

RLFP

Revista
Latinoamericana de
Filosofía
Política

Centro de Investigaciones Filosóficas

ISSN 2250-8619 • Vol. XIV • N° 11 • 2026 • Buenos Aires • Argentina

**RISK-CENTERED ORDINAL CONSEQUENTIALISM:
A DEFENSE OF DEMOCRACY**

Mariano Garreta Leclercq

RISK-CENTERED ORDINAL CONSEQUENTIALISM: A DEFENSE OF DEMOCRACY

MARIANO GARRETA LECLERCQ

Universidad de Buenos Aires (Argentina)

Consejo Nacional de Investigaciones Científicas y Técnicas (Argentina)

mgarretaleclercq@yahoo.com

“Me sé del todo indigno de opinar en materia política,
pero tal vez me sea perdonado añadir que descreo de la
democracia, ese curioso abuso de la estadística”.

JORGE LUIS BORGES, 1976¹

“...Hemos aprendido a la luz de las trágicas experiencias
de los años recientes que la democracia es un valor más
alto que una mera forma de legitimidad del poder, porque
con la democracia no sólo se vota, sino que también se
come, se educa y se cura”.

RAÚL ALFONSÍN 1983²

ABSTRACT

Why should we value democratic institutions? How does the number of individuals whose preferences are satisfied contribute

1. “I know I am completely unworthy of having an opinion on political matters, but perhaps I can be forgiven for adding that I do not believe in democracy, that curious abuse of statistics.” Final lines of the prologue to *La moneda de Hierro*.

2. “[W]e have learned in light of the tragic experiences of recent years that democracy is a higher value than a mere form of legitimacy of power, because with democracy one not only votes, but also eats, is educated, and is healed.”—President Alfonsín, inaugural speech to the Legislative Assembly of the Argentine Republic (1983).

to the quality of decisions made through this procedure? Should political decisions be entrusted to experts instead of allowing majorities to have decisive influence?

Suppose we set aside the egalitarian values inherently present in democratic institutions and their supposed epistemic potential. We argue that increasing the satisfaction of the free and informed preferences of those affected—whether by generating the greatest possible consensus around binding decisions or by creating spaces of partial deregulation—contributes crucially to decision quality. This is because relevant political decisions are invariably made in contexts of risk. Risk acts as the non-epistemic bridge between the number of individuals supporting a decision and the quality of the expected outcome.

In contexts marked by a significant probability of error and a high cost of error to individual well-being, decision-making should be conducted through an open and perpetually revisable system. This system must combine majority rule with the effective possibility of contestation at the individual or collective level, seeking to incrementally accommodate individual preferences. Rather than justifying governmental decisions solely on the grounds that they are “correct” or that decision procedures track the truth, the system’s objective should be to remain sensitive to the range of acceptable gains and tolerable losses for each affected individual. The fundamental problem is not agent fallibility itself—the making of factual or normative errors—but rather the exposure of affected individuals to risk levels incompatible with the sustainability of their life projects.

We offer an instrumental, non-epistemic justification of a republican conception of democracy. To this end, we develop a novel consequentialist normative theory, which we term *risk-centered ordinal consequentialism*. This conception aims to account for defining aspects of political normativity. In this framework, the criterion of correctness—without relinquishing the pursuit of objective and impartial reasons—is the degree to which decisions approach the twin ideals of fully free and informed consensus and deregulation without interpersonal domination. In a republican democracy, the maximization of expected good—or utility—and of freedom as non-domination are extensionally equivalent.

Keywords: Risk, Contestatory Democracy, Epistocracy, Consequentialism, Ordinal Expected Utility Maximization, Partiality.

RESUMEN

¿Por qué deberíamos valorar las instituciones democráticas desde un punto de vista instrumental? ¿Qué contribución hace el número, la ampliación de la cantidad de individuos cuyas preferencias son satisfechas, a la calidad de las decisiones tomadas a través de esa estrategia? ¿Deberían las decisiones políticas ser dejadas en manos de expertos en lugar de aspirar a que las mayorías tengan una influencia decisiva?

Supongamos que dejamos de lado los valores igualitarios inherentes a las instituciones democráticas y su supuesto potencial epistémico. Nuestra hipótesis es que la creciente satisfacción de las preferencias libres e informadas de los afectados —ya sea generando el mayor consenso posible en torno a las decisiones impuestas a todos o creando espacios de desregulación parcial— contribuye a la calidad de las decisiones, ya que las decisiones políticas relevantes se toman invariablemente en contextos de riesgo. El riesgo es el puente no epistémico entre la cantidad de individuos que apoyan una decisión y la calidad del resultado esperado. En contextos caracterizados por una probabilidad significativa de error y un alto costo del error para el bienestar individual, la toma de decisiones debe llevarse a cabo mediante un sistema abierto y perpetuamente revisable que combine la regla de la mayoría con la posibilidad efectiva de impugnación tanto a nivel individual como colectivo. Este enfoque debe tratar de acomodar gradualmente las preferencias individuales. En lugar de justificar exclusivamente las decisiones gubernamentales sobre la base de que son las opciones “correctas” o las que tomaría un agente que conociera la verdad, el objetivo del sistema debe ser permanecer sensible a la gama de ganancias deseables y pérdidas tolerables de los bienes en juego para cada individuo afectado, dada su situación específica. El problema no es la falibilidad de los agentes —cometer errores en el plano fáctico o normativo— sino exponer a los afectados a niveles de riesgo incompatibles con la sostenibilidad de sus proyectos.

Nos proponemos ofrecer una justificación instrumental, no epistémica de una concepción republicana de la democracia. Con este fin, hemos elaborado un nuevo tipo de teoría normativa consecuencialista, que denominamos *consecuencialismo ordinal centrado en el riesgo*. Esta concepción pretende dar cuenta de los aspectos definitorios de la normatividad política. En este nuevo tipo de consecuencialismo, el criterio de corrección, sin renunciar a la búsqueda de

razones objetivas e imparciales como parte del proceso, es el grado en que las decisiones se acercan a los polos opuestos del consenso libre e informado y la desregulación sin dominación interpersonal. En una democracia republicana, la maximización del bien esperado —o utilidad— y de la libertad como no dominación son extensionalmente equivalentes.

Palabras clave: Riesgo, democracia contestataria, epistocracia, consecuencialismo, maximización de la utilidad esperada ordinal, parcialidad.

1. Introduction

1.1 A Three-Dimensional Disagreement

The aim of this paper is to defend the notion that democracy, understood from a republican point of view, is superior to other forms of political decision-making, such as an elitist or purely majoritarian approach. Like epistocrats and epistemic democrats, we assume that the decisive criterion for choosing between systems of political decision-making is the quality of their outcomes, evaluated independently of any inherent procedural virtues. In our view, the proponents of what we might call *the epistemic model*—epistocrats, aggregative and deliberative epistemic democrats, and philosophers who endeavor to reduce the political phenomenon to an exclusively or predominantly moral problem—subscribe to an inadequate conception of the nature of political disagreement. The prevailing attitude among philosophers is to conceive of political disagreement as a *two-dimensional* phenomenon. Making good collective decisions would involve adequately resolving disagreements among those affected concerning 1) judgments on the facts relevant to the matter, and 2) the normative framework that should guide the choice. However, from the perspective we will defend, reasonable political disagreement

is normally a *three-dimensional* phenomenon. Making good collective decisions means resolving reasonable disagreements among the parties involved about 1) the facts relevant to the decision, 2) the values or norms involved, and 3) the attitude toward risk that should guide the choice. When there is no full certainty that decisions will produce the intended outcomes—and the cost of failure is significant for the well-being of those affected—a central feature of instrumental rationality is the requirement for adequate factual risk management.

Our hypothesis is that this third dimension, which is inherent and unavoidable in almost all significant contexts of choice, introduces the logic characteristic of republican democratic politics into the collective decision-making process. We are referring to the instrumental value inherent in five aspects of the political process: 1) the rule of the majority, 2) the expansion of free consensus among those concerned beyond the specific outcomes of majority rule—even when those outcomes diverge from the preferred normative framework accepted by each party, 3) the maximization of individual freedom, 4) the availability of mechanisms for challenging majority political decisions, and 5) the elimination of unnecessary collective regulations. The third dimension of disagreement, associated with risk management, explains why these institutions make an irreplaceable contribution to the quality of the outcomes of the political process. Risk explains both the instrumental value of numbers—the aggregation of free wills—and the irreducible value of the individual, which are characteristic of democratic politics from a republican point of view.

After sections 1.1 and 1.2, in which we present the assumptions and background of our proposal, in section 2.1 we develop an argument against epistocracy. This argument also has implications for the conceptual plausibility of epistemic conceptions of democracy. The aim of the argument is to show that even if it were possible to resolve disagreements about the factual and normative judgments relevant to decision-making, the foreseeable persistence of disagreements about the level of

acceptable risk would continue to justify reliance on political procedures—majority rule and contestatory institutions. The argument against epistocracy rests on two related foundations: the analysis of “the case of the firm”, as we term it, and the stark contrast between this scenario and decision-making situations—such as Parfit’s well-known Miners case—where a single rational attitude toward risk is uniquely determined, thus obviating the need for a political process to decide the best expected outcome. Section 2.2 examines the theoretical background of the proposal and develops the argument. Section 3 provides a detailed explanation of why an institutional framework that equally integrates majority rule and the contestatory institutions associated with the republican tradition is instrumentally superior to a system that relies exclusively on majority rule. Conceptually, the proposal we defend involves the formulation of a form of ordinal prospective consequentialism that gives a central role to divergences among individuals about the level of acceptable risk associated with decisions. We will call this new type of normative theory “*risk-centered ordinal consequentialism*.” In Section 4, we delineate the differences between our proposal and other normative consequentialist conceptions that share significant similarities. First (4.1), we examine the differences between our proposal and the cardinal prospective utilitarianism developed by John Harsanyi. This theory shares with our proposal the assignment of a key role to the potentially divergent risk attitudes of those affected. Second (4.2), we explain the differences between the account presented here and the consequentialist defense of the same republican institutional model developed by Philip Pettit, which focuses on the agent-neutral value of freedom as non-domination. In section 5, we return to the normative level and explain why our proposal is incompatible with two new possible types of epistocracy, which we call “risk-centered epistocracy” (5.1) and “maximum precaution epistocracy” (5.2). Section 6 makes a comparison between the proposed consequentialist strategy and the ordinal utilitarian conception

outlined by Arrow as part of his defense of the instrumental value of democratic political competition. Diverging from the approach found in Arrow's proposal, the framework we defend aims to clarify a structural similarity between economic market competition and political competition within the democratic process. Section 7 summarizes the arguments presented throughout the text and emphasizes the most significant conclusions.

1.2 A Paradigmatic Case of Political Disagreement

Let us imagine that in a contemporary democracy there is a debate about what the government's economic policy should be. There are three sources of disagreement among citizens. The first source of disagreement is normative. Citizens support different conceptions of social justice. We can assume, for example, that some take a utilitarian approach. They believe that the goal of economic policy should be to maximize the net amount of wealth in society, since this is equivalent to, or leads to, maximizing the welfare of those concerned. Others believe that the welfare of the less advantaged should take precedence over maximizing the overall welfare of society. We can associate this position with the theory of John Rawls (1971, 1993),³ or with a

3. When we assert that agents adopt a "utilitarian," "Rawlsian," or "prioritarian" perspective, we do not imply that they apply it with full fidelity or absolute consistency. Rather, we suggest that a specific interpretation of these ideas—perhaps partial and limited—guides their deliberations. For instance, we assume that agents identified as "Rawlsians" advocate for economic policies based on an interpretation of Rawls's difference principle. However, we do not address the more complex task of evaluating state economic policy based on Rawls's complete theory of justice. This theory includes the first principle of justice (stipulating equal civil and political liberties) and the first part of the second principle (establishing equality of opportunity). While a comprehensive Rawlsian evaluation would require considering these factors alongside the difference principle, we omit them

form of utilitarianism combined with a commitment to provide a guaranteed social minimum for the impoverished. Another position bearing significant similarities to the aforementioned view is typically referred to as prioritarianism (see Parfit 1991, 1997; Holtug 2006; Buchak 2017, 2022). Prioritarianism also puts the needs of the less advantaged first, but in some of its versions it allows for the implementation of lower levels of redistribution than those that the difference principle of Rawlsian theory seems to require. Of course, there are other conceptions of justice. However, to make the example more manageable, let us keep these options in mind.

The second source of disagreement is factual. There will be disagreement about which are the most efficient means of achieving the goals associated with the conception of justice to which each agent or group subscribes. Let us consider a simple example. Some argue that active government intervention in the economy produces significant benefits, while others maintain that it is desirable to minimize such intervention. Admittedly, technical discussions of this nature are often intertwined with normative or ideological considerations, or with the vested interests of those affected. Yet, even leaving these factors aside, substantial disagreements remain. It is quite conceivable that advocates of different conceptions of justice will endorse proposals that reflect these opposing positions. There will foreseeably be utilitarian, Rawlsian, or prioritarian proposals that envisage different degrees of government intervention in the economy.

The third source of disagreement arises from the fact that, in standard policy-making contexts, success is never guaranteed, and the cost of failure is always significant for the well-being of those affected. Any plan of action formulated in light of a normative framework and a set of relevant factual

to simplify the example. Similar simplifying assumptions apply to the other normative frameworks discussed.

judgments inherently involves an attitude toward risk. To make a decision in such circumstances, we must decide how much we are willing to risk in order to try to achieve our goals. That is, how much of what we currently have and value are we willing to put at stake in order to achieve what we want in the future. In practice, different individuals and groups exhibit diverging attitudes toward risk. While some are risk-averse to varying degrees, others are risk-seeking. Consequently, two agents sharing identical goals may reasonably disagree regarding the acceptable level of exposure required to achieve their objectives.

It is one thing to endorse a specific normative framework; it is another matter entirely—and one not necessarily determined by that endorsement—to decide how much we are willing to risk to achieve our goals. Distinguishing these two issues is crucial, as the decision involves wagering not only the goods associated with the goal itself but potentially other valued goods as well. A Rawlsian believes that we must aim to design an economic and institutional system that maximizes the position of the least advantaged. However, they recognize that if a miscalculation occurs—specifically, if the efficiency costs of these measures prove excessive—the policy’s outcome may be the opposite of what was intended. Given the magnitude of the cost of such a decision in the case of error, it may be more desirable to implement a more cautious policy that would produce a lower level of redistribution, resulting in a smaller decline in the level of total wealth in society. Decisions of this kind are unavoidable when supporting or promoting almost any specific policy measure. Furthermore, they have a significant impact on the likelihood of achieving broader consensus or, conversely, of deepening polarization among individuals who subscribe to divergent value systems or factual beliefs. For example, it seems reasonable to assume that a Rawlsian and a utilitarian who are significantly risk-averse might support more similar economic policies, reaching a broader consensus than would be possible between utilitarians and Rawlsians who

are more risk-prone. Risk-prone Rawlsians and utilitarians are likely to support policies that attempt to more radically alter the status quo in order to further the goals of the evaluative framework they support. These would be the hawks of each position. The former, are perhaps willing to drastically reduce the total wealth of society in order to improve the situation of the less advantaged group. The latter, are perhaps willing to try to maximize total wealth even if it means a very substantial deterioration in the living conditions of the less advantaged. In such a case, the disagreement about the correct attitude toward risk deepens polarization.

Attempts by governments to successfully implement their chosen policies can always fail. There is always a significant probability of this happening, however small, and the cost to the welfare of those affected by this outcome is usually substantial. In these cases, the combination of the three variables mentioned and the balance struck between them seems to play a crucial role in explaining both the agreement and the disagreement among the actors interacting in the decision-making process. Let us suppose, then, that there are various policies enjoying significant support in the hypothetical society we are considering. The desirability of each policy depends on three variables: 1) a value framework considered correct or reasonable; 2) factual estimates of the projected outcomes of their implementation; and 3) the adoption of a specific attitude toward risk.

Let us assume that there are three decision-making systems that could be adopted. The first option is to have the resolution taken by an elite of experts, agents with a level of information and epistemic competence superior to that of the average citizen. The second alternative is simply to resort to majority rule. All those concerned have one vote and one vote only, and the option that wins the support of the most citizens is implemented. The third approach is to submit the decision to a republican democratic system, which we could call, following Philip Pettit, *contestatory democracy*. In this

case the government is controlled both by the usual electoral mechanisms and by forms of contestation in which citizens, collectively or individually, have the possibility of questioning and trying to reverse decisions previously supported by the majority. This system involves the existence of a free and active public sphere, a free press; the existence of forums—more or less institutionalized—that allow citizens to deliberate; the usual checks and balances that we associate with the republican and liberal tradition; the rule of law; the separation and independence of powers; the existence of a bill of rights that allows majority decisions to be legally challenged; the existence of depoliticized decision-making mechanisms; free access to information about government actions.⁴

We will take an instrumental view of the value of the political process.⁵ The best decision-making system will be the one that, in the medium- or long-term, tends to produce the “best outcomes”, the right results according to standards of evaluation that are independent of the values that the process

4. We ground our discussion in the model of republican democracy initially developed by Philip Pettit in *Republicanism: A Theory of Freedom and Government* (1997) and further elaborated in texts such as “Republican Liberty, Contestatory Democracy” (1999) and “Democracy, Electoral and Contestatory” (2000). Although Pettit has continued to refine various aspects of this proposal, the latter article serves as a comprehensive synthesis of the specific conception we adopt as our model.

5. Traditional justifications for democracy rely on procedural or intrinsic arguments, appealing to foundational political values like freedom, equality, and fairness. The value of democratic procedures—such as universal suffrage and the right to contest decisions—is typically understood to stem from their capacity to realize these ideals. However, this approach often neglects the quality of democracy’s actual outcomes. We adopt a distinct, instrumental, and consequentialist strategy to provide an outcome-oriented rationale for democratic institutions. We advance the ambitious thesis that, in contexts characterized by risk, maximizing individual liberty and maximizing the well-being of the affected population are extensionally equivalent. Crucially, this approach complements, rather than critiques, arguments that highlight the procedural or non-instrumental value of the system.

may embody.⁶ What counts as “good outcomes” for the political system? According to Brennan and Landemore, to answer this question we must focus on values and goals that are widely shared by people around the world, such as “human rights protection, economic prosperity, peace, security, basic dignity, individual freedom, education, health, and general individual well-being” (2022, 8). The central question remains: what is the optimal model for collective decision-making among the proposed alternatives: epistocracy, majoritarian rule, or contestatory democracy?

Let us start by focusing on the epistocrats’ answer to this question. Epistocrats (see Brennan 2016 and also Malcolm 2022, Gibbons 2021, Jones 2020, Mulligan 2018, Bell 2015, López-Guerra 2014, Guerrero 2014) believe that citizens who are more knowledgeable, informed and competent should have more political influence. An epistocratic decision-making system will be superior in terms of the quality of the outcomes it produces; it will be superior from an instrumental point of view. As Brennan, one of the most influential proponents of this perspective, argues:

A system is epistocratic to the extent that greater knowledge and the good faith to act on this knowledge are *de jure*, legal prerequisites for holding power or are legal grounds for being granted greater power through law. (2017, 54)

6. Brennan and Landemore define instrumentalism as follows: 1. There are procedure-independent right answers to at least some political questions. To illustrate, think of a criminal jury trying to assess whether someone is guilty. There is an independent fact that they are trying to discover. This is in contrast with, say, a game of dice, where the “right” answer is procedurally defined: whoever rolls the highest number wins. 2. What justifies the distribution of power, or a particular way of making political decisions, is (at least in part) that this distribution or method tends to get the right answer. (2022, 7)

An extensive empirical literature on voter competence consistently demonstrates that the majority of citizens in contemporary democracies lack the basic political information requisite for informed decision-making. This central challenge is often referred to as “the problem of the political ignorance of the average voter” or “voter incompetence” (see Achen and Bartels 2016, Brennan 2016, Somin 2013, Caplan 2007, Bartels 1996, Schumpeter 1942). If the majorities affected by this phenomenon are excluded from the decision-making process and greater political influence is assigned to informed and competent agents, we can anticipate a corresponding improvement in the quality of political decisions and their eventual effective outcomes. In summary, epistocracy would be desirable because it contributes, or could contribute, more effectively than democracy to the common good. Why should the majority of citizens accept the decisions of an elite of experts? The answer is simple. Because we assume that these agents are acting in good faith, seeking the right solutions or the best decisions based on their knowledge. Of course, this does not mean that they choose policies that they believe in good faith to be best *for themselves* or to promote their *own interests*. Rather, their decisions are based on 1) beliefs about facts that are properly justified from an epistemic point of view, and 2) an adequate interpretation of the common interests of all concerned. Instead of talking about “an adequate interpretation of the common interests of all concerned,” we could say that they make their decisions in the light of judgments about justice that are true or correct, or the most reasonable conception of justice, or at least a reasonable conception of justice. In the latter case, the core premise is that this evaluative framework is not truth-apt—that is, susceptible to the predication of truth or correctness—yet it possesses specific normative force, or is at least widely accepted. Under these conditions, the decision-makers (the minority of experts) can employ the following rationale: they argue that their decisions do not reflect their particularistic

interests, but rather constitute the objectively correct choice. Hypothetically, these would be the same decisions the excluded population would have chosen if they possessed the relevant empirical information, were competent to evaluate it, and could adequately interpret the common interests at stake. Consequently, any rejection of the elite's decisions by the affected actors is viewed as a form of disagreement that lacks normative or moral relevance, dismissible as a merely "unreasonable disagreement." Let us assume that there is no difficulty in identifying those who are epistemically competent agents and in excluding from the decision-making process those who do not meet the required threshold of competence or who are morally unreasonable. We assume that the experts act in good faith and try to base their decisions on an adequate normative framework and on factual judgments that are relevant and well-founded from an epistemic point of view. Going back to the case, it is very unlikely that experts can truly establish the correct normative framework or the truth about justice. It is not clear how the experts could manage to determine, for example, whether some version of utilitarian theory or, on the contrary, a deontological conception such as Rawls' theory is the correct conception of justice. After all, philosophers are also epistemically competent and informed actors who debate in good faith, and who have not managed to come close to a definitive answer to this question in debates that have lasted for decades. But let us put this problem aside and imagine that it has been overcome or resolved. Under these conditions, would it be correct to affirm that the decisions of the experts should be accepted by all concerned, because they are the decisions that all agents would have supported in the end, if they had the necessary epistemic competence and moral reasonableness? Can we assume that the experts' decision really promotes the common good, or at least the welfare of the majority of those concerned?

2. An Instrumental Argument Against Epistocracy

2.1 *The Case of the Firm*

We will try to demonstrate that the answer to the two previous questions is categorically negative. Even if we grant the epistocrat all the assumptions on which their position depends, it still fails because it lacks the instrumental value its advocates attribute to it. To prove this statement we will analyze a case in which the factual and normative disagreements were overcome but a significant degree of reasonable controversy remains around the third source of disagreement: the attitude to be taken towards risk. This is what we will term “the case of the firm”. It will shortly be demonstrated that, in cases like this, it is possible to determine the best decision based on the accepted goals of those affected and well-justified facts that all informed and competent agents can agree on. In this scenario everyone recognizes the goal of maximizing profits and assigns the same probabilities of success/failure to the options at stake. They all describe the results of choosing each of these options and the value or desirability of these results in the same way. In the case of the firm, there is a correct decision from an objective and impartial point of view. The point is that choosing and implementing it in contexts of risk does not guarantee the quality of the result at both the individual and collective levels. The rationale for delegating decisions to an expert elite rests on their purported ability to identify the “right” decision based on objective and impartial considerations. However, the instrumental failure of these decisions—their inability to guarantee quality results in typical political contexts—fundamentally undermines the underlying justification for the epistocratic model. The same considerations apply to the thesis that the majority rule is the best decision-making system because it allows us to track down the truth—identifying the correct decision in light of objective and impartial considerations, which are supposedly equally

acceptable to all—or to increase the probabilities of identifying a decision that fits this description. We can call the advocates of this type of epistemic conception of democracy,⁷ together with the epistocrats, advocates of the epistemic model.⁸ Now, let us turn our attention to the case.

Consider the case of a small company run by 15 partners who have invested the same capital. So far, important decisions have been taken democratically: all partners have one vote and the decision that wins the support of the majority is implemented. The partners now have to choose between three different investment opportunities. The first, *N1*, could yield a very high profit, but also entails high losses if something goes wrong. Profits and losses are shared equally by all partners. Imagine that if *N1* is successful, each partner wins \$300,000. If it fails, each partner loses \$850,000. The second investment, *N2*, exhibits a smaller difference between profit and loss. If successful, each partner wins \$100,000 and if unsuccessful, each partner loses \$120,000. The probability of success is also high, 0.8. The third option, *N3*, has very different characteristics. Both the winnings and the losses are higher than in the other two options. In case of success, each partner wins \$350,000 and in case of failure each partner loses \$1,200,000. The probabilities of success are the same, 0.8. Under these conditions, there is no difficulty in making an objective comparison between the three investment alternatives. Which is the best option? The answer seems simple: the one with the highest expected monetary value, taking into account the magnitude of gains and losses and the probability of these outcomes. The expected value of *N1* is \$70,000; the expected value of *N2*, \$56,000; and the expected value of *N3*, \$40,000.

7. See Cohen 1986; Coleman and Ferejohn 1986; Nino 1996; Gaus 1997; List 2001; Estlund 2008; Landemore 2013; Goodin and Spiekermann 2018.

8. For a different critique of what we term “the epistemic model,” see Urbinati (2014, 81–127) and Lafont (2020, 75–100).

Suppose 14 of the partners temporarily lose the ability to understand the investment options. Only *S1* is qualified to assess the information relevant to decision-making, such as the potential gains and losses of the various options and the likelihood of each outcome. *S1* is also correct in believing that the appropriate normative framework for the decision is one that prioritizes profit maximization over other considerations. We are proposing a moral reading of the case of the firm. The goal accepted by all is that the firm's decisions should result in the greatest possible well-being for the partners. In other scenarios, to promote the same goal, perhaps investments that do not pollute should be given priority over the choice of polluting options, even if they are much more profitable. Perhaps in other scenarios there are different moral limits to maximizing the well-being of these agents because of their consequences for the well-being of third parties. But we assume that this is not the case in the scenario we are considering. We stipulate that the three aforementioned investment options pose no environmental harm, thereby precluding consideration of this or any other potential externality. The partners' decisions do not harm third parties. On the other hand, the solution to the distributive justice problem is much simpler and more convincing in this scenario than in the political case we presented earlier. This is because a strictly equal distribution of profits and losses among the partners seems perfectly feasible and compatible with the efficiency of the system. The relevance of theories of justice such as Rawls's or utilitarianism in real economic policy contexts depends to a large extent on the fact that the ideal of strict equality does not appear to be feasible or morally attractive. In the case of the firm the situation is more favorable for egalitarian intuitions. Everyone contributes the same capital and everyone is entitled to the same profits.

