SMARTDEMOCRACY: ARTIFICIAL INTELLIGENCE PLUS COGNITIVE DEMOCRACY AS A RADICAL PARTICIPATORY ALTERNATIVE

SMARTDEMOCRACY: LA INTELIGENCIA ARTIFICIAL SUMADA A LA DEMOCRACIA COGNITIVA COMO UNA ALTERNATIVA RADICALMENTE PARTICIPATIVA

Carlos Aguilar Blanc
Universidad de Huelva, Huelva, España
carlos.aguilar@dthm.uhu.es
Orcid id: https://orcid.org/0000-0001-6204-0911

Recibido: septiembre de 2025
Aceptado: noviembre de 2025

Keywords: SmartDemocracy, cognitive democracy, epistocracy, artificial intelligence, legitimacy, authority, constitutional design, XAI, digital rights.
Palabras clave: Democracia inteligente, democracia cognitiva, epistocracia, inteligencia artificial, legitimidad, autoridad, diseño constitucional, XAI, derechos digitales

Abstract: Liberal democracies face legitimacy crisis evidenced by epistemic, functional, and procedural deficits. This article proposes SmartDemocracy or Intelligent Cognitive Democracy (ICD), a constitutional model reconciling political equality, epistemic competence, and functional efficacy through tricameral architecture: Popular Chamber (inalienable universal suffrage), Cognitive Chamber (temporary credentials via merit/affectation/experience), and Synthetic Chamber (Explainable Artificial Intelligence XAI) under citizen oversight. The Algorithmic Tricephalous Executive implements Programmable Execution Pacts with verifiable indicators and automatic sunset clauses. Five absolute prohibitions prevent technocratic degeneration. ICD aims to surpass restrictive epistocracy and traditional representative democracy, institutionalizing knowledge under permanent democratic control through mandatory AI explainability and adversarial audits.

Resumen: Las democracias liberales enfrentan una crisis de legitimidad evidenciada en déficits epistémicos, funcionales y procedimentales. Este artículo propone como solución la SmartDemocracy o Democracia Cognitiva Inteligente (DCI), modelo constitucional que reconcilia igualdad política, competencia

epistémica y eficacia funcional mediante una arquitectura tricameral: Cámara Popular (sufragio universal), Cámara Cognitiva (credenciales temporales por mérito/afectación/experiencia), y una Cámara Sintética (Inteligencia Artificial explicable XAI) bajo supervisión ciudadana. El Ejecutivo Tricéfalo Algorítmico implementa Pactos de Ejecución Programables con indicadores verificables y cláusulas sunset automáticas. Cinco prohibiciones absolutas previenen la degeneración tecnocrática. La DCI pretende superar a la epistemocracia restrictiva y a la democracia representativa tradicional, institucionalizando un conocimiento bajo control democrático permanente mediante la explicabilidad obligatoria de la IA y auditorías críticas independientes.

1. Introduction

Liberal representative democracies face a structural legitimacy crisis that transcends cyclical fluctuations. The symptoms are evident and global: growing electoral abstentionism, citizen disaffection toward traditional institutions, the rise of populist and authoritarian movements, and the widespread perception of governmental inefficacy in the face of transnational challenges. This phenomenon, termed «democratic backsliding», does not occur through violent *coups d'état* but rather through the gradual erosion of norms and institutions, frequently perpetrated by elected leaders who dismantle checks on executive power (Levitsky & Ziblatt, 2018). The United States, Hungary, Poland, India, Brazil, and Turkey experienced setbacks characterized by electoral manipulation, judicial weakening, restriction of media freedoms, and extreme polarization through crisis narratives exploiting informational asymmetries. Contemporary democracies exhibit epistemic deficits (decisions based on deficient information), functional deficits (poor policy performance), and procedural deficits (oligarchic capture, lack of accountability). Fewer than 30% of Americans can name two First Amendment rights, and the majority lacks basic knowledge about candidates and governmental outcomes (Brennan, 2016, cap. 2, pp. 26-30). A critical aspect is the «hooligan voter»: committed but biased citizens who process information selectively to confirm preexisting prejudices (Brennan, 2016, cap. 1, pp. 4-5; cap. 2, pp. 36-43). This epistemic tribalization is structural: given that individual votes have infinitesimal probability of changing electoral outcomes, no incentive exists to process information objectively, converting politics into identity warfare where defeating the adversary matters more than solving problems.

1.1 Current State of the Debate

Facing this diagnosis, five theoretical responses have emerged that reconfigure democratic debate. Contemporary debate has transitioned from purely procedural conceptions toward approaches incorporating epistemic and instrumental considerations: the question is not simply whether democracy is fair by virtue of its procedures, but whether it can reasonably be expected to produce correct or at least not gravely erroneous decisions (Estlund, 2008, p. 98).

The deliberative and participatory approach holds that the crisis derives from

reducing participation to voting without effective deliberation. Fishkin proposes deliberative polling where informed citizens deliberate on complex policies, producing decisions superior to the aggregation of uninformed preferences (Fishkin, 2009, caps. 3-4), while Landemore defends open democracy based on sortition and inclusive deliberation, arguing that the demos's cognitive diversity systematically surpasses elite wisdom if properly organized (Landemore, 2020, cap. 3). Estlund's epistemic proceduralism establishes that democratic decisions are legitimate when procedures have better-than-random tendency to produce correct decisions, reintroducing instrumentalism without falling into pure consequentialism (Landemore, 2013, cap. 2, pp. 47-48).

The epistocratic approach questions universal suffrage, arguing that the majority lacks necessary knowledge for informed decisions. Brennan proposes conditioning voting on demonstrable political competence through examinations or plural voting where informed votes count more (Brennan, 2016, cap. 8, pp. 208-214), developing the «competence principle» according to which it is presumably unjust to forcibly deprive of life, liberty, or property as a result of decisions made by an incompetent deliberative body (Brennan, 2016, cap. 6, pp. 141-142). Caplan contends that voters are systematically irrational, embracing erroneous economic beliefs for psychological pleasure without direct personal consequences (Caplan, 2007). However, deliberative mini-publics with randomly selected citizens demonstrate capacity to deliberate competently on complex issues, suggesting that competence emerges in adequate deliberative contexts (Landemore & Brennan, 2020, pp. 167-168). Brennan's most sophisticated proposal is universal suffrage with epistocratic veto: a council of citizens with demonstrated competence that can veto incompetent legislation but not create it, preserving perceived legitimacy while improving decisional quality (Brennan, 2016, cap. 8, pp. 215-218).

The populist approach locates the problem in the exclusion of popular sectors by technocratic elites who emptied politics of substantive content. Laclau and Mouffe develop post-Marxist theory where political antagonism represents the ontological condition of the political: «the field of the social is traversed by antagonisms» that cannot be eliminated through totalizing rationality (Laclau & Mouffe, 2001, pp. 122-145). Mouffe holds that attempting to depoliticize decisions through technical consensus empties democracy of content (Mouffe, 1993, pp. 140-162), reinterpreting the Schmittian concept: antagonism must transform the enemy into agonist, legitimate opponent whose existence is democracy's condition of possibility (Mouffe, 2000, pp. 11-34). Dworkin would object that this celebration of antagonism can undermine institutional stability without a framework of shared principles (Dworkin, 1986, pp. 176-190).

The authoritarian meritocratic approach questions whether legitimacy must derive from electoral processes rather than effectiveness. Bell defends the Chinese model as a «political meritocracy» system where leaders are selected through competitive processes evaluating administrative capacity: electoral democracies face systemic weaknesses in leader selection: voters frequently lack adequate knowledge, motivation,

or time horizon (Bell, 2015, pp. 3-17). Waldron would object that meritocracy falsely presupposes an objective criterion of political competence identifiable independently of democratic processes (Waldron, 1999, pp. 101-130), while Sen argues that systems without democratic accountability lack incentives to respond to citizen needs (Sen, 1999, pp. 146-159).

The digital approach explores how decentralized technologies, artificial intelligence, and massive deliberative platforms can overcome traditional representative democracy's limitations. Weyl develops quadratic voting where individuals pay for as many votes as they wish using voice credits, but the cost is quadratic (Lalley & Weyl, 2018, p. 33), enabling efficient expression of preference intensity. Platforms like vTaiwan implement scalable deliberative architectures through machine learning: an open consultation process bringing together citizens and government to deliberate and reach rough consensus that generated decisive governmental action in 80% of 26 issues processed between 2015-2018 (Hsiao *et al.*, 2018, p. 2; p. 1). Morozov warns that technological solutionism ignores structural power and inequality dimensions (Morozov, 2013, pp. 5-15), while Noveck recognizes that technology by itself does not produce more competent citizens but must be complemented with robust deliberative processes (Noveck, 2015, p. 87).

We propose a sixth approach: SmartDemocracy or «Intelligent Cognitive Democracy», a constitutional model reconciling political equality, epistemic quality, and functional efficacy through tricameral architecture, supervised algorithmic executive, and temporary micro-credential system incorporating merit, direct affectation, and lived experience without creating permanent castes. SmartDemocracy maintains universal suffrage as ultimate source of procedural legitimacy, adding open cognitive filters improving informational quality, deliberation, and policy evaluation through explainable artificial intelligence under strong human supervision.

1.2 Diagnosis of Democratic Deficits

We identify three convergent deficits eroding democratic legitimacy. The epistemic deficit manifests in decisions based on inadequate, biased, or manipulated information. Public choice theory identifies «rational ignorance» as a structural problem: given that individual votes have infinitesimal probability of changing outcomes, it is rational not to invest in becoming informed because cost exceeds expected benefit (Downs, 1957, pp. 207-259). Empirical studies confirm alarming ignorance levels: broad majorities do not identify partisan positions on key issues, ignore basic economic data, and subscribe to systematically erroneous beliefs persisting despite contrary evidence (Caplan, 2007). During the 2000 elections, fewer than 57% of Americans knew that Gore favored greater spending than Bush, and fewer than half knew their positions on abortion or environmental regulation (Brennan, 2016, cap. 2, pp. 26-27).

