architecture for Language Understanding,” available at <https://research.google/blog/transformer-a-novel-neural-network-architecture-for-language-understanding/>.

⁵³ See Cem Dilmegani, *In Depth Guide to Web Scraping for Machine Learning in 2024*, AIMultiple: RSCH. (Jan. 2, 2024), <https://research.aimultiple.com/machine-learning-web-scraping/>; Dennis Layton, *ChatGPT—Show Me the Data Sources*, Medium (Jan. 30, 2023), <https://medium.com/@dlaytonj2/chat-gpt-show-me-the-data-sources-11e9433d57e8>; *So You’re Ready to Get Started*, Common Crawl, <https://web.archive.org/web/20230522201837/https://commoncrawl.org/the-data-get-started/> (last visited Mar. 16, 2024).

⁵⁴ See Benjamin Sobel, *A Taxonomy of Training Data: Disentangling the Mismatched Rights, Remedies, and Rationales for Restricting Machine Learning*, in *Artificial Intelligence and Intellectual Property* 221, 229–36 (Jyh-An Lee et al. eds., 2021).

used by human beings, who do enjoy the capacity for reflection on final causes and transcendent ends.⁵⁵

This analysis suggests that legal decisions made by machine learning AI alone cannot satisfy virtue jurisprudence criteria for legal justification. But this does not mean machine learning should be excluded from all processes of legal decision-making. AI is only a tool, but it can be a powerful tool. The question is whether and how the tool is subject to human agency and oversight.

Human agency and oversight are basic components of most statements of AI ethics.⁵⁶ Under the EU AI Act, human oversight is required for all “high risk” AI systems, including the ability for human deployers of AI systems to understand the system, notice biases, intervene while the system is in use, choose not use the system, or override decisions made by the system.⁵⁷ The EU AI Act encourages human oversight by design, including through dashboards and other user interfaces that allow human deployers to fulfill these obligations.⁵⁸ Legal guardrails such as these will help orient AI development towards a culture the centers human persons. Bar associations, legislative and regulatory bodies, courts, and other legal institutions will need to develop and publish professional ethics guidelines that embed similar requirements for AI tools used to assist legal decisionmakers. In this endeavor, those of us who are stewards of the rule of law must remember the humanistic focus of our enterprise. What we are doing is only justifiable as “law” if the final decisions rest with human persons who can feel, reflect, and refer to transcendent ends.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

⁵⁵ David W. Opderbeck, *AI and Creation: Some Perspectives from the Bible, Plato, Augustine, Aristotle, and Aquinas*, forthcoming in *The Journal of Law and Religion* (draft on file with author).

⁵⁶ See, e.g., Future of Life Institute, *Asilomar AI Principles* (August 11, 2017), available at <https://futureoflife.org/open-letter/ai-principles/>; Amnesty International and Access Now, *The Toronto Declaration* (May 16, 2018), available at <https://www.torontodeclaration.org/declaration-text/english/>.

⁵⁷ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, Art. 14. “High risk” AI systems cover a variety of activities, including the administration of justice. *Id.*, Annex III.

⁵⁸ *Id.*