The case of the firm appears to present the ideal scenario for an epistocratic decision system. The designated decision-maker, *S1*, possesses all the necessary information to maximize expected outcomes, ensuring a rationale that is universally

acceptable and provides an objective, impartial solution. If the political sphere were anything like this case, a minority of experts, the equivalent of agent *S1*, would be able to make good decisions from an instrumental point of view. The rest of the citizens, a majority consisting of agents who are epistemically and perhaps normatively or morally incompetent, should accept these decisions because they correspond to the choices these agents would have made had they been competent and rational. But is this conclusion accurate?

To answer this question, we must first determine the specific conditions under which delegating a political decision to a minority of experts—while excluding affected parties—is instrumentally acceptable. Specifically, when does this approach reliably and measurably enhance the quality of outcomes? The clearest case is defined by the assumption that the majority of citizens are epistemically incapable of assessing the options, and the competent agent faces no risk of error. This scenario presupposes full confidence that the expert will make the objectively correct decision and ensure successful implementation. Suppose *S1* knows with certainty at *T* that investment *N3* will be successful at *T1*. If the partners' goal is to maximize profits, then investing in *N3* at *T1* is undoubtedly the right decision, and the majority's opposition to this decision would be simply irrational. If *S1* makes the unilateral decision that the firm should invest in this option, there can be no objection from an instrumental point of view. It is undoubtedly the best option for everyone. This assertion is based on the fact that the decision successfully passes two different tests. First, the expert must justify the decision by arguing that if all partners shared the same evidence and guiding goal, they would (or ought to) arrive at the identical decision. Second, if these circumstances recur across an indefinite number of future instances, repeating the correct decision must reliably guarantee the optimal outcome.

For instance, if we know that *N3* will never produce losses and will therefore yield higher profits than the other

alternatives, we know that it is the best decision in all future situations where a choice must be made between $N1$, $N2$, and $N3$. Repeating the decision an indefinite number of times will not only produce the best result in each specific instance of the decision, but also the best overall outcome. If we are in a context where, as in the case of the firm, the probabilities of error are significant and their costs high, there is also a type of scenario that is, at least in principle, favorable to the epistocratic strategy. The well-known Miners case, formulated by Derek Parfit (2011, 159),⁹ fits this description. Let us look at the reasons.

There are 100 workers trapped in a mine. There are two mine shafts: Shaft A and Shaft B. All 100 workers are located in the same shaft, but we do not know which one. The probability that they are in either shaft is equal (0.5). There are three floodgates: 1, 2, and 3. We can close only one of the gates to prevent imminent flooding. If we close Gate 1 or Gate 2, the probability of saving all 100 workers—or conversely, of all perishing—is 0.5. On the other hand, if we close Gate 3, we ensure with certainty that 90 workers will be saved.

Almost all of us agree, without any need for calculation, that we must close Gate 3, even if it means giving up the possibility of saving all 100 workers. The calculation that supports this idea is simple. The expected utility of closing gates 1 or 2 is 50 utility units (100 utility units corresponding to saving everyone multiplied by the probabilities of success, 0.5), while the expected—and actual—utility of closing Gate 3 is 90 utility units (90 utility units corresponding to saving 90 lives multiplied by the probabilities of success, 1). This case is often used to demonstrate the necessity and feasibility of replacing

9. Another similar and frequently cited example is the medical treatment case proposed by Jackson (1991).

the classical consequentialist goal of maximizing *actual utility* with the maximization of *expected utility*.¹⁰

Let us suppose that we have to choose between 1) leaving the decision of which gate to close to an expert who knows the data we have just presented, or 2) leaving it to the miners' relatives, friends, and acquaintances. Let us assume that the majority of these people are inclined to close Gate 1 because they mistakenly believe that the probability of everyone being saved in this case is not 0.5, but 0.99. There is no evidence that this belief is true, it is just a rumor spread through social networks by some very popular and irresponsible people. Let us assume that what the majority wants is not taken into account. The expert makes the decision to close Gate 3. Under these conditions, it is not possible to object to the decision made from an instrumental moral point of view. This does not change even if, after the expert's decision, it turns out that, in this case, all 100 workers would have been saved if Gate 1 had been closed, as the majority had preferred on the basis of false judgments. The fact that the probabilities of everyone dying or being saved by closing Gate 1 were 0.5, of course, takes this possibility into account.

The expert's decision cannot be challenged because it clearly satisfies the two aforementioned tests. The expert is in a position to argue that if the miners' relatives and acquaintances recognized the priority of saving more people, and were they aware that the chance of saving everyone by closing Gate 1 was 0.5 (instead of 0.99), they would have reached the identical decision: to close Gate 3. Imagine two possible worlds, WP1 and WP2, where the mine accident occurs daily. In WP1, the

10. A further question is whether this replacement implies adopting expected utility solely as a criterion for rational decision-making—a thesis that appears highly plausible—or also as a criterion for moral correctness. Some consequentialists, such as Shaw, argue that expected utility maximization serves as the criterion for both decision-making and moral correctness (see Shaw 1999, 27–31).

authorities consistently close Gate 1; in WP2, they close Gate 3. Because of probabilistic variance, we cannot conclude that after 100 days, WP1 will have 5,000 dead miners while WP2 has only 800. For instance, like a fair coin landing on the same side multiple times in a row, it is possible that in WP1, all workers are saved consecutively many times, while at identical points on the parallel reality (WP2) timeline, 10 workers perish each time. However, given the law of large numbers, we know that the incidence of fatalities in the second possible world will tend to be drastically lower than in the first as the time sequence progresses. This confirms our intuition: choosing to close Gate 1 (or 2) instead of Gate 3 is an irrational, risk-driven attitude that leads to suboptimal outcomes in the medium- or long-term. Consistently repeating the decision to close Gate 3 represents the optimal strategy, provided that the factual conditions remain constant and we adhere to the same normative framework (namely, that all lives are of equal value and, therefore, that saving a greater number of lives is preferable to saving fewer). In cases such as the mine scenario, the optimal solution is a ‘depoliticized’ one, determined by objective reasons and the correct evaluative framework. This holds true even when we lack certainty regarding the actual outcome of the decision and the cost of error is severe.

Alternatively, suppose the majority of those affected by the accident—the friends and relatives of the trapped miners—possess superior information. If the number of people involved in the decision process is sufficiently large and the conditions for Condorcet’s Jury Theorem are met,¹¹ the majority vote will

11. As summarized by Anderson (2006, 10–11): “This theorem states that if voters (a) face two options, (b) vote independently of each other, (c) vote their judgment of what the correct solution to the problem should be (i.e., they do not vote strategically), and (d) have, on average, a greater than 50% probability of being correct, then the probability that the majority vote will yield the correct answer approaches 1 as the number of voters approaches infinity (and rapidly approaches 1 even with modest numbers of voters).”

never select gates A or B and will, therefore, converge on the correct decision: closing Gate 3. In this scenario, the resulting decision would be as depoliticized, objective, impartial, and epistemically adequate as that of the expert.¹²

In the case of the firm the situation is quite different. In scenarios with these characteristics, the expert, *S1*, cannot pass either of the two tests. Their decisions are not justified from an instrumental point of view; there is no reason to conclude that these decisions efficiently promote the interests of the firm. They do not guarantee the sustainable growth of the partners' profits or its survival in the relevant time horizon. As we saw, *S1* knows that the expected value of *N1* is higher than that of *N2* and that the expected value of this second option is higher than that of the third option, *N3*. We assume that the other partners cannot evaluate the data, that they are completely blind. While *S1* possesses greater knowledge than the other partners and appears better positioned to make a rational decision, this advantage does not necessarily extend to determining the optimal choice for the company. Their superior position is, in fact, limited to knowing what decisions best promote their *own* interests. Let us suppose that *S1* is not in a position to face the liabilities or bear the risks associated with investment in *N1* or *N3*, should such ventures prove unsuccessful. Of course, both options have a high probability of success and would yield a higher profit than *N2*. But *S1* knows that, given their previous capital, their financial situation, or their medium- or long-term projects, if something goes wrong, the result for them of the company investing in the riskiest and most lucrative options would be bankruptcy. *S1* may then conclude that the best decision, from

12. Mason plausibly defines “prospective consequentialism” as a conception in which the right action is the one that “best balances risk and good consequences” (2014, 197). In the Miners case, there is only one acceptable balance: a linear utility function corresponding to risk neutrality.

their perspective, is for the company to invest in *N2*. They are not willing to risk losing as much as they might lose by investing in *N1* or *N3*, given the probabilities of success, in order to gain more than they might gain by investing in *N2*. If *S1* suddenly realizes that they have the necessary capital to withstand the losses that would result from investing in *N1* or *N3*, or if they are not worried about the losses for some other reason related to their circumstances, they may prefer that the company invest in one of these riskier options. It could be assumed that if *S1* wants to make the best decision for the company, they could choose the optimal decision in light of objective metrics (measurable monetary gains and losses and their associated probabilities) and impartial standards (the shared goal of profit maximization). We know that *N1* is the option with the highest expected value: the best objective balance of gains and losses given the probabilities. But is it the best choice for the company? *S1* knows that *N1* is the option with the highest expected value, and may know whether it is the best decision for their interests and needs, but really knows nothing more. If *S1* were to choose one of the three options, they would not meet the requirements of the two tests that the expert could pass in a case of certainty and in the Miners case.

Let us suppose that *S1* chooses *N1*, the option that we could describe as superior from an objective and impartial point of view. In this case, however, it does not satisfy the first test because it cannot affirm that the other partners would also have chosen this option if they knew all the relevant facts and adopted the same normative framework, the goal of maximizing profits. This could only have happened if they were all risk-neutral. However, there is no reason why this should actually be the case, or why it should be required. Given the same information as *S1*, some risk-averse partners might have decided to support the investment in *N2*—as we initially assumed *S1* did—while some risk-seeking partners might have supported *N3*. However, *S1* does not know what those agents

would have decided had they possessed the information on which *S1* based their decision. Given the same information and the same normative framework, they could have made different, yet equally rational, decisions.

What distinguishes the Miners case from the case of the firm is that in the former there seems to be a single rational attitude toward risk. Under these conditions, there is a single calculation of expected utility that is equally valid for any informed agent who subscribes to the assumed normative framework. On the other hand, in the case of the firm, there are several attitudes toward risk that are equally rational for different agents. Option *N1* might be the most rational for *S1*, while *N2* and *N3* might be more rational for the other agents. The expert's decision can be rejected by the majority of those affected because it is based on the flawed premise that they would have made the same choice if they had the same knowledge and subscribed to the same normative framework. If the expert's decision, whatever it may be, is not in line with the risk attitude of the majority, the decision of this group would be different from that of the agent and equally rational.

Similar considerations apply to a situation in which the members of the majority meet the threshold of epistemic competence required by the jury theorem. The majority, if the epistemic model had justificatory force in this type of scenario, could justify its decisions simply because they are the correct ones, or the ones most likely to be correct.

For example, the majority supporting *N1* might argue that the minorities favoring *N2* or *N3* are in error. However, such a claim is untenable. In fact, all three options—provided they are adopted under conditions of adequate information—represent the optimal choice for their respective proponents.

Consequently, the members of the majority lack substantive epistemic justification for their proposal, even if their decision genuinely constituted an instance of the jury theorem. What they could argue, in instrumental non-epistemic terms, is that what benefits the majority is best for the flourishing—or at

least the survival—of the firm over time. Crucially, however, this is distinct from the arguments epistemic democrats typically offer in favor of majority rule. Nevertheless, it suffices to establish the instrumental superiority of majority rule over the decisions of a minority of experts in contexts of risk.¹³

2.2 Theoretical Background: Risk, Expected Value and Expected Utility

Before proceeding with the argument and explaining the reasons why the decisions of a minority of experts do not satisfy the second test in the corporate example, let us make some clarifications about the theoretical background that underlies the conception we are defending. The first influential theoretical approach to the role of risk attitude in decision rationality originated in the mid-18th century. It is the elegant solution to the St. Petersburg Paradox proposed by the Swiss mathematician David Bernoulli in his 1738 paper “Specimen Theoriae Novae de Mensura Sortis” (“Exposition of a New Theory of the Measurement of Risk”). The defining characteristic of the St. Petersburg game is that its expected value is infinite. If expected value were an adequate criterion for rational action, every agent should theoretically be willing to invest their entire wealth to participate. However, the typical outcome of the game is to win only small sums before

13. This claim rests on the assumption that the options favored by the majority have been assessed with epistemic competence and are supported by an adequate or reasonable normative framework—as is the case of the firm and political examples discussed earlier. In Section 5, we examine the implications of the proposed conception in contexts where majority members do not meet these standards.

being eliminated. This reality renders the decision to pay a substantial entry fee irrational.¹⁴

Daniel Bernoulli's famous solution to the paradox is that the appropriate criterion for evaluating investments in contexts of risk is not the "*expected value*" of the investment, but its "*expected utility*". In Bernoulli's model, expected utility results from a combination of 1) the value of the possible outcomes at stake, 2) the probabilities of achieving those outcomes, and 3) the risk involved, or more precisely, the agent's attitude toward

14. The St. Petersburg game operates as follows: a fair coin is tossed repeatedly until it lands on heads. If the coin lands on heads on the first toss, the player wins 2 monetary units. If it lands on heads on the second toss, the payoff is 4 units; on the third, 8 units, and so forth. The game terminates immediately upon the first instance of heads. The probability of the game ending on the first toss is 1/2; on the second, 1/4; on the third, 1/8, and so on. Consequently, the expected value of the game is infinite. This derives from the fact that the expected value contribution of each potential outcome is constant at 1 monetary unit:

$$2 \frac{1}{2} = 1; 4 \frac{1}{4} = 1, 8 \frac{1}{8} = 1; \dots$$

The sum of an infinite sequence of such units yields an infinite value. The calculation can be summarized by the following series:

$$\sum_{n=1}^{\infty} 2^n \left(\frac{1}{2^n} \right) = 1 + 1 + 1 + \dots = \infty$$

(See Samuelson 1997; Menger 1967).

Participating in the game requires an entry fee. Given that the expected value is infinite, any rational agent maximizing expected utility should theoretically be willing to invest their entire wealth to participate. However, in practice, individuals are rarely willing to pay even a moderately high amount to play. This discrepancy arises because, although the expected value is infinite, the probability of winning a substantial sum is negligible. This paradox highlights that expected value is often an inadequate criterion for evaluating the desirability of wagers or investments in contexts of risk.

risk. The calculation proposed by the Swiss mathematician to convert money into utility includes two additional factors: 1) the previous wealth of the agent and 2) the logarithmic function. The idea behind this calculation is that all agents are risk-averse to varying degrees. The logarithmic function provides a mathematical representation of this attitude, since the correlation with an increase in the money supply is a decreasing marginal growth when the logarithmic function is applied. Including the agent's prior wealth in the calculation gives mathematical expression to the idea that the degree of risk aversion is generally greater the smaller the amount of capital the agent has before making the decision. Because of their varying levels of prior wealth, different agents are positioned at different points on the concave curve produced by the logarithmic function. This is equivalent to assigning them distinct levels of risk aversion and subjective values for the involved gains and losses. Although there are exceptions, the general idea is that gains and losses do not affect people in different financial or economic situations in the same way. The loss of \$5,000 may have a minimal impact on the well-being of a high-income person, but be devastating for someone barely above the poverty line. These ideas are the logical consequence of Bernoulli's general proposition: an increase in the possession of a good, such as money, always tends to have a positive effect on the well-being of the individual, but this positive effect is marginally diminishing. Each newly acquired dollar produces less marginal utility than the one previously obtained. This characteristic defines a risk-averse agent, for whom the negative impact of a loss has a greater effect on well-being than the positive impact of an equivalent gain.

Bernoulli's model corroborates the intuition that choosing $N1$, $N2$, or $N3$ can represent an equally rational decision for each respective agent. To derive this result, it suffices to assume distinct levels of initial wealth. For instance, if we assign a total initial wealth of 16 million to $S1$, 1 million to $S2$, and

50 million to $S3$, and model utility as a logarithmic function of total wealth, we obtain a confirmation of the assumed preference scheme. Given $S1$'s level of risk aversion associated with their initial wealth, $N1$ yields a higher expected utility than $N2$ or $N3$. Specifically, the expected utility of $N1$ for $S1$ is 16.5920, whereas $N2$ yields 16.5915 and $N3$ yields 16.5898. In the case of $S2$, whose lower initial wealth implies greater risk aversion, $N2$ generates higher utility (13.8661) compared to $N1$ (13.6459) or $N3$ (8.8512). Finally, for $S3$ —who possesses greater wealth than the other agents—the riskier option $N3$ yields the highest expected utility (17.7292), surpassing $N1$ (17.7288) and $N2$ (17.7286).

Subsequent decision theories have unanimously endorsed the idea that risk attitudes play a crucial role in the rational decision-making process. However, a significant discrepancy between Bernoulli and later theorists is the adoption of what might be called permissiveness about risk attitudes. As we have seen, Bernoulli allows for a fairly wide range of risk-aversion levels resulting from the interplay between the agent's initial wealth, potential gains and losses, and the logarithmic function. However, he is not a permissivist, since according to his theory all agents are risk-averse and their degree of aversion is calculated by different applications of the same mathematical function. The problem is that it seems very difficult to deny that both risk neutrality and some degree of risk taking can be perfectly rational for different subjects in circumstances that are far from exceptional. The most historically influential approach in decision theory, Expected Utility Theory (EU), was developed in the 1940s by von Neumann and Morgenstern (see von Neumann and Morgenstern 1944).¹⁵ In EU, different agents may be risk-averse, risk-neutral or risk-seeking to different degrees. More

15. Von Neumann and Morgenstern's treatment of Bernoulli's theory is found in Chapter 3 of the work cited above.

recent approaches have emerged in an attempt to overcome the limitations and shortcomings of the classic theory. These proposals allow for a wider range of variance in risk attitudes and give the concept of risk greater weight in the rational decision-making process, specifically by involving forms of maximization weighted or constrained by agents' overall risk attitude.

Our interpretation of the case of the firm, which may seem surprising to those unfamiliar with decision theory, is not new. As we have seen, such an interpretation is possible even within the Bernoulli model, which implies that every rational agent possesses a degree of risk aversion (based on the proposed logarithmic utility function). If we move to the von Neumann and Morgenstern model, we can imagine a much larger number of variations and agents choosing different options in the light of a wide variety of utility functions associated with different degrees of risk aversion, neutrality or propensity. This scenario is deepened in more recent theories (see Kahneman and Tversky 1979, Machina 1982; Quiggin 1982; Yaari 1987). A particularly interesting example is the theory proposed by Buchak (2014), called "risk-weighted expected utility theory" (REU). Aligning with many critics of the EU, Buchak attempts to formalize the idea that subjects' attitudes toward risk play a crucial role in practical reasoning—a role that must be distinguished from both probability and utility assignments. In this theory, the axes of evaluation are not two, as in the EU, but three. A separate risk function is added to the utility and probability functions. From the REU perspective, even if two agents assign the same probability and utility to the available options, they may reach opposite conclusions if their risk functions are different. The introduction of the separated risk function allows us to express more strongly than is possible in the EU the fact that some agents give more weight to potential losses than to potential gains, while others adopt the opposite pattern. These attitudes are perfectly rational within a fairly wide range of variation. It is worth noting that there is no prior or independent criterion

of rationality to determine the right balance between prudence and risk-taking when two agents disagree. There are, of course, attitudes of extreme aversion and extreme propensity that no one would consider rational, but once they are excluded, the range of rational variation remains very wide.¹⁶

Our interpretation of the corporate example does not rely on the specific features of any of these theories. Instead, we employ specific elements of Bernoulli's theory—such as the logarithmic utility function and the premise that initial wealth influences risk sensitivity—to demonstrate that even this framework supports the intuitive conclusions outlined in the previous section. This holds true despite Bernoulli's framework assigning a less central role to risk attitude than later models. The interpretation we defend regarding the case of the firm relies on two fundamental theses endorsed by the vast majority of decision theorists, from Bernoulli to the present.

16. Buchak draws a parallel between opposing practical dispositions towards risk and William James's ideas about the balance between the epistemic desiderata of knowing truth and avoiding error. The predominance of the goal of avoiding error is an epistemic equivalent of risk aversion, while the predominance of the goal of knowing the truth is an epistemic equivalent of the disposition toward risk. Buchak points out that James's thesis is that reason cannot dictate which temperament to adopt. As she explains: "[reason] cannot tell us which temperament to have or how to balance the goals of knowing truth and avoiding error, and so we must determine this for ourselves [...] In situations in which we cannot be certain of how our acts will turn out, rationality cannot tell us how to balance the goal of allowing for the possibility of something great against the goal of securing a guarantee that things will not turn out too badly. And so, with this recognition that decision-makers have more *freedom* than we previously thought, comes a realization that decision-makers have more responsibility than we previously thought. Since there is no unique answer to the question of how a decision-maker should treat risk, an individual not only may, but must, decide for himself how he will take risk into account when making decisions" (2014, 234).

- 1) In contexts where the probability of error and its cost in terms of the well-being of those affected are significant, a rational decision is the result of taking into account simultaneously three variables: a) the probabilities of success/failure of the different options at stake, b) the positive or negative value of the possible outcomes of these decisions, and c) the rational attitude toward risk of the agent making the decision.
- 2) Different courses of action may represent the best, most rational decision for different agents, even if they fully agree about their assessment of the available options in terms of variables a) and b). It is enough that they differ in variable c), their attitude toward risk, for such an outcome to be possible or actually occur.

As we have seen, expert *S1* cannot pass the first test of instrumental rationality in the corporate example. They are unable to argue that their decision is justified by the standard criteria. This standard mandates that any agent, upon assessing the same relevant factual information and subscribing to the appropriate normative framework, would necessarily reach the same conclusion and take the same decision. In cases where there is reasonable disagreement about the most desirable level of risk, there is no single correct decision acceptable to all. On the contrary, in this type of scenario, there are several equally rational responses depending on the specific circumstances of the different agents. The idea that there is a wide range of reasonable risk attitudes in cases such as the case of the firm or the design of a state's economic, health, or defense policies is intuitively plausible. The second test, however, provides the key to confirming this intuition. In the Miners case, there are compelling grounds to conclude that a single rational risk attitude exists. Given the certainty of saving 90 of the 100 workers by closing Gate 3, an agent typically exhibits a strong aversion to the alternative options, which carry a 0.5 probability of resulting in the death of all workers. This

conclusion is intuitively compelling. But what justifies this statement? The decision to close gates 1 or 2 involves taking a considerable level of risk. However, numerous contexts exist in which an equivalent level of risk-taking would be entirely rational. Why does this logic not apply here? The answer is that closing gates 1 or 2, rather than 3, leads to a suboptimal outcome in an iterative scenario.

The number of fatalities would tend to increase dramatically over time if, instead of closing Gate 3 every day when the accident recurs, the decision-maker sometimes closes one of the other two gates. However, this phenomenon does not occur in the case of the firm. Thus, an expert would fail the second test in this context. Each of the three options available to *S1* (*N1*, *N2*, and *N3*) could either produce a favorable outcome or lead the company to failure or extinction in an iterative scenario. *N1* maximizes expected utility for a risk-neutral agent and is equivalent to maximizing expected value—i.e., the objective value at stake. *N2* is the option that maximizes expected utility for risk-averse agents, while *N3* maximizes expected utility for risk-seeking agents. In this example, we assume that the agents must, for the purposes of the model, make a joint decision: the firm will invest all agents' capital in the same option, subjecting everyone to the identical profit and loss scheme.

Consider three possible worlds. Investments must be made daily, meaning profits are realized and losses incurred every day. Suppose that, regardless of the partners' varying risk attitudes over time or their associated preferences regarding firm strategy, a manager is appointed who consistently makes the identical decision. In the first possible world (PW1), the firm always invests in *N1*. In the second (PW2), it invests in *N2*, and in the third (PW3), in *N3*. Due to the Law of Large Numbers, over a sufficiently long time horizon, the firm's total profits in PW1 will tend to exceed those in PW2 and PW3. As the number of iterations approaches infinity, the divergence in accrued benefits in favor of PW1 should increase. Since

investment $N1$ maximizes expected value, this principle implies that as the iteration count approaches infinity, the cumulative return of selecting $N1$ is statistically certain to grow, leading to a realized outcome superior to the alternatives. Expected value maximization thus rests on the guarantee that, in the long run, the average frequency of outcomes will converge with their respective probabilities.

What knowledge do we have concerning the effects of these three strategies on *expected utility* and the firm's survival or extinction over a finite time period that is relevant to the partners' well-being? In reality, as we saw in the previous section, nothing is certain. In all three cases, it is possible that the partners and firms will have gone bankrupt—that they will have ceased to exist as economic entities—long before the real benefits of maximizing *expected value* are achieved. If the manager in PW1 had decided to invest in $N1$, a losing streak would have been highly likely at least in the medium-term. Most partners at the firm would have recovered from this streak if the investment had been made in $N2$ instead. Therefore, investing in $N1$ could have caused the firm to collapse. The same could happen in PW3 if a significant number of partners could not tolerate the even higher losses resulting from the $N3$ investment. If the manager invests in $N2$ in PW2, even if things go as well as hoped, the profits may not be sufficient for the majority of the partners' projects.

As a result, the outcome for the well-being of those involved would also be suboptimal. The most probable consequence of the continuous application of *the same* investment strategy is the disappearance of companies in all three possible worlds. Just as the infinite expected value, that St. Petersburg game effectively guarantees, does not protect players from swift ruin if they bet their entire capital, so too, in the corporate example, the unwavering repetition of a single investment option cannot guarantee that the majority of the partners will avoid bankruptcy. Investing invariably in $N1$, $N2$, or $N3$ is an appropriate strategy for agents exhibiting risk neutrality,

significant risk aversion, and significant risk propensity, respectively. None of the three options, however, guarantees a better aggregate outcome over the medium- or long-term time horizons relevant to the lives of human agents. Therefore, the expert is unable to pass the second test in the case of the firm. The decision-maker, irrespective of the choice, cannot claim that it is the optimal decision on the grounds that its reiterated application in the same objective circumstances would produce the optimal cumulative result, as occurs in the Miners case (closing Gate 3). Consequently, it is not possible to assert that any of the three possible attitudes toward risk in this scenario is more rational or reasonable than the others.

It is possible to offer a simpler and more abstract model of the problem—one that illustrates the case of the firm and underlies most public policy in risk contexts. Consider the following scenario, which we term “the benefactor game.” The task is to allocate a specific sum of money to three agents. Assume we possess no information regarding their specific circumstances, such as age, gender, initial wealth, or long-term goals. Two options are available: we can either grant each agent a guaranteed \$1,000, or assign each agent the result of a lottery paying \$2,000 on “heads” and \$0 on “tails” (using a fair coin). We may apply a uniform procedure for all three agents or employ any combination of the two options. The objective of the game is to maximize expected utility—that is, to select the strategy that most efficiently promotes the future well-being of those affected. We possess all relevant facts regarding the expected value of the options and the appropriate normative framework. However, even with this information—as in the corporate example—it is impossible to determine the optimal strategy to achieve the best outcome.