The standard for measuring citizen competence proves no less problematic. Studies on voter ignorance assume certain decontextualized facts—gaps in basic judgment that experts deem

absolutely essential for possessing the political knowledge necessary to decide on political issues. However, it remains unclear whether extensive knowledge of social sciences or constitutional treaties is needed to judge whether globalization has benefited citizens. This assumption that the information measured by political questionnaires correlates with the capacity to vote correctly on matters of the common good remains undemonstrated (Landemore, 2013, ch. 7, pp. 200-201).

The problem transcends mere ignorance: many voters are irrational, embracing false beliefs because they provide psychological satisfaction without direct consequences. This rational irrationality explains the persistence of inefficient policies no expert would defend but enjoy electoral support. Voters process political information not to form accurate beliefs but to express tribal loyalties, a phenomenon termed «expressive voting» versus «instrumental voting» (Brennan, 2016, cap. 5, pp. 134-138).

The functional deficit is observed in poor policy performance, systematic nonfulfillment of electoral promises, and governmental incapacity to resolve complex problems requiring long-term vision. Democracies prioritize visible short-term benefits over deferred costs, generating excessive indebtedness, infrastructure underinvestment, and climate change inaction through democratic myopia derived from electoral incentives rewarding immediate results over sustainability.

The procedural deficit includes oligarchic capture where economic elites exercise disproportionate influence through campaign financing and lobbying (Gilens & Page, 2014), while extreme polarization blocks legislative compromises. As Duberry argues, popular sovereignty requires that citizens participate without coercion, but current digital platforms generate echo chambers: social network algorithms undermine civil society's capacity to make informed decisions (Duberry, 2022, cap. 4). (ICD) addresses this deficit through the Synthetic Chamber providing algorithmic transparency and explainability.

Empirical evidence from deliberative polling demonstrates that these deficits can be mitigated through appropriate institutional design. In Zeguo, China, participatory budgeting processes based on representative samples achieved complete budgetary transparency: participants showed significant changes in political preferences after informed deliberation, moving away from initial uninformed positions toward more coherent decisions oriented to the public good (Fishkin, 2009, cap. 4, pp. 106-111).

The convergence of these deficits erodes citizen trust. When citizens perceive that voting does not change substantive policies, that elites capture decisions independently of electoral outcomes, and that governments are incapable of solving urgent problems, disaffection becomes rational. Global democratic backsliding is a symptom of structural problems in the liberal representative model requiring profound institutional reforms.

2. Objectives, Questions, Thesis/Counter-Thesis, and Methodology

Every complex intellectual and institutional proposal demands clarity about what problems it aims to resolve

and how it proceeds to address them. Is it possible to reconcile political equality with epistemic quality without creating expert castes? Can algorithmic systems assist democratic decision-making without usurping popular sovereignty? What safeguards can prevent democratic degeneration into technocracy? This section sets out the objectives, formulates the central questions, presents «the three-dimensional legitimacy thesis» we defend, and specifies the methodology employed while acknowledging its limits.

2.1 Objectives and Questions

We pursue three interrelated objectives: to normatively justify that (ICD) offers a coherent institutional alternative surpassing identified limitations in the dominant currents of democratic reform debate, particularly in relation to epistocracy (Brennan) and deliberative democracy (Landemore, Fishkin); to specify institutional architecture with sufficient constitutional detail to evaluate its viability in western public law systems; to identify components with partial empirical validation in experiences such as deliberative polls, vTaiwan, or Citizens' Assemblies, delimiting which aspects require complementary research.

The central questions are: Is it possible to design democratic institutions reconciling political equality with epistemic quality and functional efficacy? Can cognitive filters such as temporary micro-credentials or AI-assisted deliberation improve the decisions without creating a permanent expert castes? What institutional safeguards prevent «knowledge systems» from degenerating into an oligarchic epistocracy or a captured technocracy?

2.2 Thesis/Counter-Thesis

Our central thesis maintains that the Intelligent Cognitive Democracy (ICD) model or SmartDemocracy resolves the tension between equality and competence through a three-dimensional legitimacy vector. Procedural legitimacy is preserved with the unalterable universal suffrage, where the Popular Chamber retains veto power, ratification, and political control. Epistemic legitimacy is constructed through the Cognitive Chamber (Senate of open temporary credentials) organizing specialized knowledge, lived experience, and direct affectation without coercive power, producing neutral reports and evidences syntheses. Functional legitimacy is ensured through the Synthetic Chamber (explainable artificial intelligence systems under human supervision) simulating policy impacts, identifying risks, and issuing opinions on technical coherence, without replacing democratic decision.

This three-dimensional architecture critically addresses the debate between pure epistocracies and procedural democracies. A purely epistemic conception of democracy is essentially committed to epistocracy and only contingently to democracy (Lippert-Rasmussen, 2012). Democratic legitimacy cannot depend exclusively on correct epistemic results, since this would make citizen participation dispensable if some form of non-democratic deliberation guaranteed better substantive results (Lafont, 2020, cap. 3). (ICD) avoids this trap by maintaining universal suffrage as an unbreakable foundation while adding epistemic layers that improve but never replace popular decision.

We propose operationalizing this triple legitimacy with an Algorithmic Tricephalous Executive (elected President, Rector selected by competence, Synthetic AI Council with explainability), subject to Programmable Execution Pacts with verifiable indicators, sunset clauses, and continuous evaluation. We reject both epistocracy (expert government without democratic accountability) and anti-intellectual populism (systematic contempt for evidence and specialized knowledge).

The counter-theses that will predictably be articulated from the critical literature require anticipatory refutation. The epistocratic objection (O1) will hold that any credentials system inevitably reproduces socioeconomic inequalities, benefiting educated elites (Lafont, 2015; Moraro, 2018). The governability objection (O2) will argue that tricameral architectures generate prohibitive transaction costs and decisional paralysis. The technocratic objection (O3) will assert that delegating policy evaluation to artificial intelligence, although explainable, concentrates power in those who design the algorithms and creates opportunities for opaque manipulation (Pasquale, 2015; O'Neil, 2016).

We anticipate the replies: Against O1, proposed micro-credentials are temporary (expire between 12-18 months), non-inheritable, subject to conditioned lottery to ensure pluralism, and open through three independent pathways (academic/professional merit, verifiable direct affectation, lived experience in affected communities). Against O2, modulated majorities according to risk level certified by the Synthetic Chamber accelerate urgent decisions without democratic bypass, while non-extendable maximum deadlines and sunset clauses compel periodic renewal. Against O3, the Synthetic Chamber lacks coercive power, its opinions are auditable through public source code, and every recommendation requires a meaningful explanation (XAI) plus mandatory human review by rotating citizen committees.

2.3 Combined Methodology and Traceability Criteria

We employ methodology combining three complementary approaches with differentiated scopes. Dogmatic-constitutional analysis constitutes the methodological core: we examine SmartDemocracy's institutional architecture as a constitutional draft proposal, evaluating internal coherence, norm hierarchy, competence distribution, and the conflict resolution mechanisms. We identify the structuring constitutional principles (political equality, epistemic transparency, functional accountability) and verify that proposed institutions do not contradict fundamental rights.

Comparative analysis operates selectively. We confront (ICD) with restrictive epistocracy (Brennan) and open deliberative democracy (Landemore). We extract legitimation principles from each model and evaluate whether SmartDemocracy incorporates advantages from both currents while mitigating their weaknesses. We complement with empirical evidence from partial experiences (deliberative polls, vTaiwan, participatory budgets) validating isolated components of (ICD), although no experience implements the complete model.

We apply law and economics analysis conceptually, but not formally. We invoke incentive theory (Downs's rational ignorance, Caplan's biases), transaction costs, and principles of institutional efficiency to ground the adopted design, but we do not mathematically model strategic equilibria nor quantify costs versus benefits.

Methodological transparency requires making explicit what we do not do. This article does not conduct original empirical research through surveys, experiments, or statistical analysis. We do not implement computational simulations of Intelligent Cognitive Democracy architecture nor formally model strategic equilibria. We do not perform an exhaustive historical analysis nor in-depth case studies. We do not quantify implementation costs nor calculate expected benefits in terms of policy improvements or institutional trust increases.

We recognize multiple limits. We propose a theoretical constitutional model without direct empirical testing. No complete (ICD) pilot has been implemented, although isolated components exist in experiences such as «deliberative polls» or Taiwanese digital platforms. Political viability of adoption faces formidable obstacles: established parties would lose power, political elites would have incentives to block reforms that increase the accountability, and public opinion might resist to some architectures perceived as complex or technocratic. We do not formally model strategic incentives of political actors under SmartDemocracy nor predict long-term equilibria. Finally, our proposal is framed in European constitutional traditions and might require significant adaptations for non-Western contexts.

3. Theoretical foundations and legitimacy in Smartdemocracy

To understand the institutional architecture of Intelligent Cognitive Democracy and assess its potential, we must first analyze its theoretical foundations. In a context of profound global democratic retreat, where liberal democracies face an existential legitimacy crisis, we must confront the elementary questions of political philosophy with renewed urgency: What justifies state authority in the twenty-first century? What makes a democratic system genuinely legitimate? Why would citizens have a moral duty to obey its decisions? This section unpacks the core concepts—authority, legitimacy, and obligation—that provide the legal-philosophical justification for the proposed model, presenting it as a robust response to the epistemic, functional, and procedural deficits afflicting contemporary democracies.

3.1 Core Concepts: Authority, Legitimacy, and Political Obligation

This section lays the legal-philosophical foundations upon which the institutional edifice of (ICD) rests. We will systematically analyze the distinction between authority (the right to govern and issue directives), legitimacy (the moral justification validating that right), and political obligation (the citizen's duty to obey such directives). By clarifying these pillars, we can articulate why SmartDemocracy is both an innovative institutional proposal

and a theoretically justified conception of democracy for the twenty-first century.

We define political authority as the institutional right of a system to issue binding directives and demand obedience. We adopt Joseph Raz's theory, which holds that authority functions as a «pre-emptive reason»: a reason that displaces and replaces the reasons on which it depends (Raz, 1986, ch. 3, sec. 4, pp. 41-42, 57-58). This idea crystallizes in the «pre-emption thesis»: authoritative directives do not add to an individual's balance of reasons but replace those reasons which they are meant to reflect (Raz, 1986, ch. 3, sec. 4, pp. 57-59). When a law prohibits parking in a designated zone, citizens must obey not because they have individually evaluated all underlying reasons (traffic congestion, pedestrian safety), but because the authoritative directive pre-empts their own calculation, enabling efficient social coordination. However, authority is not absolute. Raz warns about the importance of establishing clear limits, derived from the intrinsic desirability of people conducting their own life by their own lights (Raz, 1986, ch. 3, p. 56). For a directive to function as a valid pre-emptive reason without depending on mere force, it must emanate from a system citizens recognize as morally justified, thus introducing the necessity of legitimacy.