The explanation is simple. Suppose the three subjects desperately need the \$1,000 because not getting it would have a drastic negative effect on their well-being. From the perspective of the expected value, it makes no difference whether we give all three agents \$1,000 for sure or a 0.5 probability of winning

\$2,000. However, if it is vital for the three agents to have at least \$1,000 for sure there is a big difference. Similarly, if all three agents are given \$1,000 and it is vital for at least one of them to have \$2,000, while getting \$1,000 instead of \$0 does not make a significant difference to the other two, it would have been much better to subject the winnings of all three to the lottery. In this case, the agent for whom it was vital to get \$2,000 would have had a 0.5 chance of getting the amount of money they need. Of course, in other scenarios, different combinations of the two strategies may be more appropriate, depending on the divergent preferences of different agents.

Suppose we are competing against another player in the benefactor game. Our competitor consults the affected parties regarding their preferences and adopts a strategy inverse to our own: guaranteeing the first group \$1,000 while offering the second group the lottery option. The result is clear: we will lose the benefactor game. If we do not know each agent's risk attitude—the degree to which they are willing to wager current gains for greater rewards—the decision process fails to align with their subjective preferences. Without this critical information, it is impossible to determine the optimal outcome for each individual and for the group as a whole. This limitation mirrors the situation in the firm, where a manager or an external expert committee selects an investment strategy without knowing what the partners would have rationally decided had they fully evaluated the options. The same deficit arises when a committee of experts designs a nation's public policy. This holds true even in a highly idealized scenario where the committee operates solely on a correct normative framework and well-justified factual judgments. In all these cases, multiple reasonable attitudes toward risk exist. Therefore, relying on objective and impartial judgments—even those valid for all—is insufficient to identify common interests and make rational decisions.

Stated precisely, the challenge in determining the best strategy stems not from epistemic failures, but from the

existence of divergent, reasonable, and inherently subjective attitudes toward risk. Consequently, a decision-making process is required that can adequately elicit and aggregate this information to produce rational collective decisions. Our proposed theoretical framework identifies the integration of affected parties' risk attitudes as a fundamental instrumental function of a democratic system. In this regard, our proposal contrasts sharply with epistocracy.

The appeal of epistocracy lies in the possibility of reaching correct collective decisions without relying on a complex participative process. This approach assumes, for instance, that contentious political issues—such as designing a state's economic policy—can be resolved in the same manner an expert handles a problem like the Miners case. Even if citizens lack knowledge of relevant facts and hold unreasonable value judgments, experts can theoretically still determine optimal policy decisions. Epistocracy thus eliminates the need to hypothesize how the public would have decided had they been fully informed and adopted an appropriate normative framework. However, this premise fails in contexts of risk. Most significant state decisions occur in scenarios analogous to the case of the firm, not the Miners case—scenarios where a political solution is required because a purely technical or genuinely moral solution is unavailable. In such contexts, it is impossible to determine the most appropriate decision without accounting for the divergent risk attitudes of the affected parties.

Every decision made by experts inherently involves an attitude toward risk. To adhere strictly to objective and impartial considerations—specifically, the maximization of expected value—when success is uncertain and the cost of error is high, is effectively to adopt a risk-neutral attitude. Yet, there is no *a priori* reason to assume that risk neutrality constitutes the optimal strategy for all stakeholders across varying contexts. Furthermore, it is not necessarily the superior strategy within a time horizon relevant to the welfare of those affected. As

previously noted, the option that maximizes expected value guarantees the optimal outcome only in an iterative scenario where the number of iterations approaches infinity. If multiple reasonable risk attitudes exist—which is the norm—and expert decisions reflect a risk profile diverging from that of an informed majority, we cannot assume these decisions will promote the majority’s well-being in a sustainable manner within a timeframe relevant to real agents. This discrepancy between the risk attitudes of expert minorities and the affected majority is likely structural. In societies with marked inequalities in income and education, experts typically belong to advantaged social groups and, consequently, often possess different risk profiles than vulnerable agents who have more to lose. Under these conditions, it is probable that even with the best intentions, an expert government lacking democratic control will jeopardize the majority’s living conditions. This dynamic closely parallels the corporate example, where a manager makes investment decisions without adequately integrating the preferences of the affected stakeholders.

In the real world, this danger is compounded by the inevitable biases, errors, and often hidden forms of partiality that affect political decisions. If the outcome can be highly unfavorable to the majority even in an idealized scenario with an expert minority government, it is logical to anticipate far more adverse scenarios in practice. Many defenders of the democratic system interpret epistocratic or technocratic discourse as pure ideology aimed at legitimizing the rule of powerful elites. Irrespective of the motivations of its proponents—particularly within academia—there is no plausible prospect of epistocracy yielding a different long-term impact. This conclusion derives from the logic and requirements of risk management examined previously. Crucially, this result persists even if experts were to adopt the appropriate normative framework and provide a fully competent epistemic evaluation of the relevant empirical information.

3. Risk-Centered Ordinal Consequentialism as a Foundation for Contestatory Democracy

In scenarios such as the case of the firm or government economic policy, we cannot achieve the same degree of unequivocal correctness that is possible in situations like the Miners case. As demonstrated previously, even if a small group of experts identifies the optimal decision based on well-founded factual judgments and an appropriate normative framework, that decision may still be validly challenged by those affected, including the potential majority.

This challenge arises because the decision might not constitute the most instrumentally desirable outcome at the collective level. Indeed, it is often impossible to determine *a priori* whether the course selected by experts will contribute positively or, conversely, prove detrimental to the welfare of those affected in the relevant time horizon. However, this uncertainty does not imply the absence of a rational strategy for collective decision-making in such contexts. We now turn to examining the most effective methods for achieving optimal outcomes in these situations.

Revisiting the simplified model of the benefactor game, we face a central question: how can we maximize the welfare of those involved? The solution is evident, provided we assume stakeholders possess informed preferences (i.e., preferences whose satisfaction is likely to contribute to their effective well-being). In this context, the optimal strategy is to consult the affected parties and enable autonomous decision-making. It is reasonable to posit that agents are best situated to determine their admissible risk levels—specifically, whether to accept a certain \$1,000 or opt for the lottery with a potential \$2,000 payoff. Granting each individual the agency to manage their own risks maximizes expected utility at both the individual and collective levels. Furthermore, compelling reasons exist to anticipate that this approach will yield superior realized outcomes, particularly when the strategy is iterated over time

(diachronically) or across agents (synchronously). While agents are fallible, such errors are likely exceptional. Systematically allowing agents to align decisions with their risk profiles is expected to generate the optimal aggregate outcome over a sequence of simultaneous or successive decisions.

In summary, mirroring the principle of an ideal democratic system, the expected quality of outcomes is directly proportional to the control exercised by those affected by the decisions. It is crucial to emphasize that distinct groups of individuals may rationally arrive at different conclusions. Consider two extremes of the risk spectrum. If all members of a group are risk-averse, they will unanimously prefer a guaranteed payment of \$1,000. Conversely, in a group where every agent is risk-seeking, each will prefer the lottery with a potential \$2,000 payoff. Where unanimous support is unattainable and individual preference-based choices are unfeasible, we can posit that the optimal mechanism is the application of majority rule. While the majority in one group may select the safe option, another group's majority may opt for the lottery. This demonstrates that no single decision is globally optimal. Furthermore, regarding underlying factual knowledge (including probability assignments) and the evaluative framework, there is no inherent basis to value one decision over the other. In the benefactor game—unlike the firm scenario—the expected monetary value of all possible decision combinations is mathematically identical (always totaling \$3,000 for the three agents). However, the expected utility derived from these options can vary drastically across different groups.

Let us posit that the three agents in the group have a critical need for \$1,000, such that the possibility of winning \$2,000 offers no significant marginal advantage. Suppose we lack this information and decide to allocate funds via a fair coin toss (\$2,000 or \$0). As previously noted, from an expected value perspective, this strategy is equivalent to providing the guaranteed \$1,000. The mathematical expected value of both

decisions is identical. However, from the perspective of the affected agents, the expected utility of the riskier option is significantly lower.

To optimize social welfare in the benefactor game, we require a mechanism that accounts for the subjective valuations of the affected parties by consulting their actual preferences. This logic supports the conclusion that if agents possess divergent preferences due to distinct risk profiles, satisfying the majority will always represent an alternative superior to minority rule. If two agents require the guaranteed \$1,000 while the third prefers the lottery—and assuming we aggregate preferences from a purely ordinal perspective, disregarding intensity—the expected utility of guaranteeing \$1,000 to all is demonstrably greater than imposing the lottery on the dissenting majority.

The same considerations apply to the case of the firm. However, in this scenario, sophisticated strategies exist that transcend mere reliance on majority rule for optimizing expected outcomes. How do we obtain the maximum expected utility while incorporating maximum information about the stakeholders' attitudes toward risk?

The maximum possible expected utility—whether evaluated from a cardinal or ordinal point of view—is also achieved at two opposite ideal poles. The first is a unanimous, free, and informed consensus around a single decision. The second is what we can term a scenario of “deregulation without domination.” Instead of imposing a uniform decision on all parties, individuals are empowered to make their own choices. Using the best information and a sound set of values, each person can select the option that best fits their personal risk profile. Consider two companies, A and B. In Company A, all partners are highly risk-averse and prefer investment *N2* to the others. In Company B, the partners have divergent risk profiles, organized into three distinct groups. Some partners prefer *N1*, others *N2*, and still others *N3*. They all agree to dissolve the company, perhaps temporarily, so that each partner can join the subgroup where the investment most closely matches their risk profile.

In both scenarios (Company A and Company B), no partner is compelled to assume risks incompatible with their objectives. Consequently, each individual optimizes their exposure to potential outcomes—whether profits or losses—ensuring the results remain desirable or, at minimum, tolerable. For instance, if a risk-averse agent preferring *N2* is forced to invest in *N1* or *N3*, they may incur losses incongruent with their future stability, potentially leading to irreparable consequences. Conversely, an agent seeking returns exceeding the capacity of *N1* or *N2*, yet constrained by a conservative majority, may find the realized yields insufficient for their projects. The defining characteristic of this scenario is the convergence of maximum expected utility and maximum individual freedom. Agents can act in accordance with their risk attitudes only when free from coercion or arbitrary external control. Curtailing individual freedom undermines the system's instrumental potential, as it deprives the decision-making process of critical information: the authentic preferences of agents, which remain obscured under external pressure.

When decisions must be made under risk, maximizing expected utility structurally converges with maximizing individual freedom. In such cases, the optimal outcome is one where decisions are consistent with the risk profiles of all stakeholders, achieved either through unanimous consensus or through deregulation, rather than the imposition of a uniform solution.

This logic stands in sharp contrast to scenarios such as the Miners case, or any context where the premises of epistocracy—specifically, the capacity of experts to unilaterally reach rational decisions—are warranted. In such instances, while consensus remains valuable, its role is subsidiary or derivative. If an expert determines that a specific action, such as closing Gate 3, is mandatory because it maximizes expected utility, it is naturally preferable that the stakeholders understand the rationale and support the decision. However, the overriding concern remains the implementation of the optimal outcome.

Consequently, the expert has no justification for sacrificing the objectively superior course of action—or selecting a suboptimal one—merely to broaden consensus. Moreover, were deregulation possible in this context, it would prove counterproductive. Unlike the firm scenario, granting agents the license to make decisions that are suboptimal from an objective and impartial standpoint inherently lacks instrumental value.

Consider a scenario where two identical accidents occur simultaneously in separate mines. In the first instance, the expert decides to close Gate 3 (the safe option). In the second, the agent in charge adopts a high-risk strategy: closing Gate 1 in an attempt to save everyone, despite the risk of saving no one. The decision made by the second agent yields a suboptimal expected outcome. This demonstrates that, in such contexts, maximizing expected utility and exercising individual freedom do not always coincide. However, the situation differs fundamentally when the probability of error is significant, the cost of error is high, and no single rational attitude toward risk prevails. In such a scenario, it is highly probable that a diverse range of actions will maximize the expected utility for different agents, depending on their respective risk profiles. Consequently, maximizing *global* expected utility requires approximating a condition of free, informed consensus or, alternatively, deregulation without domination.

As illustrated by the case of the firm, the repeated implementation of decisions misaligned with the risk profile of the majority of those affected will predictably lead to institutional deterioration or collapse. In such scenarios, restricting an individual's freedom to assume the level of risk they are willing to bear proves detrimental to the quality of the aggregate outcome. As demonstrated by the benefactor game, when neither consensus nor full deregulation is attainable, majority rule functions as the second-best alternative. Conversely, the imposition of minority preferences upon the majority yields the worst expected outcome. Given that normative validity is anchored at the poles of informed

consensus and deregulation, any partial advancement toward either ideal—moving beyond the baseline of majority rule—is desirable and yields a significantly positive expected impact. Furthermore, it is possible to pursue progress toward both poles simultaneously.

Consider the following illustration. Imagine a scenario in which a majority of the firm's stakeholders initially supports the decision to invest in option *N1*. Subsequent deliberation, however, identifies an alternative investment: option *N1.1*. This new option presents a modified risk profile: it yields a lower profit than *N1* if successful, but entails a smaller loss if unsuccessful. Crucially, the overall utility profile associated with *N1.1* proves superior to that of option *N2*. Assume that the stakeholders who originally favored *N1* are not materially disadvantaged by adopting *N1.1*, as they retain the opportunity to achieve the returns necessary for their projects. Concurrently, a subset of the partners who supported *N2* (and rejected *N1*) deem the potential losses from *N1.1* to be tolerable. If *N1* was the initial selection, revising that decision in favor of *N1.1* becomes entirely reasonable. This revision—whether facilitated by renewed debate or alternative means of contestation—is justified by the demonstrable expansion of consensus. The broader majority support for *N1.1* thus reflects an increase in the aggregate expected utility of the chosen option.

However, the utility of broadening consensus is rendered irrelevant if *N1* represents the correct option in the sense of a settled factual or moral judgment. In such an instance, shifting from *N1* to *N1.1* would necessarily decrease, rather than increase, expected utility. This is analogous to the Miners case: if closing Gate 1 is known to save 90 workers, selecting a fourth option, Gate 4—which saves only 80—would demonstrably reduce utility. While Gate 4 is preferable to 1 or 2, there is no justification for selecting it over 3.

Under what conditions would investing in *N1* be superior to investing in *N1.1*—the option generating the most consensus—

just as closing Gate 3 is objectively superior to closing a hypothetical Gate 4? This would occur only if investing consistently in *N1* over time tended to produce the optimal aggregate expected result. In other words, the situation would have to mirror the repeated Miners case scenario, where the optimal response is invariable: the decision a risk-neutral agent would make (closing Gate 3). However, this is precisely what fails to occur in the case of the firm. In such a scenario, the invariable repetition of any single option never guarantees the best aggregate expected outcome relevant to the well-being of those affected within a meaningful timeframe. On the contrary, the optimal strategy is to adapt to the shifting free and informed preferences of those affected, even if such shifts depend solely on a change in the prevailing risk profile. If *N1.1* can generate a broader consensus than *N1* and other options at a given moment, it represents—at that particular juncture—the best decision. However, neither *N1* nor *N1.1* represents an intrinsically correct decision. A different alternative may generate greater support in a future scenario, even if the relevant facts and the accepted normative framework remain unchanged. In a system where no single decision is intrinsically correct, this consensus-driven process—governed by majority rule yet open to future correction and contestation—is an adequate mechanism to approximate the best overall result.

A comparable enhancement of expected utility can be achieved through a system of partial exemptions—often termed a strategy of “partial deregulation.” This approach aims to protect minorities whose risk profile differs from that of the majority, particularly when achieving a high level of consensus proves infeasible. Crucially, if the minority position is factually incorrect and the majority’s is correct, there is no instrumental justification for such “protection,” as it would be counterproductive to the quality of the outcome. However, in contexts of risk, this form of deregulation holds significant instrumental value. If *N1* is replaced by *N1.1*, yet agents remain who prefer *N2* or *N3*, this strategy suggests that a minority

could be exempted from the full cost burden of the majority decision. The objective is neither dissolution nor institutional schism, but rather the compensation of the most vulnerable partners for probable losses incurred when the majority pursues risks that exceed their tolerance thresholds. This compensation could take the form of an insurance mechanism, financed via a tax on prior company profits, designed to prevent the bankruptcy of these agents should the majority's high-risk decisions fail. While this mechanism may reduce short-term aggregate profits, its implementation is expected to enhance the long-term survival and stability of the group. Conversely, at the other extreme, a mechanism may be required to mitigate the lost opportunity costs for a minority willing to accept greater risks for potentially greater rewards. For instance, while the business may require the collective capital contribution of the majority for a primary investment, the arrangement could grant a minority the autonomy to pursue alternative, high-yield investments without dissolving the central enterprise.

As previously argued, this framework would lack justification if the invariable repetition of a single decision reliably produced the optimal aggregate outcome over time—that is, if there were truly “correct” decisions. Such a premise would imply that, in scenarios where the probability of error is non-trivial, a single reasonable attitude toward risk exists, alongside well-founded empirical information and an adequate normative framework. Conversely, where this condition does not hold, the quality of the overall expected outcome is substantially enhanced by adopting decisions that command the broadest possible support, or by exempting minorities from the costs or risks of majority decisions—provided this is feasible without substantially compromising the utility of the options favored by the majority.

The same logic extends to the political arena, constituting one of its defining features. To illustrate, consider four hypothetical societies (A, B, C, and D) whose citizens must decide on an economic policy. Assume the initial conditions across these four

societies are substantially identical: population size, level of industrial and technological development, access to natural resources, and educational standards. In short, they share a comparable objective baseline for evaluating the available policy options. Suppose that in each society, citizens must choose between three economic policies guided by a Utilitarian normative framework ($UEP1$, $UEP2$, $UEP3$) and three policies guided by a Rawlsian framework ($REP1$, $REP2$, $REP3$). The three options associated with each framework are analogous to the firm's investments ($N1$, $N2$, $N3$). In both cases, the first option ($UEP1/REP1$) represents the choice a risk-neutral agent would make; the second ($UEP2/REP2$) corresponds to a risk-averse agent; and the third ($UEP3/REP3$) reflects the preference of a risk-seeking agent. We can assume that $UEP3$ and $REP3$ present the sharpest contrast, representing the most divergent outcomes. Proponents of these measures are willing to incur significantly greater risks to depart radically from the status quo in pursuit of their normative ideals. Conversely, regarding the distribution of economic resources, $UEP2$ and $REP2$ represent the most cautious or moderate reform projects aimed at realizing a just society from their respective perspectives.

Suppose that in Societies A and B, a majority of citizens adhere to a utilitarian framework. In Society A, the majority demonstrates a higher tolerance for potential losses in pursuit of maximized net income, compared to Society B, where a more risk-averse attitude prevails. Consequently, the majority in Society A prefers the utilitarian economic policy $UEP1$ over the more conservative $UEP2$ (assuming $UEP3$ is rejected). Conversely, the majority in Society B is unwilling to incur the costs associated with the failure of a policy like $UEP1$. This distinction parallels the case of the firm: for Agent $S1$, the expected utility of investment $N1$ was maximal, whereas for $S2$ —the most conservative agent—option $N2$ yielded the highest expected utility. Similarly, the majority in Society A assigns a higher expected utility to $UEP1$ than to $UEP2$,

whereas the preference ordering is reversed for the majority in Society B. Therefore, *UEP1* represents the most rational policy choice for the majority in Society A, while *UEP2* constitutes the most rational choice for the majority in Society B.

Let us posit that a majority in Societies C and D subscribes to a Rawlsian conception of justice.¹⁷ Consistent with this framework, the three available economic policies (*REP1*, *REP2*, *REP3*) are designed to secure an improved position for the least advantaged compared to potential utilitarian outcomes, albeit at the expense of reducing overall societal wealth. However, the two societies differ regarding their tolerance for the risks associated with applying the Difference Principle. In Society C, the majority supports a strongly redistributive policy, *REP3*, which carries a higher risk of compromising the system's overall efficiency. In contrast, the majority in Society D adopts a more cautious attitude, supporting policy *REP2*, which entails a lower level of redistribution. The policies are defined by their specific risk profiles:

17. As previously clarified, rather than strict “Rawlsians,” we refer to agents who subscribe to the difference principle to evaluate their society's economic order. We set aside the complexities involved in a consistent application of Rawls's complete theory, which would necessitate applying the first principle of justice (regarding civil and political liberties) and the principle of fair equality of opportunity. We are not evaluating Rawls's actual position, nor the various forms of utilitarianism or prioritarianism, as theories of justice. We remain neutral regarding their ultimate moral correctness. Instead, we accept the hypothesis that appealing to such normative frameworks is an essential part of the political process. Adopting a plausible interpretation of one of these positions is arguably a necessary condition for high-quality political outcomes. However, our central contention is that these theories do not, in themselves, offer adequate justification for state decisions in contexts of risk. Even when combined with accurate factual judgments, they cannot guarantee the quality of political results. Achieving this requires an adequate risk management system—one that introduces the logic of consensus-seeking and the maximization of individual freedom characteristic of successful constitutional democracies.

- *REP1* (The Baseline): This policy is projected to generate the maximum sustainable redistribution (the standard Difference Principle optimum). However, there exists a low but significant probability that this projection is incorrect. The failure of *REP1* would impose a high cost on the welfare of the least advantaged, leaving them vulnerable to economic crisis.
- *REP2* (The Conservative Option): While this option generates less redistribution than *REP1*, the contingent cost of failure to the well-being of the least advantaged is commensurately lower, thereby better preserving the security of the most vulnerable.
- *REP3* (The High-Risk Option): If successful, this policy generates a higher level of redistribution than both *REP1* and *REP2*. However, in the event of failure, the resulting cost to the least advantaged is the most severe among the three options.

The Rawlsian normative framework does not specify the appropriate level of risk necessary to achieve its goals; therefore, it does not dictate a definitive choice among these three policy options. The framework itself is accepted by the majority, as it mandates prioritizing the improvement of the living conditions of the least advantaged (i.e., adherence to the Difference Principle). The key locus of decision, therefore, is determining the acceptable risk tolerance for maximizing the benefit to this group, rather than the acceptance of the conception of justice itself. On this specific matter, however, there persists a deep and reasonable disagreement.

The normative logic inherent in the epistemic model is flawed in contexts involving risk. Consequently, implementing any of the aforementioned policies (*REP1–3* or *UEP1–3*) solely based on their purported “correctness”—even when supported by duly justified factual judgments and operating within the appropriate normative framework—constitutes an unwarranted gamble if those affected are not included

and lack a meaningful voice in the decision-making process. Implementing such a strategy without participatory input jeopardizes the long-term prosperity, or even the survival, of the political system. Specifically, if the implemented policies either exceed the level of tolerable risk or fall below the risk threshold that the majority of citizens in the affected societies are willing to accept over the medium-to-long term, a significant deterioration of their situation will result. If a decision were genuinely the right one for all stakeholders—such as closing Gate 3 in the Miners case—it should be applied consistently in all contexts, under the assumption that it will produce the best expected aggregate outcome. However, this condition of universal applicability is precisely what fails in the case of the firm, the benefactor game, and the most consequential political decisions.¹⁸ Consequently, a majoritarian strategy is

18. For this reason, it is rarely viewed as problematic that different democracies make different decisions on almost all relevant issues: health, education, defense policy, immigration, etc. Such divergence is not necessarily a sign of error or systemic failure. Admittedly, many of these differences can be explained by varying objective circumstances and value-based disagreements. However, our point is that even if it were possible to overcome the disagreements arising from these factors, significant disagreements would predictably remain. Consequently, it would be necessary to appeal to a political logic of decision-making distinct from an epistemic or moral one—regardless of whether moral judgments possess epistemic status. Schamberger (2023) formulates a meta-ethical dilemma regarding the justification of majority rule via Condorcet's Jury Theorem (CJT). He points to the fact that the political decisions of advanced democracies tend to diverge drastically over time (diachronically) within the same society, and synchronically when comparing different nations. This appears incompatible with a realist conception of moral judgment. As Schamberger argues, if we adopt a realist meta-ethical position, we must recognize that: "(1) moral facts do not change constantly, and (2) they do not differ from country to country, or from society to society. Accordingly, under the assumption of moral realism a striking congruence of democratic decisions is to be expected if the conditions for the application of the CJT [Condorcet's jury theorem] are met. In actual fact, majority views on most political issues are subject to constant change, which is well recorded by

required: the effective aggregation of informed preferences that reflects the risk attitudes of those affected, while remaining permanently open to contestation. This openness to revision serves not merely the epistemic goal of correcting factual or normative errors (though it certainly performs that function); rather, it accommodates the subjective, partial, and evolving preferences of the individuals involved—specifically, their divergent attitudes toward risk, which cannot be reduced to a single common or generalizable interest.

However, majority rule constitutes only a preliminary mechanism. Whenever it is feasible to approximate the normative ideal poles of free, informed consensus and deregulation without

opinion polls and changing election outcomes” (7–8). This seems to suggest a dilemma: either contemporary democracies fail to satisfy the conditions for the CJT—and thus cannot be presumed to approximate the truth—or these concepts are incompatible with a realist perspective. Our approach does not require adopting a realist moral position. Rather, it demonstrates that moral realism is compatible with the hypothesis that voters possess the average epistemic competence required by the CJT ($p > 0.5$), despite the widespread disagreement observed diachronically (within the same community over time) and synchronically (across different communities). Even if participants base their decisions on correct or justified moral judgments, changes in the risk attitudes of those affected can lead to very different decisions over time, or to divergent decisions being made simultaneously by different communities. This does not explain—nor does it attempt to explain—all disagreements observed in this area (which often depend on other factors). However, it does account for a significant number of disagreements that might otherwise call into question the epistemic quality of democratic decisions. Furthermore, it helps to reconcile diversity, disagreement, and freedom with the instrumental value of the democratic system. In contrast, the epistemic model—and conceptions that reduce politics to morality—are characterized by an oscillation between two extremes. This ranges from (1) an authoritarian inclination to assert that only a very limited set of correct political responses exists (towards which all democracies ought to converge), to (2) a swift and profound skepticism regarding the rational pursuit of the common good in contexts of risk—a skepticism that arises once the implausibility of the aforementioned authoritarian goal becomes apparent.

domination, there exist decisive instrumental reasons for doing so. Consider Society B, characterized by a risk-averse utilitarian majority. Suppose it is possible to broaden consensus on the implemented policy by combining the goal of maximizing total societal wealth with the provision of a guaranteed minimum income for the most vulnerable sectors. A policy with these characteristics—designated here as *UEP2+M*—would possess a superior expected utility profile compared to *UEP2*, the option previously favored by the utilitarian majority. This is because the new proposal could garner support from members of the minority who subscribe to a Rawlsian or prioritarian conception of justice. While it is possible that *UEP2+M* yields a lower aggregate net wealth than *UEP2*, it represents a more desirable outcome regarding the quality of the expected result: it remains acceptable to the utilitarian majority while successfully attracting the endorsement of the Rawlsian or prioritarian minority, who perceive the guaranteed social minimum as a significant advancement in the pursuit of justice.¹⁹ Less advantaged citizens, unwilling to risk falling below a minimum standard of well-being solely to maximize societal welfare, possess compelling prudential grounds to support this proposal. This policy enhances the expected outcome by reducing the risk exposure for a demographic broader than that encompassed by the initial majority preference. While this implies a conceptual compromise—or a less rigorous application of the underlying normative framework—it should not be viewed as a degradation. Although it may appear as a suboptimal deviation from a strictly epistemic or moral perspective, we contend that, under the consequentialist conception defended here, it represents rational progress

19. Rawls recognizes the similarity—at least in practice—between his theory of justice and a utilitarian theory that constrains the principle of utility by providing a social minimum for the least advantaged (see Rawls 2001, 127–9).

toward the best expected aggregate outcome. We posit that this consequentialist approach captures distinctively political features that differentiate it from the mere pursuit of epistemic truth or moral correctness.