Why does a particular political system deserve recognition as valid authority? Democratic theory has traditionally offered two answers that, as Estlund notes, are insufficient separately: instrumentalist theories justify democracy by its capacity to produce good results; proceduralist theories justify it by the intrinsic value of its procedures, mainly equal participation, regardless of result quality (Estlund, 2008). Both positions, when absolutized, reveal fundamental shortcomings. Pure instrumentalism could justify efficient but oppressive regimes, while pure proceduralism cannot respond to the objection that a system with formally equal participation but systematically erroneous decisions eventually erodes its own popular support base.

Facing this false dichotomy, (ICD) postulates a three-dimensional legitimacy model requiring simultaneous satisfaction of three mutually reinforcing conditions: procedural legitimacy (decisions must emanate from procedures treating all citizens as political equals, respecting universal suffrage and participation rights without exclusions); epistemic legitimacy (decisions must systematically incorporate the best available evidence and specialized knowledge to produce informed political judgments); and functional legitimacy (public policies must produce verifiable results improving citizens' living conditions, with institutional mechanisms for correction and learning). This three-dimensional model reconciles demands for equality, knowledge, and efficacy, creating a robust legitimacy standard adapted to the complexity of contemporary challenges.

We define political obligation as the moral duty to obey laws and collective decisions, distinguishing it from obedience through mere prudence or unreflective habit. We reject two extreme positions: voluntarism, which holds that only explicitly consented obligations are morally binding; and natural authoritarianism, which claims certain institutions possess inherent authority, contradicting the democratic principle of equal moral dignity. From

SmartDemocracy's framework, we contend that political obligation emerges directly from triple legitimacy. When a system simultaneously satisfies procedural, epistemic, and functional conditions, it generates a strong presumption of obedience that can only be defeated by extraordinary reasons, such as manifest injustice. This conception explains why civil disobedience may be justified before regimes persistently failing in any of the three dimensions without collapsing into anarchism that rejects all forms of political obligation.

3.2 The Legitimacy Triangle: A Three-Dimensional Model for Twenty-First Century Democracy

We can now deepen into the normative heart of (ICD): the three-dimensional legitimacy vector. This section will theoretically ground each of the three dimensions—procedural, epistemic, and functional—demonstrating that they are complementary and interdependent.

Procedural legitimacy constitutes a fundamental and irreducible normative pillar of SmartDemocracy. It is grounded in the principle of equal moral dignity of all citizens: every adult citizen has an inalienable right to participate in collective decisions affecting them, regardless of knowledge level, wealth, or social status. This conception is rooted in a solid democratic tradition: Jefferson's principle of radical political equality, for whom government is only secure if all the people participate in ultimate authority (Jefferson, 1785/1982, Query XIV); Paine's argument (1776) that representation is a practical necessity and not a concession to elite superiority; and Madison's warnings (1788) about the risk that representatives develop interests alien to those they represent.

However, in the twenty-first century digital context, formal procedural equality is insufficient. A system where citizens vote with equal power but base their decisions on information manipulated or filtered by opaque algorithms does not fully realize the democratic ideal. As Duberry argues, popular sovereignty is threatened by digital platforms' «echo chambers» and filter bubbles, which together with the erosion of «informational privacy» undermine society's capacity to make genuinely informed decisions (Duberry, 2022, cap. 4). For formal equality to be effective, it must be complemented with a dimension guaranteeing information quality.

Epistemic legitimacy responds to the requirement that collective decisions incorporate the best available knowledge about facts, causal relationships, and foreseeable policy consequences. This is not about seeking absolute truth, but designing institutional procedures facilitating access to quality evidence and promoting informed deliberation. We adopt Estlund's concept of «epistemic proceduralism», according to which democracy is justified not only by the intrinsic equality of its procedure, but also because it tends to produce epistemically superior decisions to politically acceptable alternatives (Estlund, 2008). This last clause is crucial: it discards pure epistocracy because such experts' authority can be reasonably rejected by citizens. Therefore, knowledge must be institutionalized through transparent

mechanisms that organize expertise without creating permanent castes.

The tension between equality and competence traverses political philosophy. Epistocratic vision has great representatives: Plato distrusted the demos for its inability to discern the common good (Platón, ca. 380 a.C./1992); Mill suggested giving more votes to those demonstrating greater education (Mill, 1861/1991); while Brennan updates classical epistocracy by arguing that democracy fails when ignorant citizens decide on complex policies, proposing filters that exclude or weight votes according to demonstrable knowledge (Brennan, 2016). However, as Moraro warns (2018, p. 209), although education improves the capacity to govern wisely, it is not unreasonable to suspect that the educated group will present other biases—of class, race, gender—that could cause more epistemic harm than their education contributes.

The Cognitive Diversity perspective led by Landemore responds from another angle: the demos's cognitive diversity—when organized deliberatively through appropriate institutions—systematically surpasses any homogeneous elite, because it aggregates perspectives, experiences, and mental models that no isolated expert possesses (Landemore, 2020, cap. 4, pp. 89-112). Crowd wisdom is not magical but results from specific conditions: judgment independence, genuine diversity, appropriate aggregation mechanisms, and absence of informational cascades that prematurely homogenize opinions through mimetic contagion (Sunstein, 2019).

(ICD) resolves this secular tension not by choosing one pole over the other, but by integrating them into a superior architecture. The institutional architecture channels specialized knowledge, the experience of those directly affected, and perspectives of marginalized communities into deliberative instances without coercive power. These bodies produce informative syntheses that guide, but never replace, the citizenry's final decision. Thus, the Kantian ideal of public use of reason is institutionalized and the «computational deficit» that Duberry (2022, cap. 5) identifies in current democracies is addressed.

Functional legitimacy is the model's consequentialist dimension: a political system must produce results that verifiably improve citizens' living conditions. It is insufficient for decisions to be egalitarian and informed if they systematically fail to resolve the problems they intend to address. Functional legitimacy requires implementing institutional mechanisms to systematically track policy performance, learn from errors through continuous evaluation, and correct ineffective policies before they petrify through inertia. Traditional democracies often lack these architectures for adaptive learning. SmartDemocracy materializes this dimension through continuous evaluation systems and automatic «sunset clauses» that trigger review or termination of policies failing to meet their objectives.

The three dimensions of legitimacy are equally necessary and exist in a relationship of interdependence. None can be subordinated to the others without provoking democratic system collapse. Sacrificing equality for efficiency inevitably degenerates into authoritarian technocracy, emptying popular sovereignty of content. Ignoring knowledge for equality

produces mediocre decisions based on «rational ignorance» (Downs, 1957) or «systematic irrationality» (Caplan, 2007), feeding citizen disaffection. Neglecting results for procedural purity generates cynicism and erodes trust in democratic institutions. (ICD) operationalizes this triple legitimacy through tricameral architecture that functionally separates instances responsible for preserving popular sovereignty, organizing knowledge, and evaluating results, coordinating them through procedures ensuring no instance usurps others' authority. This functional separation updates Montesquieu's principle that power must check power (Montesquieu, 1748/1995) and Madisonian reciprocal controls (Madison, 1788/2003), extending them to the epistemic and functional realms.

3.3 SmartDemocracy's Cognitive Model: How Information Improves Collective Decisions

The central argument of Intelligent Cognitive Democracy is that greater quality and accessibility of relevant information improves democratic decisions, but only under specific institutional conditions designed to mitigate biases and disinformation. Our cognitive model rests on three fundamental premises: political decisions involve predictions about future consequences; better information tends to generate better predictions, as empirical studies on deliberative polls demonstrate where citizens accessing balanced information and deliberating among peers change their opinions toward positions more coherent with evidence, as shown in the participatory budgeting process in Zeguo, China (Fishkin, 2009, cap. 4, pp. 95-125); and cognitive diversity, well organized, surpasses concentrated knowledge, since a diverse group often surpasses individual experts because it aggregates perspectives no individual possesses alone (Landemore, 2020, cap. 4, pp. 89-112).

Condorcet's Jury Theorem formalizes this intuition: if voters have a probability greater than 50% of finding the correct answer and vote independently, the probability that the majority is correct approaches 100% as the group grows. However, this theorem is vulnerable in the real world. Sunstein has identified the danger of «informational cascades» a phenomena in which individuals violate the independence condition by ignoring their own information to follow others' judgment (Sunstein, 2019), which can lead to rapid convergence toward false beliefs, effects that social networks amplify through «echo chambers» (Centola, 2022, pp. 932-934). (ICD)'s institutional architecture is designed precisely to counteract these cascades and biases, ensuring structural epistemic diversity and providing access to quality information to leverage collective wisdom while mitigating its vulnerabilities.

This section has presented the normative foundations upon which Intelligent Cognitive Democracy is constructed. We have defined a robust conception of three-dimensional legitimacy, requiring simultaneous satisfaction of procedural (political equality), epistemic (informed knowledge), and functional (effective results) criteria. This triad surpasses traditional theories' limitations and offers a more complete justification for democratic authority in the twenty-first century. This model does not sacrifice political equality

on the altar of expertise, but strengthens it. Facing pessimistic diagnoses viewing citizen incompetence as an intractable problem, SmartDemocracy postulates a more optimistic vision: equipping citizens with institutional tools so that popular power is not only formally sovereign, but also collectively intelligent and effective. Upon this solid theoretical foundation, it is now possible to construct and detail (ICD)'s institutional architecture.

4. Institutional architecture of SmartDemocracy (ICD model)

The tricameral architecture we propose updates constitutional principles whose efficacy has been historically proven. Montesquieu (1748) formulated the separation of powers; Madison (1788) developed the system of checks and balances; Kant (1784) and Rawls (1993) grounded public reason; Dworkin (1977) established rights as insurmountable limits to power. Our innovation lies in operationalizing these principles through three contemporary instruments: The Explainable Artificial Intelligence (XAI), temporary micro-credentials that certify relevant knowledge, and continuous evaluation of results. The question is whether these new instruments can materialize promises that classical constitutionalism formulated but could not fully guarantee.