Consider a case of partial deregulation in Society C, where a Rawlsian majority imposes a stringent redistribution program. Assume that specific economic agents can be exempted from the tax burden of this program (Policy *REP3+E*). These agents are willing to accept elevated risks in exchange for potentially higher returns, yet their investments appear unattractive if yields are significantly attenuated by the prevailing tax regime. Granting this form of partial—and potentially temporary—exemption affords this minority the freedom to pursue their ventures. Provided this exemption remains consistent with the system's broader redistributive goals—for instance, by enhancing overall economic efficiency—it is instrumentally desirable because it maximizes aggregate expected utility. If such exemptions are viable and compatible with the system's continuity, a greater number of individuals gain the economic freedom to pursue projects informed by their respective risk attitudes. Thus, *REP3+E* is superior to *REP3*—the option initially favored by the majority—in terms of maximizing global expected utility. The realization of such corrected policy outcomes—whether shifting from *UEP2* to *UEP2+M* (classic utilitarianism with a social minimum), or moderating a Rawlsian regime to foster market freedom (*REP3* to *REP3+E*)—can be attributed to different political processes. These include corrections driven by electoral cycles, the iterative application of majority rule, or recourse to mechanisms of political and judicial contestation.

In contexts of risk, achieving optimal outcomes depends upon the continuous availability of mechanisms for reviewing and adapting decisions. This dynamic process is crucial, driven by the evolving, informed preferences of stakeholders, as well as by open political competition. Undoubtedly, the complexity of the political landscape far exceeds that of the idealized case

of the firm. However, the fundamental logic remains entirely applicable. This logic necessitates two key protective measures:

- 1) Establishing safeguards for the most vulnerable partners against majority decisions favoring ambitious investment or political programs (assuming these partners cannot simply withdraw their capital).
- 2) Developing mechanisms to mitigate or compensate for the opportunity costs imposed by a conservative—risk-averse—majority policy.

These practices, aimed at reducing disagreement or expanding the scope for deregulation, are ubiquitous in actual democratic politics. As noted, their value may appear dubious from a strictly epistemic or moral perspective. Nevertheless, for the reasons discussed, they possess decisive instrumental value: they represent the most effective strategy for enhancing the quality of the expected political outcome in contexts of risk. Crucially, managing reasonable disagreement in collective decisions under risk entails a distinctively democratic political logic. This logic is fundamentally different from that which characterizes attempts to address disagreement in the domains of factual knowledge, morality, or justice. Indeed, it would be fair to say that these are opposing or antagonistic logics in some respects.

When individuals disagree on a decision—either due to conflicting beliefs about matters of fact or incompatible moral judgments—the quality of the outcome depends entirely on whether the true judgment prevails. Ascertaining the truth in these cases appears to be both necessary and sufficient for making a rational decision. This perspective is shared by both epistocrats and epistemic democrats. For instance, when two individuals must decide whether to consume a certain food, and one is convinced it is perfectly safe while the other believes it is poisoned, the quality of the outcome depends entirely on whose judgment is correct. Individual autonomy in decision-

making does not inherently possess instrumental value in such instances. This principle similarly applies in the moral realm, assuming the existence of true or false judgments in this domain. The desirability of the outcome is contingent solely upon the prevalence of the true judgment. The Miners case illustrates this logic within a standard consequentialist framework. Let us consider the perspective of one of the ten workers who will die if Gate 3 is closed.

Assume that this agent is fully aware of the consequences: they know that the cost of the correct policy—saving 90 miners with certainty—is their own death. From a prudential point of view, this agent might prefer that, instead of closing Gate 3, the authorities close one of the other two gates. This would increase their probability of survival from 0 to 0.5. However, the worker knows that they have no moral right to demand such an alternative; they must accept the sacrifice for the greater good. Their objections to the authority's decision to close Gate 3 would carry no moral weight. This logic of sacrificing individual interests—and even the individual's existence—for a greater common good (or for “truth” or “what is right”) stands in antithesis to the logic that governs democratic politics. While this consequentialist logic may be appropriate for situations like the Miners case, it is ill-suited to contexts like the case of the firm or the analogous domain of economic policy.

When a decision must be made in a context where the probability of error is significant, the cost to the well-being of those affected is high, and reasonable disagreement exists regarding the most rational attitude toward risk, maximizing individual freedom constitutes the strategy that yields the best expected outcome. The logic of democratic politics, viewed as a strategy for achieving the optimal medium- to long-term result, is to avoid—or at least minimize—the sacrifice of individuals or their partial interests. Democratic politics constrains—through institutional design, primarily majority rule and contestatory institutions—the practical power of narratives that seek to morally justify the opposing strategy. In scenarios such as the

case of the firm or the debate on economic policy—unlike the Miners case where ten workers are destined to perish—those whose partial interests are endangered, or who face a non-consensual risk, always retain the right to voice dissent and persist in political contestation. The option to reopen political debates and the potential for contestation must remain perpetually available. This accessibility serves not only to correct factual or moral errors but, more importantly, to prevent harm or the unacceptable marginalization of individuals and their interests to a degree that the mere pursuit of truth or moral correctness—divorced from the political process—cannot guarantee, or can only achieve contingently. When there is reasonable disagreement about the level of risk to be faced, no impartial moral or objective scientific discourse is sufficient to justify sacrificing anyone. Every decision is inherently partial; without a political system that applies majority rule and possesses institutions oriented toward consensus-building and deregulation, the system’s medium- or long-term outcome will be suboptimal. The appeal to truth—whether religious, moral, philosophical, or scientific—as the ultimate framework for political decision-making is not only a tendency toward totalitarianism, as Arendt argued,²⁰ or a form of sectarianism,

20. Arendt believed that appeals to truth were incompatible with the nature of politics. Political discourse—especially in a democracy—is dominated by persuasion, by the search for voluntary agreement and consent. However, the kind of validity to which an agent commits themselves when claiming that their position is true (and not just a well-founded opinion) would involve “an element of coercion” with “tyrannical tendencies” (see Arendt 1977, 239). This is because recognizing that a belief or statement is true is tantamount to placing it “beyond agreement, dispute, opinion, or consent”; it is tantamount to closing the debate and abandoning the field of politics (Arendt 1977, 240). From Rawls’s perspective, appealing to truth as a justification for a common coercive norm is likewise inadequate. Given the burdens of judgment, public justification of such an appeal is unfeasible. Those affected could always reject the norm or policy in question on the basis of reasonable arguments, rendering its implementation arbitrary. As

as Rawls pointed out, but also a primitive and irrational strategy from an instrumental point of view.²¹

Rawls argues, when people claim that their opinions should be imposed on others simply because they are true—knowing they cannot publicly prove this claim—this insistence can only be interpreted as an attempt to arbitrarily impose their will on others (see Rawls 1993, 61).

21. According to Landemore, this type of approach renders political practice—as we know it—meaningless. As she notes, the fact that “[...] there are better and worse answers and that those can be known is already validated by the fact that there would be no point in reasoning about politics if it wasn’t the case”. Indeed, the existence “of a procedure-independent standard and the possibility of knowing it are inscribed in, and presupposed by, the very nature of our discursive exchanges and, specifically, our ways of engaging each other’s reason in political arguments” (Landemore 2017, 284; see also Martí 2006). The proposals of Rawls and Arendt have been the object of similar criticism by Cohen (2009) and Beiner (2008), respectively. According to Beiner, “Arendt was incontestably right to put the emphasis that she did upon the diversity of human opinions. But at the same time, we should be careful not to fetishize cognitive diversity. As much as people differ in the content of their judgments, they are drawn together in political argument and rational debate because of what they *share*: a common aspiration to apprehend moral and political truth. [...] Our opinions divide us, but we wouldn’t be motivated to formulate our opinions, articulate them, and thus share them with others unless we were participants at a deeper level in a shared quest for moral and political truth that draws us back together as a community of truth-seekers. Exercising political judgment means that holding differing opinions never separates us sufficiently to cancel out this community of shared aspiration” (2008, 132). Similarly, Cohen challenges the idea—attributed to Rawls—that the concept of truth has no place in public reason: “[T]hat claim makes the idea of public reason unnecessarily contentious, as if it were committed to the view that the truth about justice does not matter. It also makes the idea of public reason hard to understand, because it proposes to leave the concept of truth behind while preserving notions of belief, assertion, judgment, reason, and objectivity, all of which are essential to an idea of public reason” (2009, 5). While these criticisms may be misplaced regarding Arendt and Rawls (a question we set aside here), they fundamentally do not apply to the approach we are developing. The conception we defend shares with Arendt and Rawls the thesis that appeals to factual or moral truth have no autonomous or self-sufficient justificatory force in politics. Arendt’s response is to replace discourse centered on the search for truth with “representative thinking,”

Promoting the instrumental quality of the political process presupposes a reduction of moral and factual disagreement; it necessitates an endeavor to approximate truth or a standard of correctness in both areas. To achieve the best expected outcome, one must indeed approximate factual and moral truth within the process. However, when there is disagreement about the level of risk to be taken, this logic must be subordinated to the logic of freedom. This third source of reasonable disagreement—disagreement about the level of risk in collective decisions—introduces the distinctively political logic. This logic mandates that the ultimate framework for decision-making be defined by the goal of identifying majority preferences, broadening consensus by modifying options, and identifying spaces for partial deregulation or special protection for minorities in a dynamic process always open to revision. The instrumental quality of the entire process depends on the adequate and simultaneous management of all three sources of disagreement.

centered on opinion formation in the political arena. For Rawls, the strategy is to replace the appeal to truth with the practical—not merely epistemic—concept of the reasonable. Our proposal does not seek to replicate these strategies or to adopt a revisionist stance toward political practice. Rather, we recognize that the appeal to truth—the search for objectivity and impartiality—while ineradicable and essential to the meaning of political practice, cannot possess full justificatory force within the democratic process. Within this process, the final word and direct decision-making power are reserved for electoral procedures (dominated by majority rule) and institutions that favor the normative ideals of free consensus-building and deregulation—even when these specific decisions conflict, in the view of some affected agents, with the decision that would result from a direct appeal to factual or moral truth. The ultimate goal of the political process is to maximize freedom. From the perspective of a purely epistemic logic, these institutions and practices may appear incoherent. However, in light of our framework, their priority over direct appeals to truth depends not on their intrinsic value, but on their instrumental value. As we have seen, in risky contexts, maximizing the expected quality of outcomes by rational means requires maximizing the influence of the free and partial will—specifically, the attitude toward risk—of those affected in the decision-making process.

Failure in any one dimension is sufficient to compromise the instrumental value of the whole decision-making system.²²

A final clarification is in order. We do not suggest that the issue of risk should play a central role in the public sphere or in the deliberations taking place within political institutions. To assume so would be a misinterpretation of our proposal. Political debates typically center on disagreements concerning factual and normative (prudential or moral) judgments. We do not posit that this practice should be modified to make the risk involved in policy proposals an explicit topic of debate. While such discussions may arise spontaneously, they are not a requirement of our framework. Our hypothesis is that divergent attitudes toward risk among agents possess a key instrumental value in political activity, just as they do in economic activity. This subjective and dispersed information must be incorporated and assessed in the decision-making process to ensure optimal outcomes. However, as previously stated, the sphere for processing this information lies primarily within political practices and institutions characterized by majority rule and mechanisms of contestation. Disagreements regarding the acceptable level of risk in joint decisions are not necessarily best addressed through a process aimed at the search for truth or correctness. Rather, they are processed through aggregative mechanisms guided by a distinct normative logic. These mechanisms point toward the two key

22. Contrary to a pervasive common-sense illusion—one likely inherent to the participant's understanding of the practice—the moral point of view and the search for objective truth are inadequate compared to the strategy of republican democratic politics (when conducted under appropriate conditions). These direct mechanisms are insufficient for regulating interpersonal interaction in complex decision-making contexts characterized by uncertainty and risk. The central thesis defended here is that the superiority of democratic politics is instrumental: it provides the necessary indirect framework for collectively and rationally pursuing the same ends—truth, moral correctness, and justice—that actors seek through direct discourse.

ideal normative poles: expanding free and informed consensus and deregulation without domination.

In the political process, the risk attitude of those affected is expressed through their decisions, preferences, and support for (or rejection of) initiatives—often irrespective of the explicit public justification offered. We adopt the predominant view in decision theory that an agent's attitude toward risk is an irreducibly subjective phenomenon—one typically associated with, but not determined by, their particular situation. Attitudes toward risk are agent-relative phenomena that resist translation into agent-neutral considerations. Consequently, disputes regarding the definition of a reasonable risk level are ill-suited for a public debate aimed at achieving universally applicable, objective, and impartial justifications. It is possible that the political institutions we defend also possess intrinsic moral value and epistemic instrumental value. However, the purpose of our proposal is to emphasize that, even absent these values, democratic institutions play a key instrumental role by serving rational risk management. This approach provides an instrumental justification of democracy that is conceptually independent of purely epistemic and moral considerations.

4. A Comparison Between Risk-Centered Ordinal Consequentialism and Other Known Consequentialist Conceptions

4.1 Harsanyi's Risk-Centered Cardinal Utilitarianism

By comparing the proposal under discussion to two consequentialist theories with which it shares several points of contact, we can better unpack its normative and conceptual implications. We begin with John Harsanyi's utilitarianism (see Harsanyi 1955, 1975, 1977, 1979, 1982). Like our proposal, Harsanyi's framework maintains that information about the involved parties' typically divergent attitudes toward risk is

crucial for determining which decision maximizes expected utility. For Harsanyi, an agent's attitude toward risk is equivalent to the intensity of their preferences—that is, the subjective importance they attach to promoting or avoiding certain states of affairs resulting from individual or collective decisions. The theory of expected utility developed by von Neumann and Morgenstern allows for cardinal measurement of individual agents' preferences at the *intrapersonal* level. Harsanyi posits (1) that these cardinal comparisons reflect the different intensities of agents' preferences, and (2) that by introducing additional premises, the von Neumann-Morgenstern theory can be extended to enable interpersonal comparisons of well-being. This cardinal assessment of expected utility is critical for utilitarianism. For utilitarians, utility maximization acquires moral status only because it moves beyond the mere satisfaction of the ordinal, informed preferences of the greatest number. If utility maximization were merely ordinal—as it is in the proposal we defend—it would lack this moral status, remaining tied instead to agent-relative and prudential considerations. From a utilitarian perspective, the moral correctness of a decision depends strictly on the net maximization of the good, which may conflict with the informed preferences of the majority of affected individuals.

Suppose ten agents are affected by a decision between two courses of action. Six prefer action A, while the remaining four prefer action B. This information is insufficient to determine the correct decision from a utilitarian perspective. If the intensity of preferences is identical or very similar (assuming preferences accurately reflect the anticipated well-being for each agent), then the superior option in an *ordinal* evaluation will coincide with the superior option in a *cardinal* evaluation. However, if the minority's preference for B is significantly more intense than the majority's preference for A, the latter option—despite maximizing ordinal utility—will not necessarily maximize cardinal utility. Consequently, it would not be the morally correct decision from a utilitarian standpoint. The

cardinal character of actual (or expected) utility maximization is essential to achieving the level of impartiality that defines utilitarianism as a moral theory.²³ It is not simply a question of how many people desire a particular outcome, but of identifying the right outcome from an impartial perspective. The correct decision is the one that produces the greatest net good, regardless of its distribution among individuals. Admittedly, in decisions affecting large populations, we can assume that the will of a large majority (provided they understand what contributes to their well-being) will likely coincide with the right course of action; the greater intensity of a minority's preferences can rarely compensate for the sheer weight of numbers. Thus, utilitarianism appears to offer a limited argument in favor of democracy. Yet, this argument is indeed limited, as utilitarianism also readily supports epistocracy or elitist conceptions. Given that determining the impact of decisions on well-being requires extensive empirical knowledge—and identifying the utility-maximizing choice is computationally difficult—it often seems wiser to entrust this task to experts.²⁴

In Harsanyi's theory, the fact that individual preferences in typical contexts (where the probability and cost of error are significant) are inseparable from the adoption of a risk attitude supports the plausibility of the orthodox utilitarian interpretation of moral correctness. This fact theoretically enables the measurement of the cardinal expected utility maximization on which this interpretation depends. In the proposal we defend, the inseparability of individuals' preferences—conceived ordinally—from the adoption of an attitude toward risk has the opposite effect. We argue that

23. See Bykvist 2014 on the importance of adopting a cardinalist conception in contemporary utilitarianism.

24. For an insightful analysis of the tensions between utilitarianism and democracy, see Woodard (2019, 164–167).

any collective decision that maximizes ordinal expected utility, yet falls short of free and fully informed consensus or deregulation without domination, lacks full moral status. In other words, such a decision remains partial or agent-relative. Consequently, minorities affected by these decisions retain the right to contest them and advocate for reforms that foster outcomes more favorable to their preferences. The legitimacy of this process rests on the principle that it is always desirable—when feasible—to modify adopted policies to broaden consensus or create spaces for deregulation (e.g., through protections for vulnerable minorities or exemptions for those willing to take greater risks). One might argue that if it were possible to identify the decision maximizing cardinal expected utility—and if utilitarianism truly offered the best interpretation of impartiality—then the logic of approximating the poles of consensus and deregulation in our proposal would be rendered superfluous.²⁵

25. Our approach aligns with Jonathan Riley's (1990) argument that ordinal utility maximization—assuming those affected have informed preferences—is equivalent to the outcomes of majority rule and, more broadly, to the functioning of constitutional democratic institutions. Riley views democratic institutions as the application of utilitarian logic within contexts of limited information. In such contexts, applying cardinal utility directly is either impossible or prohibitively costly; thus, relying solely on ordinal preference information becomes the rational alternative. As Riley explains: "[I]f the acquisition and use of any utility information richer than purely ordinal utility information involves more costs than benefits to society in at least some situations, then utilitarianism itself properly recommends reliance on purely ordinalist decision-making procedures rather than a full-blown utilitarian calculus in those situations. Our hypothesis is that these purely ordinalist utilitarian institutions can be seen as democratic institutions. Utilitarianism recommends democracy in this case because there is no ideal agent to perform without cost the many tasks implicit in the notion of the utilitarian calculus. Democratic government is seen as a set of procedures for calculating the general welfare in environments where a full-blown utilitarian calculus is too costly to establish and run. In particular, democratic decision making takes account of the fact that

As in the epistemic model, if we are certain that we have made the right decision, broadening consensus is incidental or, in some cases, dangerous—particularly when it involves modifying the decision to increase the level of support it can obtain. Similarly, deregulation only opens up the possibility of a deterioration in the quality of the outcome. Consider the case of firm example: if investing in *N1* at *T1* is truly the option that maximizes expected cardinal utility, then for a utilitarian, the subjects who reject this decision (because they wish to avoid the associated risks) have no legitimate claim. Investing in *N1* at *T1* would be the morally correct decision. Consequently, such individuals ought to sacrifice their partial preferences to promote the greater expected good, even if it is to their detriment; indeed, they should be willing to jeopardize their own survival as economic agents. The decision-making process in the firm would thus adopt the moral logic of sacrificing individual interests, mirroring the Miners case. There would be no moral reason to attempt to broaden consensus, deregulate, or protect minority interests. On the contrary, sacrificing the interests of minorities becomes a moral obligation.

personal utility information is initially dispersed and costly to retrieve, compare, and aggregate” (342–343). The utilitarian defense of democracy proposed by Riley shares with ours the adoption of a consequentialist and ordinalist perspective. However, the most salient differences between the two proposals are as follows. First, our proposal is prospective; it centers on the notion of expected utility rather than actual utility. Second, it focuses on the disagreement regarding the attitude toward risk that affected parties (as decision-makers) should adopt. Neither of these issues is addressed by Riley. Third, our justification for democratic institutions does not depend on the cost or impossibility of obtaining the information required to identify which decision maximizes cardinal utility. Given reasonable disagreement about the level of risk that collective decisions should involve, the normative goal of our prospective consequentialist approach is not the maximization of cardinal utility (expected or otherwise). Rather, taking majority rule as a baseline, the goal is to approximate the normative poles of free and informed consensus and deregulation without domination.

However, for well-known reasons, Harsanyi's utilitarianism does not pose a significant obstacle to the consequentialist republican conception we defend. First, there is broad agreement that Harsanyi's attempt to solve the problem of cardinal interpersonal utility comparisons is highly problematic (see Sen 1970, 1979; Weymark 1991; Roemer 1996; Narens and Skyrms 2020). Second, the premise that utilitarianism offers the appropriate interpretation of impartiality has been the target of powerful objections. As Amartya Sen (1970, 142–3) observes, even sophisticated versions of utilitarian theory—such as Harsanyi's—could justify a social system in which a minority is subjected to slavery. For instance, if slaves constitute a minority, and their intense suffering produces (or is compatible with) an increase in the overall welfare of society—even if the welfare gain for each member of the majority is modest—there would be no moral basis within the theory to question the institution. Even if such a result remains merely a conceptual possibility—unlikely to materialize in practice, as utilitarians often argue—it constitutes a significant theoretical flaw. Agents subjected to slavery would ostensibly bear a moral obligation to accept their condition, possessing only partial and selfish prudential reasons to oppose it. Consequently, a normative theory leading to such conclusions appears highly implausible.

Nevertheless, it is equally clear that the ordinal maximization of expected utility produced by majority rule—although it reveals the predominant risk attitude within the group—does not preclude outcomes that are intuitively as unacceptable as those resulting from the cardinal utility principle. Fundamentally, in the consequentialist ordinalist proposal we advocate, the moral acceptability of majority rule does not hinge on strict impartiality aimed at satisfying the cardinal aggregation of expected well-being. Instead, its legitimacy rests on continuous, ever-evolving progress toward impartiality. This progress is fostered by majority rule when implemented against a backdrop of practices and institutions that strive

to broaden free and informed consensus while phasing out coercive regulations. This is particularly crucial in scenarios where agents have the capacity to make divergent decisions reflecting the risks they are willing to take, without morally compromising the ability of others to do the same. Almost invariably, when centralized collective decisions must be made via majority rule, politics entails sacrificing the preferences or projects of some individuals for the sake of others. Inevitably, some prevail while others do not, and not everyone achieves their desired outcomes to the same degree.²⁶

26. From the perspective of our proposal, Bernard Williams is correct to state that we should not view political engagement simply as arguing with those who disagree. As he observes: “[W]e should not think that what we have to do is simply to argue with those who disagree: treating them as opponents can, oddly enough, show more respect for them as political actors than treating them simply as arguers—whether as arguers who are simply mistaken, or as fellow seekers after truth. A very important reason for thinking in terms of the political is that a political decision—the conclusion of a political deliberation which brings all sorts of considerations, considerations of principle along with others, to one focus of decision—is that such a decision does not in itself announce that the other party was morally wrong or, indeed, wrong at all. What it immediately announces is that they have lost.” (2005, 13) Politics is not merely applied moral theory or applied science; it possesses its own logic which, as we have seen, is distinct from the logic of searching for impartial moral correctness and objective knowledge. In the political decision-making process, the dominant normative goal is not merely to approximate truth or moral correctness, but to approximate the twin normative poles of broadening free and informed consensus and deregulation without domination. In politics, ensuring the quality of outcomes in contexts of risk requires engaging—simultaneously and sometimes sequentially—in two very different games. This dual practice is equally necessary to pursue that goal, though it can create instability and confusion. First, as proponents of the epistemic model and moral conceptions of politics emphasize, the public justification of any actor’s proposals must seek to provide objective and impartial grounds. This goal is immanent to political practice; in a significant part of the process, we must treat others as “arguers” or “fellow seekers of truth.” This constitutes the first game and is a prerequisite for the system’s instrumental value. Crucially, however, a commitment to objectivity and impartiality is insufficient to guarantee the

From the perspective we defend, however, for this partial outcome to be morally acceptable, it must meet certain conditions. It must occur:

1. In a choice between options that have been competently evaluated from an epistemic point of view, based on a reasonable normative framework; and
2. Against a backdrop of institutions that minimize the magnitude of, this partial outcome (i.e., the fact that some win while others lose) by attempting:
 - a. To broaden the consensus around the options at stake, even at the expense of the strict integrity of the prevailing normative framework; or
 - b. To limit or avoid unnecessary regulation, always allowing for the possibility of re-examining and challenging majority decisions.

In light of this model, it can never be argued that the decisions of the majority are morally superior to those defended by the minority. Both the majority's decisions and the minority's alternatives are imbued with a degree of partiality incompatible with the kind of legitimation that moral or justice-related discourse entails. While such discourse is necessary as a basis for the quality of political decisions—and remains an integral part of the political process—it is always insufficient in

quality of results. The epistemic-moral logic possesses instrumental value only when subordinated to the process of political competition—a process that is aggregative, always open to revision, and irremediably marked by partiality and risk (see Section 6). This is the second game. Ultimately, in the final stage of decision-making, other citizens (or their representatives) are not merely debaters or truth seekers whose proposals are rejected because they are proven wrong, or accepted because they are proven right. Rather, they are opponents who have temporarily lost or won in the competitive process.

contexts of risk to provide an adequate or complete justification for such decisions.

In summary, there are several possible candidates for characterizing the best outcome of a decision process from a consequentialist moral point of view:

1. Actual utility maximization.
2. Expected value maximization.
3. Cardinal expected utility maximization, distinguishing between two scenarios:
 - a. Cases where there is a single reasonable attitude toward risk (usually neutrality, as seen in the Miners case).²⁷
 - b. Cases where different subjects can adopt different and equally reasonable attitudes toward risk (aversion, neutrality, or propensity, as in the firm example).
4. Ordinal expected utility maximization.
5. Ordinal expected utility maximization corrected by institutions and practices of contestation and deliberation. These practices aim to expand free and informed consensus and create instances of non-regulation free from domination. (This option is equivalent to implementing institutional arrangements such as contestatory democracy or other republican or liberal models).

27. Although it would require more detailed analysis, it seems plausible to assume that in the Miners case, a cardinal and an ordinal assessment of the preferences of those affected should yield the same result. One might assume that some miners have a stronger preference for survival than others, and this may indeed be true. However, it seems appropriate to act as if everyone assigns the highest possible intensity to preserving life. This is because, first, accessing such information in this context is impossible. Second, accounting for differences in intensity seems morally controversial when the very existence of the agents is at stake. In this scenario, the famous dictum that John Stuart Mill attributes to Jeremy Bentham—“everybody to count for one, nobody for more than one”—seems to render differences in the intensity of preferences irrelevant.