4.1 Popular Chamber (Congress): Direct Sovereignty

The Popular Chamber constitutes the core of democratic sovereignty, integrating all adult citizens through verified digital identity. Each citizen obtains a cryptographic credentials that guarantee vote authenticity and ballot secrecy, resolving the technical contradiction between verifiability and anonymity.

Membership in the Popular Chamber is automatic and unconditional, with no education, wealth, or competence requirements. This decision responds to Brennan's epistocratic debate (2016, ch. 2, pp. 26-30), who documents that fewer than 30% of Americans can name two First Amendment rights and proposes conditioning suffrage on demonstrated political competence. However, as Landemore (2013, ch. 7, pp. 200-201) notes, the standard for measuring citizen competence proves problematic: it is unclear that one needs extensive knowledge of social sciences or constitutional treaties to judge whether policies benefit citizens.

The Popular Chamber exercises four cardinal functions. First, legislative initiative: any proposal backed by 0.5% of the census automatically enters the process. Fishkin's experience (2009, ch. 4, pp. 106-111) in Zeguo, China, demonstrates that participatory budgeting can achieve complete transparency and effectively implemented decisions, overcoming limitations of traditional representation.

Second, final ratification through direct voting eliminates the Madisonian risk that representatives develop interests distinct

from the represented (Madison, 1788). No law takes effect without explicit approval from the Popular Chamber, transforming Jefferson's concept "all power derives from the people" (Jefferson, 1801) from abstract principle into concrete practice. As Paine (1776) argues, representative government is necessary only due to practical limitations, not superiority over direct participation.

Third, veto power: the Popular Chamber can veto any decision from the other chambers through referendum convened when 1% of the census requests it, introducing what Rosanvallon (2008) terms «counter-democracy» as citizens' capacity to resist and correct deviated institutional decisions. Veto approval requires simple majority, recognizing the epistemic value of status quo as accumulated knowledge according to Habermasian communicative rationality (Habermas, 1996).

Fourth, political oversight through constructive vote of no confidence: an absolute majority can remove the Executive President with a reasoned motion, while simultaneously electing a successor until mandate completion.

Majorities are differentiated according to decision type and risk level across three tiers. Routine decisions require simple majority without quorum, recognizing Downs's «rational ignorance» (1957, pp. 207-259): given the cost of becoming informed, abstention on low-impact issues is rational. Important decisions (military intervention, debt exceeding 3% of GDP) demand 60% with 30% quorum, preventing «informational cascades» where decisions are adopted through contagion rather than deliberation (Sunstein, 2019). Critical decisions or structural reforms (fundamental rights modification, binding treaties, constitutional changes) require 70% with 40% quorum, responding to the Rawlsian concept of overlapping consensus: fundamental changes require agreement transcending the particular comprehensive doctrines (Rawls, 1993).

Three complementary participation mechanisms operationalize this design. First, secure electronic voting with cryptographic credentials guaranteeing authenticity, secrecy, and traceability, allowing voting from any device during minimum 72-hour periods. Second, an assisted digital deliberation: before each vote, the platform presents neutral synthesis from the Cognitive Chamber (maximum 2,000 words), impact simulations from the Synthetic Chamber, and factually verified arguments. Third, rotating deliberative mini-publics: a lottery selects representative samples of 500-1,000 citizens who deliberate in person with expert access, publishing recommendations alongside official synthesis.

This design maintains effective popular sovereignty while avoiding two problematic extremes. It avoids «delegative democracy» where representatives exercise discretionary power without continuous accountability. It also avoids «manipulable referendism» where loaded questions, strategic timing, or disinformation campaigns distort popular will.

4.2 Cognitive Chamber (Senate): Specialized Knowledge

The Cognitive Chamber represents an institutional innovation that organizes specialized knowledge, lived experience, and direct affectation without usurping the Popular Chamber's decision-making authority. Its exclusive function is producing neutral evidence syntheses, drafting legislative projects at the Popular Chamber's request, evaluating technical coherence of proposals, and facilitating informed deliberation. This functional separation resolves the contradiction between democratic equality and epistemic quality.

The Chamber comprises citizens bearing «temporary microcredentials» that accredit a relevant knowledge, verifiable direct affectation, or lived experience in impacted communities. Three independent pathways grant microcredentials, responding to three legitimate forms of epistemic authority. The Academic/professional merit pathway (M) contributes up to 3 points, recognizing knowledge acquired through formal education, scientific research, or accredited professional experience. The Direct Impact pathway (A) contributes up to 3 points, recognizing the special interest of those who will be directly affected by the rules to be drafted. The Lived experience pathway (E) adds up to 3 points, recognizing situated knowledge from historically marginalized communities whose perspective is epistemically valuable but frequently ignored.

These three pathways capture three distinct but equally legitimate forms of knowledge relevant to political decisions. Academic merit provides theoretical understanding of causal mechanisms and aggregated empirical evidence. Direct affectation provides knowledge of concrete consequences and preferences of those living policies daily. Lived experience provides perspectives challenging implicit assumptions of privileged majorities and preventing systematic epistemic blindness.

Conditional lottery operates as an anti-capture mechanism ensuring substantive pluralism. When qualified applicants exceed available seats (variable composition between 150-500 members depending on topic complexity), selection operates through weighted lottery guaranteeing a minimum representation from each M/A/E pathway (minimum 25% each), disciplinary diversity within M, territorial and socioeconomic diversity. This lottery prevents academic elites or organized interest groups from capturing deliberation through overrepresentation.

Strict temporality establishes that all microcredentials expire automatically after 18 months without possibility of automatic renewal. This temporality prevents formation of permanent castes of «certified experts» monopolizing specialized deliberation.

The Cognitive Chamber produces three document types. First, neutral evidence syntheses (maximum 2,000 words) presenting the state of knowledge, identifying robust consensuses, legitimate disagreements, and gaps. Second, technical coherence evaluations identifying internal inconsistencies in legislative proposals, unanticipated impacts, and implementation risks. Third, reports that do not determine whether proposals are politically good or bad but

whether they are technically viable and coherent with declared objectives.

4.3 Synthetic Chamber: Explainable Artificial Intelligence

The Synthetic Chamber constitutes an Explainable Artificial Intelligence system (XAI) under permanent human supervision, with exclusive function of simulating impacts of alternative policies, monitoring real-time implementation, and issuing risk opinions on legislative proposals. "Incorporating advanced computational capacity into democratic architecture does not constitute a type of algorithmic technocracy but rather the amplification of collective capacity to predict the possible consequences of norms still in the drafting phase.

The Synthetic Chamber lacks coercive power. It cannot veto laws, impose policies, or substitute human decision. Its authority is exclusively epistemic: when models predict that policy will produce effects significantly different from those promised, it issues reasoned alert that deliberators can accept, reject, or weigh. This architecture preserves human sovereignty while leveraging advantages of automated processing: speed, consistency, capacity to integrate massive evidence volumes, and absence of motivated biases.

«Risk Coherence Opinions (RCO)» evaluate legislative proposals before final vote, classifying them into three categories: Low Risk (green) for policies with predictable effects based on robust evidence, quickly reversible, with limited impact; Medium Risk (yellow) for policies with partially predictable effects, mixed evidence, moderate reversibility, significant distributive impact; High Risk (red) for policies with highly uncertain effects, contradictory or absent evidence, potentially irreversible consequences, systemic impact.

«Executive Performance Opinions (EPO)» monitor implementation of approved policies, comparing observed results with objectives declared in Programmable Execution Pacts. When deviation between goal and reality exceeds predefined thresholds (15% for ordinary policies, 10% for structural reforms), SmartDemocracy issues alert activating mandatory review. Three consecutive alerts in quarterly periods trigger automatic sunset clause: policy expires unless the Popular Chamber explicitly renews it after deliberation on failure causes.

This modulation of majorities according to certified risk level reconciles decisional agility versus protection against costly errors. Low-risk policies (green RCO) are approved with simple majority, facilitating experimentation and iterative learning. Medium-risk policies (yellow RCO) require qualified majority of 55%. High-risk policies (red RCO) require reinforced majority of 60% plus extended deliberation period (minimum 90 days), ensuring that commitments with potentially irreversible consequences have robust support and sustained deliberation.

Five algorithmic transparency requirements guarantee auditability and democratic control. First, meaningful Explainability Artificial Intelligence (XAI): every opinion includes explanation in natural language about which factors most influenced prediction, what assumptions underlie it, and what limitations exist.

Second, open and audited source code: algorithms, training data, and methodologies are public, allowing verification of bias absence. Third, mandatory human oversight: no opinion is published automatically. A rotating committee of 21 citizens selected by lottery reviews each opinion, can request additional explanations, and demands corrections if errors are detected. Fourth, periodic adversarial audits: quarterly, external teams attempt to identify biases through «algorithmic red teaming». Findings are published entirely and generate correction obligation within 60 days. Fifth, absolute prohibition of political scoring: (ICD) cannot evaluate individual citizens, parties, or social movements. Its exclusive object is abstract public policies.

4.4 Algorithmic Tricephalous Government: Democratic Legitimacy and Technical Competence, Assisted by Explainable Artificial Intelligence

The tricephalous executive power reconciles tension between democratic legitimacy (elected President) and technical competence (selected Rector), with advice from the Synthetic Council (SynC)—revolutionary component of the ATG—that materializes constant XAI support through explainable algorithmic systems under rotating human supervision.

Presidential election through two rounds guarantees what Madison (1788) considered essential: that whoever governs has explicit majority support, not merely relative plurality. The three-year mandate responds to Jeffersonian balance between governmental continuity and frequent democratic renewal (Jefferson, 1785/1982). The prohibition of immediate reelection prevents gradual democratic erosion through power perpetuation.

The Rector selection process through competition evaluated by the Cognitive Chamber introduces what King (2003, p. 27) describes as «competence without epistocracy». The criteria—40% technical competence, 40% executive experience, 20% social understanding—respond to the criterion that democratic legitimacy cannot depend exclusively on correct epistemic results. The requirement of presentation before the Popular Chamber with possibility of 60% veto adds democratic filter to avoid pure technocracy.

Functional division responds to Montesquieuian separation of powers (Montesquieu, 1748), applied internally to the executive. The President exercises functions where direct electoral legitimacy is crucial, while the Rector manages where technical competence is determinant. This division prevents tribal decisions based on identity rather than knowledge.

The co-decision mechanism for strategic matters, with conflict resolution through Cognitive Chamber reports and Synthetic Chamber simulations, and final decision by the Popular Chamber, operationalizes Rawlsian «public reason»: decisions must be justified with reasons all reasonable citizens can accept (Rawls, 1996). When President and Rector disagree, the process clarifies whether conflict concerns values (where democratic legitimacy prevails) or efficacy (where technical evidence is relevant).

Structural integration of XAI does not constitute a technocracy but a radical democratization of complex analysis. SynC lacks definitive coercive power—the Popular Chamber has final decision—but introduces mandatory epistemic friction preventing impulsive or captured decisions. Mandatory explainability transforms algorithmic «black box» into mass civic pedagogy tool.

The ATG eliminates rigid ministries, substituting them with «program managers» charged with specific missions. These managers are appointed by consensus of President and Rector from winning slates in competence tournaments supervised by the Cognitive Chamber, with variable mandates (6-24 months) not synchronized with ATG cycle to maintain operational flexibility.

States of exception require ATG unanimity for provisional activation for maximum 72 hours; after this period, they require immediate Cognitive Chamber review on necessity and proportionality, while (ICD) simulates alternative scenarios and sets maximum temporal ceiling based on historical evidence. Extension proceeds only with reinforced 2/3 majority in the Popular Chamber. The system guarantees that no state of exception can extend beyond two partial renewals (16 months absolute maximum). All measures are recorded in immutable blockchain for subsequent audit.

Budget abandons departmental logic for allocation to missions with verifiable objectives. Public procurement operates through transparent auctions with blind tests, complete digital traceability, and «most favored citizen» clause: when SynC detects more advantageous conditions granted to any supplier, these apply retroactively to all similar contracts, preventing hidden favoritism.

Differentiated revocation recognizes distinct sources of legitimacy. The President can only be recalled popularly, preserving what Jefferson (1801) considered sacred: direct popular will. The Rector can also be removed for technical noncompliance (fewer than 50% of indicators for 6 months), recognizing that technical competence requires verifiable results, not just credentials.

In summary, the institutional architecture we propose institutionalizes Aristotelian prudence, Montesquieuian separation, and Dworkinian limits. The Popular Chamber exercises ultimate sovereignty with rules making political equality effective. The Cognitive Chamber organizes knowledge without creating castes. SmartDemocracy subjects AI to public reason and protects rights as absolute limits. The Tricephalous Executive transforms government into measurable missions with timely corrections. Mechanisms of transparency, continuous evaluation, and revocation ensure that power is exercised communicatively, not strategically.

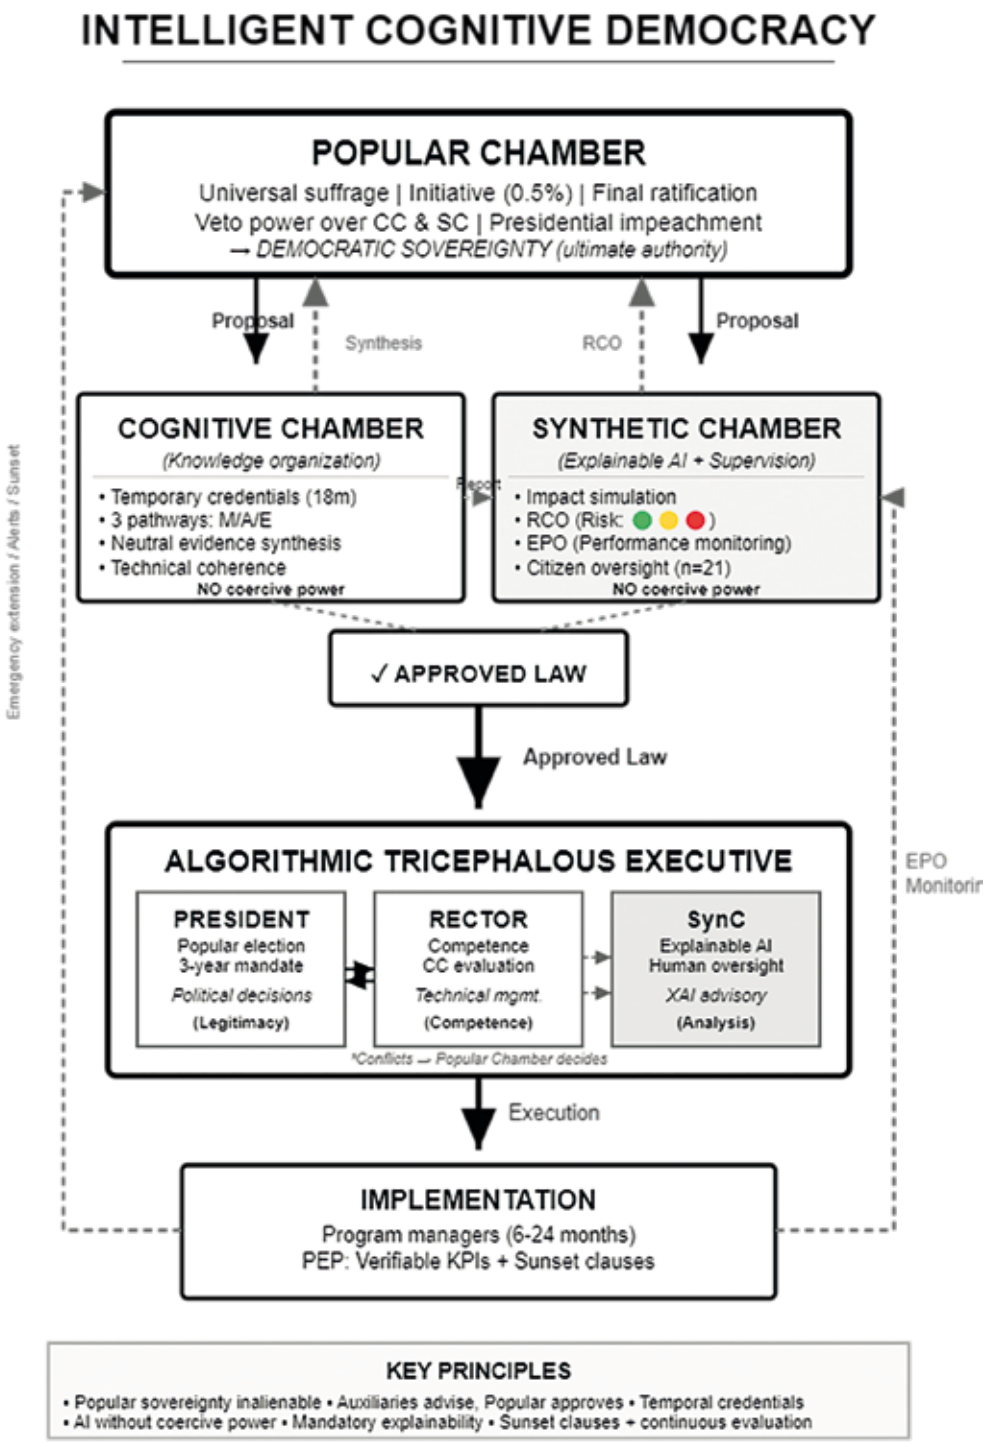
INTELLIGENT COGNITIVE DEMOCRACY
POPULAR CHAMBER
Universal suffrage | Initiative (0.5%) | Final ratification
Veto power over CC & SC | Presidential impeachment
→ DEMOCRATIC SOVEREIGNTY (ultimate authority)
Proposal
Synthesis
RCO
Proposal
COGNITIVE CHAMBER
(Knowledge organization)
• Temporary credentials (18m)
• 3 pathways: M/A/E
• Neutral evidence synthesis
• Technical coherence
NO coercive power
Report
SYNTHETIC CHAMBER
(Explainable AI + Supervision)
• Impact simulation
• RCO (Risk:)
• EPO (Performance monitoring)
• Citizen oversight (n=21)
NO coercive power
Emergency extension / Alerts / Sunset
✓ APPROVED LAW
Approved Law
ALGORITHMIC TRICEPHALOUS EXECUTIVE
PRESIDENT
Popular election
3-year mandate
Political decisions
(Legitimacy)
RECTOR
Competence
CC evaluation
Technical mgmt.
(Competence)
SynC
Explainable AI
Human oversight
XAI advisory
(Analysis)
*Conflicts → Popular Chamber decides
EPO
Monitorir
Execution
IMPLEMENTATION
Program managers (6-24 months)
PEP: Verifiable KPIs + Sunset clauses
KEY PRINCIPLES
• Popular sovereignty inalienable • Auxiliaries advise, Popular approves • Temporal credentials
• AI without coercive power • Mandatory explainability • Sunset clauses + continuous evaluation

5. Fundamental rights and algorithmic layer: guarantees and limits

The incorporation of artificial intelligence in democratic processes generates an unprecedented constitutional tension: how to preserve fundamental rights when political decisions incorporate algorithmic systems that operate through opaque logics for ordinary citizens? This section establishes the normative framework that articulates this apparently difficult coexistence. First, we will specify structural principles and specific digital rights adapted to algorithmic risks. Second, we will define mandatory standards for explainability and meaningful human review of AI. Third, we will establish absolute prohibitions as a hard and inviolable core. Finally, we will design the complete traceability architecture that transforms artificial intelligence from a potential threat into a tool for democratic deepening.

5.1 Classical and Digital Rights in the ICD Architecture

Intelligent Cognitive Democracy recognizes the complete catalog of fundamental rights from liberal constitutional tradition with immediate direct efficacy. Incorporating artificial intelligence into institutional architecture, particularly through the Synthetic Chamber, does not weaken these rights but demands their reinforcement through specific guarantees adapted to algorithmic risks.