As demonstrated by the preceding considerations, the best alternative available in most political contexts is Option 5. This option is equivalent to implementing the set of institutions that characterize contestatory democracy. Normally, when political decisions must be made, we cannot know which option will effectively maximize actual utility (1). Furthermore, the idea of maximizing expected value (2) lacks practical relevance and is conceptually problematic. We cannot make a rationally grounded prediction about the impact on the welfare of those affected simply by multiplying the value of possible outcomes by their probability of occurrence. This approach is valid only if a single rational attitude toward risk exists—specifically, neutrality (which corresponds to a linear utility function)—as seen in the Miners case.²⁸

28. When analyzing examples like the Miners case, philosophers tend to present them as an application of von Neumann and Morgenstern's theory to the field of normative ethics. Horacio Spector (2016) rightly points out that this is a mistake. The moral plausibility of these cases depends on the assumption that only a linear utility function is acceptable—implying that there is only one correct attitude towards risk: neutrality. On the contrary, one of the defining features of expected utility theory is that the utility functions of different subjects can vary significantly according to their attitudes toward risk. Agents can be assigned concave, convex, and linear utility functions, corresponding to risk aversion, risk-seeking behavior, and risk neutrality, respectively. These divergent functions can be equally rational. As we have seen, in Bernoulli's theory, all subjects' utility functions are concave—taking the form of a logarithmic function. For Bernoulli, we are all risk-averse to varying degrees. Consequently, the conception of expected utility often found in philosophical literature implies a return to the pre-Bernoullian notion of expected value calculus originally formulated by Fermat and Pascal. In reality, due to a content-contingent feature of the cases invariably invoked in this literature, any deviation from risk neutrality is deemed intuitively unacceptable. For instance, a risk-seeking agent might argue that the expected utility of saving 100 miners with a 0.5 probability of success is greater than the expected utility of saving 90 with certainty. This is analogous to claiming that the expected utility of earning \$100 million with a 0.5 probability of success exceeds that of earning \$90 million with certainty. The expected *value* of the first option is significantly

In the context of the Miners case, we are justified in assuming that the decision maximizing expected value is the optimal decision. If the decision could be repeated indefinitely,

lower than that of the second (50 million versus 90 million). However, let us return to the case of the firm. If the survival of the company—or other objective or subjective factors—comes into play, the preference ordering may be reversed regarding the expected utility that decisions provide to the partners. If earning 90 million is insufficient to avoid bankruptcy, the rational course of action may indeed be to risk everything to earn 100 million. The problem with this logic is that it is completely unacceptable when applied to human lives. Suppose we are willing to risk the lives of 90 workers who could certainly be saved by closing Gate 3, in order to try to save 100 workers by closing gates 1 or 2—even though there is a 0.5 chance that all 100 workers will die as a result of this gamble. This implies that some individuals' lives hold more value than others'. Specifically, it suggests that the potential survival of the 10 additional workers (at a 50% probability) is weighted more heavily than the certain survival of the 90. This is, of course, morally unacceptable. While each additional million dollars gained by taking greater risks may be worth more than the previous million (as in the firm example), such marginal valuation cannot be applied to human existence. The plausibility of examples like the Miners case depends on these kinds of substantive moral considerations. Governments must sometimes make decisions such as those illustrated by the Miners case—decisions that should not be subject to political logic, but rather to a combination of technical, scientific, and moral considerations. The problem, however, is that these cases represent only a small subset of the decisions required in both the private sphere (individual and associative) and the political sphere. When the issue is not life itself but the *quality* of life—or other goods distributed by politics such as opportunities, income, wealth, rights, and freedoms—the divergences regarding the most appropriate attitude toward risk can vary drastically across subjects. Crucially, these divergences can be perfectly reasonable. Harsanyi's utilitarian theory represents the most consistent and comprehensive application of expected utility theory in normative ethics. This theory is potentially applicable to all types of decisions, not merely those where risk neutrality is the only rational attitude. The challenge, however, is that it relies on a cardinal interpersonal evaluation of utility. Despite the efforts of Harsanyi and other scholars, this requirement appears infeasible and susceptible to the same conceptual or normative difficulties plaguing other utilitarian theories. The normative conception we defend, risk-centered ordinal consequentialism, aims to avoid both difficulties.

the overall outcome would be clearly superior to that resulting from any other strategy. However, this possibility is limited to a relatively small and perhaps exceptional subset of cases. The case of maximizing cardinal expected utility, exemplified by Harsanyi's utilitarian theory, avoids the aforementioned scope restriction but suffers from the normative and feasibility challenges previously discussed. Philosophers and economists have not been able to offer a convincing solution to the problem of interpersonal comparisons of well-being. Moreover, the maximization of cardinal expected utility appears to justify the violation of basic minority rights in certain scenarios—even if these are unlikely in practice. Our proposal demonstrates that contestatory democracy avoids all these difficulties. The mere ordinal maximization of expected utility—majority rule applied to options competently evaluated from factual and normative standpoints—does not, by definition, require interpersonal comparisons of well-being. Consequently, there are no significant epistemic or practical obstacles to determining which option possesses majority support. Crucially, however, the critical issue remains: the rule of the majority in itself lacks any independent justificatory moral force.

We have argued that compelling reasons exist to consider decisions made by a majority superior to those made by a minority, primarily because majority decisions reflect the group's overall attitude toward risk. However, majority rule can impose unacceptable costs on minorities—a consequence as objectionable as those arising from the unrestricted utility principle. Therefore, the moral acceptability of majority rule depends on its embedding within a framework of institutions. These institutions must aim to approximate the twin normative ideal poles of free and informed consensus and non-regulation without domination—scenarios in which the maximization of ordinal and cardinal expected utility coincide. At both ideal poles, the informed preferences of all affected parties are satisfied, ensuring there is no possibility of divergence between a cardinal and an ordinal assessment.

It is important to note that ordinal expected utility maximization is not necessarily equivalent to maximizing actual utility in any particular instance. Furthermore, it does not necessarily coincide with maximizing expected value or with the outcome that maximizes cardinal expected utility. When applying majority rule over time, the partners in the firm might choose to invest in $N2$ at $T1$, $N1$ at $T2$, $N2$ again at $T3$, $N3$ at $T4$, or any other combination. Similarly, these decisions could be complemented at different times with exemptions that allow some agents to assume more risk while providing insurance for the most vulnerable. For instance, if the majority of partners chooses to invest in $N1$ at $T2$, it is entirely possible that the option maximizing actual utility at that moment is one of the alternatives. This discrepancy can occur with any majority decision. Likewise, even if an interpersonal cardinal evaluation of well-being were possible, the option maximizing cardinal expected utility might differ from the one maximizing ordinal expected utility. In the time horizon relevant to the well-being of those affected, the only viable strategy for achieving the best overall outcome is a flexible form of ordinal maximization. This approach requires constant revision and correction, guided by the twin ideal poles of reaching a free and informed consensus and promoting deregulation without domination. It is impossible to ascertain whether we are maximizing actual utility in every decision, just as it is impossible to determine cardinal expected utility with certainty. We also lack a definitive understanding of whether actual or expected utility is the correct standard of moral validity. Nevertheless, the republican democratic political strategy maximizes the probability of individual and collective survival and flourishing in contexts of risk. This strategy involves approaching an equivalence between maximizing the expected good (however defined) and securing individuals' freedom to choose their own risks without interference or domination.

If a democratic republican system functions adequately for a considerable period, we can hypothesize that it will

tend to produce results approximating those derived from the principle of cardinal expected utility. While the application of majority rule ensures only ordinal utility maximization and fails to reflect the intensity of individual preferences, we can assume that the institutions of a stable, functioning contestatory democracy will tend to progressively reflect this dimension. Let us accept Harsanyi's equivalence between preference intensity and agents' attitude toward risk: the degree to which a subject is willing to risk losing a good is directly proportional to the intensity of their preference for obtaining it. Conversely, risk aversion captures the intensity with which an agent seeks to avoid potential losses of assets in exchange for future acquisitions. The ideal poles that justify the value of republican institutions represent scenarios where all subjects see their preferences satisfied. These scenarios are characterized by uncoerced, informed consensus and deregulation devoid of interpersonal domination. In such contexts, agents decide freely making choices guided by their respective risk attitudes. At these poles of full consensus and deregulation, ordinal and cardinal utility maximization are equivalent. Consequently, the approximation to these ideal poles constitutes an approximation to the cardinal maximization of expected utility. Our hypothesis is that the republican system facilitates a progressive approximation to these ideal poles, even if they are never fully realized. Majority rule, which is already a closer approximation to the ideal poles than a decision supported only by a minority, is merely a first step. The republican model always permits the review and correction of collective decisions through institutions and practices that seek to broaden consensus beyond a simple majority. Citizens have the right to question and challenge initial decisions through various channels of contestation. These channels can lead to the establishment of special insurance systems, protections, and partial exemptions. If the republican system is capable of producing this approximation to the ideal poles, it serves as a path to the cardinal maximization of expected

utility, incorporating the preference intensity of those affected throughout the political process.

It is not possible to directly calculate the decisions that maximize cardinal expected utility, due to the problem of interpersonal comparisons of well-being and the high epistemic complexity of the issue. However, it is possible to predict that, in the medium- and long-term, institutions that foster approximation toward the poles of full consensus and deregulation will also bring us closer to a scenario of maximizing cardinal expected utility. The republican model, and its underlying justification—risk-centered ordinal consequentialism—pursues a normative goal compatible with Harsanyi's prospective utilitarianism. As with Harsanyi's position, it aims to be applicable in the paradigmatic political context: scenarios characterized by reasonable disagreement regarding the correct or most appropriate attitude towards risk. However, our proposal has two significant advantages over Harsanyi's theory and other variants of the utilitarian position. The first advantage concerns the tractability of its application in practice. As previously noted, directly calculating the decision that maximizes expected utility does not seem feasible. Despite its undeniable complexity, it is much simpler and more feasible to identify which decisions are supported by a majority, and to determine whether that initial consensus can be broadened through political competition, the reopening of debates, and ongoing instances of individual and collective contestation. In other words, it is possible to assess whether there has been a progressive approximation to the poles of free and informed consensus and deregulation without domination.

The second advantage concerns the affinity of our proposal with moral convictions regarding the inviolability of individuals' fundamental rights—convictions embodied in both moral intuition and normative deontological theories. Our proposal does not make the protection of individual fundamental rights contingent upon a controversial moral principle. Instead, it is grounded in the institutions of constitutional democracies.

These liberal and republican institutions have historically proven to be the most effective means of guaranteeing such rights, regardless of the moral disagreements inherent in the public culture of these societies. In the approach we defend, the institutions of a contestatory democracy take the place of the principle of utility (or any other attempt at social coordination based exclusively on moral considerations). While utilitarianism remains the prevailing consequentialist theory in academic philosophy, democratic institutions represent a consequentialist model forged by the most successful social practices. Consistent with our framework, there is no conceptual space to justify or legitimize the violation of individual rights or the sacrifice of one group's interests for those of others. Should such violations occur, they must be recognized as a failure of the system rather than part of its proper functioning. This is because the best expected outcome and the maximization of individual freedom—achieved at the poles of fully free consensus and total deregulation—can never differ. The replacement of the principle of utility with majority rule, combined with deliberative practices and institutions for reviewing and contesting decisions, aims to broaden free and informed consensus beyond the initial threshold. This strategy minimizes the discrepancy between actual procedural results and the ideal poles, where all differences between maximizing expected utility and maximizing individual freedom are dissolved.

Utilitarianism can be regarded, in some respects, as an attempt to assign meaning and value to the pain and suffering of individuals in the name of a presumed collective benefit. It seeks to fictionally transform a tragedy into a moral triumph. In contrast, the political consequentialism we defend seeks to create institutions that avoid the sacrifice of individuals. This approach rests on the premise that sacrifice is never rational or intrinsically valuable, even if it is inevitable. The core argument posits that in contexts of risk, actual agents take precedence over both empirical facts and normative frameworks, as neither is

sufficient to yield universally justified decisions. Consequently, the partial will of those who impose their preferences becomes decisive. Thus, the sacrifice of an individual's projects or well-being is a regrettable outcome for which specific agents bear responsibility. Such an outcome lacks complete justification based on reasons acceptable to all and can never claim legitimacy derived from a purported higher authority. Given the inherent partiality of any decision-making process involving risk, the decisions of authorities must be subject to continuous review and contestation. Such determinations must be guided by a political logic that transcends the mere pursuit of factual truth and moral correctness. When faced with risk, the supposed value of sacrificing individual interests is often a product of ideology or the imposition of one group's power over another. Without democratic politics, consequentialism is ideology.

4.2 The Consequentialist Defense of Republicanism by Philip Pettit: Freedom as Non-Domination and Contestatory Democracy

Just as Harsanyi represents the closest theoretical analogue to our proposal within utilitarianism, the republican political conception developed by Philip Pettit offers the strongest link to our framework. The primary connection between our proposal and Pettit's has already been explicitly noted at the outset. From a normative standpoint, the aim of this paper is to defend the specific interpretation of contestatory democracy (cf. Pettit 1997, 1999, 2001; Maynor 2003) outlined by Pettit, distinguishing it from both majoritarian (or populist) proposals and epistocratic (or elitist) conceptions. Pettit's proposal centers on the ideal of freedom as non-domination and remains consistently consequentialist. He argues that freedom consists neither in self-mastery nor merely in the absence of interference by others, but in "the absence [...] of mastery by others"—that is, the absence of domination. According to Pettit, an agent or group is free to the extent that no other person or

group possesses “the capacity to interfere in their affairs on an arbitrary basis” (1999, 165; cf. Pettit 1997, 2001, 2012, 2014). While freedom as non-domination is identified as the supreme political value, Pettit insists that it should not be viewed as a side-constraint that limits the maximization of other goods (as in a deontological framework), but rather as a goal to be promoted or maximized.²⁹ To quote Pettit

I maintain that non-domination is a goal which such institutions should seek to promote, not a constraint that they have to honour in the pursuit of other goals; I defend a consequentialist version of republicanism. This republican doctrine, as we shall see, is a consequentialism with a difference: it allows us to say that the institutions which promote people’s freedom as non-domination go to constitute that freedom, not to cause it; the doctrine does not countenance any temporal or causal gulf between civic institutions and the freedom of citizens. (1997, 81)

Like our proposal, Pettit’s theory aims to neutralize the dangers to individuals’ basic rights and freedoms associated with other consequentialist positions, such as utilitarianism. Similarly, it is a prospective theory in which the promoted goal is expected value, not actual value. Rational risk management is a relevant element of this conception. In its early formulations, the need to reduce uncertainty is a prominent feature in explaining the value of freedom as non-domination and its superiority over the mere interpretation of freedom as non-interference (see Pettit 1997, 83–90). Our proposal aligns with Pettit and the wider republican tradition by assigning a central role to freedom as non-domination. This framework is particularly illuminating when applied to the case of the firm. Rational collective decision-making requires the accurate aggregation of stakeholders’ risk attitudes. However, such information is

29. For an example of a deontological conception of the ideal of non-domination, see Forst 2013.

inaccessible in the presence of arbitrary domination. If certain individuals possess the power to threaten or sanction others, the victims are likely to endorse options incompatible with their actual risk profiles—accepting volatility that disrupts their life plans. This dynamic undermines the quality of decision outcomes. For instance, if a majority of partners prefers investment $N2$ but votes for $N1$ due to intimidation by $N1$'s proponents, the firm fails to capture true preferences. Over time, a sequence of decisions distorted by such domination can lead to the company's decline, insolvency, or bankruptcy. Notwithstanding these theoretical convergences, we must address the crucial conceptual differences between Pettit's republicanism and our own framework. The fundamental divergence lies in Pettit's reliance on orthodox consequentialist moral philosophy. In works subsequent to *Republicanism* (1997), Pettit maintains that consequentialism is an inescapable normative perspective, particularly in the political realm. His definition of this perspective clarifies the distinction we aim to highlight. Defending an agent-neutral conception of moral rightness, Pettit posits that "the right alternative in any choice is a promotional function of the agent-neutral good" (2012b, 42). As he elaborates:

The neutral good is the sort of good that can be identified in common terms from any perspective. It might also be described as the non-indexical good, since the main mark is that its expression does not require the use of an indexical such as 'I' or 'mine', 'we' or 'ours', 'here' or 'now'; it contrasts with agent-relative goods such as the success of my projects, the welfare of my children, the prosperity of my country (2012b, 42).

When Pettit argues that the primary goal of a political system is to promote freedom as non-domination, he defines it as an agent-neutral good. Freedom as non-domination must be promoted because it is the highest agent-neutral political value, taking precedence over alternative agent-neutral goods whose maximization might conflict with it. Our proposal

aims to defend a contestatory conception of democracy via a consequentialist strategy focused on maximizing the expected good. We define this good as the maximization of the freedom of those affected—interpreted not as mere non-interference, but as the institutional protection against arbitrary domination. However, the key idea of our proposal (and the point of divergence from Pettit’s republicanism) is the recognition that, in contexts of risk, there is no direct way to justify collective decisions based exclusively on agent-neutral goods—whether freedom as non-domination, global welfare (as in utilitarianism), or liberal conceptions of equality.

As we have seen, in contexts where the probability of error is significant and the costs substantial, the adoption of an irreducibly subjective and partial attitude toward risk is inescapable—and indeed desirable within the context of adequate institutions. Consequently, individual freedom cannot be restricted by appealing directly to welfare maximization or equality—that is, to agent-neutral values. Nor can such restrictions be justified by appealing merely to some supposed moral or factual truth, even if we accept that such truth is available. Furthermore, any progress toward the best expected aggregate outcome in the medium- or long-term requires a political decision-making system that combines majority rule with a constant effort to approach the normative poles of uncoerced and informed consensus and deregulation free from domination. Even if one rejects the specific republican thesis prioritizing freedom as non-domination over other agent-neutral goods, one must recognize that the best expected outcome cannot be pursued without a republican political system. In such a system, the maximization of the expected good is extensionally equivalent to the maximization of freedom as non-domination.

To return to our earlier example, consider the choice between economic policies derived from divergent normative frameworks—such as utilitarianism, Rawlsian or prioritarian conceptions, or a republican conception centered on non-

domination—and varying risk profiles. It is necessary to subject these competing proposals to a decision-making system designed to maximize the political agency of those concerned. This system must integrate majority rule with consensus-seeking practices, opportunities for deregulation, and mechanisms for individual or collective contestation of the majority will. When applied in contexts where the probability of error is significant and the cost to affected individuals' well-being is high, all decision-making procedures are insurmountably agent-relative. This result holds true even if the options at stake emerge from an epistemically competent evaluation of facts and an adequate normative framework. This is because no decision can be made without adopting an attitude toward risk; that is, a purely subjective factor always plays a decisive role in the choice. Consequently, the following procedures all produce different levels of maximization of the agent-relative expected good:

1. Minority rule;
2. Majority rule;
3. Majority rule applied against the backdrop of practices or institutions seeking free and informed consensus or deregulation without domination;
4. Ideal scenarios, such as:
 - 4.1 Unanimous consensus on a single decision, derived from identical risk attitudes subscribed to by all affected; or
 - 4.2 A situation where agents choose different options appropriate to their divergent risk profiles.

In all these cases, the result remains the same: the enforcement of the preferences of a specific group of individuals—described indexically—even if the decisions contingently express the will of all (as in cases 4.1 and 4.2). In the presence of risk, the optimal achievable outcome is a set of decisions aligned with the will of each affected agent,

rather than with an objective “common good.” The outcome is never purely objective or impartial. As we have established, the maximum possible expected utility in contexts of risk is achieved equally when all affected parties provide free and informed support for the same decision, or when they can freely and deliberately make different decisions based on their divergent but reasonable attitudes toward risk.

If the option that received unanimous support in the first scenario truly represented the common good, it would be inconsistent to assert that the second scenario (if feasible) is as desirable as the first. If we knew with certainty which decision was the correct one, from a consequentialist perspective it would be optimal to impose it on everyone, rather than allowing free choice. The point is that, as we have seen, in contexts of risk it is entirely rational to assign equal value to the twin poles of consensus and deregulation. This is because, in contexts of risk, the maximum that can be achieved is a sub-moral status. We are not in the realm of morality, but rather, the realm of politics.

To invoke Rousseau’s well-known distinction, the will of all is not identical to the general will. In contexts of risk, neither a general will nor a determinate common good exists. All we can attempt is the search for an adequate normative framework and correct factual judgments relevant to the decision. This process occurs before and after decisions made by majority rule—which must always remain open to contestation—and relies on institutions aimed at broadening free consensus and avoiding unnecessary regulation. This represents the best way to rationally approach the true common good, even if that goal can never be fully or directly achieved. Each step in this iterative approximation to the ideal poles is inseparable from the maximization of individual freedom. No given decision possesses full moral status. Consequently, the purported moral correctness of a particular decision does not justify the violation of political rules or the sacrifice of individual interests. This stance does not derive from a fetishistic commitment to freedom or an affirmation of its intrinsic superiority; rather, it stems from

the fact that the only rational guarantee of achieving the best expected outcome in the medium- and long-term is the stability and proper functioning of a system constituted by republican democratic institutions. The individual is not protected from arbitrary power merely by identifying the supreme political value or the correct interpretation of the common good. In contexts of risk, it is necessary to acknowledge that even a correct understanding of the common good cannot justify individual sacrifice. Consequently, the best outcome is always inextricably linked to the maximization of individual freedom. Falling short of this goal—a frequent occurrence—holds no redeeming moral value; it is always, invariably, a regrettable failure.³⁰

5. Contestatory Democracy vs. Two Variants of Risk-Centered Epistocracy

5.1 Risk-Centered Epistocracy

Addressing a possible objection to our proposal serves to clarify its implications. It could be argued that the framework we defend does not preclude the conclusion that, in practice, a government of experts (or epistocracy) might be preferable to both majority rule and contestatory democracy on instrumental grounds. Beyond possessing accurate factual information

30. A further distinction between our proposal and Pettit's is that, despite adopting a consequentialist position, Pettit does not take an instrumental approach to the value of the political system. He does not commit to the hypothesis that there are standards of political correctness independent of democratic procedures (see Pettit 2015, 82). Our proposal accepts this hypothesis from the outset. We seek to demonstrate that, in contexts of risk, the rational promotion of such independent standards—however defined—is inseparable from a republican political process such as the one we have described. As we have seen, at the ideal poles of this process, the maximization of the expected good is extensionally equivalent to the maximization of individual freedom.

and selecting an appropriate normative framework, experts might conceivably discern the risk attitudes of those affected. Furthermore, they might predict the outcome of a political process in which decisions are based on:

1. well-justified factual judgments;
2. a reasonable conception of justice;
3. prevailing attitudes toward risk; and
4. an attempt to broaden consensus around binding decisions, while eliminating unnecessary regulation of individual and collective action.

This form of elitist governance would be undesirable if we believe the political process is functioning adequately, or if we expect it to do so. However, the proposal gains relevance if traditional epistocrats are correct that the political ignorance of the majority is an established and insurmountable fact. Yet, experts cannot make optimal decisions solely on the basis of well-founded factual judgments and universally valid normative principles. As we have seen, in addition to making an adequate factual and normative assessment, they must ascertain how much *each* subject is willing to risk—specifically, the reasonable risk attitude each agent will adopt in a given context. This is because a general or global risk attitude cannot simply be ascribed to subjects to predict their behavior. It is well established that a single person can exhibit profoundly different risk attitudes across different decision-making contexts. This phenomenon underlies the insurance-gambling paradox: the same agent may be highly risk-averse in one context (e.g., purchasing expensive fire insurance for their home) while concurrently choosing extremely risky options in the stock market.³¹ Furthermore, temporal variations in an

31. A classic approach to this phenomenon was developed in a seminal paper by Milton Friedman and Leonard Savage (1948).

individual's attitudes are often significant. These variations appear to stem from a complex set of factors: psychological dispositions, prior wealth, available resources and capacities, life circumstances, religious or moral beliefs, and broader conceptions of the good.

If policy experts knew the risk attitudes that affected parties would adopt—assuming those parties based their decisions on well-founded empirical information and a reasonable normative framework—they could make rational decisions without relying on a political system to reveal these preferences. This would enable a novel form of elitist government, which we might term “risk-centered epistocracy.” The fundamental difficulty with this position, however, is that there are compelling grounds to doubt that experts can possess this kind of knowledge. Consider a scenario requiring the management of assets for multiple agents who currently lack decision-making capacity. We are tasked with allocating their capital in financial markets. In the case of agent *S1*, we possess a record of their previous investments. Analyzing this record enables us to deduce *S1*'s prior risk tolerance and infer the strategy they would endorse were they not incapacitated. While hypothetical, this inference may sufficiently justify decisions made by their representatives. In contrast, consider agent *S2*. Despite possessing substantial capital, *S2* has either been incapacitated since birth or has a history of economically irrational behavior stemming from factual errors or incoherent value judgments. If *S2* lacks a history of competent agency, how might we determine what their “reasonable” risk attitude would be in the present? Furthermore, what decisions would *S2* have accepted had they operated as a rational and informed agent?

If the political competence of the citizenry aligns with the diagnosis offered by epistocrats, these citizens are functionally equivalent to agent *S2*, rendering the authentic representation of their interests unfeasible. Consequently, a risk-centered epistocracy appears untenable. Broadly, this mode of governance faces objections analogous to those raised

in the socialist calculation debate. Hayek famously argued that no central planning body can access the dispersed, often subjective information³² required to make efficient economic decisions. Without market mechanisms aggregating the choices individuals make based on their localized beliefs and projects, central experts operate blindly. The challenge is not merely gathering scattered propositional knowledge; it is the impossibility of centralizing tacit knowledge—the “know-how” guiding agent behavior that cannot be fully articulated in conceptual form (Lavoie, 1986).

Entrepreneurs and investors possess deep, context-specific knowledge regarding their projects and past experiences. This knowledge is critical for evaluating risks and projecting success, yet agents are often unable to articulate it propositionally. Because this “know-how” is tacit, it cannot be transmitted as accessible data to an external expert observer. Instead, it is the market price of the resulting goods and assets that translates this tacit understanding into public information. Absent the free interaction of the market, price signals do not emerge, rendering centralized economic calculation impossible. As

32. In this context, it is worth invoking Hayek’s well-known observation regarding the progress of economic theory: “...it is probably no exaggeration to say that every important advance in economic theory during the last hundred years has been a further step in the consistent application of subjectivism” (1955, 31). Hayek was referring primarily to the impact of the marginalist revolution and the replacement of objectivist theories of value—such as Adam Smith’s cost-of-production theory or Marx’s labor theory of value—by subjectivist theories. Our proposal attempts to take a parallel step in democratic theory: reclaiming the role of subjectivity and partiality by recognizing the third dimension of political disagreement—specifically, the structural function played by the risk attitudes of those affected by the political process. We advocate for a subjectivist theory of political value. This does not imply skepticism regarding the impartiality or correctness of moral judgments in collective decision-making, nor does it dismiss the objectivity of factual judgments. Instead, it acknowledges that the need to adopt an attitude toward risk in most political decisions introduces an irreducible subjectivity and partiality to the entire process.