We establish five structuring principles articulating the relationship between classical rights and the algorithmic layer. First, the «principle of democratic non-regression»: no automation can reduce the level of participation, transparency, or citizen control existing in an analog system. Second, the «principle of reinforced transparency»: when a political decision incorporates algorithmic analysis, the required transparency level exceeds that of purely human decision, because technical opacity could conceal biases evident in human deliberation. Third, the «*pro persona* and *pro democratia* principle»: facing normative ambiguity regarding rights' scope against algorithmic systems, the most favorable interpretation to person and democratic process prevails. Fourth, the «principle of universal revisability»: every decision affecting rights must be revisable through procedure allowing argument presentation, contrary evidence contribution, and reasoned resolution from independent human instance. Fifth, the «principle of strict proportionality»: restrictions on rights through algorithmic systems require more rigorous justification than traditional restrictions, because automation magnifies systematic error risks and facilitates mass surveillance.

These principles prevent two symmetric risks: naive techno-utopianism assuming automation as neutral, ignoring its potential to amplify biases and enable authoritarian surveillance through massive political data processing (Veale & Zuiderveen Borgesius, 2021); and paralyzing techno-pessimism rejecting any use of artificial intelligence in governance, depriving democracy of tools that could improve decisional quality and strengthen evidence-based accountability.

We specify a catalog of specific digital rights complementing classical rights. The

right to verified but pseudonymous digital identity guarantees political participation authenticity while protecting privacy through dissociation between civil identity and specific political actions, implemented through unique cryptographic credentials. The right to meaningful algorithmic explanation requires that when a political decision incorporates Synthetic Chamber analysis, citizens access comprehensible explanation about influential factors, underlying assumptions, processed data, and recognized limitations, using interactive visualizations and verifiable natural language. The right to participation data portability allows complete download of democratic participation history in interoperable standard format. The right to modulated political rectification and erasure operates in nuanced manner: algorithmic predictions cannot be based indefinitely on obsolete data (we establish five-year expiration for data used in predictive modeling), but historical memory of adopted political decisions is preserved for democratic accountability. The right to neutral comparative information guarantees access to Cognitive Chamber neutral syntheses, presenting arguments for and against with equal depth, both subjected to equivalent fact-checking, preventing manipulation through asymmetric presentation.

Luciano Floridi has argued that the regulation of artificial intelligence must be grounded in an 'infosphere' that integrates human values into technical design (Floridi, 2021). We agree with the cited author, but we maintain that this requires going a step further through concrete institutional specification that translates abstract principles into verifiable procedural guarantees and effective remedies. For this reason, we have proposed this catalog of rights that addresses the particular risks of algorithmic governance.

[5.2] Mandatory Explainability and Meaningful Human Review

The incorporation of artificial intelligence into democratic processes requires a fundamental guarantee: no algorithmic system can influence political decisions without explaining its recommendations in a manner comprehensible to citizens without specialized technical training. This principle responds to the risk of algorithmic «black boxes», where decisions emerge from opaque processes that make it difficult to detect biases, errors, or manipulations (Pasquale, 2015). Consider an example: if the Synthetic Chamber recommends rejecting a social housing policy by arguing that "historical data show low effectiveness," but fails to reveal that these data come exclusively from deliberately underfunded programs, technical opacity conceals political bias under the appearance of scientific neutrality.

We propose three progressively detailed levels of explainability. Basic explainability consists of a summary of no more than five hundred words that answers four elementary questions: What question did the system attempt to answer? What information did it use? Which factors weighed most heavily in its recommendation? What limitations does it have? If the system evaluates a pension reform, it must specify whether it used data from ten or thirty years, whether it included crises or only growth periods, and whether it assumed that life expectancy will continue to increase.

These specifications have direct political consequences: projecting thirty years using only growth data will produce recommendations radically different from including cyclical recessions. Technical explainability provides complete documentation—the model architecture, datasets with dates and sources, source code—and must be accessible within five business days upon request by any citizen. If a trade union suspects that the algorithm underestimates the impact of a labor reform on specific sectors, it can commission an independent audit that identifies biases in the data or in the model's assumptions. Counterfactual explainability offers, for contested decisions, explanations of the type "if conditions had been X instead of Y, the recommendation would have been Z." If the algorithm predicts the failure of a rural education policy, counterfactual explanation may reveal that the determining factor was low digital connectivity, not the characteristics of the population; this distinction is crucial because the infrastructural problem is solvable through public investment.

However, explainability requires meaningful human review: oversight exercised by persons with real authority to reject, modify, or overturn algorithmic recommendations. We propose three institutional mechanisms. First, a rotating committee of twenty-one citizens selected by lottery reviews each opinion before publication, with partial quarterly renewal. To verify that it exercises real oversight, if the rejection rate falls below five percent for two consecutive quarters, an external audit is triggered that examines whether the committee developed excessive deference toward algorithmic recommendations. Second, economic incentives to detect algorithmic errors: creation of a civic bug bounty program where any citizen who identifies a confirmed methodological error receives financial compensation (v.gr., from ten thousand to five hundred thousand euros/dollars) according to the magnitude of the detected error, transforming algorithmic auditing into a shared responsibility that exponentially multiplies detection capacity. Third, public registry of disagreements between algorithm and humans with biannual evaluation: each time the committee modifies or rejects a recommendation, the disagreement is recorded along with the human justification. Audit teams analyze these disagreements to identify systematic biases and propose methodological corrections.

5.3 Absolute Prohibitions as Constitutional Hard Core

Five absolute prohibitions delimit insurmountable limits of artificial intelligence use, operating as unmodifiable constitutional hard core because they protect the conditions of possibility themselves of democratic regime. These prohibitions apply absolutely to SmartDemocracy and to any algorithmic system that might be incorporated in future into (ICD) institutional architecture.

First prohibition: political scoring of citizens. No algorithmic system can assign scores, ratings, or political profiles to individual citizens based on their expressed opinions, voting history, participation in deliberations, or any manifestation of political preferences. Empirical evidence from Chinese social credit system demonstrates that political scoring inevitably derives into authoritarian

control tool punishing legitimate dissent (Liang *et al.*, 2018).

Second prohibition: exclusively automated decisions on fundamental rights. No fundamental right can be restricted through decision adopted exclusively by algorithmic system without substantive human intervention with authority to modify or override decision. Automation can assist preliminary evaluation, but final decision on rights restriction must always rest with identifiable human authority, accountable and subject to judicial review. This prohibition materializes the principle of human dignity: persons have right to have crucial decisions affecting their life made by other human beings capable of empathy, contextual judgment, and consideration of particular circumstances that current algorithms do not capture (Hildebrandt, 2015).

Third prohibition: use of artificial intelligence for deliberative emotional manipulation. Algorithmic systems cannot be designed to exploit psychological vulnerabilities, cognitive biases, or emotional states with objective of tilting democratic deliberation toward predetermined conclusions. This includes microtargeting techniques presenting personalized arguments calculated to maximize emotional persuasion without consideration for veracity, and recommendation systems creating filter bubbles reinforcing preexisting beliefs while hiding contradictory evidence.

Fourth prohibition: deliberate opacity of critical governance systems. Developers of algorithmic systems intervening in democratic processes cannot intentionally design incomprehensible architectures, obfuscate code with purpose of hindering democratic audit, or employ obfuscation techniques preventing verification of undue bias absence. This prohibition does not demand that all code be trivially comprehensible but prohibits deliberate opacity: when methodological alternatives exist with equal predictive performance but greater interpretability, choice of opaque architectures must be publicly justified. Burden of proof rests on whoever designs opaque systems, not on citizens attempting to audit them (O'Neil, 2016).

Fifth prohibition: retroactive use of artificial intelligence to rewrite political history. Algorithmic systems cannot modify, delete, or alter historical records of democratic deliberations, votes, opinions, or political decisions once published. Cryptographic versioning with verifiable timestamps through public blockchain, which we have already discussed for states of exception and which we now generalize to the entire Synthetic Chamber through immutable recording with SHA-256 cryptographic hashes, ensures the immutability of historical records, preventing Orwellian revision of the past. This guarantee protects democratic memory as public good: capacity to evaluate past government performance, learn from institutional errors, and hold decision-makers accountable depends on reliable historical records (Duberry, 2022).

These prohibitions trace the red line between legitimate algorithmically assisted democracy and illegitimate technocratic authoritarianism. Democracy can and should harness artificial intelligence to improve information quality and facilitate mass deliberation without crossing the threshold where automation usurps popular sovereignty, suppresses dissent, or manipulates citizens' will.

5.4 Algorithmic Reinforced Justification

The principle of reinforced justification inverts the intuitive logic: when a political decision incorporates artificial intelligence analysis, the standard of justification required must be higher than that of a purely human decision. This inversion responds to three philosophical-political foundations. First, structural epistemic asymmetry: when a human legislator proposes a policy, citizens can directly evaluate their arguments; when an algorithmic system recommends, technical complexity blocks this ordinary democratic scrutiny, demanding compensatory transparency superior to the standard. Second, magnitude of systemic impact: a legislator who commits an error produces bounded consequences; an algorithm that replicates a bias operates at massive scale. If an official discriminates against social housing applicants, they commit localized injustices; if an algorithm incorporates variables that correlate with ethnicity, it replicates that discrimination systematically across hundreds of thousands of cases without specific human review. Third, prevention of technical capture: requiring the making explicit of normative assumptions prevents technical complexity from functioning as a shield against democratic scrutiny. When an opinion recommends fiscal austerity, it must make explicit whether it assumes that cuts will generate business confidence that will compensate for contractionary effects, or whether it privileges budgetary balance over economic stabilization.

Reinforced justification is operationalized through three mandatory components. The technical-methodological component specifies what analysis was performed, what information was processed (dates and verifiable sources), and how reliable the predictions are using comprehensible probabilities. Instead of stating that "tax reform will increase revenue by eight percent," it must specify "there is a seventy percent probability that revenue will increase between five and eleven percent, assuming that taxpayers do not substantially modify their behavior." The substantive-evaluative component translates technical results into political implications, explicitly identifying the normative judgments that mediate between prediction and recommendation. If the model predicts that a policy will reduce unemployment by two points but increase regional inequality by fifteen percent, it must make explicit: "This recommendation privileges aggregate efficiency over territorial equity. Those who prioritize territorial cohesion could reasonably reject it." The comparative component contrasts the method employed with reasonable alternatives. Consider the decision about where to build a hospital: one method might calculate with statistical precision how many lives would be saved but without revealing its reasoning; another would give less precise estimates but clearly show which factors it considers and how it weighs them. The first method is like a sophisticated calculator that gives the final result without showing the operations; the second like solving the problem on a blackboard. If we value maximizing numerical accuracy, the first is superior; if we value democratically auditing the decision, the second is preferable. This choice is not technical but political: it affects which errors we tolerate and who can challenge decisions.