Elizabeth Anderson notes, “[s]ocially dispersed information can be transmitted in three forms: talk, votes, and market prices. Markets respond primarily to price information; democratic states primarily to talk and votes” (2006, 8–9).

It seems reasonable to posit that the limitations of central economic planning apply equally to political governance. Just as a committee of experts cannot efficiently plan an economy due to dispersed propositional and tacit knowledge, a political body operating as an external observer cannot access the localized information that drives the actual political process. Consequently, it lacks the data essential for making instrumentally rational decisions. The argument for free markets and the democratic political process is strengthened by analyzing the central role of risk in both contexts. In neither the economic nor the political realm is it feasible to attain optimal outcomes solely through objective assessments of facts, values, and norms. Rather, when the probability of error is significant and the costs high, rational decision-making requires adopting—both individually and collectively—an appropriate risk management strategy. Beyond the dispersed, articulable, or tacit factual knowledge central to Hayek’s argument, there exists another source of dispersed knowledge: the assessment of risk. This form of knowledge is often tacit, inarticulate, and irreducibly subjective. This reality applies equally to choosing an investment and implementing public policy. Consequently, it is essential to understand the risk attitude each affected party adopts in light of the specific facts and values guiding the process.

How could experts determine *ex ante* the level of risk an agent is willing to incur to secure a future benefit from a specific policy? The answer is that such knowledge is inaccessible to a central planner. Risk attitudes and effective preferences can only be revealed through the operation of the political process itself. Mechanisms such as deliberation, electoral competition, and institutional contestation function as a discovery procedure, much like the price system in a free market reveals willingness

to invest and relative value. Thus, risk-centered epistocracy suffers from the same epistemic deficits as a centrally planned economy.

5.2 Epistocracy of Maximum Reasonable Precaution

It might also be objected that we overlook a possible variant of epistocracy that is more compatible with the logic of our proposal. Suppose it proves impossible to ensure that the majority of citizens in a political community attains an adequate minimum threshold of epistemic competence and moral reasonableness. Our defense of contestatory democracy depends on the premise that majorities can reach this threshold at a reasonable cost. If this is not feasible in a given community, what recourse remains? Suppose, as we have argued, that the expert elite cannot predict the risk attitude citizens would have adopted had they been properly informed. Nevertheless, it could be argued that this elite might still be capable of making instrumentally sound decisions by adopting an *appropriate* attitude toward risk on their behalf.

Buchak argues that when we are compelled to make a decision for another agent and lack knowledge of the risk attitude they would adopt, we must choose “under the assumption that he has the most risk-avoidant attitude within reason” (2017, 632). Buchak applies this principle to a different context; her aim is to defend a conception of justice she terms “relative prioritarianism” as an alternative to both utilitarianism and Rawlsian theory.³³ However, the criterion is highly relevant to our current context. It serves as the foundation for a novel

33. According to Buchak, the principle behind her proposal—the “Risk Principle”—can be formulated as follows: “When making a decision for an individual, choose under the assumption that he has the most risk-avoidant attitude within reason unless we know that he has a different risk attitude, in which case, choose using his risk attitude” (2017, 632).

form of epistocratic governance: the “epistocracy of maximum reasonable precaution.” This conception does not aim to ascertain what the majority *would* decide if it could make informed decisions, nor does it attempt to deduce the specific risk attitude agents would adopt based on an epistemically and normatively competent assessment.

Instead, experts would identify options that are epistemically and normatively well justified, attributing maximum reasonable risk aversion to the citizenry. This strategy minimizes the likelihood that the majority is exposed to unmanageable losses or serious harm. We can expect this approach to yield superior social outcomes compared to classical epistocracy, which neglects the role of risk and consequently leaves society vulnerable to intolerable consequences. Unlike “risk-centered epistocracy,” the “epistocracy of maximum reasonable precaution” enables decision-making without the need to predict citizens’ behavior under idealized conditions.

The problem, however, is that the appeal of the “epistocracy of maximum precaution” seems very limited. First, it is difficult to predict the medium- to long-term outcomes of such a strategy.³⁴ A plausible assumption is that it would be far

34. If an analogy is drawn between the success of a political community and the evolutionary success of species, an epistocracy based on maximum reasonable precaution might appear promising. Scenarios chosen by risk-averse agents likely possess a significant evolutionary advantage over those favored by risk-neutral or risk-seeking agents. As Okasha argues: “[I]n any realistic biological scenarios... [n]atural selection will favor risk-averse organisms over their risk-seeking counterparts... For the former have lower variance in reproductive output, which is often evolutionarily advantageous” (2007, 229; see also Okasha 2011). Okasha hypothesizes that this evolutionary phenomenon explains the prevalence of risk aversion among agents—a subject widely studied empirically in economics. However, it is reasonable to assume that this analogy, while defensible, remains significantly limited. Notably, these considerations highlight, from a novel perspective, the inadequacy and danger of classical epistocracy. In the present context, this position implies either neutrality regarding risk or the imposition of the preferences of minorities—who, as we have seen, are

less effective in promoting the welfare of those affected than a strategy sensitive to the diverse risk attitudes of informed agents. Admittedly, this model seems more promising than simply imposing the preferences of a small minority of experts—regardless of whether those experts are guided by a search for truth or are biased toward their own risk attitudes, cloaking their decisions in the rhetoric of objectivity and impartiality. A further, crucial problem is that this form of epistocracy is likely unrealistic and unsustainable in practice. No political subject possesses a compelling reason to support its implementation.

A “maximum precaution epistocracy” would prioritize the interests of both the most vulnerable and the majority—groups that frequently coincide. Yet, since expert elites often belong to or are beholden to dominant social sectors, they lack the compelling motivation to enact such a strategy. Classical epistocracy, by contrast, aligns with the interests of powerful sectors. It presents policies that are structurally biased in favor of privileged minorities as impartial and objective decisions—even when those policies are otherwise justified by reasonable factual and normative considerations. Furthermore, if we posit that the majority lacks epistemic competence and holds unreasonable moral views, we can infer that they would find no reason to support the decisions of a maximum precaution epistocracy. Instead, they would predictably prefer projects that are significantly more speculative, yet ultimately counterproductive to their own interests. It must be emphasized that failing to achieve a robust competitive democracy—characterized by a developed political culture, an active public sphere, and established institutions of electoral and non-electoral competition (such as the separation of powers, the rule of law, a bill of rights, and checks and balances)—invariably results in poor prospects for

likely to be less risk-averse than majorities. Viewed through the lens of the evolutionary analogy, both options appear to be self-defeating strategies.

the majority. Without a well-ordered democratic republican system, the majority is condemned to endure suboptimal long-term political outcomes.

We can now establish an order of preference among the various decision-making systems considered. In contexts of risk, two options—if feasible—should produce identical results and rank highest in the order of preference: contestatory democracy and risk-centered epistocracy. In both cases, the best expected outcome, or the maximization of expected ordinal utility, is achieved. Risk-centered epistocracy aims to replicate the results of a well-functioning contestatory democracy, but it attempts to do so by bypassing the actual political process. However, the primary limitation of this second option is that, as we have argued, it does not appear feasible. The same epistemic arguments against the possibility of an efficient planned economy, compounded by the role of risk, seem to apply equally to its practical implementation. The second option in the order of preference is majoritarian democracy. Provided the majority of citizens reaches an adequate minimum threshold of epistemic and moral competence, majority rule guarantees that collective decisions are drawn from options capable of effectively promoting the well-being of the agents involved. Furthermore, it ensures that the gains and losses associated with those decisions remain within a range tolerable to the largest group. This makes majority rule superior to classical epistocracy, as noted in Section 2. Crucially, if these levels of competence are achieved, there is no compelling reason to stop at the majoritarian model rather than advancing to the republican contestatory model. Admittedly, such a reason exists within the logic—albeit misguided—of the epistemic model. For instance, if majority rule is viewed strictly as an application of Condorcet’s Jury Theorem, where decisions are recognized as objectively “correct,” then the normative core of the republican model—the expansion of free and informed consensus and deregulation without domination—would lose

its *raison d'être* or potentially be viewed as counterproductive³⁵ Given the shortcomings of the epistemic model considered, we must conclude that the preference for majority rule over the republican model lacks plausibility in contexts of risk.

However, if classical epistocrats are right about the fundamental incapacity of majorities to make rational decisions, then the next option in the hierarchy is not classical epistocracy, but an epistocracy of maximum reasonable precaution. The limitations of this option have already been noted: the quality

35. Anderson offers a similar critique—albeit from a purely epistemic perspective—regarding attempts to ground the value of majority rule in Condorcet's theorem. As she argues: "Democratic decision-making needs to recognize its own fallibility, and hence needs to institute feedback mechanisms by which it can learn how to devise better solutions and correct its course in the light of new information about the consequences of policies. Periodic elections are one critical feedback mechanism of this sort. The Condorcet jury theorem does not represent the necessity of such mechanisms. Since it suggests that majorities are nearly infallible from the start, why would they ever need to correct their initial decision?" (2006, 12). Political institutions require feedback and revision mechanisms to correct mistakes and respond to changing circumstances. The epistemic value of such mechanisms is undeniable given the prospective nature of political decisions—even if Condorcet's Theorem implies they are superfluous. Consequently, for epistemic reasons, majority rule alone is insufficient for an instrumentally successful political system. The quasi-infallibility promised by the theorem would eliminate both the need for error-correction mechanisms and the instrumental value of the poles of free consensus-building and deregulation. However, as we have seen, these are distinct requirements—both of which are desirable in the proposal we are defending. John Dewey's pragmatist epistemic conception of democracy—the alternative favored by Anderson—centers on error-correction mechanisms. Yet, it offers no explanation for the instrumental normative value of the poles of free and informed consensus and deregulation without domination. As previously argued, the impossibility of overcoming the provisional and always contested status of political decisions depends on the instrumental value of these normative poles, not merely on the epistemic limitations of the system. Even if we knew the best decision in light of an appropriate epistemic and normative evaluation, the instrumental value of the democratic institutions that serve to approximate these twin political normative poles would be preserved.

of its results would likely be inferior to those of the republican model, and there appears to be no political actor motivated to implement or defend such a system. The last position, the least desirable tier, is shared by classical epistocracy and a democracy in which the majority of citizens are incapable of making a minimally competent epistemic and normative assessment of the available political options. Admittedly, classical epistocrats might challenge this ranking, insisting that an epistocratic government led by experts genuinely committed to the common good is superior to a democracy dominated by the political ignorance of the majority. However, even granting an idealized scenario where experts are shielded from the influence of powerful minorities without the vigilance of an informed majority, confidence in the instrumental value of this model seems plausible only in the short term. In the medium- to long-term, both political paths—one leading to a democracy dominated by irrationality, the other to a classical epistocracy—appear to condemn the majority to the same fate: a progressive deterioration of their living conditions.

6. Democracy, Market and Partiality

The proposal we defend is supported by a structural similarity between the market and the democratic political process that is often overlooked. Neglecting this parallel can lead to confusion and misinterpretation. To avoid this, it is useful to examine how Kenneth Arrow conceived the analogy between the market and democracy, and to distinguish his approach from the one we defend. Arrow addresses this question in a particularly interesting way in a critical review of John Rawls's *A Theory of Justice*. One of his central concerns in that work centered on an issue that received only marginal attention in Rawls's proposal: the role that justice as fairness assigns to the democratic political process.

Arrow criticizes the idea, which he attributes to Rawls, that the main function of majority rule is merely to resolve disagreements between conflicting interpretations of the theory of justice. While Arrow accepts the limitations Rawls places on drawing a close parallel between economic and political competition, he insists that the similarity is more relevant and profound than Rawls—and most political philosophers—recognize. According to Arrow

Clearly, there is something to Rawls's position, which indeed he shares with many political philosophers, as he notes. A political system in which there is no other-regardingness will not function at all. Further, Rawls is right in saying that the analogy with the market is imperfect. In the market, he agrees that selfish behavior is socially correct, but holds that the political process can never lead to perfect justice if based on self-seeking behavior. But I would argue that the analogy, though imperfect, is not completely wrong either. Political competition does serve some of the same functions in its sphere as economic competition. Further, the expression of one's own interests in voting seems to me an essential part of the information process needed for voting. Unless voters express their interests, how is anyone going to know if the ends of justice are in fact being carried out? "If I am not for myself, then who is for me?" said Hillel, though he continued in more Rawlsian terms, "and if I am not for others, then who am I?" (1973, 258).

Political competition in a democratic system does not simply serve the function of settling disagreements between different interpretations of a shared impartial normative framework, as Rawls seems to assume in *A Theory of Justice*. Rather, it is a process in which agents strive to advance their partial interests, even when acting within a normative framework oriented toward the common good; it always involves an element of self-interested behavior. This renders the functioning of democratic politics structurally similar to that of the market. Significantly, however, for Arrow this partiality is not only ineradicable in both spheres but also desirable. In both politics and economics,

partiality helps achieve the inherent goods of each practice, however distinct they may be. The competition between agents seeking to advance their partial interests is a primary driver of the market economy's capacity to generate economic value (wealth, development) and democratic politics' capacity to produce political value (the common good, however it may be defined).

Arrow's idea is best illustrated by the economic policy debate referenced throughout this paper. Consider two societies, A and B, both of which adopt the Rawlsian theory. Both subscribe to the difference principle, which mandates that inequalities in income and wealth are acceptable only if they result in the greatest benefit to the least advantaged. However, these societies apply the principle differently. Suppose the magnitude of resource transfers from the middle and upper sectors to the least advantaged is significantly greater in Society A than in Society B. This variance reflects their distinct demographic structures. In Society A, the least advantaged sector is numerous; by allying with other sectors—such as a large lower-middle class—it attains significant political influence and secures the implementation of demanding redistributive policies, the costs of which fall mainly on the upper strata. Conversely, in Society B, the middle class is much larger. While the majority of citizens accept that redistributive policies are fair in principle, they promote their own partial interests by supporting less radical redistribution to minimize their own tax burden. Arguably, the fact that the democratic process produces different policies in societies A and B is a desirable outcome, demonstrating the value of this decision-making system in adapting to diverse circumstances. This system is well-suited to ensuring the stable and efficient functioning of societies with different compositions.

In both scenarios, decisions are grounded in a reasonable and impartial normative framework aimed at the common good. However, the dynamics of political competition lead to divergent interpretations of this framework, prioritizing

the distinct partial interests of the majority in each society. This balance and integration between impartiality and partiality appears reasonable and desirable. It serves as an apt illustration of the position Arrow defends. The same logic extends to other impartial normative frameworks, such as utilitarianism or prioritarianism.

Nonetheless, these ideas are subject to a salient objection. It can be argued that partiality in politics does not yield desirable results from the standpoint of justice. It may be true that the most efficient way to generate wealth in economic activity is through competitive processes where agents seek to maximize their own returns—a context in which self-interested behavior is both acceptable and conducive to the goals of the practice. Yet, if the political system is highly susceptible to the partial interests of individuals or groups, this appears to threaten the realization of justice, even if such partiality is prudentially desirable for the majority. For example, if the difference principle is accepted, but a lower level of wealth redistribution is implemented than what is economically feasible because it is prudentially advantageous for the middle-class majority, this decision undermines the goals of justice conceived from an impartial perspective.

A similar dynamic emerges if we adopt a utilitarian normative framework. If a policy exists that maximizes total expected utility evaluated in cardinal terms, there is no reason to prefer a decision that has more votes—that is, one superior in an ordinal evaluation—but which yields less total net utility. Gains in the partial dimension appear to inevitably come at the expense of impartiality; indeed, this statement borders on the tautological. Consequently, the maximization of ordinal utility generated by majority rule holds a merely prudential status and must be subordinate to the moral point of view. Justice always takes precedence over individual partial preferences. If prioritizing the majority's partial preferences implies a decline in the quality of outcomes when viewed through purely impartial standards, it lacks justification. This characterizes

the classical perspective of philosophers who view politics as a mere proxy for morality, or simply an imperfect application of moral principles.

While we reject this philosophical logic, the proposal we defend is not vulnerable to the common objections raised against Arrow's position. We fully agree with Arrow that political competition does serve some of the same functions in its sphere as economic competition, and that the partial interests of individuals provide information essential to the quality of decisions revealed by an effective political process—provided reasonably favorable conditions are met. This principle extends to the outcomes of the contestation mechanisms characteristic of the republican model we advocate. In both cases—whether decisions are reached via majority rule or institutional contestation—the will or partial interests of those affected are inevitably expressed, however much the discourse may center on impartiality and the common good. Crucially, however, the reference to “partial interests” differs between our proposal and Arrow's. In our view, partiality persists as a structural component of decision-making, even were it possible for agents to purge all self-interested appeals and limit their justifications to plausibly impartial considerations.

As previously noted, our proposal highlights that in standard political and everyday contexts, agents must adopt a specific attitude toward risk to make any decision.³⁶ This

36. Despite Arrow's influential work on the role of risk in the economic process (see Arrow 1951, 1953, 1965, 1970, and 1996), he never integrated risk into the specific logic of democratic politics. His ordinalist utilitarianism—unlike the republican ordinalist consequentialism we defend—does not center on risk. Our proposal essentially synthesizes Arrow's ordinalism with Harsanyi's focus on the relevance of risk to rational decision-making processes. However, the abandonment of cardinalism reveals a distinct form of normativity, one separate from impartial morality and particularly suited to conceptualizing political processes. Indeed, Arrow was highly critical of introducing risk considerations into moral reasoning and welfare economics (see 1963, 10). Harsanyi, who defended the role of risk in both fields, explicitly

step is an essential, unavoidable, and irreducibly subjective part of the process. The partiality to which Arrow refers—termed Partiality 1—can be eliminated or neutralized by a moral discourse appealing to “objective and impartial” reasons presumed to be universally valid. This is the typical strategy of the epistemic model and of a strong tradition in political philosophy that conceives politics as a mere proxy for morality. However, the partiality introduced by risk—Partiality 2—is ineradicable in most contexts involving relevant collective choices. Crucially, Partiality 2 persists even in the presence of impartial justifications—that is, even after Partiality 1 has been eradicated. This constitutes the primary difference between the two positions: Partiality 2 is largely ineradicable

confronted Arrow’s position (see Harsanyi 1979, 298–300). Conversely, John Rawls endorsed Arrow’s stance in *A Theory of Justice* (see Rawls 1971, 323–325). Our proposal need not adjudicate this debate. As we have seen, the focus our proposal places on risk is part of a political conception, not merely a moral or epistemic one. Perhaps Arrow and Rawls were correct to suggest that when risk becomes a factor in our deliberations, we depart from the domain of pure morality. The inescapable necessity of taking a position on risk in nearly all relevant political decisions is a central feature of our proposal. This necessity, in turn, precludes specific outcomes from being legitimately presented as morally correct or obligatory. Consequently, every decision remains inherently revisable—a condition that does not depend exclusively on agent fallibility. The direct normative goal of the system is not moral correctness, but the expansion of informed consensus and deregulation without domination. In the form of indirect consequentialism we defend, what holds indirect moral value is the continuity of the system, which offers the greatest probability of producing the best expected outcome from a moral point of view in the medium- to long-term. Our proposal therefore recognizes the indispensable role of risk attitudes, partiality, and individuality in collective decisions, without renouncing the commitment to the common good. Risk, channeled through political competition within democratic republican institutions, neutralizes the autonomous legitimating capacity of moral discourse—that is, the possibility of justifying decisions outside these institutions—while contributing to the overall moral quality of outcomes in a timeframe relevant to the well-being of those affected.

and effectively neutralizes the independent justificatory potential of “objective” or “impartial” moral discourse.

A second distinction arises regarding instrumental value. In contrast to Partiality 1 (simple self-interest), Partiality 2 (a reasonable risk attitude) possesses distinct instrumental utility within moral discourses aimed at the common good. Conversely, the accommodation of Partiality 1 within collective decision-making is typically regarded—contra Arrow—as a normative defect. Impartialist moral theories, whether utilitarian or Rawlsian, negate the validity of claims predicated solely on selfish considerations. While agents retain partial motivations, they are ethically obligated to subordinate them to the common good wherever these diverge. If an agent subscribes to utilitarianism, for example, they must be prepared to sacrifice fundamental personal projects should they impede the maximization of utility. In such contexts, claims grounded in Partiality 1 lack standing, as their satisfaction is incompatible with the collective good.

From the perspective of morality or justice, Partiality 1 can and should be eliminated or overcome; failure to do so results in a normatively inferior outcome. However, attempting to implement decisions based on objective or impartial discourse in contexts of risk does not eliminate Partiality 2; rather, it implies that the legitimizing power of such discourse is never conclusive in these contexts. Conversely, rather than worsening the expected outcome in terms of moral standards, incorporating this information—Partiality 2 of those involved—has, as demonstrated, precisely the opposite effect. Increasing the number of individuals whose preferences are satisfied—the ordinal maximization of expected utility in a process that is always open and revisable—is the only rational method to achieve the best substantive outcome in the medium- to long-term in contexts of risk. There is no opposition or incompatibility between the goal of discourses centered on impartiality and Partiality 2—or, at minimum, it does not present the type of conflict that exists between such discourses and Partiality 1.

We have attempted to show that the adequate representation and channeling of Partiality 2 is one of the great virtues of an adequate political system. This approach, which we term republican ordinalist consequentialism, creates a system of risk management in which individuals play an irreducible role. We argue that this is a decision-making framework superior to any direct, self-sufficient appeal to moral or factual truth. Ultimately, it is superior because it better promotes the instrumental goal of producing the best possible outcomes in the relevant time horizon, judged by independent factual, prudential, and moral standards. Our proposal, like Arrow's, emphasizes a fundamental structural similarity between market competition and democratic political competition. However, this similarity does not rest on Partiality 1 (narrow self-interest). Rather, it rests on Partiality 2: the adoption of reasonably divergent risk attitudes by the agents involved. This function, on non-epistemic grounds, strengthens the case for both the free market in the economic sphere and contestatory democracy in the political arena.

The goal of agents in the market is profit maximization. However, in standard market contexts, agents rarely accumulate significant wealth by pursuing this goal directly. Success instead requires recognizing the importance of subjective maximization: the pursuit of profit within reasonable margins of risk—margins that vary substantially between different economic agents. In economics, risk is a primary variable driving the subjective determination of value. As Daniel Bernoulli theorized in the early 18th century, foreshadowing the marginalist revolution, identical goods can possess vastly different levels of utility for individuals with divergent needs or plans. This subjective nature of value, combined with the epistemic problem of dispersed knowledge (both articulated and tacit), strengthens the case for the freedom of economic agents. Just as Hayek argued regarding epistemic limitations, the existence of risk transforms individual freedom into a condition for the efficiency of the economic system.

Although it may contradict the intuition of many philosophers, adhering to objectivism in the economic field constitutes a denial of reality. On the contrary, the principles of subjectivity and freedom form the true basis of economic prosperity. By the same token, objectivism in politics functions as both a denial of reality and an ideology in itself. In the political sphere, risk subjectivizes and partializes normative frameworks that claim to provide a universally valid definition of the common good—whether those frameworks appeal to moral, ideological, or scientific authority. The critique here is not that these discourses are merely cloaks for egoism (Partiality 1), but rather that they lack the justificatory power they claim in contexts of risk. The rationale for this is clear: under conditions of risk, every decision is inherently partial. It constitutes a specific stance imposed by one agent upon another, and thus lacks universal validity. Consequently, risk renders republican democratic politics necessary as a means of producing value in collective action. This system relies on decision-making procedures based on the ordinal maximization of the informed preferences of those affected. While the common good is pursued through discourses aspiring to impartiality, it is conceptually and practically essential that real individuals, with their divergent goals, retain the final word. There can be no authority above a process that is intrinsically open and revisable. Within this framework, majority rule justifies decisions against a background of institutions oriented toward two normative poles: expanding uncoerced consensus and minimizing centralized regulation, provided such authority is not required to eradicate domination.³⁷

37. One might question the analogy drawn between the case of the firm and that of economic policy. Although, as established, our analysis of the case of the firm includes moral considerations, a key distinction remains: in the business context, partners may appeal to purely egoistic considerations. They arguably possess the right to defend their own interests without reference to the common good. In contrast, regarding economic policy, citizens must

Arrow concludes his review of *A Theory of Justice* by noting that insofar as individuals are truly individuals—“each an autonomous end in himself”—they “must be somewhat mysterious and inaccessible to each other.” Consequently, “there cannot be any rule that is completely acceptable to all. There must [...] be the possibility of unadjudicable conflict, which may show itself logically as paradoxes in the process of social decision-making” (Arrow 1973, 263). This perspective is consistent with theoretical frameworks where impartiality relies on the aggregation of preferences. However, unlike cardinal utilitarianism, which presupposes access to preference intensity, Arrow’s ordinal framework implies that individuals remain autonomous ends in themselves precisely due to structural and insurmountable information deficits. Because full impartiality is precluded by the opacity of individual preferences, insoluble conflicts are inevitable. Logically, these manifest as paradoxes within collective decision-making, most notably in the Impossibility Theorem formulated by the author (Arrow 1963).

Arrow posits that the fundamental opacity of other individuals, and the inability of any system of rules to impartially settle the conflicts arising from this inaccessibility, serve as the foundation for autonomy. These limitations guarantee that individuals remain ends in themselves and possess an ineliminable moral standing. This represents a paradoxical and inverted “Kantian consequentialism,” primarily because it constitutes the antithesis of a central tenet in Kant’s moral philosophy. From a Kantian perspective,

appeal to the common good in light of a conception of justice centered on impartiality—such as utilitarianism, prioritarianism, or Rawlsian theory. However, from the perspective of our approach, this difference is irrelevant due to the underlying structural analogy. This analogy rests on the fact that, in both cases, risk introduces subjectivity, partiality, and generates reasonable disagreement *within* the respective normative frameworks—whether based on self-interest or impartiality.

the notion that the same set of rules applies to all rational beings is foundational; moral laws—the laws of reason—are universal, a fact inseparable from individual autonomy. Arrow, conversely, suggests that the inviolability of the individual and their equal dignity derive precisely from the fact that there can be no set of rules common to all, equally acceptable to all, or rational for all.

From a deontological perspective, one might argue that Arrow's idiosyncratic conception—despite its Kantian echoes—remains viable only within the conceptual limits of utilitarianism. Such an approach errs by attempting to establish impartiality through the aggregation of partial and subjective preferences. Arrow's proposal thus represents a distinctive attempt to secure the inviolability of the individual within a framework that tends to have precisely the opposite effect: undermining it. However, this conception of autonomy is rendered meaningless if one abandons the erroneous assumption that impartiality can be constructed by aggregating partial preferences. From this deontological perspective, nothing is gained by maximizing the weight of affected partial perspectives (or Partiality 1) in the decision-making process; indeed, the result is quite the reverse. Such attempts yield only the paradoxes of social choice and a degradation of the moral standpoint—a likely outcome when politics is organized through democratic institutions without the reintroduction of the appropriate normative framework from the outside or above.