This comprehensive architecture we propose aims to transform artificial intelligence from a potential threat into a tool for democratic deepening. Extreme transparency seeks to prevent technocratic capture. Rotating citizen oversight democratizes control over technical systems. Economic incentives mobilize distributed collective intelligence. Public registry of disagreements enables institutional learning. Reinforced justification compels making explicit the assumptions that in human decisions might remain implicit, incorporating specialized technical knowledge without sacrificing citizen political control over the fundamental directions of the political community.

6. Theoretical comparison and positioning

Every institutional proposal must position itself critically against existing alternatives and justify why it merits serious consideration. This section positions SmartDemocracy on the conceptual map of contemporary democratic theory. First, we contrast Intelligent Cognitive Democracy (ICD) with epistocracy and deliberative democracy, demonstrating how it overcomes the false dichotomy between epistemic exclusion and naive participatory optimism. Second, we document that components of ICD have been empirically tested through successful democratic experiments. Third, we defend gradual implementation that respects Tocquevillian prudence and distinguishes ideal normative justification from non-ideal institutional engineering. ICD is not invented *ex nihilo* but rather systematically integrates dispersed lessons from political philosophy and historical democratic practice.

6.1 Contrast with Epistocracy and Open Democracy

SmartDemocracy challenges centuries of representative democratic theory while recovering proven intuitions such as direct Greek voting and updating them with rigor. Where Brennan identifies insurmountable ignorance and proposes exclusion, (ICD) recognizes the deficit but organizes knowledge without amputating rights. Where Landemore celebrates radical openness and trusts pure cognitive diversity, SmartDemocracy accepts that certain technical expertise or lived experiences deserve specific deliberative channels.

The tension between political equality and epistemic quality has traversed political philosophy since Plato (Mladenovi-, 2020, p. 172). Plato distrusted the «demos» for its inability to discern the common good and proposed educated guardians to govern (Platón, 1992 [ca. 380 a. C.]). Mill suggested giving more votes to those who demonstrated greater education (Mill, 1861/1991). Jefferson, in contrast, defended radical political equality: influence over government must be shared among all the people, for if each individual participates in ultimate authority, government will be secure (Jefferson, 1785/1982, Query XIV). Brennan updates epistocracy by arguing that democracy fails when ignorant citizens decide on complex policies; his solution: filters that exclude or weight votes according to demonstrable knowledge (Brennan, 2016). Moraro argues that political competence alone

does not guarantee reliable voter behavior, since a selfish and competent voter might vote worse than a selfish and incompetent one by better identifying which candidate favors their particular interests over the common good (Moraro, 2018, pp. 208-209). Landemore responds that the cognitive diversity of the demos—when organized deliberatively—surpasses any homogeneous elite (Landemore, 2020).

(ICD) reconciles these positions through three-dimensional legitimacy: it preserves political equality in the Popular Chamber, channels expert knowledge and lived experience by fusing human deliberation with explainable artificial intelligence in the Cognitive Chamber, and evaluates risks and policy impacts through the Synthetic Chamber which lacks coercive power (Rask & Shin, 2024, pp. 2-6). Knowledge improves decisions without becoming permanent privilege.

This reconciliation recovers classical constitutional principles. Montesquieu taught that separating functions protects liberty: legislating, executing, and judging must be in different hands so that power checks power (Montesquieu, 1748/1995). Madison refined the idea through reciprocal controls between branches that prevent ambition from running wild (Madison, 1788/2003). SmartDemocracy extends this logic: it separates deciding, gathering knowledge, and translating complexity (Dorf & Sabel, 1998, pp. 267-268). Kant demanded that all authority justify its acts with public reasons that anyone can understand (Kant, 1784/1996); Rawls specified that these reasons must be acceptable to reasonable citizens even when they disagree (Rawls, 1996). The Cognitive Chamber operationalizes this ideal: its reports make evidence explicit, signal uncertainties, and separate facts from values. Habermas would add that rational deliberation requires communicative symmetry (Canale, 2022, p. 47; Habermas, 1996); (ICD) responds with procedural rules guaranteeing equivalent time for rebuttal and mandatory exposure to dissent.

Dworkin insisted that individual rights act as insurmountable limits against aggregated utility (Dworkin, 1977). The Synthetic Chamber incorporates this constraint: no technical opinion can endorse policies that injure the fundamental rights, and any alert of discrimination or arbitrariness activates institutional pauses.

6.2 Empirical Experiences: Deliberative Polling, vTaiwan, Participatory Budgets

Intelligent Cognitive Democracy does not emerge from a theoretical vacuum. Empirical experience shows that some isolated components already function correctly; what is lacking is their integration into a coherent constitutional framework (King, 2003, p. 27). Fishkin demonstrated through deliberative polls that ordinary citizens, when they receive balanced information and time to deliberate, shift their preferences toward more informed and nuanced positions (Fishkin, 2009). Representative groups deliberate for days, consult with experts who hold opposing positions, and produce considered judgments. These judgments subsequently influence actual policies, such as wind energy expansion in Texas or health infrastructure in China (King, 2003, p. 27; Roberts, 2023, p. 9). The Cognitive Chamber institutionalizes

this logic: temporary micro-credentials select participants with intellectual merit, direct stake, or lived experience. These participants are provided with verified information, and their syntheses guide the Popular Chamber.

Taiwan, with vTaiwan and the Pol.is platform, represented an experiment in massive digital deliberation: algorithms mapped emerging consensus among thousands of participants without forcing binary polarization (Moats *et al.*, 2024). Where social networks amplify divides, Pol. is identified areas of surprising agreement and recognized proposals formulated by minorities outside mainstream thinking as valuable contributions (Boucher *et al.*, 2020, p. 120; Roberts, 2023, pp. 19-51). DCI adopts this philosophy in its Synthetic Chamber: explainable AI simulates policy impacts, evaluates risks, and issues opinions on technical coherence, always with human oversight and public traceability.

Participatory budgets democratized local fiscal decisions: neighbors prioritize projects, oversee execution, and correct deviations (Menendez-Blanco & Bjørn, 2022; Fung & Wright, 2001, pp. 5-8). (ICD) scales these lessons: each Programmable Execution Pact includes verifiable KPIs, sunset clauses, and continuous audits by the Synthetic Chamber. The difference from isolated experiments: SmartDemocracy constitutionalizes institutional learning, converting correction into obligation rather than voluntary exception.

6.3 Gradual Constitutionalization and Theory/Implementation Separation

Epistemic concerns distinguish AI innovation from other ethically complex fields such as biotechnology, especially with opaque AIs like current deep neural networks (Buhmann & Fieseler, 2023, p. 150). Constitutionalizing (ICD) requires prudence. Tocqueville warned that institutions take root when born from local customs, not abstract impositions (Tocqueville, 1835-1840/2000). The sensible path would be to implement controlled pilot programs in sectoral policies, evaluate them subsequently and independently with clear metrics, and proceed to scale them up only if data support them. This democratic experimentalism—proposed by Sabel and Dorf—enables error correction without compromising the legitimacy of the entire system (Dorf & Sabel, 1998, pp. 267-270). Precedents exist: universal suffrage was gradual conquest; social rights were constitutionalized after decades of struggles and incremental reforms.

Rawls distinguished ideal theory (principles of justice under optimal conditions) from non-ideal theory (how to advance from reality toward justice) (Rawls, 1999). This article operates in the ideal register: it normatively justifies SmartDemocracy's three-dimensional legitimacy, its design principles, and its safeguards. Detailed institutional engineering—how to implement micro-credentials in diverse legal systems, what technical standards for explainable AI, how to design platforms for massive digital deliberation—corresponds to

complementary research that addresses specific contexts. Separating justification from execution is not evasion; it is methodological rigor: first, demonstrate that DCI is desirable and compatible with democratic values; then, build operational manuals adapted to different national realities.

7. Conclusions

This final section synthesizes the argument developed throughout the article. First, we recapitulate the main thesis: Intelligent Cognitive Democracy reconciles procedural, epistemic, and functional legitimacy through tricameral institutional separation. Second, we respond to the research questions posed by showing how ICD dissolves the false dichotomies between equality and knowledge. Third, we specify our original contribution compared to epistocracy and deliberative democracy. Finally, we acknowledge the study's limitations and identify the necessary future lines of research to move from theoretical principles to verifiable practical implementation.

7.1 Main Thesis Recapitulation

Intelligent Cognitive Democracy reconciles three dimensions of legitimacy that contemporary democracies treat as mutually exclusive. Procedural legitimacy is preserved through unalterable universal suffrage: the Popular Chamber retains ultimate decision-making power, veto, and political control without epistemic exclusions. Epistemic legitimacy is operationalized through the Cognitive Chamber, which organizes specialized knowledge, direct affectation, and lived experience via temporary micro-credentials (expire at 18 months), open through triple pathways (merit/affectation/experience), and subject to conditioned lottery preventing the formation of permanent castes. Functional legitimacy materializes through the Synthetic Chamber: explainable artificial intelligence systems simulate policy impacts, issue auditable risk opinions, and modulate majorities according to certified uncertainty, but lack coercive power. The Algorithmic Tricephalous Executive translates these legitimacies into governance through verifiable missions: Programmable Execution Pacts with public indicators, continuous evaluation, and sunset clauses that automatically terminate policies that fail to achieve their stated objectives. SmartDemocracy does not dilute popular sovereignty nor create oligarchic epistocracy: it institutionalizes knowledge under permanent democratic control.

7.2 Answers to Research Questions

First question: Is it possible to reconcile political equality, epistemic quality, and functional efficacy? The tricameral architecture demonstrates that it is. Institutional separation of functions (deciding/informing/evaluating) dissolves the false dichotomy between equality and knowledge. Every citizen votes with equal power in final decisions; some citizens rotate temporarily in specialized deliberation; no algorithm decides without mandatory human review.

Second question: Can cognitive filters improve decisions without creating castes?

Temporary micro-credentials (expire after 18 months) are non-hereditary, open through three pathways, and possibly subject to conditional lottery—all of which prevent aristocratic consolidation by design. Empirical evidence analyzed shows that informing citizens improves decision-making quality without eliminating pluralism.