Our proposal is intended to provide an answer to that line of reasoning. It attempts to draw both utilitarians and deontologists into the political logic implied by the conceptual inversion found in Arrow's conception. In practice, moral degradation also occurs when one attempts to place morality or knowledge above republican institutions. Individual freedom and autonomy do not depend solely on the success of discourses centered on impartiality and objectivity. In decision-making scenarios involving risk, the quality of outcomes relies on an institutional framework that incorporates subjective

information—specifically, the partiality and risk attitudes of those affected—through both aggregative and competitive mechanisms. Correctly interpreted, Arrow’s ordinalist utilitarianism intended to be a form of consequentialism articulated through democratic political institutions. Within this framework, the individual is an irreducible entity—an end in themselves—whom the system exists to serve and who cannot be sacrificed to theories of the common good. Consequently, the republican ordinalist consequentialism we defend links the maximization of the good (or expected utility, however defined) to the maximization of individual freedom, mediated by reasonable disagreement regarding risk. Ultimately, this seeks to generalize Arrow’s idea, validating it against any competing conception of impartiality, justice, or objectivity.

The proposal we defend, like Arrow’s, posits an intrinsic link between the irreducible value of individuals—specifically, their lack of subordination to any higher authority—and the claim that no system of impartial rules can adequately resolve their disagreements. When agents hold diverging but reasonable attitudes toward risk and their wills conflict, no higher authority can adjudicate the disagreement without relying on consultation, aggregation mechanisms, and the expansion of free consensus. As in most political contexts, no external observer can determine *ex ante* what the optimal collective decision is regarding the quality of the outcome. Such determinations require a genuinely democratic process in which affected parties freely express their informed preferences. This mirrors the ideal operation of a free market, where economic agents reveal investment preferences and risk attitudes through a dynamic process that cannot be anticipated by an expert planner. Thus, in contexts of risk—both political and economic—the maximization of individual freedom holds fundamental collective value; as we have argued, it has a decisive positive impact on the quality of expected outcomes.

Contrasting our proposal with Arrow’s framework also enables us to distinguish it more effectively from the defense

of democracy articulated by Pettit in his more recent work. In his essay, “The General Will, the Common Good, and a Democracy of Standards,” Pettit addresses the problem of political domination. He asks: under what conditions can a state impose its decisions on dissenting citizens without constituting arbitrary domination? Even if the state’s action constitutes interference, when does it not imply a threat to the system’s prevailing goal of promoting freedom as non-domination? Pettit’s answer is that domination does not exist within what he calls a “democracy of standards.” Essentially, the political system avoids generating arbitrary domination when it institutionalizes the outcome of discursive processes characterized as “acceptability games”—a concept Pettit juxtaposes with “acceptance games.”

Both types of games involve multiple parties negotiating a solution to a problem and determining a joint course of action. The operating assumption is that any solution is better than none; however, none of the available options garner universal support. The aim of acceptance games is to reach a common decision through a process where parties appeal exclusively to self-interested reasons to defend their preferred solution. A typical example is the negotiation between a seller and a buyer regarding price. Each agent begins the process with an opening offer. If the buyer’s offer falls short of the seller’s asking price, the two may negotiate toward an acceptable equilibrium, even if the final agreement deviates from their initial expectations. Even in this context, agents may introduce “commonly relevant reasons... [asking] the other to recognize the fairness of their offer” (Pettit 2019, 33). However, doing so transforms the interaction into acceptability games, in which the dominant goal is the pursuit of the common good through reasons that are valid or relevant to all. As Pettit states, in an acceptability game

...the parties do not make self-interested bids in the hope of gaining the acceptance of others. Rather what they each do

is to identify their preferred solution and make a case for it on the basis of reasons that are purportedly relevant, if not given the same weight, in every perspective. They argue that in virtue of the reasons adduced in its support, the solution is acceptable from every point of view; it is a solution that each has a reason to prefer. The parties may argue on this pattern in the hope of securing an eventual consensus on a solution that is supported in the view of all. But they may also argue on the assumption that at the end of the exchange there will be a number of candidate solutions standing and that differences have to be resolved by recourse to voting or some such device. And in that case, of course, they will look for a device that itself passes muster under the acceptability game. This might involve any of a range of voting measures, referral to a select committee, recourse to an outside arbitrator, or even the toss of a coin. (2019, 33–34)

According to Pettit's hypothesis, if political institutions operate according to the logic of acceptability games, the authorities' decisions will not result in the arbitrary domination of citizens. If the policies implemented are based on reasons that are recognized as valid and weighty by all, even those citizens who support or prefer solutions rejected in the process must admit the legitimacy of the policies enacted. In such scenarios, agents whose preferences are not adopted would not be subject to the arbitrary enforcement of decisions that ignore their interests and preferences. They would experience interference from political authorities, but not arbitrary domination. Pettit's strategy is another example of the moralized impartialist logic that dominates political philosophy. In essence, it is the same logic to which Arrow objected in Rawls's *A Theory of Justice*. From our perspective, the problem with Pettit's strategy is that "acceptability games" do not eliminate arbitrary domination in contexts of risk. Freedom as non-domination is central to our proposal, incorporating both agent-relative and agent-neutral dimensions. In contexts of risk, when one group imposes a decision on another, arbitrary domination may persist regardless of the quality of the impartial reasons invoked. If

the first agent imposes their risk profile on the second, this constitutes arbitrary domination. Consequently, an appeal to reasons that are universally valid must be complemented by an adequate risk management process that strives to broaden the acceptability of decisions. This expansion of consensus should occur through majority rule, open to contestation, and with special protections for minorities—such as insurance and spaces of deregulation. Otherwise, there is no medium-term guarantee that decisions will benefit the group or maximize its chances of joint survival.

As previously argued, when a minority imposes its risk profile on the majority for an extended period, the group's survival is threatened irrespective of the objective and impartial reasons supporting the decisions. In the case of the firm, *S1* can correctly assert that *N1* represents the best objective balance (the one that any risk-neutral agent would choose) between probabilities of success and the variation in potential outcomes. However, this does not prove that investing in *N1* is the best decision at any specific juncture or as a long-term strategy. This holds equally true when evaluating the variation between gains and losses in the implementation of a conception of justice whose goal, rather than promoting individual profits, is the common good. In contexts of risk, merely satisfying the requirements of an acceptability game is insufficient to eliminate domination. Nor does it serve as a reliable indicator of the instrumental value of the decision-making system.

Pettit and the deliberative democrats (whose ideas he adopts as a starting point)³⁸ rightly emphasize that political institutions must be subject to the logic of acceptability games, rather than functioning merely as acceptance games—as suggested by purely aggregative or “economic” conceptions of

38. Pettit makes explicit references to Habermas (1984; 1996), Cohen (1986; 1989; 2010b), and Gutmann and Thompson (1996).

democracy. The quality of decisions depends on an epistemically adequate assessment of facts valid for everyone, and on a normative framework recognized as reasonable by all. However, as we have seen, meeting these standards is a necessary but not sufficient condition for high-quality outcomes. To achieve this result, we must employ the most inclusive and participatory political process possible, ensuring decisions reflect the real, subjective, and partial will—risk attitudes—of those affected. There is no room here for external arbiters, expert committees, or other epistocratic strategies as direct decision procedures. The logic here mirrors the free market. A market approximates an ideal by reflecting the risk attitudes and subjective preferences of the largest possible number of participants (investors, entrepreneurs, and consumers) in a way that a committee of expert planners cannot emulate. Similarly, an inclusive democratic process must be deployed to generate successful outcomes regarding the normative goals that guide it. To maintain legitimacy and promote freedom as non-domination, the political process must embody both majoritarian elements (regular elections) and protections for minorities, guaranteeing avenues for contestation and the permanent revisability of decisions. This process involves competition and aggregation, resembling a game of acceptance played against the backdrop of games of acceptability.

In order to function properly from an instrumental perspective, the political system must integrate successive layers of acceptance and acceptability games. The games of acceptance on which effective decisions depend in each specific instance, the need for political competition laden with partiality even in highly idealized situations, and the appeal to the strength of numbers are not mere concessions to the imperfection of the world. They possess instrumental value and constitute the correlate of the irreducible value of individual freedom. Aligning with Arrow's perspective—though distinct in our focus on the role of risk—we maintain that at least part of the irreducible value and dignity of the individual

does not depend on their ability to appeal to universally valid reasons, but rather on the impossibility of making rational decisions derived exclusively from those reasons. This leads us to recognize that the collective decision-making system must also reflect, to the greatest extent possible, reasons valid only for each particular individual.

7. Synthesis of Arguments and Conclusions

One way to synthesize our instrumental defense of contestatory democracy, as developed throughout this paper, is to distinguish among three possible scenarios.

First scenario: *Instrumental justification of collective decisions under conditions of certainty.*

Premise 1: Decision-maker *S1* knows with certainty at time *T* that actions *A1*, *A2*, and *A3* will produce states of affairs *E1*, *E2*, and *E3*, respectively, at time *T1*.

Premise 2: *S1* adopts a correct or reasonable normative framework (*NFC/R*) that enables them to establish a ranking of the aforementioned states of affairs that is valid for all agents affected by the decision. For instance: *E3* is better than *E1*, and *E1* is better than *E2* (i.e., $E3 > E1 > E2$).

Premise 3: *S1* knows the relevant facts, adopts *NFC/R*, and has the capacity to perform action *A3* at *T*, which will certainly produce at *T1* the most desirable state of affairs for all concerned: *E3*.

Example: *S1* is the decision-maker in a modified case of the firm and knows with certainty at *T* that *A3* (investing in *N3*)—the option that generates the most profits in case of

success and the greatest losses in the event of failure—will be successful at $T1$, producing the best outcome: $E3$.

Conclusion: Given **Premise 3**, $S1$ is justified in making a binding collective decision, through the unilateral execution of action Ax at T (in this case $A3$), irrespective of any opposition from the other affected parties, because this decision impartially promotes the interests of all. $S1$ is justified in making a binding collective decision because it is the “right” or “correct” one. Under these conditions, no one can reject the agent’s action on the grounds of valid, acceptable, or relevant instrumental arguments.

Second scenario: *Instrumental justification of collective decisions under conditions of significant probability of error, involving considerable welfare costs and a single rational attitude towards risk.*

Premise 1: Decision-maker $S1$ is justified in believing that, at T , actions $A1$, $A2$, and $A3$ will produce states of affairs $E1$, $E2$, or $E3$ with a specific probability P , or the states of affairs $E1f$, $E2f$, or $E3f$ with the corresponding complementary. The second set of states represents the foreseeable outcomes resulting from the failure to realize the intended states in the first set.

Premise 2: $S1$ adopts NFC/R , which enables them to establish a ranking of preferability among the aforementioned states of affairs. This ranking is derived by balancing—in light of NFC/R —the probabilities of success against the consequences of failure (obtaining negative or undesirable outcomes). This calculation assumes a single rational risk attitude. Given these three elements— NFC/R , the probabilities of the different outcomes, and a singular rational risk attitude—a unique ranking is established.

Premise 3: Since the Order of Preferability of Actions (*OPA*) is based on objective, factual, and impartial normative considerations, and since there is only one rational risk attitude in this context, the ranking is instrumentally valid—at least from an ideal perspective or under appropriate conditions—for all affected parties.

Premise 4: *S1* can establish an *OPA* given their knowledge of the relevant facts, their ability to appropriately assign probabilities to various possible outcomes, their adherence to *NFC/R*, and their rational risk attitude. In light of the *OPA*, performing action *Ax* is the best decision. This decision satisfies two tests of instrumental correctness:

- **First test:** If the agents rejecting the decision were not making an epistemic or normative error and subscribed to the only rational attitude toward risk, they would reach the same conclusion: performing *Ax*. Therefore, they would support the same decision as *S1*.
- **Second test:** In an iterated scenario, *S1*'s decision to perform *Ax* will predictably produce the best aggregate outcome within a timeframe relevant to the welfare of those affected.

Example: The Miners case. Closing Gate 3 (action *A3*, which entails a probability of 1 of producing *E3* and 0 of producing *E3f*) yields an expected (and actual) utility of 90 units. In contrast, closing gates 1 and 2 (actions *A1* and *A2*, which entail a probability of 0.5 of producing *E1* or *E1f* and *E2* or *E2f*, respectively) each yields the same result: 50 units of expected utility (100 units weighted by a probability of 0.5). In this scenario, the only acceptable attitude toward risk is neutrality. The possibility of a risk-prone agent assigning an expected utility greater than 90 units to closing gates 1 or 2—which introduce the chance of either saving everyone or everyone perishing—is excluded.

Conclusion: Given **Premise 4**, *S1* is justified in making a binding collective decision, through the unilateral execution of action *Ax* at *T*, which will produce outcome *Ex* at *T1* with probability *P*, or outcome *Exf* with the corresponding complementary probability, irrespective of any opposition from the other affected parties. This justification holds because the decision satisfies the two tests of instrumental correctness, thereby maximizing the net expected utility. Consequently, the decision is justified as the “right” or “correct” one. Under these conditions, no one can reject the agent’s decision on the grounds of valid instrumental arguments.

Third scenario: *Instrumental justification of collective decisions under conditions of significant probability of error, involving considerable welfare costs and legitimate disagreement among those affected regarding the most rational / reasonable level of risk.*

Premise 1: Decision-maker *S1* is justified in believing that, at *T*, actions *A1*, *A2*, and *A3* will produce states of affairs *E1*, *E2*, or *E3* with a specific probability *P*, or the states of affairs *E1f*, *E2f*, or *E3f* with the corresponding complementary. The second set of states represents the foreseeable outcomes resulting from the failure to realize the intended states in the first set.

Premise 2: *S1* adopts *NFC/R*.

Premise 3: *S1* knows that, under normal circumstances, reasonable disagreements about the acceptable level of risk when making a decision are to be expected. These disagreements may arise even among agents who accept the same factual description, assign the same probabilities to outcomes, and evaluate them in light of the same *NFC/R*. Given the disagreement among those affected about the

acceptable or desirable level of risk, the choice of any of the three options (A1, A2, or A3) does not satisfy the two tests of instrumental correctness.

- **First test:** *S1* cannot argue that the agents opposing the decision would have endorsed it if they had conducted a factually and normatively competent assessment. Disagreements about risk attitude justify dissent irrespective of those variables, as exemplified by the case of the firm, the benefactor game, and debates on economic policy design.
- **Second test:** *S1*'s decision cannot be characterized as "correct" because, in an iteration scenario, there is no guarantee that it will produce the best aggregate expected outcome. Repeating the same decision offers no guarantee of improving the quality of the outcome within a timeframe relevant to the real interests of those affected.

Premise 4: Disagreement about the level of acceptable risk is instrumentally relevant. The risk attitude of those affected provides dispersed information—both objective and subjective—that is decisive for the quality of expected medium- or long-term outcomes, which affect welfare at both individual and collective levels. This information concerns the tolerable variance in gains and losses of the goods at stake across available actions. Without processing the information about the divergent risk profile of each individual, an adequate assessment of the available courses of action from epistemic (factual) and normative (prudential or moral) points of view—considerations supposed to be equally valid for all affected parties—is insufficient to rationally promote the quality of the outcomes of the decision-making process.

Premise 5: *S1* has five options:³⁹

1. **Neutral Balance:** Implement a course of action based on an objective balance between the factual and normative judgments involved, adopting a risk-neutral attitude.
2. **Decision-maker preference:** Implement a course of action superior in light of *S1*'s own risk aversion or propensity, without considering the preferences of those affected.
3. **Majority Rule:** Implement the option supported by the majority of the group. This option is presumed to reflect the group's prevailing attitude toward risk.
4. **Expanded Consensus or deregulation (Two institutional variants):**
 - 4a: Implement options that attract a higher level of support than that generated by the initial application of majority rule. This includes mechanisms like insurance for risk-averse agents, partial exemptions for risk-prone agents, deliberative institutions, or contestation of collective decisions. These measures move beyond the initial threshold to favor decisions that expand free and informed consensus or "deregulation without domination."
 - 4b: *S1* acts as an expert, identifying the options described in 4a without developing the full participatory process. *S1* determines which decisions *would have* gained majority and subsequent support if strategies to broaden voluntary consensus had been utilized.
5. **Maximum Precaution:** Proceed under the assumption that those affected are incapable of a minimally competent

39. We decided not to discuss strategies derived from different variants of utilitarianism, based on the reasons examined in section 4.1. If utilitarianism were a consistent or viable option, *S1* could attempt to maximize either actual utility or expected cardinal utility.

epistemic evaluation, subscribe to an unreasonable normative framework, or cannot form an instrumentally rational risk attitude. In this case, since it is impossible to know what decision they would have made, *S1* combines an epistemically and normatively competent assessment with an attitude of maximum reasonable risk aversion on behalf of those affected.

Premise 6: Given **Premise 4**—which establishes the instrumental value of the affected parties’ risk attitudes—the ultimate normative goal of the collective decision-making system (i.e., obtaining the best possible expected outcome) is identified with a process in which all those affected freely decide on their preferred option. This decision must be made in light of: a) a competent epistemic evaluation of the relevant facts, b) an adequate assignment of probabilities of success and failure to the available courses of action, c) the adoption of an NFC/R, and d) alignment with the risk attitude of as many of those affected as possible (or, ideally, all of them). There are two ways in which all those affected can freely decide on their preferred option: either when a unanimous, unforced, and informed consensus is reached on a single decision, or when each agent assumes their own risks by making divergent informed decisions. The ideal goal, once the requirements of epistemic and normative competence have been met, is not only “truth” or “correctness” but an approximation to the normative poles of free and informed consensus and deregulation without domination.

Examples: The case of the firm developed in sections 2.1 and 2.2, the benefactor game discussed in section 2.2, and the example of choosing an economic policy guided by a Rawlsian, utilitarian, or prioritarian normative framework, examined in section 3. In all these cases, maximizing expected utility involves processing information about the

divergent preferences of those affected in light of their risk profiles as efficiently as possible.

Conclusion: In contexts of risk, an adequate assessment of the available courses of action from epistemic (factual) and normative (prudential or moral) points of view never directly and sufficiently justifies the imposition of binding collective decisions. Given **Premises 4** and **6**, *S1* is justified in making a binding collective decision, through the execution of action *Ax* at *T*, which will produce outcome *Ex* at *T1* with probability *P*, or outcome *Exf* with the corresponding complementary probability, if and only if *S1* has managed to minimize opposition to *Ax* by those affected (optimally, eliminating it completely) by approaching the ideal normative pole of free and informed consensus. In contexts of risk, binding collective decisions are never instrumentally justified merely because they “track the truth” or because they are “correct,” without the existence of an institutional framework that reduces the third source of political disagreement through an adequate risk management system. On the other hand, if feasible in practice, the converse strategy is equally desirable from an instrumental point of view. Instead of seeking free and informed agreement on a single course of action or a common norm, the goal can be a scenario of deregulation without domination in which those affected make divergent and informed decisions aligned with their attitude toward risk.

Based on **Premise 6**, we derive the following ranking for the options in **Premise 5**:

- **Best Strategy:** Option 4. This involves (4a) establishing a system progressing toward informed consensus or deregulation without domination (Contestatory Democracy), or (4b) the speculative anticipation by experts of the outcomes of such a system (Risk-centered Epistocracy).

- **Second Best:** Option 3 (Majority Rule), assuming decisions are made among competently evaluated options that reflect the group's prevailing risk attitude.
- **Third Best:** Option 5 (Maximum Reasonable Risk Aversion), applicable when affected parties lack competence.
- **Excluded:** Options 1 and 2, as they lack instrumental value in this context by ignoring the risk factor.

Recognizing the distinctions among these three scenarios clarifies the limited relevance of the debate between epistocrats and epistemic democrats. It also demonstrates the insufficiency of these perspectives in addressing the problem of the instrumental value of collective political decision-making processes. The epistocratic perspective can be summarized as follows:

Premise 1: The quality of the outcomes produced by a political system is the decisive criterion for evaluating its merits, independent of the intrinsic value of the procedures through which decisions are made.

Premise 2: The group whose members have superior collective epistemic competence (assuming they are genuinely committed to the common good) will make the best decisions.

Premise 3: If the goal of the political system is the common good and Premise 2 is true, then the group whose collective epistemic competence is superior should make the decisions or exercise greater political influence.

Premise 4: In contemporary democracies, the vast majority of citizens lack the information—or the competence to evaluate the information—necessary to make intelligent or rational decisions. Majorities are characterized by the

phenomenon of political ignorance. Consequently, a minority of experts or agents with higher individual epistemic competence than the average citizen will also possess higher collective epistemic competence.

Conclusion: Assuming they employ their knowledge to promote the common good, the group of experts with higher collective epistemic competence should make the decisions or possess greater political influence than the majority.

Epistemic democrats also adopt an instrumental conception of the value of the political system in their argument for majority rule; however, they reject Premise 4 of the preceding argument and, consequently, its conclusion. Epistemic democrats claim that, despite the lower individual competence of their members, the collective epistemic competence of majorities is superior to that of any minority of experts. Consequently, they argue that majorities should make political decisions and have more political influence than minorities. The debate between epistocrats and epistemic democrats would be settled decisively in favor of the institutional design defended by each group if at least one of two conditions were met: 1) political authorities had full certainty regarding the outcomes of their decisions; or 2) the failure of implemented policies did not impose a significant cost on the welfare of the people affected. Another possibility is that the expected utility maximization calculation, as illustrated by the Miners case, could be performed—meaning policy decisions would admit only one rational or reasonable attitude toward risk.

Nevertheless, none of these options is available. The probability of error in relevant policy decisions is always significant—often high—and the cost of error is usually substantial, often dramatic. Reasonable disagreement about the most desirable levels of risk seems to be as much a characteristic feature of political competition as disagreement at the normative level. The epistemic model would provide

an adequate framework for analyzing the instrumental value of various systems of collective decision-making if either the first or the second scenario actually provided an accurate characterization of political reality. However, this is rarely the case. Political decisions are usually made under conditions similar to those of the third scenario. The acknowledgment of risk—the necessity of considering the consequences of a plan's failure—inherently implies that our available knowledge and information remain inadequate for identifying the uniquely correct decision. If the adopted normative framework and the evaluation of relevant factual information were sufficient to determine the single correct course of action, then admitting that risk is involved would be meaningless.

In politics, as in a wide range of important decisions made by institutions and individuals, it is necessary to decide whether one is willing to sacrifice something valuable at time T for the possibility of obtaining something even more valuable at $T1$. Insisting on purely epistemic or moral solutions in this type of scenario simply demonstrates a failure to understand the phenomenon at hand, or arguably amounts to illusion.

As in business, a rational agent in politics who aspires to achieve their goals without relying exclusively on luck must simultaneously satisfy three conditions:

1. obtain all possible empirical information relevant to the problem at hand;
2. adopt an appropriate normative framework; and
3. manage risk exposure—either by hedging (through insurance, diversification, etc.) or by consciously accepting higher risk to maximize potential gains.

These three conditions are critical: the failure to satisfy even one is enough to guarantee a suboptimal outcome. Premise 2 of the argument, on which the plausibility of both epistocracy and epistemic conceptions of democracy depends, is false in contexts of risk. There is no guarantee that the decisions of

the group with the greatest collective epistemic competence constitute the optimal course of action from an instrumental perspective, or that they are superior to those made by other groups.

Pursuing the optimal expected outcome requires identifying the prevalent risk attitude and initiating a perpetually open, revisable process. This mechanism should encompass competition for majority formation, strategies for extending free and informed consensus, and robust contestation mechanisms designed to safeguard minorities. In most political decision-making scenarios, estimating the best expected outcome is impossible without specific information regarding individual preferences—dispersed knowledge that includes risk attitudes. This information cannot be reduced to objective factual or impartial normative judgments equally valid for all. In scenarios such as the firm, the benefactor game, or the design of a State's economic policy, this particular, subjective, and partial information is directly relevant to the quality of the aggregate expected outcome over time. Consequently, disagreement among those affected is reasonable, precisely because the information driving that disagreement determines the quality of the result (as stipulated in Premise 3 of the *Third Scenario*). Therefore, rational collective decision-making precludes strategies that insist on a universally valid or "correct" solution, impose such solutions on dissenters, or suppress disagreement in the name of a fixed common good. Fundamentally, this represents a significant conceptual pitfall for epistocrats and epistemic democrats alike: both subscribe to an inadequate conception of the normativity governing collective decision-making in contexts of risk—particularly when assessed from an instrumental perspective.

In contexts involving risk, direct appeals to truth and moral correctness, unmediated by an adequate political process, can reveal their ideological and destructive potential, posing a threat to individual freedom and the quality of the system's outcomes. This is our consequentialist reinterpretation of

Arendt's and Rawls's analyses of the dangers of relying on truth as the ultimate criterion for justifying political decisions. In a decision-making system suited to contexts of risk, the final authority rests with political institutions adopting a strategy that navigates between the normative poles of expanding free and informed consensus and deregulation without domination. In essence, a process of ordinal utility maximization is required, ideally realized either through full consensus on a single decision, or through complete deregulation without the domination of any agent or group. This conclusion constitutes Premise 6 of the *Third Scenario*, derived from Premises 3 and 4. The process, which is anchored in real individuals whose partial preferences are not neutralized or superseded by the common good, becomes morally valuable from an instrumental point of view when risk is introduced. Risk is the bridge between the number—the quantity of individuals consulted—and the quality of outcomes. In order to rationally pursue the best expected outcome, it is essential to obtain and process the widest possible array of information regarding the desirable and tolerable variation in gains and losses of the goods at stake for each affected individual.

In a hypothetical scenario where the future is certain or where we can calculate the best expected outcome based on objective data and an accepted normative framework—assuming a single rational attitude toward risk, as in the Miners case—the number of individuals supporting a decision does not directly influence the quality of the outcome. On the contrary, when there is risk, the number of participants plays a decisive role in the quality of the expected result and the subsistence and flourishing of individuals over time, under the presupposed conditions of epistemic quality and moral reasonableness. As we have seen, if the information about the level of risk that each individual is prepared to face is not taken into account, the flourishing and perhaps the survival of most members of the group will be compromised in the medium- and long-term. They will face unfavorable scenarios stemming from

excessive losses or insufficient gains, situations from which they are not equipped to recover. This will ultimately cause a drop in living conditions or, in extreme cases, systemic collapse. This is the conceptual basis of the instrumental value of republican democracy as a procedure that reveals information about the will of those affected and their attitudes toward risk, distinct from its possible epistemic potential and intrinsic value. It is also the foundation of the irreplaceable value of the individual and their freedom in collective decisions from a consistent consequentialist point of view.

At the level of collective political decisions, the conclusion of our work can be summarized with a concise maxim: “republican democracy or decadence,” or more briefly, “republic or decadence.” As we have seen, even if a minority of experts is committed to a plausible conception of the common good, the majority of those affected can expect a progressive deterioration of their living conditions if decisions are made by that minority. In the absence of an adequate risk management system, expert decisions can be categorized in three ways:

1. risk-neutral decisions;
2. decisions that reflect the risk attitude adopted by a minority (neutrality, aversion, or propensity); and
3. decisions that reflect the risk attitude of both the minority and the majority.