Third question: What safeguards prevent technocratic degeneration? Five absolute prohibitions (political scoring, exclusively automated decisions on fundamental rights, deliberative emotional manipulation, deliberate opacity, retroactive rewriting of political history) plus mandatory explainability, meaningful human review, quarterly adversarial audits, and public source code guard that artificial intelligence assists without usurping sovereignty. The Popular Chamber retains power to revoke defective models, demand new audits, or reject opinions through democratic majority.

7.3 Original Contribution and Scope

SmartDemocracy as (ICD) aims to surpass both epistocracy (Brennan) and pure deliberative democracy (Landemore, Fishkin) through verifiable institutional synthesis. Where Brennan proposes excluding incompetent voters through restrictive examinations, SmartDemocracy organizes knowledge without amputating rights. Where Landemore trusts cognitive diversity without formal epistemic structure, (ICD) channels complementary expertise through auditable credentials. Innovation resides in operationalizing Aristotelian prudence, Montesquieuian separation, and Dworkinian limits through constitutional design that translates classical principles into explainable algorithmic architecture. We do not invent epistemic governance: we update the republican tradition with tools that Montesquieu, Madison, and Tocqueville would probably have employed if transparent artificial intelligence, digital lotteries, and automated evaluation of public policies had been available to them.

7.4 Limitations and Future Lines

This study establishes the theoretical constitutional framework without direct empirical testing. We design the complete institutional architecture of Intelligent Cognitive Democracy (micro-credential rubrics, pilot phases, indicator templates); its translation into operational protocols for specific sectoral implementation constitutes the object of complementary research in development. Future lines of research include computational simulations of the tricameral system, controlled deliberative experiments, comparative analysis of political feasibility in different constitutional contexts, and municipal pilots with rigorous evaluation.

Author's note: The English version of this manuscript was produced via an initial translation using Claude (v. Sonnet 4.5), followed by full human review by the author.

References

Bell, D. A. (2015). The China Model: Political Meritocracy and the Limits of Democracy. Princeton University Press.

Boucher, S., Hallin, A., Paulson, G., Dukát, M., Rask, M., & Shin, B. (2020). Collective intelligence, democracy, and governance: Handbook. Edward Elgar Publishing.

Brennan, J. (2016). Against democracy. Princeton University Press.

Buhmann, A., & Fieseler, C. (2023). Deep learning meets deep democracy: Deliberative governance and responsible innovation in artificial intelligence. Business Ethics Quarterly, 33(1), 146-179.

Canale, D. (2022). Cuando los expertos crean derecho: deferencia, opacidad y legitimidad. Discusiones, 30, 37-60.

Caplan, B. (2007). The myth of the rational voter: Why democracies choose bad policies. Princeton University Press.

Centola, D. (2022). Network science of collective intelligence. Trends in Cognitive Sciences, 26(11), 923–941.

Dorf, M. C., & Sabel, C. F. (1998). A constitution of democratic experimentalism. Columbia Law Review, 98(2), 267-473.

Downs, A. (1957). An economic theory of democracy. Harper & Row.

Duberry, J. (2022). Artificial intelligence and democracy: Risks and promises of AI-mediated citizen-government relations. Edward Elgar Publishing.

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

Dworkin, R. (1986). Law's empire. Belknap Press of Harvard University Press.

Estlund, D. (2008). Democratic authority: A philosophical framework. Princeton University Press.

Fishkin, J. S. (2009). When the people speak: Deliberative democracy and public consultation. Oxford University Press.

Floridi, L. (2021). The European legislation on AI: A brief analysis of its philosophical approach. Philosophy & Technology, 34(2), 215-222. https://doi.org/10.1007/s13347-021-00460-9

Fung, A., & Wright, E. O. (2001). Deepening democracy: Innovations in empowered participatory governance. Politics & Society, 29(1), 5-41. https://doi.org/10.1177/0032329201029001002

Gilens, M., & Page, B. I. (2014). Testing theories of American politics: Elites, interest groups, and average citizens. Perspectives on Politics, 12(3), 564-581.

Habermas, J. (1996). Between facts and norms: Contributions to a discourse theory of law and democracy. MIT Press.

Hildebrandt, M. (2015). Smart technologies and the end(s) of law: Novel entanglements of law and technology. Edward Elgar Publishing.

Hsiao, Y.-T., Lin, S.-Y., Tang, A., Narayanan, D., & Sarahe, C. (2018). vTaiwan: An empirical study of open consultation process in Taiwan. SocArXiv. https://osf.io/preprints/socarxiv/xyhft_v1

Jefferson, T. (1982). Notes on the State of Virginia (W. Peden, Ed.). University of North Carolina Press. (Obra original publicada en 1785).

Jefferson, T. (1984). First inaugural address. In M. D. Peterson (Ed.), Thomas Jefferson: Writings (pp. 492-496). Library of America. (Original work published 1801).

Kant, I. (1996). An answer to the question: What is enlightenment? En M. J. Gregor (Ed. y Trad.), Practical philosophy (pp.

15-22). Cambridge University Press. (Obra original publicada en 1784).

King, L. A. (2003). Deliberation, legitimacy, and multilateral democracy. Governance, 16(1), 23-50.

Laclau, E., & Mouffe, C. (1985). Hegemony and socialist strategy: Towards a radical democratic politics. London: Verso.

Laclau, E., & Mouffe, C. (2001). Hegemony and socialist strategy: Towards a radical democratic politics (2nd ed.). Verso.

Lafont, C. (2015). Deliberation, participation, and democratic legitimacy: Should deliberative mini-publics shape public policy? Journal of Political Philosophy, 23(1), 40-63.

Lafont, C. (2020). Democracy without shortcuts: A participatory conception of deliberative democracy. Oxford University Press.

Lalley, S. P., & Weyl, E. G. (2018). Quadratic Voting: How Mechanism Design Can Radicalize Democracy. AEA Papers and Proceedings, 108, 33–37. https://doi.org/10.1257/pandp.20181002

Landemore, H. (2013). Democratic reason: Politics, collective intelligence, and the rule of the many. Princeton University Press.

Landemore, H. (2020). Open democracy: Reinventing popular rule for the twenty-first century. Princeton University Press.

Landemore, H., & Brennan, J. (2020). Debating democracy: Do we need more or less? Oxford University Press.

Levitsky, S., & Ziblatt, D. (2018). How democracies die. Crown Publishers.

Liang, F., Das, V., Kostyuk, N., & Hussain, M. M. (2018). Constructing a Data-Driven Society: China's Social Credit System as a State Surveillance Infrastructure. Policy & Internet, 10(4), 415-453.

Lippert-Rasmussen, K. (2012). Estlund on epistocracy: A critique. Res Publica, 18, 241-258. https://doi.org/10.1007/s11158-012-9179-1

Madison, J. (2003). The Federalist No. 51. En A. Hamilton, J. Madison, & J. Jay, The Federalist Papers (pp. 318-322). Bantam Dell. (Obra original publicada en 1788).

Menendez-Blanco, M., & Bjørn, P. (2022). Designing digital participatory budgeting platforms: Urban biking activism in Madrid. Computer Supported Cooperative Work, 31(4), 567-601. https://doi.org/10.1007/s10606-022-09443-6

Mill, J. S. (1991). Considerations on representative government. Prometheus Books. (Obra original publicada en 1861).

Mill, J. S. (1991). Considerations on representative government. Prometheus Books. (Obra original publicada en 1861)

Mladenovi- , I. (2020). Democracy, truth, and epistemic proceduralism. Croatian Journal of Philosophy, 20(59), 171-182.

Moats, D., Tang, A., & Chiu, J. (2024). Digital democracy in Taiwan: vTaiwan and the Pol.is platform. Technology in Society, 68, 101-115.

Montesquieu, C. L. (1995). Del espíritu de las leyes (M. Blázquez y P. de Vega, Trads.). Tecnos. (Obra original publicada en 1748).

Moraro, P. (2018). Against epistocracy. Social Theory and Practice, 44(2), 199-216. https://doi.org/10.5840/soctheorpract20185835

Morozov, E. (2013). To save everything, click here: The folly of technological solutionism. PublicAffairs.

Mouffe, C. (1993). The return of the political. Verso.

Mouffe, C. (2000). The democratic paradox. Verso.

Noveck, B. S. (2015). Smart citizens, smarter state: The technologies of expertise and the future of governing. Harvard University Press.

O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishers.

Paine, T. (1986). Common sense. Penguin Books. (Original work published 1776)

Pasquale, F. (2015). The Black Box Society: The Secret Algorithms That Control Money and Information. Harvard University Press.

Perfecto, has identificado correctamente la obra. Aquí está la referencia en formato APA para Mouffe, 2000:

Platón. (1992). La República (C. Eggers Lan, Trad.). Gredos. (Obra original publicada ca. 380 a.C.).

Rask, M., & Shin, B. (2024). Integrating paths: Enhancing deliberative democracy through collective intelligence insights. Societies, 14(270), 1-27.

Rawls, J. (1996). Political liberalism (Expanded ed.). Columbia University Press. (Original work published 1993)

Rawls, J. (1999). A theory of justice (Rev. ed.). Belknap Press of Harvard University Press. (Original work published 1971)

Raz, J. (1986). The morality of freedom. Oxford University Press.

Roberts, I. (2023). How can collective intelligence enhance democracy? An investigation towards human-centred design principles for empowering people engaging with deliberative technologies [Tesis de maestría, Northeastern University London].

Rosanvallon, P. (2008). Counter-democracy: Politics in an age of distrust (A. Goldhammer, Trans.). Cambridge University Press.

Sen, A. (1999). Development as freedom. Oxford University Press.

Sunstein, C. R. (2019). Conformity: The power of social influences. New York University Press.

Tocqueville, A. de. (2000). Democracy in America (H. C. Mansfield & D. Winthrop, Trads.). University of Chicago Press. (Obra original publicada en 1835-1840).

Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU artificial intelligence act. Computer Law Review International, 22(4), 97-112.

Waldron, J. (1999). Law and disagreement. Oxford University Press.

Warren, M. E., & Gastil, J. (2015). Can deliberative minipublics address the cognitive challenges of democratic citizenship? The Journal of Politics, 77(2), 562–574. https://doi.org/10.1086/680078