Due to the disparity in living conditions between the most powerful and influential minorities and the majority of citizens, the third option—convergence of the risk profiles of both groups—is highly improbable, especially in the medium- and long-term. The first option—maintaining an attitude of neutrality in the face of risk—is theoretically optimal in an infinite time horizon, but it represents a speculative gamble that could prove irrational for finite agents operating within limited time frames to realize their projects. The second option is for minorities to impose their predominant risk profile on

everyone. This alternative only increases the chances of success for minority decision-makers. Crucially, this outcome remains favorable for minorities only under a specific condition: that the elites remain insulated from the deterioration of the majority's welfare—a deterioration caused by exposure to unaffordable risks or losses incompatible with their projects. Ultimately, neither of these options is adequate for the majorities. For them, the choice is simple: republic or decadence.

The model we term “risk-centered epistocracy,” which aims to transcend the epistemic model's limitations by addressing the three primary sources of political disagreement, is theoretically nonviable. The arguments against it mirror those adduced against the planned economy: just as a minority of economic experts cannot predict the price signals generated by market mechanisms, expert political planners cannot access the dispersed, evolving knowledge—both objective and subjective—required to replicate the outcomes of a robust political system. A functioning political system involves an intricate process wherein epistemic competence and normative reasonableness are achieved, and a reasonable diversity of risk attitudes is integrated. This integration is best managed inclusively through free electoral institutions and mechanisms of contestation. Consequently, it appears infeasible to artificially simulate the outcomes that this complex, emergent process would produce.

The other epistocratic alternative we have imagined—in which experts adopt maximum reasonable aversion to risk—is less theoretically problematic. One might even defend the hypothesis that such a system would efficiently guarantee the group's subsistence. However, the doubts regarding its viability are primarily practical. Neither the privileged elites nor the vulnerable and misinformed majorities—if we accept the characterization based on sociological research usually cited by epistocrats—have any incentive to support such a system. The remaining option is a majoritarian democracy. This model treats majority decisions as final, premised on the assumption

that the majority possesses a higher level of collective epistemic competence than any minority of experts (and thus tracks the truth). However, this option is also flawed because it fails to account for the phenomenon of risk. From an instrumental perspective, when risk is present, the dominant normative logic is not the search for truth or moral correctness. Rather, it is the maximization of freedom through electoral and contestatory institutions that aspire to broaden free, informed consensus and guarantee the existence of spaces of deregulation without domination. If a majority of citizens can indeed reach a threshold of epistemic competence and normative reasonableness, there is no justification for restricting the decision-making process solely to majority rule. Under these conditions, it is feasible and desirable from an instrumental perspective to develop the full set of institutions that characterize the republican model.

Conceptually, our proposal represents a novel normative framework: risk-centered ordinal consequentialism. Its unique characteristics are best clarified by contrasting it with the three consequentialist views previously discussed. From Pettit's republican theory, we adopt the contestatory conception of democracy and the centrality of freedom as non-domination. In our proposal, as in Pettit's, the freedom at stake is not merely freedom from interference. The mere existence of domination, even when not actively exercised, can suppress the critical information necessary for making rational individual and collective decisions in contexts of risk. When individuals lack protection against arbitrary interference by others, they cannot freely form or express preferences consistent with their risk attitudes. Consequently, the dispersed and subjective information required for adequate individual and collective risk management is lost. As previously noted, the contrast between the two proposals lies in Pettit's treatment of freedom as non-domination as an exclusively agent-neutral good. Consistent with this perspective, he argues that avoiding arbitrary domination requires political decisions based on reasons acceptable to all—thereby relying on a plausible

interpretation of the common good. Pettit does not subscribe to the epistemic model but embraces an impartialist conception, akin to that found in Rawls and the dominant tradition in political philosophy. At the level of political legitimacy, non-domination is realized by practices and institutions guided by the logic of what Pettit calls “acceptability games.” As discussed above, these are discursive processes that aspire to reach agreements—or make decisions—based on acceptable reasons recognized by all affected parties. This transcends the partial and self-interested perspective of “acceptance games.”

We defend a proposal that defines freedom as non-domination as a good that is simultaneously agent-neutral and agent-relative. However, in our proposal, it is this second dimension that plays a decisive role in favor of democratic institutions. Acceptability games and their facilitating institutions do not completely eliminate arbitrary domination in contexts of risk. Even policies based on objective and impartial judgments—though equally valid for everyone—may impose a pattern of risk (a variation in possible losses and gains) that threatens the well-being and survival of those affected. Divergences between risk patterns cannot normally be resolved by acceptability games, since they express subjective and agent-relative considerations, inseparable from the distinctions between individuals, rather than their common features. Ultimately, politics is necessary not merely to discover reasons that are valid for all, but because such reasons—even if they were available, which is always questionable—are insufficient for making rational decisions in contexts of risk.

Addressing disagreements about acceptable risk levels requires combining attempts to broaden free and informed consensus—as achieved through acceptability games—with preference aggregation procedures. In other words, it requires an ordinal strategy that maximizes expected utility. We adopt the ordinalism of the utilitarian conception outlined by Arrow, along with his critique of purely impartialist conceptions that assimilate politics to the logic of moral discourse. Against

the background of political competition—where acceptability games play an ineliminable role in broadening consensus and generating legitimacy—the final decision rests on the aggregation of wills. These procedures are necessitated by disagreements and forms of partiality that are insurmountable by discursive means. In the final analysis, the decision relies on the free and irreducible will of individuals and their attitudes toward risk (what we term Partiality 2). Consequently, following the acceptability games, a return to acceptance games is required to make effective decisions. When democratic politics functions effectively, it operates as a chain in which both types of games alternate in a continuous cycle.

As we have seen, the aim is not to make the issue of risk a central topic of public debate. Nor is it to address risk-related disagreements using the logic of acceptability games—that is, seeking reasons valid for all and a single common solution. The instrumental success of such a strategy would require dissolving the disagreement over which levels of risk are most appropriate or desirable. Yet, in most cases, this is unfeasible. We are dealing with irreducibly different subjects—possessing distinct resources, capabilities, goals, and projects—who face decisions where the probability of error is significant and the cost to their well-being is high. Democratic institutions allow for persistent disagreement by grounding decisions in the actual preferences of the majority of those affected—formed in light of a competent assessment of the relevant facts and a reasonable normative framework—rather than simply in what everyone should want if they knew the truth. Crucially, these institutions always leave open the possibility for minorities to dispute and attempt to modify these decisions, regardless of the epistemic and moral quality of their foundations. In doing so, these institutions preserve the irreducible weight of individuality and disagreement regarding desirable levels of risk. They incorporate and reactivate the logic of acceptance games, albeit at a new level of complexity and legitimacy.

Acceptance games—when reflecting the pure self-interest of individuals, as in many economic interactions—constitute a form of collective coordination that is often considered normatively inferior from the standpoint of morality or impartiality. However, the dynamic differs significantly in the case of republican political institutions. Majority rule and instances of individual and collective contestation must be developed against an institutional backdrop that incorporates acceptability games: deliberative mechanisms aimed at achieving a competent assessment of facts and developing a reasonable normative framework. This is essential to maintain the presumption of their instrumental value. Although these processes ultimately manifest as acceptance games—forms of competition laden with partiality involving winners and losers—they represent a superior form of collective coordination. Specifically, they surpass models that rely directly on acceptability games as the basis for decision-making, where an epistemic authority takes total precedence over the will of the individuals who constitute the political community. Pure acceptability games do not represent the ideal goal to strive for in contexts of risk, wherein acceptance games are viewed as a mere concession to unfavorable circumstances. In such scenarios, the ideal goal must inherently include a combination of the two types of games in order to be instrumentally adequate. Only this specific combination of acceptability and acceptance games enables us to simultaneously address the three sources of political disagreement. It facilitates decisions grounded in well-justified factual judgments, a reasonable moral framework, and a risk management system guided by the normative poles of free and informed consensus and deregulation without domination.

Whereas Arrow's model prioritizes the controversial instrumental value of self-interest (Partiality 1), the risk attitude of those affected (Partiality 2) is a fundamental determinant of collective decision quality. Unlike Partiality 1, Partiality 2 subjectivizes normative frameworks without

contradicting their presuppositions and goals. As previously noted, a risk-averse Rawlsian will likely support policies that differ from those chosen by risk-neutral or risk-prone agents who subscribe, with equal honesty and consistency, to the same normative framework.

This stands in stark contrast to a scenario where a wealthy Rawlsian advocates for less demanding redistributive policies simply to avoid the personal cost of higher taxes. In such a case, the agent is not genuinely promoting a conception of justice; rather, they undermine it to advance their own self-interest. In both instances, the subject's partial interests play a key role in decision-making. However, while disagreements arising from Partiality 1 are questionable from an impartialist normative standpoint and lack instrumental value, disagreements arising from Partiality 2 exhibit the opposite characteristics. These are reasonable disagreements that, when channeled through appropriate institutions, hold decisive instrumental value. These considerations apply, equally, to any plausible normative framework, including utilitarianism or prioritarianism.

Arrow did not consider the instrumental value of the risk attitudes of those affected at the collective political level. This oversight likely stemmed from his interpreting this factor exclusively through the lens of Harsanyi's reaction against the ordinalist revolution in welfare economics. However, risk constitutes the pivotal element. It confers instrumental value upon the ordinal maximization of expected utility and links the number—the quantity of individuals consulted—to the quality of the outcome, distinct from purely epistemic considerations—provided a threshold of epistemic competence and normative reasonableness is achieved. Ultimately, risk transforms the satisfaction of individual partial preferences, majority rule, consensus-oriented institutions, and the creation of deregulated spaces into strategies that are necessary for collective survival and flourishing.

While our proposal recognizes Harsanyi's insight into the central role of divergent risk attitudes in rational collective decision-making, we fundamentally reject his reliance on utilitarian cardinalism. We view his strategy as another facet of the impartialist model—one that erroneously reduces politics to the logic of morality—and we resist this tendency. Our proposal is a form of political consequentialism. Risk-centered ordinal consequentialism takes a fundamental step to avoid utilitarianism's core problems: it substitutes the cardinal utility principle with a system of electoral and contestatory institutions. These institutions ensure the open and continuous process of ordinal maximization of expected good or utility. The key insight is that, regardless of how this good is characterized, every approach toward its maximization is extensionally equivalent to maximizing freedom as non-domination for actual individuals. The classic criticism that utilitarianism morally justifies sacrificing the basic rights or well-being of minorities to advance majority well-being does not apply to our proposal.

The severe crisis currently confronting democracies is a multifaceted phenomenon. Two key factors play an important role: technological transformations in mass communication and the concentration of power within wealthy, influential minorities whose lives are insulated from the collective fate. This context is defined by a lack of reliable information, factual errors, and the resurgence of unreasonable normative frameworks—ranging from xenophobia and racism to the demonization of minorities, adherence to implausible conspiracy theories, and the rejection of scientific knowledge. These tendencies, embodied by populist political leaders, exemplify the most problematic features of pseudo-democratic politics. Consequently, this scenario generates doubts across the ideological spectrum regarding the value of majoritarian participation in political decision-making. The urgency of the situation incentivizes the search for expedient, authoritarian solutions. Crucially, such factual and normative distortions can

produce disastrous short-term outcomes and fundamentally degrade the political culture of democracies. The importance of improving the epistemic quality of public debate and decision-making processes is undeniable. However, the argument developed in this paper leads to a clear normative conclusion: enhancing the quality of collective decision-making is incompatible with both reduced majority participation in politics and diminished public accountability, especially over an extended time horizon. Two strategies lack instrumental value: relying on a select minority of experts to make decisions, and attempting to artificially cultivate competence within a select group external to the actual democratic process. Even under ideal conditions, such strategies offer no guarantee of quality outcomes for the majorities excluded from the process. Suppose, as we have assumed throughout the text, that it were possible to identify those who reach a threshold of adequate competence, and that these individuals were genuinely committed to the pursuit of the common good. Suppose further that these agents could be shielded from the influence of economic elites and other powerful groups. Even under these stringent conditions, the development of such a system over time would threaten the welfare or subsistence of the majority. As in other human activities, avoiding such a result in politics requires reliable factual information, a reasonable normative framework, and adequate risk management strategies. No human activity can prosper over time without accounting for this third dimension, barring fortuitous success. At a political level, rational risk management requires a specific institutional foundation: electoral and contestatory practices must center on achieving and broadening free consensus, promoting deregulation, and maximizing individual freedom. Crucially, this necessitates a high level of citizen participation and influence in the decision-making process, thereby satisfying one of the defining characteristics of democratic politics informed by republican institutions.

Acknowledgements

We would like to thank the many colleagues and friends who provided critical comments throughout the extended process of developing this proposal, especially Julio Montero, Eleonora Cresto, and Martín Oliveira. Discussions at various academic conferences and research seminars informed the initial formulations of this project. We are especially grateful to Claudio López-Guerra, Sebastián Linares, Eduardo Rivera López, Francisco García Gibson, Ezequiel Monti, Luciano Venezia, Luis García Valiña, Moisés Vaca, Miguel Duranti, María Victoria Costa, Pablo Gilabert, José Luis Martí, Marcelo Alegre, Nahuel Maisley, Jeffrey Howard, Juan Espíndola, Leticia Morales, Martín Farrell, María Victoria Kristán, and Nicolás Alles for their comments on earlier versions of this proposal.

REFERENCES

- Achen, Christopher H., and Larry M. Bartels. 2007. *Democracy for Realists*. Princeton, NJ: Princeton University Press.
- Anderson, Elizabeth. 2006. "The Epistemology of Democracy." *Episteme: A Journal of Social Epistemology* 3 (1-2): 8–22.
- Arendt, Hannah. 1977. "Truth and Politics." In *Between Past and Future: Eight Exercises in Political Thought*, 227–64. New York: Penguin.
- Arrow, Kenneth J. 1951. "Alternative Approaches to the Theory of Choice in Risk-Taking Situations." *Econometrica* 19 (4): 404–37.
- . 1953. "The Role of Securities in the Optimal Allocation of Risk-Bearing." *Review of Economic Studies* 31: 91–96.
- . 1963. *Social Choice and Individual Values*. New York: Wiley.
- . 1965. *Aspects of the Theory of Risk-Bearing*. Helsinki: Yrjö Jahnssoninsäätiö.

- . 1970. *Essays in the Theory of Risk-Bearing*. Amsterdam: North-Holland.
- . 1973. “Some Ordinalist-Utilitarian Notes on Rawls’s Theory of Justice.” *The Journal of Philosophy* 70: 245–63.
- . 1996. “The Theory of Risk-Bearing: Small and Great Risks.” *Journal of Risk and Uncertainty* 12: 103–11.
- Bartels, Larry. 1996. “Uninformed Votes: Information Effects in Presidential Elections.” *American Journal of Political Science* 40: 194–230.
- Beiner, Ronald. 2008. “Rereading ‘Truth and Politics.’” *Philosophy & Social Criticism* 34 (1-2): 123–36.
- Bell, Daniel A. 2015. *The China Model: Political Meritocracy and the Limits of Democracy*. Princeton, NJ: Princeton University Press.
- Bernoulli, Daniel. 1738. “Specimen Theoriae Novae de Mensura Sortis.” *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5: 175–92.
- . 1954. “Exposition of a New Theory on the Measurement of Risk.” Translated by Louise Sommer. *Econometrica* 22 (1): 23–36.
- Brennan, Jason. 2016. *Against Democracy*. Princeton, NJ: Princeton University Press.
- . 2017. “Does the Demographic Objection to Epistocracy Succeed?” *Res Publica* 24 (1): 53–71.
- Brennan, Jason, and Hélène Landemore. 2022. *Debating Democracy: Do We Need More or Less?* Oxford: Oxford University Press.
- Buchak, Lara. 2014. *Risk and Rationality*. Oxford: Oxford University Press.
- . 2017. “Taking Risks behind the Veil of Ignorance.” *Ethics* 127: 610–44.
- . 2022. “Relative Priority.” *Economics & Philosophy* (2022): 1–31.
- Bykvist, Krister. 2014. “Utilitarianism in the Twentieth Century.” In *The Cambridge Companion to Utilitarianism*,

- edited by Ben Eggleston and Dale Miller, 103–24. Cambridge: Cambridge University Press.
- Caplan, Bryan. 2007. *The Myth of the Rational Voter*. Princeton, NJ: Princeton University Press.
- Cohen, Joshua. 1986. “An Epistemic Conception of Democracy.” *Ethics* 97 (1): 26–38.
- . 1989. “Deliberation and Democratic Legitimacy.” In *The Good Polity: Normative Analysis of the State*, edited by Alan Hamlin and Philip Pettit, 17–34. Oxford, UK: Blackwell.
- . 2009. “Truth and Public Reason.” *Philosophy & Public Affairs* 37 (1): 2–42.
- . 2010. *The Arc of the Moral Universe and Other Essays*. Cambridge, MA: Harvard University Press.
- Coleman, Jules, and John Ferejohn. 1986. “Democracy and Social Choice.” *Ethics* 97: 6–25.
- Condorcet, Marquis de. 1995. “An Essay on the Application of Analysis to the Probability of Decisions Rendered by a Plurality of Votes.” In *Classics of Social Choice*, edited and translated by Iain McLean and Arnold B. Urken, 91–112. Ann Arbor: University of Michigan Press.
- Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Elster, Jon, and John Roemer, eds. 1991. *Interpersonal Comparisons of Well-Being*. Cambridge: Cambridge University Press.
- Forst, Rainer. 2013. “A Kantian Republican Conception of Justice as Nondomination.” In *Republican Democracy: Liberty, Law, and Politics*, edited by Andreas Niederberger and Philipp Schink. Edinburgh: Edinburgh University Press.
- Friedman, Milton, and Leonard J. Savage. 1948. “The Utility Analysis of Choices Involving Risk.” *The Journal of Political Economy* 56 (4): 279–304.

- Gaus, Gerald. 1997. "Does Democracy Reveal the Voice of the People? Four Takes on Rousseau." *Australasian Journal of Philosophy* 75 (2): 141–62.
- Gibbons, Anthony. 2021. "Political Disagreement and Minimal Epistocracy." *Journal of Ethics and Social Philosophy* 19: 192–201.
- Goodin, Robert E., and Kai Spiekermann. 2018. *An Epistemic Theory of Democracy*. Oxford: Oxford University Press.
- Grofman, Bernard, and Scott Feld. 1988. "Rousseau's General Will: A Condorcetian Perspective." *American Political Science Review* 82 (2): 567–76.
- Guerrero, Alexander A. 2014. "Against Elections: The Lottocratic Alternative." *Philosophy & Public Affairs* 42: 135–78.
- Gutmann, Amy, and Dennis Thompson. 1996. *Democracy and Disagreement*. Cambridge, MA: Harvard University Press.
- Habermas, Jürgen. 1984. *The Theory of Communicative Action*. Cambridge, UK: Polity Press.
- . 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Translated by William Rehg. Cambridge, MA: MIT Press.
- Harsanyi, John C. 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *Journal of Political Economy* 6 (3): 309–21.
- . 1975. "Nonlinear Social Welfare Functions." *Theory and Decision* 6: 311–32.
- . 1977. *Rational Behavior and Bargaining Equilibrium in Games and Social Decisions*. Cambridge: Cambridge University Press.
- . 1979. "Bayesian Decision Theory, Rule Utilitarianism, and Arrow's Impossibility Theorem." *Theory and Decision* 11: 289–317.
- . 1982. "Morality and the Theory of Rational Behavior." In *Utilitarianism and Beyond*, edited by Amartya Sen and Bernard Williams, 39–62. Cambridge: Cambridge University Press.

- Hayek, F. A. von. 1945. "The Use of Knowledge in Society." *American Economic Review* 35: 519–30.
- . 1955. *The Counter-Revolution of Science: Studies on the Abuse of Reason*. Glencoe, IL: The Free Press.
- Holtug, Nils. 2006. "Prioritarianism." In *Egalitarianism: New Essays on the Nature and Value of Equality*, edited by Nils Holtug and Kasper Lippert-Rasmussen, 125–56. Oxford: Oxford University Press.
- Hong, Lu, and Scott E. Page. 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences* 101: 16385–89.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101 (3): 461–82.
- Jones, Geoff. 2020. *10% Less Democracy: Why You Should Trust Elites a Little More and the Masses a Little Less*. Stanford, CA: Stanford University Press.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47: 263–91.
- Lafont, Cristina. 2020. *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy*. Oxford: Oxford University Press.
- Landemore, Hélène. 2012. "Democratic Reason: The Mechanisms of Collective Intelligence in Politics." In *Collective Wisdom: Principles and Mechanisms*, edited by Hélène Landemore and Jon Elster, 151–89. Cambridge: Cambridge University Press.
- . 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- . 2017. "Beyond the Fact of Disagreement? The Epistemic Turn in Deliberative Democracy." *Social Epistemology* 31 (3): 277–95.

- Lavoie, Don. 1986. "The Market as a Procedure for Discovery and Conveyance of Inarticulate Knowledge." *Comparative Economic Studies* 28: 1–19.
- List, Christian. 2001. "Epistemic Democracy: Generalizing the Condorcet Jury Theorem." *Journal of Political Philosophy* 9 (3): 277–306.
- López-Guerra, Claudio. 2014. *Democracy and Disenfranchisement*. Oxford: Oxford University Press.
- Machina, Mark J. 1982. "'Expected Utility' Analysis without the Independence Axiom." *Econometrica* 50: 277–323.
- Malcolm, Finn. 2022. "Epistocracy and Public Interests." *Res Publica* 28: 173–92.
- Martí, José Luis. 2006. "The Epistemic Conception of Deliberative Democracy Defended." In *Deliberative Democracy and Its Discontents: National and Post-national Challenges*, edited by Samantha Besson and José Luis Martí, 27–56. Aldershot, UK: Ashgate.
- Mason, Elinor. 2014. "Objectivism, Subjectivism, and Prospectivism." In *The Cambridge Companion to Utilitarianism*, edited by Ben Eggleston and Dale Miller, 177–98. Cambridge: Cambridge University Press.
- Maynor, John W. 2003. *Republicanism in the Modern World*. Cambridge, UK: Polity Press.
- Menger, Karl. 1967. "The Role of Uncertainty in Economics." Translated by W. Schoellkopf and W. G. Mellon. In *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, edited by Martin Shubik, 211–31. Princeton, NJ: Princeton University Press.
- Mill, John Stuart. 1991. *Considerations on Representative Government*. New York: Prometheus Books. Originally published in 1861 by Parker, Son, and Bourn.
- Mulligan, Thomas. 2018. "Plural Voting for the Twenty-First Century." *The Philosophical Quarterly* 68: 286–306.
- Narens, Louis, and Brian Skyrms. 2020. *The Pursuit of Happiness: Philosophical and Psychological Foundations of Utility*. Oxford: Oxford University Press.

- Nino, Carlos S. 1996. *The Constitution of Deliberative Democracy*. New Haven, CT: Yale University Press.
- Okasha, Samir. 2007. "Rational Choice, Risk Aversion, and Evolution." *The Journal of Philosophy* 104 (5): 217–35.
- . 2011. "Optimal Choice in the Face of Risk: Decision Theory Meets Evolution." *Philosophy of Science* 78 (1): 83–104.
- Parfit, Derek. 1991. *Equality or Priority? The Lindley Lecture*. Lawrence, KS: University of Kansas.
- . 1997. "Equality and Priority." *Ratio* 10: 202–21.
- . 2011. *On What Matters*. Oxford: Oxford University Press.
- Pettit, Philip. 1997. *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press.
- . 1999. "Republican Freedom and Contestatory Democratization." In *Democracy's Edges*, edited by Ian Shapiro and Casiano Hacker-Cordon, 163–190. Cambridge: Cambridge University Press.
- . 2001. *A Theory of Freedom: From the Psychology to the Politics of Agency*. Oxford: Oxford University Press.
- . 2012a. *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.
- . 2012b. "The Inescapability of Consequentialism." In *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams*, edited by Ulrike Heuer and Gerald Lang, 41–70. Oxford: Oxford University Press.
- . 2014. *Just Freedom: A Moral Compass for a Complex World*. New York: W. W. Norton & Co.
- . 2015. "On the People's Terms: A Reply to Four Critiques." *Symposium: Republicanism between Justice and Democracy. Philosophy and Public Issues (New Series)* 5 (2): 79–97.
- . 2019. "The General Will, the Common Good, and a Democracy of Standards." In *Republicanism and the Future of Democracy*, edited by Yiftah Elazar and Geneviève

- Rousselière, 13–49. Cambridge, MA: Cambridge University Press.
- Quiggin, John. 1982. “A Theory of Anticipated Utility.” *Journal of Economic Behavior and Organization* 3: 323–43.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: The Belknap Press of Harvard University Press.
- . 1993. *Political Liberalism*. New York: Columbia University Press.
- . 2001. *Justice as Fairness: A Restatement*. Cambridge, MA: The Belknap Press of Harvard University Press.
- Riley, Jonathan. 1990. “Utilitarian Ethics and Democratic Government.” *Ethics* 100: 335–48.
- Roemer, John E. 1996. *Theories of Distributive Justice*. Cambridge, MA: Harvard University Press.
- Samuelson, Paul A. 1977. “St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described.” *Journal of Economic Literature* 15 (1): 24–55.
- Schamberger, Christopher. 2023. “The Metaethical Dilemma of Epistemic Democracy.” *Economics & Philosophy* 39: 1–19.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper.
- Sen, Amartya. 1970. *Collective Choice and Social Welfare*. Amsterdam: North Holland.
- . 1979. “Interpersonal Comparisons of Welfare.” In *Economics and Human Welfare*, edited by Michael J. Boskin, 183–201. New York: Academic Press.
- Sen, Amartya, and Bernard Williams, eds. 1982. *Utilitarianism and Beyond*. Cambridge: Cambridge University Press.
- Shaw, William H. 1999. *Contemporary Ethics: Taking Account of Utilitarianism*. Cambridge, MA: Blackwell Publishers.
- Somin, Ilya. 2013. *The Problem of Political Ignorance*. Stanford, CA: Stanford University Press.
- Spector, Horacio. 2016. “Decisional Nonconsequentialism and the Risk Sensitivity of Obligation.” *Social Philosophy and Policy* 32 (2): 91–128.

- Urbinati, Nadia. 2014. *Democracy Disfigured: Opinion, Truth, and the People*. Cambridge, MA: Harvard University Press.
- Von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Weymark, John. 1991. "A Reconsideration of the Harsanyi-Sen Debate on Utilitarianism." In *Interpersonal Comparisons of Well-Being*, edited by Jon Elster and John Roemer, 255–320. Cambridge: Cambridge University Press.
- Williams, Bernard. 2005. *In the Beginning Was the Deed: Realism and Moralism in Political Argument*. Princeton, NJ: Princeton University Press.
- Woodard, Christopher. 2019. *Taking Utilitarianism Seriously*. Oxford: Oxford University Press.
- Yaari, Menahem E. 1987. "The Dual Theory of Choice under Risk." *Econometrica* 55: 95–115.

Received: December 19, 2024.

Accepted: November 8, 2025.